Slicing WiFi WLAN-Sharing Access Infrastructures to Enhance Ultra-Dense 5G Networking

Maxweel Carmo[1,2] Sandino Jardim[1,2] Augusto Neto[1,3], Rui Aguiar[3] Daniel Corujo[3] Joel J. P. C. Rodrigues[4,5,6]
[1]Universidade Federal do Rio Grande do Norte (UFRN), Natal-RN, Brazil
[2]Universidade Federal de Mato Grosso (UFMT), Barra do Garças-MT, Brazil
[3]Instituto de Telecomunicações, Aveiro, Portugal
[4]Instituto de Telecomunicações, Portugal
[5]Instituto Nacional de Telecomunicações (Inatel), Santa Rita do Sapucaí, MG, Brazil
[6]Universidade de Fortaleza (UNIFOR), Fortaleza, CE, Brazil
max, sandino}@ufmt.br, augusto@dimap.ufrn.br, ruilaa@ua.pt, dcorujo@av.it.pt, joeljr@ieee.org

***Abstract*—The ascending of *5th Generation* based network systems provide the prospect for an unprecedented technological revolution in different aspects of current network infrastructures to satisfy the high demands of smart space fully. This work addresses the challenges that raise to exploiting the potentials of WLAN-sharing technology in Ultra-Dense Networking (UDN) Fifth Generation (5G) futuristic use cases. We investigate new complementary aspects of emerging 5G technologies, namely Network Function Virtualization (NFV), Fog computing, Software-Defined Networking (SDN) and others to design a unique WLAN-sharing ecosystem fully *softwarized* to allow complying with 5G UDN critical requirements in a carrier-grade basis. In the resulting approach, we empower WLAN-sharing infrastructures with fog computing technology, so that expanding its installations by wireless access carrier cloud facilities, and enabling carrier-grade *softwarized* netprogramability at the runtime. Moreover, Fog-enabled WLAN-sharing nodes follow a slice-defined approach, in the prospect to provide differentiated services, at unprecedented levels, on top of the same infrastructure through customized, isolated and independent digital building blocks. Moreover, we also enable slices to accommodate applications besides networking functions, so as to obtain ultra-low latency rates leveraging direct linkage to data producer things. For proof of concept of our approach, called — WISE (WLAN slIcing as a SErvice), we carry out experiments in a real testbed, allowing insights in its feasibility and suitability.**

I. INTRODUCTION

The ascending of 5th generation (5G) based systems [1] provides the prospect to an unprecedented technological revolution in different aspects of current network infrastructures, to fully afford the very high smart space networking demands. In smart spaces, the vast amount of heterogeneous network connected moving things (e.g., people, sensors, vehicles, and other mobile nodes) and mobile applications drive the variability and dynamics of service requirements. These requirements can vary from simple environmental sensor data gathering to high complex real-time digital video processing of in-vehicle video cameras [2].

By 2020 5G-based systems will prevail to enable the emerging of new innovative and futuristic use cases [3], including enhanced Mobile Broadband (eMM), massive Machine-type Communications (mMTC), Ultra-Reliable and Low-Latency Communications (URLL). On the one hand, these futuristic 5G use cases will give rise to the emergence of innovative and high value-added social applications, but on the other hand, they will require a very radical view of what the 5G infrastructure should become.

In addition to eMM, mMTC and URLL futuristic use cases, the 5G-based ecosystems entail the convergence of the wide variety of networks (that currently run separately) into a powerful core, thus forming an Ultra-Dense Networking (UDN) scenario. The UDN is a promising facility in 5G-based ecosystems to deal with the demand for ubiquitous availability of reliable and high data rate mobile services [4]. The spatial spectrum reuse, via low power small cells (such as femtocell and picocell), will drastically improve the coverage and capacity of available wireless access networks. Moreover, dense small cells can also offload the wireless data traffic of user equipment attached to macrocells, especially at indoor environments where more than 80% of the data traffic occurs. In this work, we believe that the WLAN-sharing technology arises as an asset to enable the realization of UDN-based 5G use cases, by complementing cellular networks with broadband wireless access.

A. Problem statement

In general terms, the WLAN-sharing technology stands to a community of wireless access networking infrastructures that agreed to share a part of their bandwidth as a WiFi signal, so that authenticated devices can connect to other members' hotspots. The Fon Wireless Ltd. [http://fon.com] is a pioneer and leading company in residential WiFi sharing technology that manages connectivity between over 21 million federated hotspots worldwide, seamlessly. A Customer-Premises Equipment (CPE) is designated as a WLAN-sharing hotspot when embedding facility sets that prepares the WiFi router to aggregate residential (private) and prime public WiFi footprints, making all the incoming wireless traffic subject to underlying networking approach.

Our studies to exploit the WLAN-sharing technology in the hypothesis to expand UDN 5G use cases reveal critical challenges, mainly by lacking methods for provisioning resources and facilities required to satisfy the high variability and dynamics of the service requirements associated with mobile data that moving things produce in smart space [5]. In fact, typical WLAN-sharing solutions provide two separate and independent WiFi networks for private (WiFi owner) and public (federate members) type independent access. However, all the incoming wireless traffic is aggregated and subject of the same underlying off-the-shelf CPE's networking capabilities. Our theoretical analysis reveal that a WLAN-sharing capable CPE must deploy a network-level service differentiation approach, to enable handling the different requirements, patterns and scope of the incoming traffic with isolation and independence, as well as provisioning customized service transport. For instance, real-time content (e.g., video data) must be subject of advanced networking service sets, whereas best-effort type networking facilities should handle packets encapsulating scalar data (e.g., raw data gathering).

The network-level service differentiation approach we advocate means deploying multiple network instances that run locally, and independently, at the WLAN-sharing capable CPE. Each network instance must feature customized networking services, specifically designed and tailored to specific scenarios and applications needs. Available CPE-tailored OS (e.g., OpenWrt, Mikrotik Router OS, and others) supports virtualization facilities that allow to create multiple independent WiFi networks. However, we highlight that all incoming packet aggregates will be handled by the same network service approach, without either any service differentiation nor application-level knowledge. Indeed, world-leading commodity-based CPE suppliers add additional networking facilities (in relation to traditional ones such as firewall, DNS, and NAT), such as Quality of Service (QoS) control, network optimization, wireless mesh connectivity, etc. However, it is the responsibility of the WiFi owner to configure all of these features by themselves, which is not minimally reasonable since mostly of them lack any specialized knowledge in network management and control. Lastly, but not least, CPEs should allow network provider to deploy software applications to run locally, seeking to afford both eMM and URLL ultra-low latency requirements through direct wireless linkage with data producers. This is currently possible only in open CPEs (e.g., Android OS compatible), requiring direct WiFi owner intervention for this, not friendly as well. Despite increasing seamless wireless broadband connectivity opportunities, the WLAN-sharing technology must evolve a lot to cope with UDN 5G use cases, mainly in regards to deploying methods for provisioning required resources and facilities to suite emerging mobility scenarios.

Based on all aforementioned concepts, their benefits and corresponding issues/challenges, this paper aims to evolving current WLAN-shared technology to address the challenges of efficiently affording the rising of the massive mobile data demand in UDN 5G use cases, as well as efficiently satisfying their high resource demands. We investigate new complementary aspects of emerging 5G technologies, namely Network Function Virtualization (NFV) [6], Fog computing [7], Software-Defined Networking (SDN) [RFC 7426], and others to design a unique and innovating WLAN-sharing ecosystem. In the resulting approach, we expand the capabilities of the WLAN-sharing CPEs by the wireless access carrier cloud computational facilities, through applying Fog computing technology. A full *softwarized* approach is adopted to enable a flexible carrier-grade netprogramable system at the runtime. Moreover, the WLAN-sharing Fog nodes follow slice-based definitions, in the prospect to provide differentiated services on top of the same infrastructure through customized, isolated and independent digital building blocks. Finally, we enable slices to accommodate applications in addition to network functions, so as to obtain ultra-low latency rates by direct linkage to data producer things.

In addition to the WISE (WLAN slIcing as a SErvice) proposal — a Fog-enabled slice-defined softwarized WLAN-sharing architecture — initial impact assessments are carried out through a preliminary prototyping evaluation for proof of concept. Thus, we allow to obtain insights in the feasibility and suitability impact associated to the WISE. The outcomes show that the Fog-enabled slice-defined softwarized WLAN-sharing solution can operate satisfactorily, at least at small-scale scenarios.

The remainder of this paper is organized as in the following. Section II presents related work and provides comments that summarize the technologies used to design our proposal. Section III introduces the WISE proposal, while preliminary results regarding the prototype assessment are discussed in section IV, followed by the conclusion and future works.

II. RELATED WORK

The ability to dynamically instantiate network and application services at CPE is a key WISE's requirement. For its realization we employ virtualization technologies, as they provide convenient encapsulation and isolation of services, allowing them to be easily created, terminated, and migrated. Although initial NFV [6] efforts have focused on the use of fully-fledged virtual machines (VM), the nature of our context, focused on WLAN enabler CPEs, led us to employ lightweight virtualization instead. Indeed recent works, e.g. [8], have considered the use of lightweight virtualization as OS tynification [9] and containers to perform slicing tasks. Lightweight virtualization presents a smaller memory footprint, scale better, allows a greater number of VNFs to operate simultaneously, are quick to deploy, and allow fast migration.

Performance studies such as the presented in [10] show the feasibility of running containers at CPEs, but they focus on devices' performance evaluation and don't propose any system-level solution to offer services closer to users. In [11] authors propose the deployment of VNFs in resource-constrained CPEs without the support of containers or VM technologies. Each VNF is implemented as a Linux namespace and rely on software modules and applications already available at the device's Operating System, e.g. iptables, and virtual bridges. The drawback is that a device can run only a restricted set of VNFs (as it depends on software components previously installed), jeopardizing its usage in dynamic scenarios where a variety of network functions and services continuously need to be instantiated/terminated. Also, as the solution lacks resource isolation among VNFs it makes it difficult to support multi-tenancy.

The work carried out in [12] advocates the advantages and identifies open challenges in running VNFs at Fog nodes. There is no assumption about the network technology or hardware constraints and the focus is on analyzing the impact of migrating VNFs between Fog nodes.

In [13] authors employ high-volume servers to provide virtual wireless networks at the urban landscape through VM technologies. Each virtual network is created as a VM or a group of them. In contrast, our work leverages the existing WLAN infrastructure by considering the use of resource-constrained Fog nodes to support the deployment of virtual WLANs while leaving the task of providing rich resources at higher levels of the hierarchy in case they are needed. In this way, we are able to offer dense (virtual) WLANs at a lower cost. Besides, the referred work lacks flexibility as the virtual networks are statically configured and deployed, i.e., there is no softwarized programmability in managing them.

It becomes evident that none of the related works described hereinabove are able to provide an efficient WLAN-sharing approach that is able to satisfy the rigorous requirements of UDN 5G use cases, for the reasons described in the following. The limitations of the related work motivate carrying out our work, that leverage the WLAN infrastructure currently existing on densely populated urban centers (both commercial and residential) for the realization of 5G-capable smart space scenarios, by complementing connectivity opportunities through broadband wireless systems.

III. THE DESIGN

The WISE system aims at enabling service-oriented WLANs, allowing such networks to provide a myriad of differentiated services to hosts and things via network resources slicing. A slice has the property of running as customized (offering its all set of services), isolated, and independent network instances. WISE introduces Fog computing at the WLAN domain to provide flexibility in services offering, allowing, for instance, the deployment of network services requiring low latency close to end-users. The proposal leverages SDN and virtualization technologies to allow provisioning, controlling and managing of the Fog resources in a programmable way.

A key requirement of the approach is to provide a low-cost solution able to integrate seamlessly with the current WiFi networks, as such domains usually don't cope with expensive equipment. In this way, we employ lightweight Fog nodes with limited computing and storage capacity compared with high-end servers usually employed at the ISP premises or Cloud.

Fig.1 depicts the high-level architecture of a WISE-enabled node. Besides providing typical WLAN services such as attachment/connectivity for mobile devices, the node provides locally-deployed custom procedures encapsulated as slice abstractions. Through WiFi interface virtualization the node presents a number of different virtual WLANs (V-WLANs), each one associated to a particular slice. In this way, packets from different V-WLANs are treated differently at the local network before leaving towards the Internet. For instance, in a simple WLAN-sharing scenario the WISE node could be configured to provide two types of services via two distinct V-WLANs: one for residential users and another for IoT support.

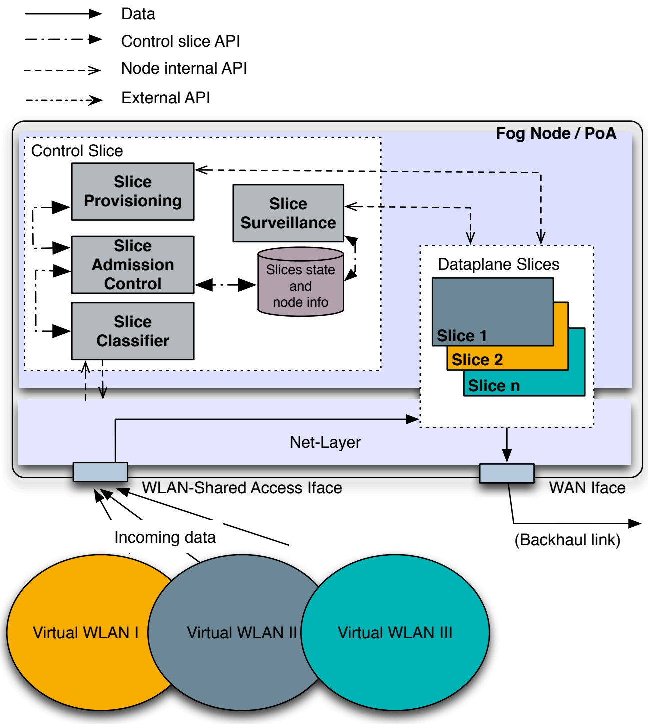

Fig. 1: WISE high-level architecture.

The architecture is structured around two main blocks: control slice and data plane slice. The latter is in charge of processing data packets from mobile nodes according to the slice offered services, while the former runs the modules responsible for the life-cycle management (creation, termination, maintainability) of each data plane slice instance. In essence the control slice identifies the right slice for a new data flow (slice classifier module), verifies whether new flows can be admitted or not (adequacy to node resources, policies, and so on), makes the provisioning of resources for data plane slices, and monitors their performance.

The behaviour of slice admission control is controlled by an entity that is external to the node (an orchestrator located at the cloud, for instance) and is in charge of configuring, for instance, the policies the module will follow in admitting flows. Also, the slice surveillance communicates with external elements to inform about the state of the node (for instance, node resources usage history, how the slices perform, etc).

The sequence diagram depicted in Fig.2 illustrates the steps needed to admit a new data flow in the particular case

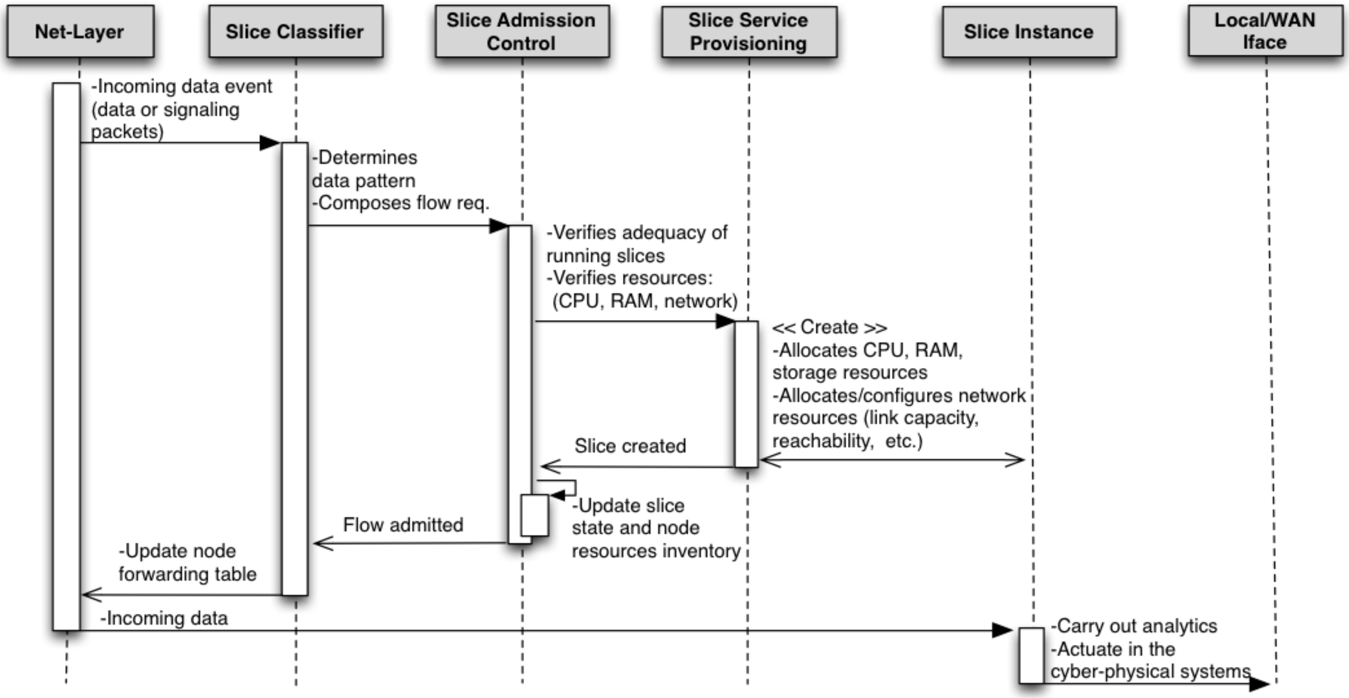

Fig. 2: Sequence diagram for slice creation in response to flow admission.

where the required slice is not instantiated. The Net-Layer, which represent the node's network stack, signalizes the slice classifier module whenever a new flow arrives from any of the available virtual WLANs. An arriving flow in our context can be both application data or signalling messages (SIP, for instance) asking a server on the Internet for a new session establishment. By performing a cross-layer inspection of the packet's headers the classifier determines the data pattern of the flow (e.g., a video session) and composes the flow network requirements. Such information together with the flow identification is delivered to the slice admission control that verifies the existence of a running slice that meets the flow requirements and the availability of resources (e.g., cpu, memory, bandwidth, storage) to accommodate the new flow. In the case the slice is unavailable but there are enough resources, the service provisioning will create a new slice by allocating to it appropriate node resources and starting the services that represent the slice workflow. Next, the slice state and the node resources inventory are updated and the node's forwarding table configured to allow the flow's packets to be forwarded to the new slice.

In case the flow is not admitted due to lack of resources, the mobile node can try to connect to a virtual WLAN belonging to another Fog node at its reach where the admission process described above would start again. Naturally, network assisted methods, where the system can assist the mobile node to connect to a more appropriate V-WLAN to increase the chances of the flow admission can be employed, but they are out of the scope of the present work.

IV. PROTOTYPE EVALUATION

To evaluate the present work we deployed a low-cost Fog solution where the resources of a standard WLAN access point were extended by attaching to it an inexpensive single-board computer — a Raspberry Pi 3 model B. The AP is in charge of providing radio connectivity, while the SBC the computing resources to bring applications and network services closer to customers. The inflexibility and varying capabilities of the APs software stack led us to employ the OpenWRT (https://openwrt.org) — an Operating System supporting a broad range of commercial switching devices — together with the Open vSwitch — OVS [14] kernel module (to steer, via OpenFlow, packets to the WISE node). Such a solution abstracts software differences among vendors, exposing a homogeneous device with appropriate flexibility.

The WISE node deploys an OVS to redirect the incoming traffic to the right software entity for processing. For instance, upon identifying a new flow the WISE's OVS contacts the slice classifier module (see Fig.1) that, in response (in case the flow is admitted), installs OpenFlow rules back at the switch to forward the packets to the appropriate slice.

For this particular implementation we used the Docker technology (https://docker.com) to manage the containerized slices services. Each slice is formed by one or multiple Docker containers embodying network procedures (VNFs) or application layer services.

The Alpine Linux (https://alpinelinux.org), a tiny distribution, was used as the base for the containers, allowing an implementation as small as 11 MB for the application services we used on the experiment. The image size is a key parameter

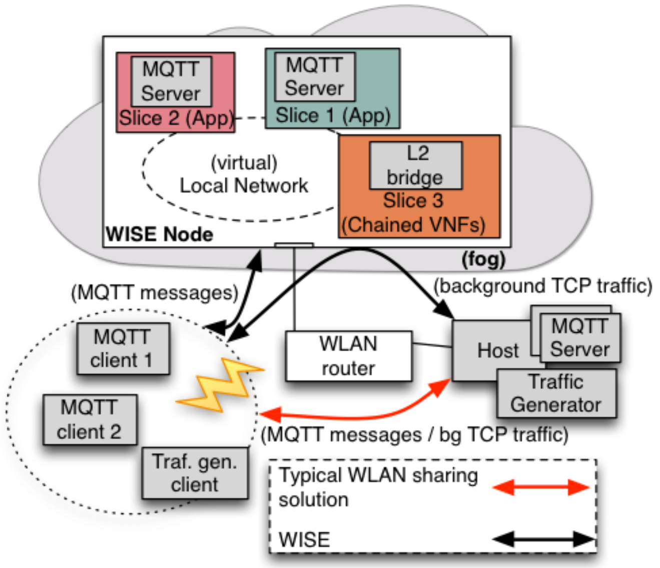

Fig. 3: Evaluation model.

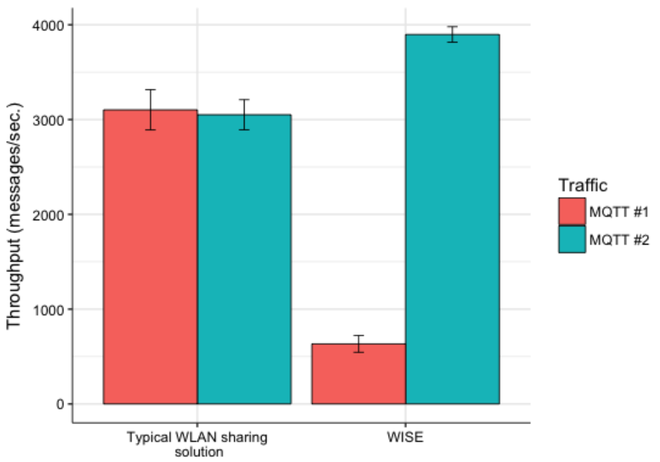

Fig. 4: Bandwidth isolation between services by limiting the slices' CPU cycles.

due to the storage and memory node constraints, and the network delay required to transmit the image from a repository to the edge of the network where the WISE node is located.

Fig.3 depicts the testbed used to perform proof-of-concept performance comparison between two solutions: a typical WLAN sharing deployment and the WISE proposal. As the former does not employ local processing capabilities, the communication among clients is mediated sole by the WLAN router. In both cases, two mobile nodes establish communication with MQTT brokers (https://mosquitto.org) — a popular publish/subscribe messaging protocol suitable for constrained environments like IoT—, while a third node communicates with a iperf traffic generator (https://iperf.fr) to produce TCP background traffic at the local links. For the WISE scenario the MQTT brokers are located in the Fog node, while for the typical WLAN sharing experiment the brokers are deployed at a distinct dedicated machine. In a production environment such machine would be located at the Internet (e.g., in the cloud), but, for sake of simplicity, in our experiments it is deployed at the local network. Besides hosting the two slices that offer the MQTT broker service, the WISE node also implements a third one that acts as a middlebox for processing every packet involving the communication between iperf end-points. The slice deploys three containers where each one implements a simple two-port L2 bridge which sole function is to forward any packet it receives. The containers are chained in a way that the packets cross the three bridges in sequence before getting back to the WLAN router.

In the first proof-of-concept experiment we make use of both MQTT brokers to compare the typical WLAN sharing deployment behaviour against WISE's. Each MQTT client was configured to generate 15K messages of 100 bytes each towards its respective MQTT broker. At the typical WLAN sharing solution, as there is no differentiation of services, both MQTT instances compete against each other for the available bandwidth, resulting in a fair sharing of the resources. Fig.4 shows that the average throughput for each application is approximately 3100 messages per second. As aforementioned, at the WISE solution the MQTT services are encapsulated in two different slices. For this experiment the WISE architecture leveraged the Linux cgroups mechanism to limit the Raspberry CPU cycles per slice in order to offer differentiated response rates. The two MQTT slices were set to use 20% and 80% of the available host's CPU, respectively. Reflecting such distribution, the average throughput for the first slice decreased to 633 messages/sec, while for the second it reached 3897 messages/sec.

Naturally, for emerging scenarios, where the demands for diverse network requirements and the need for dynamic network reconfiguration is of utter importance, the behaviour of the typical WLAN sharing scenario is far from being effective.

Although Docker includes mechanisms to perform CPU and memory resources isolation, it may not be enough in certain configurations. For instance, the slice 3 implements L2 bridging service by making use of a conventional Linux bridge. Although for the viewpoint of the WISE system the slice is a set of manageable Docker containers similar to the ones implementing application services, they operate at the kernel level, making it difficult to limit their CPU cycles. To exemplify the problem Table I summarizes the results of a experimentation set where slice 1 and 3 are supposed to have 80% and 15% of the available CPU cycles, respectively. Although slice 1 reaches an average throughput of 3.26 Mbps, it decreases to 1.23 Mpbs in presence of background traffic at slice 3.

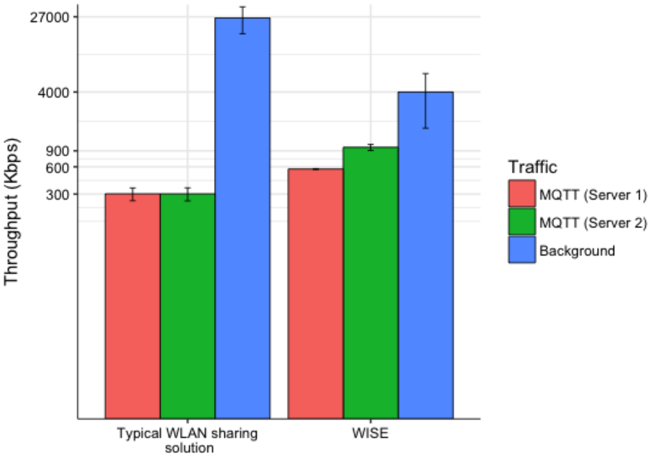

Fig. 5: Node resources limitation per slice through network traffic control.

TABLE I: Insufficiency of limiting CPU cycles to enforce slice isolation

Slice	Scenario 1	Scenario 2	CPU share
Slice 1 (MQTT service)	3.26 Mbps	1.23 Mbps	80%
Slice 3 (L2 VNF service)	– Mpbs	72.8 Mbps	15%

Naturally, the straightforward solution for the above problem is to implement L2/L3 related network functions (NAT, firewalls, etc.) at the user space. But, as WISE works upon resource-constrained devices, such solution could reduce the system performance. As an alternative simple solution we deploy a traffic shaping mechanism to enforce bandwidth limits for the incoming traffic. In this way WISE indirectly controls the CPU resource used by each slice, including kernel-level network functions. Fig.5 depicts another proof-of-concept comparison between a typical system and WISE to illustrate the effectiveness of the solution. Services throughput are shown in a log scale. For the typical WLAN-sharing scenario the background TCP traffic causes the MQTT throughput to stay around 302 Kbps in average, for each instance. At the WISE experiment the incoming traffic is limited at different rates (1, 2, and 5 Mbps for slices 1, 2, and 3, respectively). As a result the WISE system manages to increase slice 1 throughput to 566.97 Kbps and slice 2 to 984.80 Kbps, in average, by throttling packets before they are processed by slices. It is worth noticing, Nonetheless, new studies are required to analyze how the additional scheduling impacts the overall system performance in face an increasing traffic volume

Fechamento

V. CONCLUSION AND FUTURE WORK

In this paper we proposed the WISE, a Fog-enabled slice-defined softwarized WLAN-sharing, which aims at evolving current WLAN-shared technology to address the challenges of efficiently affording the rising of the massive mobile data demand in UDN 5G use cases, as well as efficiently satisfying their high resource demands. Through the WISE approach, differentiated services can be provided at unprecedented levels on top of the same infrastructure, allowing customized, isolated and independent building blocks (i.e., slices). Moreover, we also enable slices to accommodate applications besides networking functions, so as to obtain ultra-low latency rates leveraging direct linkage to data producer things. The feasibility of the proposal is assessed via experiments in a real testbed, allowing insights in its proof of concept.

For the next steps, we intend evolve the WISE proposal with offloading control facilities, in order to improve the system scalability capacities through leveraging a Fog-cooperative computing approach. Moreover, we will submit the WISE ecosystem to different 5G use cases, in order to assess the impact of the offloading control facilities in terms of scalability and performance.

ACKNOWLEDGMENT

This work was partially supported by the FCT – Fundação para a Ciência e a Tecnologia through the UID/EEA/500008/2013 Project, and Finep with resources from Funttel, Grant No. 01.14.0231.00, under the Centro de Referência em Radiocomunicações - CRR project of the Instituto Nacional de Telecomunicações (Inatel), Brazil. We thank the Coordination for the of Personnel in Higher Education (CAPES), the National Council for Scientific and Technological Development (CNPq) and the *Fundação de Amparo à Pesquisa do Estado de Mato Grosso* (FAPEMAT).

REFERENCES

[1] "5g vision: The 5g infrastructure public private partnership: the next generation of communication networks and services," 5G Public-Private-Partnership (5G PPP), Tech. Rep., Aug 2017. [Online]. Available: https://5g-ppp.eu/wp-content/uploads/2015/02/5G-Vision-Brochure-v1.pdf

[2] Q. Zhang and F. H. P. Fitzek, *Mission Critical IoT Communication in 5G*. Cham: Springer International Publishing, 2015, pp. 35–41. [Online]. Available: https://doi.org/10.1007/978-3-319-27072-2_5

[3] "Imt vision: Framework and overall objectives of the future development of imt for 2020 and beyond," International Telecommunication Union (ITU), Tech. Rep. M.2083, Set 2015. [Online]. Available: https://www.itu.int/dms_pubrec/itu-r/rec/m/R-REC-M.2083-0-201509-I!!PDF-E.pdf

[4] C. Niu, Y. Li, R. Q. Hu, and F. Ye, "Fast and efficient radio resource allocation in dynamic ultra-dense heterogeneous networks," *IEEE Access*, vol. 5, pp. 1911–1924, 2017.

[5] F. S. D. Silva, A. V. Neto, D. Maciel, J. Castillo-Lema, F. Silva, P. Frosi, and E. Cerqueira, "An innovative software-defined winemo architecture for advanced qos-guaranteed mobile service transport," *Comput. Netw.*, vol. 107, no. P2, pp. 270–291, Oct. 2016. [Online]. Available: https://doi.org/10.1016/j.comnet.2016.04.019

[6] "Network functions virtualisation (nfv); architectural framework," ETSI Industry Specification Group (ISG), Tech. Rep. ETSI GS NFV 002 V1.2.1, Dez 2014.

[7] S. Yi, C. Li, and Q. Li, "A survey of fog computing: Concepts, applications and issues," in *Proceedings of the 2015 Workshop on Mobile Big Data*, ser. Mobidata '15. New York, NY, USA: ACM, 2015, pp. 37–42. [Online]. Available: http://doi.acm.org/10.1145/2757384.2757397

[8] F. Ramalho and A. Neto, "Virtualization at the network edge: A performance comparison," in *World of Wireless, Mobile and Multimedia Networks (WoWMoM), 2016 IEEE 17th International Symposium on A*. IEEE, 2016, pp. 1–6.

[9] S. Natarajan, R. R. Krishnan, A. Ghanwani, D. Krishnaswamy, P. Willis, A. Chaudhary, and F. Huici, "An Analysis of Lightweight Virtualization Technologies for NFV," Internet Engineering Task Force, Internet-Draft draft-natarajan-nfvrg-containers-for-nfv-03, Jul. 2016, work in Progress. [Online]. Available: https://datatracker.ietf.org/doc/html/draft-natarajan-nfvrg-containers-for-nfv-03

[10] A. Krylovskiy, "Internet of things gateways meet linux containers: Performance evaluation and discussion," in *Internet of Things (WF-IoT), 2015 IEEE 2nd World Forum on*. IEEE, 2015, pp. 222–227.

[11] R. Bonafiglia, S. Miano, S. Nuccio, F. Risso, and A. Sapio, "Enabling nfv services on resource-constrained cpes," in *Cloud Networking (Cloudnet), 2016 5th IEEE International Conference on*. IEEE, 2016, pp. 83–88.

[12] A. Manzalini, R. Minerva, F. Callegati, W. Cerroni, and A. Campi, "Clouds of virtual machines in edge networks," *IEEE Communications Magazine*, vol. 51, no. 7, pp. 63–70, 2013.

[13] O. Braham and G. Pujolle, "Virtual access point to the cloud," in *Cloud Networking (CLOUDNET), 2012 IEEE 1st International Conference on*. IEEE, 2012, pp. 206–208.

[14] B. Pfaff, J. Pettit, T. Koponen, E. J. Jackson, A. Zhou, J. Rajahalme, J. Gross, A. Wang, J. Stringer, P. Shelar *et al.*, "The design and implementation of open vswitch." in *NSDI*, 2015, pp. 117–130.