

Resource Slicing in Virtual Wireless Networks: A Survey

Matías Richart, Javier Baliosian, Joan Serrat, and Juan-Luis Gorricho

Abstract—New architectural and design approaches for radio access networks have appeared with the introduction of network virtualization in the wireless domain. One of these approaches splits the wireless network infrastructure into isolated virtual slices under their own management, requirements, and characteristics. Despite the advances in wireless virtualization, there are still many open issues regarding the resource allocation and isolation of wireless slices. Because of the dynamics and shared nature of the wireless medium, guaranteeing that the traffic on one slice will not affect the traffic on the others has proven to be difficult. In this paper, we focus on the detailed definition of the problem, discussing its challenges. We also provide a review of existing works that deal with the problem, analyzing how new trends such as software defined networking and network function virtualization can assist in the slicing. We will finally describe some research challenges on this topic.

Index Terms—Wireless network slicing, wireless network virtualization, wireless resource management, slice isolation, 5G, LTE, WiFi.

I. INTRODUCTION

PRESENTLY, one of the major concerns of wireless networks comes from the spectrum scarcity in face of a constantly increasing demand of traffic from the end users. This challenge has led to the consideration of new access technologies or to improve the efficiency of existing ones. Paradigms such as heterogeneous networks, the combination of different Radio Access Technologies (RATs), the use of differentiated services or the cognitive radio concept have appeared as candidate alternatives to increase the efficiency of wireless networks. On the other hand, these paradigms will potentially increase the costs of network operators, making network management and operation more

complex, and consequently, requiring the deployment of more infrastructure.

Recently, the Wireless Network Virtualization (WNV) concept has appeared as a new alternative to help on the achievement of this efficiency goal, reducing capital and operational costs. The term WNV covers a wide variety of virtualization flavors, similar to how network virtualization applies to cabled networks. An accurate definition of network virtualization is given in [1]:

Network virtualization is any form of partitioning or combining a set of network resources, and presenting (abstracting) it to users such that each user, through its set of the partitioned or combined resources has a unique, separate view of the network. Resources can be fundamental (nodes, links) or derived (topologies), and can be virtualized recursively. Node and link virtualization involve resource partition/combination/abstraction.

In summary, the goal of network virtualization is to create logical partitions of some existent physical network resources in an efficient manner. This partitioning is also known as resource *slicing*, which becomes, as we will see throughout the paper, a complex research problem in the wireless domain.

This paper focuses on the slicing problem: how it could be implemented, its challenges, and what is still missing to achieve a complete virtualization approach that could slice the wireless medium. We review recent works on two important aspects of slicing: **resource allocation and isolation**. Slicing implies the allocation of the necessary resources to satisfy independent service requests, but, on dealing with wireless resources, and due to the particularities of the wireless medium, assuring slice isolation becomes a difficult task, even more when Quality of Service (QoS) or Service Level Agreements (SLAs) constraints come into play. In contrast to previous surveys on wireless virtualization [2]–[4], our interest focuses on the problems derived from wireless slicing and on the existing techniques that deal with these specific problems.

The rest of the paper is organized as follows: In Section II we introduce the concepts of Wireless Network Virtualization (WNV), Software Defined Networking (SDN) and Network Function Virtualization (NFV). In our opinion, these concepts are fundamental enablers for the implementation of slicing and in Section II we describe their relationship to wireless slicing presenting some existing solutions. In Section III we thoroughly explain the definition and motivations of wireless slicing. In Section IV we present some

Manuscript received February 26, 2016; revised June 17, 2016 and July 29, 2016; accepted July 31, 2016. Date of publication August 2, 2016; date of current version September 30, 2016. This work has been supported in part by FLAMINGO, a Network of Excellence project (318488) supported by the European Commission under its Seventh Framework Programme, by Comisión Académica de Posgrado, Universidad de la República through the program *Becas de Apoyo a Docentes para Estudios de Posgrado* and by the project TEC2015-71329-C2-2-R (MINECO/FEDER). The associate editor coordinating the review of this paper and approving it for publication was F. De Turck.

M. Richart is with the School of Engineering, University of the Republic, Montevideo 11300, Uruguay, and also with the Network Engineering Department, Universitat Politècnica de Catalunya, 08034 Barcelona, Spain (e-mail: mrichart@fing.edu.uy).

J. Baliosian is with the School of Engineering, University of the Republic, Montevideo 11300, Uruguay.

J. Serrat and J.-L. Gorricho are with the Network Engineering Department, Universitat Politècnica de Catalunya, 08034 Barcelona, Spain.

Digital Object Identifier 10.1109/TNSM.2016.2597295

background of the main wireless technologies and we describe the problems of resource allocation and isolation in the context of those technologies. We review the current efforts for implementing the slicing approach in wireless networks in Section V and, based on the analysis of these works, we identify the challenges and research directions in Section VI. Finally, we give some concluding remarks in Section VII.

II. RELATED CONCEPTS

In this section we briefly describe three new concepts (WNV, SDN and NFV) we consider fundamental enablers for the wireless slicing purpose. We also review some existing works on these topics, showing their contribution to the implementation of slicing.

A. Wireless Network Virtualization

WNV aims to share a common network infrastructure, including the radio resources, among different virtual networks. We can identify five different goals behind this paradigm:

- The definition of an abstraction layer to simplify the provisioning of wireless access from heterogeneous networks.
- High-level management and programmability of wireless networks.
- Network slicing by service, user or application.
- Infrastructure sharing.
- Radio spectrum sharing.

Wireless virtualization, in comparison with wired network virtualization, encompasses the virtualization of specific wireless hardware and the radio spectrum as well. The virtualization of the wireless medium introduces a number of challenges that do not exist in the wired domain, e.g., the signal propagation, the interference, the user mobility or the considered radio access technology. All these particularities will be the focus of WNV.

Virtualization of a wireless network can be applied at different layers and degrees, from only virtualizing the core network to virtualizing the radio spectrum and physical layer of base stations. Even more, the motivations for virtualizing a wireless network can be very diverse: from enabling the infrastructure sharing among several operators, to offering a layer of abstraction in order to simplify the network management. There is an extended bibliography devoted to WNV, treating the subject under different perspectives, tackling a specific problem or using a particular technology. Wen *et al.* [2] and Liang and Yu [3] offer a comprehensive view on WNV and present existing works on this field.

WNV as an Enabler for Slicing: Virtualization and slicing are two concepts so coupled that virtualization becomes the principal technology enabler for slicing. Nowadays, all slicing proposals consider each slice as some kind of virtual network in order to achieve the objectives behind wireless network virtualization.

Some frameworks for virtualizing wireless networks have been proposed in the last years [5]–[8]. In general, these proposals do not provide details about their implementation, but

they present candidate design guidelines. These works provide the foundations for a wireless network slicing design. Common to all proposals, there are two requirements wireless virtual networks will need to satisfy:

- the coexistence of different virtual networks mapped onto the same physical network,
- the isolation of the virtual networks so as to avoid conflicts between coexistent virtual networks.

These issues are discussed in this paper as the *Slicing* problem (see Sections III and IV).

B. Software Defined Networking

The main idea of SDN is to decouple the data and control planes, moving the control plane from the network devices to a central location [9]. Then, the forwarding devices (switches or routers) just apply the forwarding rules programmed by a controller element.

These SDN ideas imply a separation between the network's policies definition, their implementation in hardware, and the forwarding of data. With this separation, considerable flexibility is achieved, which allows a simpler management of the network [10].

Applying SDN to the wireless domain, some works have proposed different designs, frameworks and tools [4], [11]. Many of the SDN proposals have concentrated on decoupling the management from the hardware and the technology, to be able to give a unique interface to control a heterogeneous network. As we will show next, several works have focused on abstracting from the wireless technology and on allowing network programmability, but do not consider the isolation and coexistence of virtual networks over a shared infrastructure. However, both paradigms, SDN and virtualization, seem to be necessary for achieving a complete cross-layer solution to manage, program, share and slice a heterogeneous wireless network.

SDN as an Enabler for Slicing: Deploying and managing a sliced wireless network is a complex task if it is not handled correctly. In our opinion, SDN is the necessary tool for easing this task, and it is crucial for achieving the needed flexibility and programmability a sliced wireless network will need. Even though SDN does not appear as the solution for the slicing problem itself, we mention some existing ideas of how SDN would be helpful for wireless network slicing.

A good example of how SDN can be helpful is FlowVisor [12]. This slicing tool is designed and used in wired networks to achieve slicing and flow isolation. Although in a wireless scenario the problems are different, the main ideas of *flow-based slicing* and *control message isolation* can be used in the wireless domain.

In the cellular domain, SoftCell [13] focuses on the core network of cellular providers. It proposes the use of the SDN paradigm at the providers network by using common switches and middle-boxes instead of specific proprietary hardware. This approach tackles particular problems of this kind of networks, such as scalability and high bandwidth requirements. In a similar way, MobileFlow [14] proposes an architecture for deploying SDN onto mobile operators, in particular

with 3GPP infrastructure. The architecture is called Software-Defined Mobile Network (SDMN) and its main idea is “to provide maximum flexibility, openness, and programmability to future carriers without mandating any changes in user equipment”. In this architecture, the data and control planes are decoupled and the functions of each plane are virtualized.

These works are the most direct application of the SDN paradigm on a cellular network. Although, not the focus of our work, the slicing of the backbone of a cellular network is an important aspect of any slicing approach. Increasing the flexibility and programmability of this part of the network will be essential.

A more ambitious approach is proposed in SoftRAN [15]. It defines a *virtual big-base station* that logically groups geographically close physical base stations. The idea is that these physical base stations can be centrally managed to facilitate the radio resource allocation and the interference mitigation. For this, the authors propose an abstraction of the radio resources through the virtualization of the physical resources. Also, an Application Programming Interface (API) is proposed to export the state of the network to an external manager which can program the control plane. In this case, the control plane of the wireless devices is decoupled from hardware.

The previously mentioned ideas of wireless resource abstraction can be used in a slicing approach to set-up slices and to specify their resources. Also, the decoupling of control from hardware and the centralized management are necessary to develop different control planes for the different slices.

For WiFi systems, some works have been proposed in the last few years, introducing the SDN paradigm to wireless LANs. Odin [16], [17] is a framework for enabling the programmability of a wireless network. The most interesting idea is the definition of an abstraction called Light Virtual Access Point (LVAP). This formalizes a logical connection between a client and an AP to maintain the status of the association. The idea is that these LVAPs can be allocated on any kind of hardware (WiFi AP or LTE eNodeB). Then an API is defined so as to access network parameters and to reconfigure the network. The objective of this approach is to build high-level applications to manage the network. Similar works that extend, use and improve these ideas are: Empower [18] and AeroFlux [19], [20].

In our opinion, these approaches are essential for implementing slicing on WiFi devices. For example, the LVAP idea could be extended or modified to be used in slicing, having an LVAP per slice and so, easing the wireless configuration of each slice.

C. Network Function Virtualization

The main idea behind NFV is the decoupling of network functions from the physical network equipment where they run on [21]. This is achieved by removing their execution from specific hardware and, by means of virtualization, run on standalone hardware, with the additional possibility to be deployed on any location.

Hence, a network service can be decomposed into a set of network functions, which are then virtualized and executed on

general purpose servers. This way, the Virtualized Network Functions (VNFs) can be easily created, moved or destroyed, anywhere and at anytime, giving flexibility and lowering costs to the network operator. With NFV, more dynamic and service-aware networks are possible with lower operating and capital expenses [21].

NFV as an Enabler for Slicing: As already mentioned, we see NFV as an enabler for slicing a wireless network. It will make the creation and management of slices easier to perform if some functions can be taken from proprietary hardware, virtualized and run centrally. In the following, we describe some works where we found ideas that can be applied to slicing a wireless network.

The approach of Software Defined Radio (SDR), where signal processing functions are run in a centralized manner by general purpose hardware is clearly a NFV approach. CloudRAN [22], [23] is a SDR architecture with the objective of taking away the signal processing functions from the Base Stations to put them in the cloud. This way, these heavy processing functions can be run on general purpose hardware, and therefore this solution reduces capital costs and promotes the deployment of new technologies. Another example is the work in [24] where the authors apply NFV to the Evolved Packet Core (EPC) of a LTE network. EPC is the core network for LTE systems consisting on a number of entities in charge of functions such as: mobility, routing and forwarding, access control, pricing, etc. The objective of this work is to virtualize the functions of all these entities on the cloud, but grouping some of the functions on the same server to reduce transactions over the network.

CloudMAC [25] proposes to move all MAC processing functions currently run by WiFi APs to the cloud. In this proposal, the functionality of APs is limited to forwarding frames, consequently all the processing is done on a central server where Virtual APs are running as virtual machines. To connect the physical AP (now called Wireless Termination Point, WTP) to the Virtual Access Point (VAP), tunnels and OpenFlow switches are used.

Having network functions decoupled from hardware and grouped at a single location could be a solution to many of the slicing problems we will describe later on this paper. For example, with an SDR approach or with the idea of a decoupled MAC from CloudMAC, it could be possible to modify the MAC-layer or PHY-layer implementation to adapt it to new requirements (e.g., prioritizing some traffic over other). Also, with NFV at the core of the network, it would be possible to assign to each slice specific network functions and to remove others so as to tailor the slice for a specific scenario.

III. THE SLICING APPROACH

In the following section we define in detail what is understood by a *Slice* and we also introduce some of the motivations for slicing a wireless network.

A. Definition

The concept of slicing in the context of network virtualization is multifaceted. For the most general definition, a slice can be considered as a set of flows belonging to different *end*

users (mobile clients of the wireless service in our case). Then, a slice supports flows of multiple end users, but at the same time, an end user can participate in multiple slices. A flow (stream of packets) is an atomic entity in this approach, it can have specific QoS requirements and is a member of a single slice. For example, a flow in an IP network could be defined by the tuple composed by source and destination IP addresses and ports.

Furthermore, a slice can be defined as a subset of network resources allocated to a *tenant* (virtual operator or service provider), with complete control over those resources. An important aspect of the slicing design approach is the delivery of customization and programmability tools to the tenant.

Examples of slices can be: all the flows whose source or destination is a given type of device such as sensors; or all the flows from a VoIP service; or, all the flows with source or destination the end user of a given operator. Depending on the specification of a slice, an end user can participate on different slices but, slices are always independent between each other. In the context of wireless networks, we envision two big scenarios for using slices:

- *Quality of Service Slicing*: the idea is to create slices to offer different services and assure some type of QoS within the slice. For example, a slice can be created to give service to a specific group of devices with the same requirements (sensors or smartphones) or by type of application (e.g., a slice for multimedia services).
- *Infrastructure Sharing Slicing*: this is the traditional idea of network virtualization applied to the wireless domain. There is a tenant (e.g., Mobile Virtual Network Operator), which is given a slice of the network. The tenant has complete control over the network infrastructure and functions within the slice.

An example scenario for applying the QoS Slicing is given in [26]. This scenario consists of a future 5G network operator offering differentiated types of services depending on the specific use case. For example, a high-throughput service for smart-phones, a low-rate non-critical service for Internet of Things (IoT) or Machine to Machine (M2M) communications and a low-latency service for critical real-time communications. So, the scenario is a combination of these use cases, each one with its specific requirements, and the operator has to provide service and management for all of them jointly. To cope with these requirements, isolated slices are defined, each one giving service to a specific group of users or devices (see Figure 1).

Another aspect of the definition of a slice is related to where in the network or up to what level slicing should be applied. A good classification of this facet of slicing is given in [3] where different levels of slicing are detailed:

- *Spectrum-level slicing*: The spectrum can be sliced by time, space or frequency multiplexing, or by an overlaid access. It can be considered as link virtualization.
- *Infrastructure-level slicing*: It is the slicing of physical network elements, such as: antennas, BSs, processors, memory. It is accomplished mostly by virtualization.
- *Network-level slicing*: It is the slicing of all the network infrastructure.

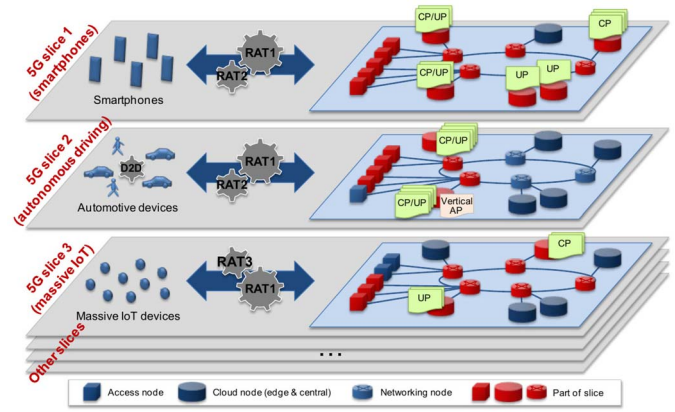


Fig. 1. Example of slices in a 5G scenario. From [26].

B. Motivation

We have already mentioned some of the motivations for implementing slicing on the wireless domain. In what follows, we thoroughly detail the major benefits the slicing approach will bring to wireless networks.

1) *Heterogeneous Service Differentiation*: In the current context, where there is a wide variety of services and devices that wireless networks have to deal with, slicing becomes a way to isolate and accomplish different requirements simultaneously. On sharing resources, slicing will enable the creation of customized services with fine control features of QoS [27]. The idea is to divide the network into slices made of different resources and capacities so as to offer differentiated services for heterogeneous use cases. Even more, in [28], slicing is presented as one of the key enablers of future 5G Systems to manage the expected heterogeneous requirements.

It is also possible to define slices for specific applications, which may require customized network capabilities [29]. With virtualization, a layer of abstraction over the slice can be defined so as to control the network as a *black box* to easily specify application requirements. Another possible approach is to have customized slices per type of device or per type of client requirement. Network slices will offer efficient resource utilization as each slice can be customized for a specific service and on a dynamic on-demand way. This dynamism, is the key difference from existing similar proposals as VPNs.

2) *Network Management*: As said in [30]: “The management of different applications with contradictory requirements on a common infrastructure can be performed via separated network slices”. Slicing the network will allow to individually configure the networks *edge-to-edge* and define specific functions for each case, while sharing the same infrastructure and avoiding higher costs. For example, slicing will allow to allocate only the necessary functions and to reserve resources on the entire path of the communication, allowing the network configuration to be tailored for each case.

Slicing will also provide flexibility to dynamically create and destroy slices depending on the operators policies, with the help of NFV and SDN. The objective is to virtualize as many functions as possible, and those that cannot be virtualized should be programmable and configurable [26]. Even

more, in the case of slices defined per type of service or device, as it is known which service each slice is servicing, the network can be simplified by removing functions that are not necessary. For example, if a slice is giving access to static sensors, mobility management can be reduced to a minimum. This way, management is simplified, becoming easier to develop autonomic management for each specific slice.

3) *Heterogeneous Radio Access Technologies*: Slicing can also help on the management of networks using heterogeneous Radio Access Technologies (RATs). It is becoming more common to have different RATs working on the same network as a way to alleviate the spectrum scarcity problem. For example, WiFi has become an important player on the mobile business as a way to *offload* data transmissions of mobile devices like smart-phones or tablets. This way, end users avoid the extra cost penalty when exceeding the contracted data usage limit. For instance, the Office of Communications from U.K. (Ofcom) reported that in the U.K. 81% of mobile consumers use WiFi at some point [31]. Resource allocation from different technologies can be handled from a slicing perspective where, depending on parameters such as: throughput, user location or costs, the best RAT is assigned to each slice.

The spectrum efficiency can also be improved with slicing, as it is possible to match different requirements to the best available radio resources [32]. To allow this possibility, the network must encompass virtualized or programmable wireless interfaces, as well as different wireless technologies, as expected in future wireless networks. Then, as predicted in [2], the future leads to the coexistence and convergence of different wireless technologies composing a service-oriented infrastructure. Slicing appears as one possible solution to allow this coexistence by simplifying the management.

4) *Infrastructure Sharing*: Another important motivation for slicing is infrastructure sharing. It is similar to the service differentiation concept but, in this case, each slice can be used by a different operator offering its own services. For example, there are in the U.K. 41 mobile virtual operators, which are customers of mobile infrastructure providers [31]. Most of them offer similar services of voice, SMS and data as the incumbent operator. Slicing will facilitate the infrastructure management and will provide isolation between different operators.

From a different point of view, the idea of sharing the infrastructure will give operators more flexibility to change their logical network and efficiently use their resources [33]. This idea is also backed by the Telemanagement Forum [34], which quoted: "The expectation is that 5G will offer multiple virtual networks with different cost/performance characteristics across a shared infrastructure".

5) *Flexibility for New Services and Business Models*: From a business point of view, network slicing will promote the introduction of new use cases without increasing costs thanks to the ability to share the infrastructure by different slices. This can allow to provide service to devices with low traffic demands on highly dense areas (e.g., IoT) without increasing costs, as 5G will need to do. Besides, as a standardized API for programming the network could be offered, slicing will

leverage the Everything as a Service (XaaS) business model and allow third parties to explore new opportunities.

IV. PROBLEM DESCRIPTION

The most difficult problem of slicing a wireless network is found at the base stations and the associated wireless links. Achieving effective slicing is challenging mainly because of the variability of wireless links' capacity and because of limited resources. In wireless communications, the capacity of the link depends on the Signal-to-Interference and Noise Ratio (SINR) and on the available bandwidth of the link. Additionally, the SINR is variable over time due to, for example, the distance between the transmitter and receiver, the location of the interferers or the obstacles in between the communicating nodes. Besides, the available radio spectrum is regulated and bounded, without the possibility to be increased, in contrast to the usual deployment of wired networks.

Having in mind these issues, in this section we identify the problems of resource allocation and isolation in wireless slicing and show the complexity of tackling them. As the problems are closely related to the radio access technology and, in particular, to the medium access technique used, we first give a brief introduction to this matter.

A. Medium Access Techniques

All wireless technologies include medium access control functions to decide when each of the devices can transmit. As the two major technologies used nowadays are the 3GPP LTE standard and the IEEE 802.11 standard, we briefly explain the way each of these technologies access the medium, to easily identify the derived problems from wireless slicing.

1) *Medium Access Control in LTE*: In LTE, medium access is performed by Orthogonal Frequency Division Multiple Access (OFDMA) on the downlink and by Single Carrier - Frequency Division Multiple Access (SC-FDMA) on the uplink. OFDMA is based on dividing the available bandwidth into orthogonal frequency sub-carriers, assigning a set of sub-carriers to each user. OFDMA uses sub-carriers distributed through the available spectrum while SC-FDMA can only use adjacent sub-carriers. The multiple access technique works this way: each time interval, called Transmission Time Interval (TTI), an assignment decision is made and each user is assigned a certain amount of radio blocks in the time and frequency domains. This task of assigning resources to users is called *scheduling*.

In the time domain, LTE supports two types of frames of TTIs: one for Frequency Division Duplex (FDD) mode and one for Time Division Duplex (TDD) mode. In FDD mode, the frame consists of 10 TTIs, which are divided in time slots of 0.5 ms and can carry 7 OFDM symbols (in the default configuration with short cyclic prefix). Two consecutive time slots define a sub-frame of 1 ms and, in general, the scheduling works on a sub-frame level. In TDD, which is basically designed for coexistence with legacy systems, the frame is divided into two half-frames of 5 ms.

In the frequency domain, the sub-carriers are grouped into sub-channels of 180 kHz (12 sub-carriers) and the number

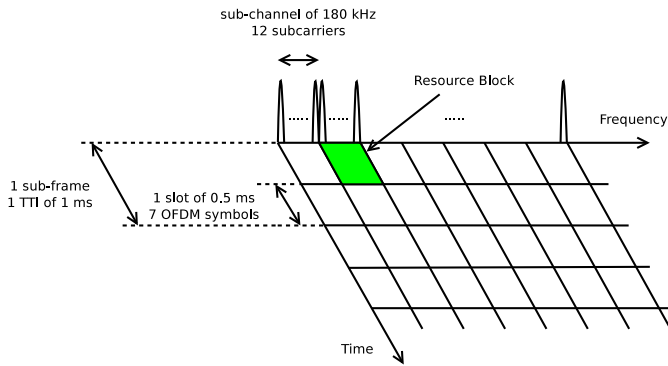


Fig. 2. LTE time-frequency frame.

of sub-channels depends on the available bandwidth. The radio resource unit consisting on 1 time slot and 1 sub-channel is called a Resource Block (RB) or Physical Resource Block (PRB) and it is the basic resource unit of allocation (see Figure 2).

In summary, the responsibility of the scheduler is to decide how to assign PRBs among users taking into account the channel conditions and QoS requirements. This complex task presents several challenges for which various scheduler designs and implementations have been proposed [35].

2) *Medium Access Control in WiFi*: The IEEE 802.11 standard [36] defines four coordination functions (or methods to arbitrate the access to the medium): the *Distributed Coordination Function* (DCF), the *Point Coordination Function* (PCF), the *Hybrid Coordination Function* (HCF) (which uses two mechanisms EDCA and HCCA) and the *Mesh Coordination Function* (MCF). Most devices, when working in infrastructure mode, use DCF or EDCA by default.

DCF uses *Carrier Sense Multiple Access with Collision Avoidance* (CSMA/CA) to regulate the access to the medium. In this access method a device must sense the medium (physical carrier sense) before starting to transmit. If the medium is not busy, the device is able to transmit. More specifically, a device must sense the medium is idle for a period of time (called DIFS) before transmitting. If the medium is busy (on a transmission attempt), the device waits for the current transmission to end. Then, before attempting to transmit again, the device waits for a random backoff time while the medium is idle and, then, it transmits its frame (see Figure 3). This backoff time is selected randomly from the interval $[0, CW]$ (the Contention Window). CW is a variable parameter, which is duplicated every time the device tries to transmit, and can take values between the limits CW_{min} and CW_{max} .

In EDCA, some of the access medium control parameters are adaptable, so as to prioritize the access and to provide QoS. These parameters are:

- The Arbitration Inter-Frame Spacing (AIFS). It defines the time between two frame transmissions.
- The backoff variables CW_{min} and CW_{max} .
- The Transmission Opportunity (TXOP) Limit, which defines the period of time a device can use the medium after gaining access.

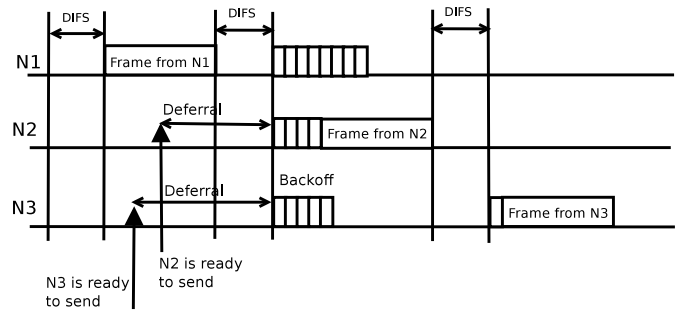


Fig. 3. Distributed Coordination Function backoff procedure.

Then, for providing QoS, different Access Categories (AC) are defined and each one has different configurations of these parameters so as to prioritize the access to the medium.

B. Resource Allocation

On implementing slicing on a wireless network, the main issue is how to assign resources to the different slices. This is known as the resource allocation problem and is well studied in computer science and operations research [37], [38] for other research fields.

For the wireless scenario, the resource allocation problem considers these aspects:

- 1) Definition of what the resources are.
- 2) If necessary, definition of a model for representing the resources.
- 3) Selection of a way to partition the resources.
- 4) Definition of a way for modeling the request of resources.
- 5) Design of a mechanism for assigning the resources to the slices.
- 6) Design of a mechanism to keep or update the assignment in case of changes.

Regarding aspects 1 and 2, in the wireless domain, what the resources are and how they are modeled can vary depending on the point of view. The resources can be the available radio spectrum (divided also in time, frequency or space), the available transmission time or the capacity of the medium.

For points 3 and 4, current works propose different models of the allocation problem, from low-level and highly technology-dependent models (resource-based models) to more general high-level models (throughput-based models) [39]. For example, for the LTE case, several works propose resource requests consisting of a number of PRBs to be assigned to each slice [40]. This model has the advantage that the requests are given in allocation units, facilitating the implementation of the allocation. However, such low-level requests are difficult to handle from high-level management entities such as operators or service providers. A more high-level model considers a fraction of resources per request. In this case, there is not an exact demand of a number of PRBs but a relative percentage of resources. In the case of LTE, this would translate easily to PRBs, but this model could also be used for other technologies where resources are quantified differently. Another approach is to partition the time instead of the actual

resources. The idea is that, given an interval of time, split this interval into fractions to be used by each slice. It is similar to a Time Division Multiplexing (TDM) approach, but the fractions of time used by each slice do not have to be of the same size or periodic. All previously mentioned models share their dependency on the channel conditions, and consequently the resulting throughput for each slice is not guaranteed and can be time varying. More strict but complex models propose to make reservations according to specific QoS parameters such as throughput, bandwidth or minimum data rates. For this model, a translation to resource assignments has to be done and a dynamic control is necessary to update those mappings over time. As channel capacity varies, the necessary resources to accomplish the same request may vary as well.

After deciding what are the resources, how they are modeled and how to model the requests, the next step is to define a mechanism to assign the resources so as to fulfill the requests (points 5 and 6 of the list). Because of the particularities of wireless transmissions, applying the decided allocations onto the wireless hardware is not trivial. Hence, the resource allocation problem also includes the development of techniques and mechanisms that translate the decisions into actions over the resources. The type of actions and the actual implementation will depend mostly on the used wireless technology, but also on the models defined in previous steps. Current proposals on this issue are given in Section V.

C. Isolation

Controlling that the resource allocation (and so the slice specification) is not violated over the time is another difficulty for any slicing approach, we call this the *isolation problem*. The fundamental idea of isolation is to prevent the deterioration on the performance of one slice because of any change on another slice (like the number of end users, flows or channel conditions) or because of the removal or set-up of slices.

The complexity of assuring isolation in wireless networks appears because of the high variability of the channel conditions and because of users' mobility. The capacity of a wireless link can vary significantly depending on several factors: the distance between the client and the AP or BS, the interference, the environment (indoor, outdoor, surrounding objects), the radio access technology. All these considerations make isolation (and also resource allocation) very difficult to implement.

The problem of isolation can be more or less difficult to solve depending on how resources are modeled and on the used wireless technology. Because of the different models, isolation can be interpreted as the maintenance of the assigned resources or as the maintenance of the requested throughput or bandwidth, despite of changes on any other slice. Also, in some slicing approaches, isolation is implicit to the resource allocation solution, for example, if while assigning resources it is guaranteed that there are no resource overlapping.

In summary, an important aspect of wireless resource slicing is how to keep satisfying the requests in spite of this variability. Nevertheless, the isolation issue is the less treated aspect in the literature in spite of being one of the major challenges of

research in this field. Dynamic resource allocation techniques which use the information of the channel conditions to take fast and accurate decisions will be needed.

V. EXISTING APPROACHES FOR WIRELESS SLICING

Facing the issues and complexities we have mentioned regarding the implementation of an efficient slicing, some ideas and mechanisms have been proposed in the last few years. In this section, we review these proposals and explain their characteristics, advantages and drawbacks. A summary of these works is given in Table I. In this table we classify the works based on the considered technology and summarize their main characteristics and objectives.

A. A Classification of Current Solutions

Current proposals for resource allocation and isolation of virtual slices in wireless networks are significantly dependent on the wireless technology, focusing on 3GPP LTE or IEEE 802.11 standards.

For the case of LTE (or other cellular technologies such as WiMAX) the vast majority of approaches modify the frame scheduler to assign PRBs to the slices (*PRB scheduling*). Some works, trying to avoid such a low level strategy, propose mechanisms which schedule the use of resources between slices in a higher layer. This approach is generally done at the MAC-layer or at the Network-layer, we call it *Slice scheduling*. The third category is *Traffic shaping*, controlling the traffic (packets) that are sent to the scheduler, with the implicit intention to modulate the schedulers behavior. So, for LTE three strategies prevail:

- *PRB scheduling*
- *Slice scheduling*
- *Traffic shaping*

For WiFi, similar approaches are proposed. In our review, we identify three predominant strategies:

- *EDCA control*: This strategy modifies the EDCA parameters (CW, AIFS and TXOP) to prioritize the access to the medium of the different slices.
- *Slice scheduling*: The idea is to have each slice as a virtual machine over the physical AP and schedule the use of transmission resources between the slices (e.g., in a TDM-like approach).
- *Traffic shaping*: In this case, the traffic coming from different slices is shaped before sending the data to the MAC-layer.

In Table I we explicitly mention which of these classifications are used by the reviewed works along with their main objectives. In fact, some works only focus on achieving some grade of slicing, others on resource allocation or resource embedding while others only on isolation. Also, the earliest works on the area of slicing were focusing on experimental testbeds, hence, the use cases and objectives of these early works are different from present approaches.

B. LTE and WiMAX Context

In the following, we review the main proposals to the slicing problem in the context of cellular networks. We focus

TABLE I
SUMMARY OF WIRELESS RESOURCE SLICING PROPOSALS

	Technology	Objective	Resource Allocation and/or Isolation Mechanism		Implementation
			Uplink	Downlink	
Zaki <i>et. al.</i> [40], [41]	LTE	Assign resources based on predefined contracts	None	Scheduling PRBs	OPNET simulation
Li <i>et. al.</i> [42]	LTE	Load balancing in multi-cell slicing.	None	Scheduling PRBs	OPNET simulation
Karnaugh-map embedding [43]	LTE	Embedding requests	None	Scheduling PRBs	None
Dynamic embedding[44]	LTE	Dynamic embedding	None	Scheduling PRBs	Simulation
Kamel <i>et. al.</i> [45]	LTE	Optimal resource allocation. Fairness among users	None	Scheduling PRBs	Matlab simulation
Virtual Basestation [46]	IEEE 802.16 (WiMAX)	Slice isolation	None	Traffic shaping	Real, WiMAX prototype
VNTS [47]	IEEE 802.16 (WiMAX)	Air-time fairness	None	Traffic shaping	Real, WiMAX prototype
NVS [39]	IEEE 802.16 (WiMAX)	Slicing framework	Slice and flow scheduling	Slice and flow scheduling	Real, WiMAX prototype
CellSlice [48]	IEEE 802.16 (WiMAX)	Slice isolation	Slice scheduling and traffic shaping	Sustained rate control	Real, WiMAX prototype
C-VAP [49]	IEEE 802.11 (WiFi)	Slicing Air-time fairness	EDCA parameters	None	Matlab simulation
Nakauchi <i>et. al.</i> [50]	IEEE 802.11 (WiFi)	Slicing. Air-time control.	EDCA parameters	EDCA parameters	Simulation in QualNet
ViFi [51]	IEEE 802.11 (WiFi)	Slicing. Air-time control.	EDCA parameters	Slice scheduling and traffic shaping.	Real, WiFi prototype
Derakhshani <i>et.al.</i> [27]	IEEE 802.11 (WiFi)	Air-time control	EDCA parameters	None	Matlab simulation
Smith <i>et. al.</i> [52]	IEEE 802.11 (WiFi)	Isolate experiments in testbed	Slice scheduling	Slice scheduling	Real in ORBIT testbed
Mahindra <i>et. al.</i> [53]	IEEE 802.11 (WiFi)	Testbed. Multiple concurrent experiments	None	Traffic shaping	Real in ORBIT testbed
Virtual WiFi [54]	IEEE 802.11 (WiFi)	Client virtualization	None	Slice scheduling	Real, WiFi prototype

on the two most popular technologies in the literature: LTE and WiMAX. Although WiMAX has lost attention and LTE is expected to evolve with the introduction of 5G systems, these solutions are an interesting approach to the problem. After the description of the proposals we briefly compare the different approaches.

1) *PRB Scheduling*: Zaki *et al.* [40], [41] present a framework for LTE virtualization. The authors propose an architecture for virtualizing the LTE Base Stations (called eNodeBs (eNB) in LTE architecture) with the objective of having different operators sharing the same physical resources. The solution is based on a Hypervisor (as in CPU virtualization), which hosts different virtual nodes, allocates the resources and is responsible of the spectrum sharing and data multiplexing. The Hypervisor will accomplish two tasks: (i) it will host several virtual eNBs onto a physical eNodeB, scheduling the physical resources among them; (ii) it will schedule the wireless resources among the different virtual eNodeBs. For this second task, the solution uses the Physical Resource Block (PRB) as the minimum resource granularity that can be allocated, and assigns them among the different virtual nodes, instead of among the users (as done typically by a scheduler). The PRBs are scheduled to the different virtual eNodeBs based on previously arranged contracts, which specify different guarantees for the operator owning a virtual eNodeB. The contracts can set different PRBs allocation policies:

- a fixed amount of PRBs,
- a maximum amount of PRBs to be allocated dynamically according to the current estimated demand or,
- a best effort allocation with no guarantees.

After the Hypervisor allocates PRBs to the virtual eNBs, each virtual eNB allocates the PRBs to its users.

In this work, a comparison between the fixed and dynamic approaches is done using the OPNET simulator. For showing the benefits of the approach it is assumed that multiple operators have their traffic peaks at different moments of time. In our opinion, this assumption is unrealistic or not well founded. Also, although the scheduler handles the coexistence of different slices over a shared physical eNB, the mechanism does not offer explicit isolation. If the demand exceeds the available resources, the assignment is reduced proportionally.

In [42], the framework from [40] is used and extended through a more detailed algorithm for scheduling PRBs for the virtual nodes. The objective of the solution is to dynamically allocate PRBs based on the estimated demand of the slices. The demand is estimated separately for Guaranteed Bit Rate (GBR) traffic and non-GBR traffic. The allocation goal is to satisfy the GBR demands and then to allocate PRBs to non-GBR traffic proportionally. If the total demand overloads the amount of available resources, the assignment is done proportionally to each slice demand. An interesting aspect of this solution is the addition of a load balancing mechanism to distribute the load of a slice among different eNodeBs (in a multi eNB scenario). The idea is that if a eNB is overloaded and a neighbor eNB has available resources, a user is selected to be migrated to the unloaded eNB.

The Karnaugh-map-like Embedding Algorithm (KEA) considered in [43] deals with the problem of embedding virtual wireless networks (slices) requests onto the physical wireless resources. Specifically, this work concentrates on the allocation of resources to slices when the requests come along the time (on-line requests). This approach presents some obvious drawbacks when compared to resource allocation with all requests arriving at once (off-line requests). To handle this dynamic scenario, requests are grouped within a time window, and the

embedding is done at the end of each window. The spectrum is modeled as a two-dimensional (frequency and time) grid of assignable resources and an algorithm based on Karnaugh-map is used to make the assignment. In the article there is no explanation on how this mechanism would be implemented in real hardware or in current wireless technologies. For example, how this can be implemented by the LTE scheduler seems complicated.

In [44] an extended version of the previous mechanism [43] is presented. The proposal in this work considers *dynamic embedding* to avoid requests rejections due to topological constraints. This type of rejections happen when, sufficient resources (time-frequency slots) are available for a request, but due to the arrangement of existing assignments in the 2-D grid of time and frequency, there is not enough contiguous space for that request. The mechanism of dynamic embedding rearranges the assignments on each time slot, giving priority to already assigned requests. Through simulations it is shown that with dynamic embedding the number of rejections is significantly reduced in comparison to a static embedding. Although an interesting approach, it suffers the same problems as in [43]. Even more, this mechanism can seriously affect the scheduling times if too elaborate calculations are needed every time.

Kamel *et al.* [45] propose a scheduling mechanism to slice an LTE network into several virtual networks owned by different Service Providers (SP). For each SP, a contract is agreed, which defines the minimum amount of PRBs that will be assigned. Differently from previous works, in this case the scheduler assigns PRBs to users (as LTE generally does) but it is modified to follow a specific optimization strategy to allow slicing. The optimization problem objective is to maximize the transmission rate obtained by each user subject to a set of constraints: not to exceed the total BS power, not to assign the same PRB to more than one user, to assign at least the minimum agreed PRBs to each SP and to keep certain fairness among users. The solution is numerically evaluated by a Matlab simulation, which shows the effectiveness of the proposed heuristic when compared to other solutions. However, the lack of more realistic simulations or deployments, with variable traffic and channel conditions, make the proposal difficult to compare to others.

2) *Slice Scheduling*: Network Virtualization Substrate (NVS) [39] proposes an architecture and algorithms for slicing a WiMAX (IEEE 802.16) network. The main contribution of this work is a mechanism for scheduling slices, which guarantees the requested resources or bandwidth demand while keeping isolation between slices. In this case, the scheduling is implemented by modifying the WiMAX flow scheduler which is located at a higher level than the PRB scheduler. The scheduler decides at each time interval which slice should use the transmission resources. Then, each slice can decide how to schedule its own flows with the given resources following different allowed strategies to finally send the packets to the frame scheduler. Therefore, the frame scheduler is not changed in the way PRBs are assigned, however, modifications at the MAC layer of the base station will be needed. This causes the approach to face similar deployment constraints as

the PRB Scheduling proposals, as in general, this software is proprietary and manufacturer dependent.

3) *Traffic Shaping*: Virtual Basestation [46] is an architecture for the virtualization of WiMAX base stations (BTS) to achieve resource sharing and isolation between multiple virtual network slices. The proposal adds a new layer over the WiMAX network called virtual BTS substrate. This substrate acts as a virtualization layer and provides a platform where virtual machines (VMs) are created and executed for each slice. These VMs operate as virtual BTSs, and emulate an isolated private BTS for each slice. This framework presents two interesting aspects: the definition of the virtual BTS as an entity separated from the physical BTS and an isolation mechanism based on traffic shaping decoupled from the BTS. This idea makes the proposal feasible and independent of hardware. However, modifications on network components such as in the ASN-GW are needed for control, data tunneling and isolation.

An implementation of the isolation mechanism is discussed in [47]. The mechanism (called Virtual Network Traffic Shaper (VNTS)) obtains information from the wireless interface about the current transmission rate and uses this value, jointly with the number of clients and the weight of the slice, to shape the traffic. The shaping is implemented outside the BTS to control the offered load to the frame scheduler and to assure the fraction of resources assigned to each slice. Some important limitations can be foreseen with this proposal: (i) the traffic shaping is done independently of the number of available resources, resulting in an inefficient use of resources if some slice does not provide traffic; (ii) the mechanism focuses on isolation and not on resource allocation, there is no explanation on how to assign resources when new slice requests arrive. It also lacks an isolation mechanism for the uplink traffic and no study is done on the latency performance.

CellSlice [48] is another resource slicing proposal which does not need to modify the scheduling algorithms at the base stations but instead proposes a shaping mechanism at the gateway. The ideas are similar to other works, the network is divided into slices and each slice specifies a reservation of a fraction of the total resources. The objective of the mechanism is to assure that the requests are satisfied, while maintaining isolation among slices and using resources efficiently. The work focuses mainly on uplink flows, which are difficult to control, as traffic originates from the clients. The method used for this control is based on the adaptation of a specific parameter of the WiMAX standard, the *maximum sustained rate*. Hence, although an interesting mechanism, it highly depends on the BS scheduler capability to control the rate of a flow through an adaptable parameter.

4) *Discussion*: In most of the PRB Scheduling proposals reviewed so far the description of the algorithms is very vague, and not explicitly shown. Also, in all cases, slicing agreements are based on the number of PRBs that are guaranteed to the operator, which, in our opinion, is a too low-level approach. This appears problematic for a slice tenant which could not have enough knowledge or information to decide the correct number of PRBs to request. Moreover, this does not guarantee a fixed performance when considering variable rates and variable channel conditions. Agreements related to more

high-level variables such as a percentage of the total resources would be more appropriate. In addition, PRB Scheduling slicing will require a modification of the scheduler, which can be a difficult task because of the complexity of the scheduling algorithms. Also, an important aspect of the schedulers are the tight time constraints that the algorithms must satisfy, an issue not considered by the proposed solutions.

On the other hand, Slice Scheduling and Traffic Shaping techniques do not need to modify the PRB scheduler and are generally easier to deploy. Also, as these are higher level mechanisms, the allocations are made on fractions of the total resources or on guaranteed bandwidths. Nevertheless, the current proposals also face some drawbacks:

- Without controlling the scheduling, it is more difficult to control the traffic coming from end users. Only [48] proposes a solution to this issue, but it highly depends on the technology and the hardware.
- Slice Scheduling at upper layers does not always guarantee that the scheduling will be kept at low layers, for example, queue buildup can happen at the MAC layer or the frame scheduler which can disrupt the scheduling.
- Traffic Shaping can increase latency if queue management is not made properly or jointly considered with the shaping.

C. IEEE 802.11 Context

In this section we detail recent works dealing with the slicing problem for the WiFi technology. As explained previously, because of the distributed nature of the medium access control, these solutions have to deal with problems different from those of cellular networks. At the end of the review we briefly discuss the advantages and drawbacks of each approach.

1) *EDCA Control*: Control-theoretic optimization of Virtual APs (C-VAP) [49] is a control-theory approach for adjusting the CW of the clients in a sliced WLAN in order to provide optimized throughput and fairness to virtual slices. The slicing mechanism uses a Proportional Integral (PI) controller adapting the CW parameter of each client. This way, the mechanism achieves the same throughput in all slices independently of the number of associated clients. Although the authors present a complete formal approach to the problem, the solution does not provides any type of guarantee to each slice, just fairness among them.

In [50] a mechanism is presented to create virtual APs (VAPs) over a physical AP, in such a way that each VAP has its own MAC transmission queue and virtual machine. Thus, there is a set of EDCA parameters specific to each queue, enabling an air-time-based mechanism to isolate the slices. The mechanism adjusts the parameters of each queue (each slice) and the parameters of the associated clients, all based on previously defined requirements for each slice and on the number of devices in the network. The mechanism allows the definition of target air-time ratios for each VAP and adjusts the CW_{min} to achieve those ratios. Although the solution is said to consider variable rates, this is not shown in the simulations done. Besides, the description of the algorithm is not clear about the control mechanism of the CW_{min} parameter.

The work in [51] (ViFi) proposes a similar idea to that found in [50] for the uplink traffic, and a slice scheduling mechanism for the downlink traffic. The uplink control mechanism configures two EDCA parameters, the CW_{min} and the transmission opportunity τ for each client. The authors argue that the joint control of both parameters provides better fine-grained tuning of the throughput. The final goal of the control mechanism is to guarantee pre-established air-times for the uplink flows of each client. For the downlink, it uses a scheduler which, in a first stage, schedules packets per-slice depending on the requirements of each slice and, in a second stage, schedules packets per-user of each slice in a round-robin manner. In this work the management of variable rates is peculiar. There is no real isolation, when the rate drops on one client, the throughput is increased on other clients of another slice. The evaluation is done on a real implementation, which is an important progress from previous works. However, more extensive evaluations, where a more dynamic scenario is tested, would be interesting. Also, an evaluation of the convergence time of the algorithm and of the air-time usage would be needed.

In [27] a mechanism for joint control of association and air-time is presented. The objective of this proposal is to maximize the total throughput of a virtual *multi-AP WiFi network* by controlling the clients associations to APs and the air-time obtained by each slice sharing the infrastructure. By analytically modeling the behavior of the clients, the authors found the optimal transmission probability that maximizes throughput and also guarantees that the air-time request of each slice is maintained. Then, a control mechanism adapts some EDCA parameters to obtain this optimal transmission probability. This proposal only controls the client's behavior but no control on the AP generated traffic is done. Besides, in the article it is not clear how the wireless capacity variability is managed. The optimization appears to be based on a static deployment of clients. Also, as it seems that the optimization algorithm is executed on a central controller, a mechanism to inform each client of the parameters to be used is needed.

2) *Slice Scheduling*: Smith *et al.* [52] propose a mechanism for isolating experiments over a shared WiFi infrastructure. The objective is to separate a testbed of APs into independent virtual testbeds so as to be able to run simultaneous and isolated experiments. The isolation is achieved by a TDM-like mechanism where different experiments are allocated in separated time slots. This approach has two major drawbacks, the synchronization for enabling and disabling all virtual nodes of one experiment, and the context switch cost at the devices when switching between different experiments.

Virtual WiFi [54] is a proposal designed for *client virtualization*. It tackles the problem where a client runs several virtual machines (VMs) on its device and these virtual machines handle independently the connections to an AP. The objective is to have many VMs within a single device with a single wireless interface connected to different networks simultaneously. The work contributes with a very valuable analysis on the problems of sharing and slicing a wireless interface: the support of all the functionalities of the wireless interface inside the VMs, and the ability for each VM to establish its own connections

with its own credentials. The proposed architecture enables the access to the physical interface from inside any VM to have the same management functions and to create isolated connections. Nevertheless, it needs to modify the device driver at the host and the firmware of the wireless card. The issues tackled in this work are important for a full wireless virtualization where slices can access low level wireless functions.

3) *Traffic Shaping*: In [53] an empirical comparison of different approaches for isolation in concurrent experiments on a shared testbed is conducted. In particular, the authors study the efficiency of space and time isolation and conclude that no mechanism gives sufficient isolation if the bandwidth of the different slices is not controlled. Then, a mechanism for traffic shaping and admission control is proposed to enforce slices to specific bandwidths. No details are given on how the bandwidths are selected or how the admission control mechanism takes decisions.

SplitAp [55] is a proposal to assure slice isolation for the uplink traffic in a virtual WiFi network. The method applies traffic shaping on the client side based on commands sent by the AP. For this, special software has to be installed on the client: a traffic shaping module and a control and reporting module. The control and reporting module is responsible of two tasks: reporting usage parameters (as the MCS and packet size) to the AP; and controlling the shaping module. The algorithm at the AP uses the information sent by the clients to estimate the uplink air-time usage of each slice, and if the predefined policies for each slice are not kept, it broadcasts a command to adjust the air-time usage at each client. However, in the article it is not explained the mechanism at the client which takes the command sent by the AP and converts it to a traffic shaping rate. In our opinion, this would be the actual slicing mechanism. Besides, as other similar works that control uplink traffic, changes are needed at the client side, an approach that currently appears unfeasible.

4) *Discussion*: Regarding the EDCA approaches, some drawbacks are identified: there is a total lack of analysis of the real feasibility to adapt the EDCA parameters on the hardware. For example, in some devices the EDCA parameters are coupled to the hardware queues, and the number of those hardware queues is fixed. Hence, the number of possible slices that can independently control the EDCA parameters is fixed. The possible values the EDCA parameters can take is limited and this is not considered, neither is the time granularity to which these values can be modified. The variability of the channel conditions is also not well considered, for example, the convergence time or accuracy of the mechanisms are important metrics to show how fast the algorithms can adapt the resource assignment to new channel conditions.

Doing the slicing at higher layers avoids some of the problems mentioned above but introduces new ones. For example, for traffic shaping mechanisms, information from lower layers is needed. This cross-layer communication is not always easy to implement if access to the firmware is not available. The technique of slice scheduling does not appear as the appropriate approach for isolation, as current solutions only tackle sharing problems.

VI. RESEARCH CHALLENGES

Because of the novelty of wireless slicing and more generally of wireless virtualization many challenges remain not addressed or at least not solved properly. In this section, we explore some of the challenges that make wireless slicing an interesting and promising research topic.

A. Isolation in Random Access Networks

For the particular case of technologies that use random access methods (e.g., WiFi) the isolation of virtual slices is complex and not fully studied. As already stated, two major challenges prevail in this technology: the randomness and the distributed nature of the access control. The most comprehensive works that deal with these problems include traffic shaping and EDCA control for the downlink and uplink slicing (Section V-C1). However, many aspects of isolation such as variable traffic, mobility and variable channel capacity are not deeply treated. Therefore, designing a mechanism for effectively slicing the uplink and downlink with strictly assured isolation is still a challenge.

B. Technology Agnostic Solutions

One of the biggest research challenges is to obtain a mechanism that could perform resource allocation and isolation of wireless slices independently of the wireless technology. The air-interface, the spectrum, the protocols for wireless and for the backbone are different for the different technologies. Then, there is not yet a unified approach dealing with any of the above mentioned factors. This technological dependency will be a problem when slicing a heterogeneous network.

When dealing with this issue, there is a trade-off between flexibility (or abstraction) and performance [29]. It seems to be difficult to use the same approach to virtualize and slice different wireless technologies without affecting the performance, as each technology has its own particularities and mechanisms for optimization. Also, for allowing virtualization to offer slicing and abstraction, a common language would have to be defined to be able to specify, manage and control the heterogeneous infrastructure. For example, assuring certain throughput to a given slice appears complex without knowing the underneath technology, or at least the channel access method used (deterministic or random access).

Ideas like modeling a “general” wireless network or developing layers of abstraction could be useful to reach this objective. For instance, approaches for network programmability [11] such as *Software-Defined Radio (SDR)* [56]–[58] or *programmable MAC protocols* [59]–[61] may be used to circumvent this challenge. In these proposals a generic API is provided to develop new wireless protocols (or modify existing ones) over generic radio transmitters or vendor-specific wireless devices.

C. Dynamics and Time Constraints

The existing proposals have not given enough attention to the impact of the adaptation of transmission parameters on the slicing techniques. As the link quality varies dynamically

over time, many wireless standards incorporate an autonomous modulation and coding scheme (MCS) adaptation mechanism and/or a transmit power adaptation mechanism, so as to select the best transmission parameters for the current conditions. These reconfigurations of parameters are done with short latency, for example, in WiFi the MCS control algorithm takes decisions within 100ms intervals. Hence, slicing mechanisms have to be efficient and fast when reacting to changes or when searching for reallocations.

Related to this, many proposals lack a detailed performance study in terms of resource consumption (processor, memory or storage) and in terms of execution times. In wireless networks, changes can happen very fast and the time interval between transmissions is of the order of milliseconds. This way, new mechanisms with better reaction to changes need to be developed for slicing. For example, feedback control theory [62] and machine learning [63] techniques should be considered. Control theory would help with the design of controllers with guarantees of performance and stability. Furthermore, machine learning can help to learn control policies without the need to model a very complex environment.

D. Real Deployments

Only few works have deployed and tested their slicing proposals on real networks. In the wireless domain, doing a real deployment is critical for the evaluation of solutions.

Furthermore, in real deployments it is common to find scenarios with multiple BSs or APs belonging to the same access network. Then, resource allocation and isolation on a multi-cell (multi-AP) network needs to be considered carefully. Deploying slices sharing multiple BSs or APs can bring new issues such as: interference between slices or load unbalance. For instance, sharing the spectrum could be accomplished cooperatively considering the interference among the cells and some load estimation mechanism. Then, resource assignment to each slice inside each cell could be more accurate [41].

Additionally, in a real deployment it becomes necessary to decide up to which level virtualization should be applied to achieve an efficient sliced solution. As we already mentioned, slicing can be done at different levels, namely: from application and flow slicing to hardware and spectrum slicing. As stated in [64], virtualization can be considered at different levels with respect to providers and operators:

- *Universal Virtualization*, where the network is viewed as a cloud of BSs where the tenant has to choose and configure all the components to provide the desired service, and is totally transparent to the resource provider.
- *Cross-infrastructure Virtualization*, the idea of this paradigm is to share resources among infrastructure providers, where there is a pool of resources, and the tenants can choose the resources which best fit their needs.
- *Limited intra-infrastructure virtualization*, is the virtualization inside a single infrastructure provider, in this case there is spectrum sharing only between tenants inside the network of the provider.

Deciding which of these approaches better fits current network deployments is an open challenge.

E. User Mobility and Interference

The mobility of users is a particular feature of wireless networks that brings new challenges to slicing. Not only because mobility generates variations in links capacity and performance, or because it makes the number of users on a network to vary significantly, but also because it adds management complexity. Wireless networks have to deal with the management of user mobility, handle handovers and assure QoS despite of the location of the user.

It is clear that new problems are introduced when allowing the user mobility in a sliced network. In this case, not only a user will change the BS or AP it is connected to, but also it could change of slice (if changing of operator or service is needed). Then, handover mechanisms to move a user across slices are necessary. As slices may be owned by different entities and can belong to totally independent virtual networks, implementing this seems complex. A possible approach could share a central mobility manager across slices, however, this may need to be a third party agent with an open interface controller. In addition, centralization would add latency in a task with strict time constraints. Alternatively, a distributed solution could also be considered. However, the distribution of mobility management can add new problems such as more signaling overhead between the management entities. In summary, a good solution to the mobility problem in a virtual sliced scenario will need to support handoffs between BSs, slices and technologies while maintaining the service quality.

F. Control of Final Users

Another major challenge in wireless resource slicing is the access control of end user devices to the medium. The complexity of this problem depends greatly on the wireless technology used. For example, the most used medium access control in the IEEE 802.11 standard is totally distributed, i.e., the AP does not have any possibility to control how and when an end user will transmit. In this case, the allocation and isolation of resources used by end users becomes complex as there is little control over users' devices. In contrast, in 3GPP LTE, the scheduler at the BS aside from scheduling the downlink traffic, it also schedules the resources for the uplink traffic, having complete control of the resources on both links. However, information from the end users is needed in order to gather knowledge about the traffic that is generated, as well as the channel conditions.

G. Complex Wireless Management Functions and Configurations

Most wireless equipment has complex management functions, and is manufacturer specific, involving the programming of drivers and low level software. When sharing a base station by multiple slices, these specific functions have to be used with care as commands from different slices can conflict with

each other. Also, in general, each wireless link has its particular configuration parameters like its frequency of operation, bitrate or transmit power, which can be very different from another link sharing the same infrastructure.

Also, in architectures with central controllers (e.g., [15] and [19]) special care has to be taken on local functions at the devices. The possible delays between a central controller and the physical devices imply that the physical devices have a more updated view of the local state. Consequently, the physical devices, in certain scenarios, can manage in a better way their resources locally. So, the controller will have to manage the network globally while each device could take local decisions, without interfering nearby devices.

H. Compatibility With Other New Technologies

To satisfy the ever increasing requirements of future wireless networks, other new technologies, aside from slicing, are being proposed. For example:

- *Extreme Densification and Offloading*, which consists on massively deploying base stations on a given area, and complementary, using offloading techniques to redirect the traffic through different networks.
- *Millimeter Wave* consists on the use of higher frequencies of the spectrum, those of millimeter wavelength, to overcome the spectrum scarcity. These frequencies are more affected by path-loss and do not have good penetration through walls, therefore, they are expected to be used for indoor communications, along with dense deployments.
- *Massive MIMO* to spatially increase the spectral efficiency by using multiple antennas of Multiple Input Multiple Output (MIMO) technology.

How these technologies interact with a sliced design of a wireless network has still to be studied. For instance, slicing a MIMO interface introduces new challenges as several transmissions can happen in parallel, and then multiple slices would be “transmitting” at the same time.

I. Security

One of the main features of slicing is the abstraction process where the slice is viewed as a whole network, and the slice tenant can manage and configure it in its own way. This flexibility introduces higher security risks to wireless networks, as the slices share the same physical infrastructure. Hence, it is of crucial importance to offer security and isolation at the configuration, management and programming levels. However, to the best of our knowledge, there is a complete lack of research in security for wireless slicing.

In addition, wireless networks offer authentication and encryption on the air-interface, but with slicing, those functions need to be splitted between the slices. Furthermore, all the security issues related to virtualization and hardware sharing, are also relevant to the slicing approach. Extensive research efforts will need to undertake these challenges with extreme care before virtualization and slicing in wireless networks can be deployed.

VII. CONCLUSION

In this paper we introduced the concept of slicing as an integral approach of wireless networks virtualization. We outlined its relation to current trends such as SDN and NFV. We also described how slicing can benefit future wireless networks to satisfy new challenging requirements, showing possible scenarios of application.

We have introduced the two major problems of slicing wireless resources: resource allocation and isolation. We defined and explained in detail these problems in the context of wireless slicing, and presented the challenges for solving them, emphasizing those challenges caused by the wireless medium variability.

Next, we presented a classification and review of existing proposals for the predominant wireless technologies. We compared the different approaches and highlighted their advantages and drawbacks. Finally, we analyzed some research challenges and improvements needed for wireless resource slicing to become a reality.

In summary, slicing, as an integral part of wireless virtualization, appears as one of the technological enablers for meeting future wireless network requirements. Although many research efforts have been taken on wireless virtualization and slicing, several challenges still remain unsolved. Our intention with this article was to briefly survey some current proposals, highlighting their contributions and envisioning prospective research directions.

REFERENCES

- [1] A. Wang, M. Iyer, R. Dutta, G. N. Rouskas, and I. Baldine, “Network virtualization: Technologies, perspectives, and frontiers,” *J. Lightw. Technol.*, vol. 31, no. 4, pp. 523–537, Feb. 15, 2013.
- [2] H. Wen, P. K. Tiwary, and T. Le-Ngoc, *Wireless Virtualization*. Cham, Switzerland: Springer, 2013.
- [3] C. Liang and F. R. Yu, “Wireless network virtualization: A survey, some research issues and challenges,” *IEEE Commun. Surveys Tuts.*, vol. 17, no. 1, pp. 358–380, 1st Quart. 2015.
- [4] M. Yang *et al.*, “Software-defined and virtualized future mobile and wireless networks: A survey,” *Mobile Netw. Appl.*, vol. 20, no. 1, pp. 4–18, 2015.
- [5] J. Sachs and S. Baucke, “Virtual radio: A framework for configurable radio networks,” in *Proc. 4th Annu. Int. Conf. Wireless Internet*, Maui, HI, USA, 2008, Art. no. 61.
- [6] M. Hoffmann and M. Staufer, “Network virtualization for future mobile networks: General architecture and applications,” in *Proc. IEEE Int. Conf. Commun. Workshops (ICC)*, Kyoto, Japan, 2011, pp. 1–5.
- [7] K. Nakauchi, K. Ishizu, H. Murakami, A. Nakao, and H. Harada, “AMPHIBIA: A cognitive virtualization platform for end-to-end slicing,” in *Proc. IEEE Int. Conf. Commun. (ICC)*, Kyoto, Japan, 2011, pp. 1–5.
- [8] J. S. Panchal, R. D. Yates, and M. M. Buddhikot, “Mobile network resource sharing options: Performance comparisons,” *IEEE Trans. Wireless Commun.*, vol. 12, no. 9, pp. 4470–4482, Sep. 2013.
- [9] F. Hu, Q. Hao, and K. Bao, “A survey on software-defined network and openflow: From concept to implementation,” *IEEE Commun. Surveys Tuts.*, vol. 16, no. 4, pp. 2181–2206, 4th Quart. 2014.
- [10] D. Kreutz *et al.*, “Software-defined networking: A comprehensive survey,” *Proc. IEEE*, vol. 103, no. 1, pp. 14–76, Jan. 2015.
- [11] J. Qadir, N. Ahmed, and N. Ahad, “Building programmable wireless networks: An architectural survey,” *EURASIP J. Wireless Commun. Netw.*, vol. 2014, no. 1, pp. 1–31, 2014.
- [12] R. Sherwood *et al.*, “Can the production network be the testbed?” in *Proc. OSDI*, Vancouver, BC, Canada, 2010, pp. 365–378.

- [13] X. Jin, L. E. Li, L. Vanbever, and J. Rexford, "SoftCell: Scalable and flexible cellular core network architecture," in *Proc. 9th ACM Conf. Emerg. Netw. Exp. Technol.*, Santa Barbara, CA, USA, 2013, pp. 163–174.
- [14] K. Pentikousis, Y. Wang, and W. Hu, "Mobileflow: Toward software-defined mobile networks," *IEEE Commun. Mag.*, vol. 51, no. 7, pp. 44–53, Jul. 2013.
- [15] A. Gudipati, D. Perry, L. E. Li, and S. Katti, "SoftRAN: Software defined radio access network," in *Proc. 2nd ACM SIGCOMM Workshop Hot Topics Softw. Defined Netw.*, Hong Kong, 2013, pp. 25–30.
- [16] L. Suresh, J. Schulz-Zander, R. Merz, A. Feldmann, and T. Vazao, "Towards programmable enterprise WLANs with Odin," in *Proc. 1st Workshop Hot Topics Softw. Defined Netw.*, Helsinki, Finland, 2012, pp. 115–120.
- [17] J. Schulz-Zander *et al.*, "Programmatic orchestration of WiFi networks," in *Proc. USENIX Annu. Tech. Conf. (USENIX ATC)*, Philadelphia, PA, USA, 2014, pp. 347–358.
- [18] R. Riggio *et al.*, "Programming software-defined wireless networks," in *Proc. 10th Int. Conf. Netw. Service Manag. (CNSM)*, Rio de Janeiro, Brazil, 2014, pp. 118–126.
- [19] J. Schulz-Zander, N. Sarrar, and S. Schmid, "Towards a scalable and near-sighted control plane architecture for WiFi SDNs," in *Proc. 3rd Workshop Hot Topics Softw. Defined Netw.*, Chicago, IL, USA, 2014, pp. 217–218.
- [20] J. Schulz-Zander, N. Sarrar, and S. Schmid, "AeroFlux: A near-sighted controller architecture for software-defined wireless networks," presented at the Open Netw. Summit, Santa Clara, CA, USA, 2014.
- [21] R. Mijumbi *et al.*, "Network function virtualization: State-of-the-art and research challenges," *IEEE Commun. Surveys Tuts.*, vol. 18, no. 1, pp. 236–262, 1st Quart. 2016.
- [22] C. Chen, "C-ran: The road towards green radio access network," China Mobile Res. Inst., White Paper 2.5, 2011.
- [23] M. Hadzialic, B. Dosenovic, M. Dzaferagic, and J. Musovic, "Cloud-RAN: Innovative radio access network architecture," in *Proc. 55th Int. Symp. (ELMAR)*, Zadar, Croatia, 2013, pp. 115–120.
- [24] H. Hawilo, A. Shami, M. Mirahmadi, and R. Asal, "NFV: State of the art, challenges, and implementation in next generation mobile networks (vepc)," *IEEE Netw.*, vol. 28, no. 6, pp. 18–26, Nov. 2014.
- [25] P. Dely *et al.*, "CloudMAC—An OpenFlow based architecture for 802.11 MAC layer processing in the cloud," in *Proc. IEEE Globecom Workshops (GC Wkshps)*, Anaheim, CA, USA, 2012, pp. 186–191.
- [26] J. Erfanian and B. Daly, "5G white paper," NGMN Alliance, White Paper, Mar. 2015.
- [27] M. Derakhshani, X. Wang, T. Le-Ngoc, and A. Leon-Garcia, "Virtualization of multi-cell 802.11 networks: Association and airtime control," *IEEE Trans. Mobile Comput.*, to be published.
- [28] "5G systems," Ericsson, Stockholm, Sweden, White Paper Uen 284 23-3244, Jan. 2015. [Online]. Available: <http://www.ericsson.com/res/docs/whitepapers/what-is-a-5g-system.pdf>
- [29] H. Wen, P. K. Tiwary, and T. Le-Ngoc, "Current trends and perspectives in wireless virtualization," in *Proc. Int. Conf. Sel. Topics Mobile Wireless Netw. (MoWNeT)*, Montreal, QC, Canada, 2013, pp. 62–67.
- [30] K. Kusume *et al.*, *Deliverable D1.5. Updated Scenarios, Requirements and KPIs for 5G Mobile and Wireless System With Recommendations for Future Investigations*, document ICT-317669, METIS Project, Apr. 2015. [Online]. Available: https://www.metis2020.com/wp-content/uploads/deliverables/METIS_D1.5_v1.pdf
- [31] "Infrastructure report 2014. Ofcom's second full analysis of the UKs communications infrastructure," Ofcom, Office Commun., London, U.K., 2014. [Online]. Available: <http://stakeholders.ofcom.org.uk/binaries/research/infrastructure/2014/infrastructure-14.pdf>
- [32] W. Tong and Z. Peiying, "5G: A technology vision," Huawei Technol., Shenzhen, China, Tech. Rep. M3-023985-20131104-C-1, 2013. [Online]. Available: <http://www1.huawei.com/en/about-huawei/publications/winwin-magazine/hw-329304.htm>
- [33] Y. Zaki, "Future mobile communications: LTE optimization and mobile network virtualization," Ph.D. dissertation, Faculty Phys. Elect. Eng., Univ. Bremen, Bremen, Germany, May 2012.
- [34] B. Graham. (Nov. 2015). *NFV and SDN Answer or Question?* [Online]. Available: <http://www.cnsm-conf.org/2015/files/sdnnfv-keynote.pdf>
- [35] F. Capozzi, G. Piro, L. A. Grieco, G. Boggia, and P. Camarda, "Downlink packet scheduling in LTE cellular networks: Key design issues and a survey," *IEEE Commun. Surveys Tuts.*, vol. 15, no. 2, pp. 678–700, 2nd Quart. 2013.
- [36] *802.11-2012—IEEE Standard for Information Technology—Telecommunications and Information Exchange Between Systems—Local and Metropolitan Area Networks—Specific Requirements—Part 11: WLAN MAC and PHY Specifications*, IEEE Standard 802.11-2012, Mar. 2012.
- [37] J. Y. Hui, "Resource allocation for broadband networks," *IEEE J. Sel. Areas Commun.*, vol. 6, no. 9, pp. 1598–1608, Dec. 1988.
- [38] L. Georgiadis, M. J. Neely, and L. Tassiulas, *Resource Allocation and Cross-Layer Control in Wireless Networks*. Hanover, MA, USA: Now, 2006.
- [39] R. Kokku, R. Mahindra, H. Zhang, and S. Rangarajan, "NVS: A substrate for virtualizing wireless resources in cellular networks," *IEEE/ACM Trans. Netw.*, vol. 20, no. 5, pp. 1333–1346, Oct. 2012.
- [40] Y. Zaki, L. Zhao, C. Goerg, and A. Timm-Giel, "LTE wireless virtualization and spectrum management," in *Proc. 3rd Joint IFIP Wireless Mobile Netw. Conf. (WMNC)*, Budapest, Hungary, 2010, pp. 1–6.
- [41] Y. Zaki, L. Zhao, C. Goerg, and A. Timm-Giel, "LTE mobile network virtualization," *Mobile Netw. Appl.*, vol. 16, no. 4, pp. 424–432, Jun. 2011. [Online]. Available: <http://dx.doi.org/10.1007/s11036-011-0321-7>
- [42] M. Li *et al.*, "Investigation of network virtualization and load balancing techniques in LTE networks," in *Proc. IEEE 75th Veh. Technol. Conf. (VTC Spring)*, Yokohama, Japan, 2012, pp. 1–5.
- [43] M. Yang, Y. Li, L. Zeng, D. Jin, and L. Su, "Karnaugh-map like online embedding algorithm of wireless virtualization," in *Proc. 15th Int. Symp. Wireless Pers. Multimedia Commun. (WPMC)*, Taipei, Taiwan, 2012, pp. 594–598.
- [44] J. van de Belt, H. Ahmadi, and L. E. Doyle, "A dynamic embedding algorithm for wireless network virtualization," in *Proc. IEEE 80th Veh. Technol. Conf. (VTC Fall)*, Vancouver, BC, Canada, 2014, pp. 1–6.
- [45] M. I. Kamel, L. B. Le, and A. Girard, "LTE wireless network virtualization: Dynamic slicing via flexible scheduling," in *Proc. IEEE 80th Veh. Technol. Conf. (VTC Fall)*, Vancouver, BC, Canada, 2014, pp. 1–5.
- [46] G. Bhanage, I. Seskar, R. Mahindra, and D. Raychaudhuri, "Virtual basestation: Architecture for an open shared WiMAX framework," in *Proc. 2nd ACM SIGCOMM Workshop Virtualized Infrastruct. Syst. Architect.*, New Delhi, India, 2010, pp. 1–8.
- [47] G. Bhanage, R. Daya, I. Seskar, and D. Raychaudhuri, "VNTS: A virtual network traffic shaper for air time fairness in 802.16e systems," in *Proc. IEEE Int. Conf. Commun. (ICC)*, Cape Town, South Africa, 2010, pp. 1–6.
- [48] R. Kokku, R. Mahindra, H. Zhang, and S. Rangarajan, "Cellslice: Cellular wireless resource slicing for active RAN sharing," in *Proc. 5th Int. Conf. Commun. Syst. Netw. (COMSNETS)*, Bengaluru, India, 2013, pp. 1–10.
- [49] A. Banchs, P. Serrano, P. Patras, and M. Natkaniec, "Providing throughput and fairness guarantees in virtualized WLANs through control theory," *Mobile Netw. Appl.*, vol. 17, no. 4, pp. 435–446, Aug. 2012. [Online]. Available: <http://dx.doi.org/10.1007/s11036-012-0382-2>
- [50] K. Nakachi, Y. Shoji, and N. Nishinaga, "Airtime-based resource control in wireless LANs for wireless network virtualization," in *Proc. 4th Int. Conf. Ubiquitous Future Netw. (ICUFN)*, 2012, pp. 166–169.
- [51] K. Guo, S. Sanadhya, and T. Woo, "ViFi: Virtualizing WLAN using commodity hardware," *ACM SIGMOBILE Mobile Comput. Commun. Rev.*, vol. 18, no. 3, pp. 41–48, 2014.
- [52] G. Smith, A. Chaturvedi, A. Mishra, and S. Banerjee, "Wireless virtualization on commodity 802.11 hardware," in *Proc. 2nd ACM Int. Workshop Wireless Netw. Testbeds Exp. Eval. Characterization*, Montreal, QC, Canada, 2007, pp. 75–82.
- [53] R. Mahindra *et al.*, "Space versus time separation for wireless virtualization on an indoor grid," in *Proc. Next Gener. Internet Netw. (NGI)*, Kraków, Poland, pp. 215–222.
- [54] L. Xia *et al.*, "Virtual WiFi: Bring virtualization from wired to wireless," *ACM SIGPLAN Notices*, vol. 46, no. 7, pp. 181–192, 2011.
- [55] G. Bhanage, D. Vete, I. Seskar, and D. Raychaudhuri, "SplitAP: Leveraging wireless network virtualization for flexible sharing of wlans," in *Proc. IEEE Glob. Telecommun. Conf. (GLOBECOM)*, Miami, FL, USA, 2010, pp. 1–6.
- [56] R. Dhar, G. George, A. Malani, and P. Steenkiste, "Supporting integrated MAC and PHY software development for the USRP SDR," in *Proc. 1st IEEE Workshop Netw. Technol. Softw. Defined Radio Netw. (SDR)*, 2006, pp. 68–77.
- [57] D. A. Burgess and H. S. Samra, "The openBTS project," 2015. [Online]. Available: <http://openbts.org/>
- [58] M. Bansal, J. Mehlman, S. Katti, and P. Levis, "OpenRadio: A programmable wireless dataplane," in *Proc. 1st Workshop Hot Topics Software Defined Netw.*, Helsinki, Finland, 2012, pp. 109–114.

- [59] M. Neufeld, J. Fifield, C. Doerr, A. Sheth, and D. Grunwald, "SoftMAC—Flexible wireless research platform," in *Proc. HotNets-IV*, College Park, MD, USA, 2005, pp. 1–5.
- [60] I. Tinnirello *et al.*, "Wireless MAC processors: Programming MAC protocols on commodity hardware," in *Proc. IEEE INFOCOM*, Orlando, FL, USA, 2012, pp. 1269–1277.
- [61] B. Jooris *et al.*, "TAISC: A cross-platform MAC protocol compiler and execution engine," *Comput. Netw.*, to be published.
- [62] J. C. Doyle, B. A. Francis, and A. R. Tannenbaum, *Feedback Control Theory*. Courier Corporat., 2013.
- [63] T. M. Mitchell, *Machine Learning*. New York, NY, USA: McGraw-Hill, 1992.
- [64] X. Wang, P. Krishnamurthy, and D. Tipper, "Wireless network virtualization," in *Proc. Int. Conf. Comput. Netw. Commun. (ICNC)*, San Diego, CA, USA, 2013, pp. 818–822.



contact point of the TM Forum with UPC.

Joan Serrat received the telecommunication engineering degree and the Ph.D. degree in telecommunication engineering from the Universitat Politècnica de Catalunya (UPC), in 1977 and 1983, respectively. He is currently a Full Professor with UPC, where he has been involved in several collaborative projects with different European research groups, both through bilateral agreements or through participation in European funded projects. His topics of interest are in the field of autonomic networking and service and network management. He is currently the



Matías Richart received the computer engineer and master's degrees from the University of the Republic, Uruguay, in 2011 and 2014, respectively. He is currently pursuing the Ph.D. degree with the University of the Republic and the Polytechnic University of Catalonia, Spain. His research interests include autonomic control and resource management of wireless networks, virtual wireless networks, heterogeneous wireless networks, and network simulation.



with the Department of Computer Science, University of the Republic.

Javier Baliosian received the computer engineer degree from the University of the Republic, Uruguay, in 1998, and the Ph.D. degree from the Polytechnic University of Catalonia, Spain, in 2005. Since then, he has been involved in several research projects with different groups such as the Computer Laboratory, University of Cambridge, the Laboratory of Communication Networks, Royal Institute of Technology, Sweden, and the Ericsson Ireland Research Centre, where he was a Researcher and a Project Coordinator until 2007. He is currently



virtualized networks and functions.

Juan-Luis Gorricho received the telecommunication engineering degree and the Ph.D. degree from the Universitat Politècnica de Catalunya (UPC), in 1993 and 1998, respectively. He is currently an Associate Professor with the UPC. His recent research interests are in applying artificial intelligence to ubiquitous computing and network management, with special interest on using smartphones to achieve the recognition of user activities and locations, and applying linear programming and reinforcement learning to resource management in