

Customisation and Trade-offs in 5G RAN Slicing

Conor Sexton, Nicola Marchetti, Luiz A. DaSilva
CONNECT - Trinity College Dublin, Ireland.

Abstract—The heterogeneity of the requirements for 5G necessitate a versatile 5G radio access network (RAN); slicing offers a way of realising a flexible RAN through customised virtual sub-networks. In this paper, we focus on how enabling lower layer flexibility in the RAN affects the development of RAN slicing, particularly in relation to ensuring isolation between RAN slices. We first examine how RAN slices may be individually tailored for different services. We follow this up with an examination of the potential time-frequency resource structure of the RAN, focusing on the trade-off between flexibility and the overhead related to ensuring coexistence between contrasting RAN slices. Based on this analysis, we suggest an approach that permits the allocation of resources to a service-type to be performed separately to resource allocation for individual services belonging to that type.

Keywords—5G; slicing; RAN; resource allocation

INTRODUCTION

5G will cater to numerous different vertical industries, which typically target the needs of a specific group of people and machines, presenting a diverse range of requirements to the network. This is a marked change from previous generations in which the network was designed with a single, dominant market in mind (such as LTE and data services). The 5G RAN will be faced with an unprecedented level of heterogeneity in the requirements it must satisfy, including high reliability, low latency, high data rates, and low energy consumption. Each of these requirements necessitate contrasting behaviour from the RAN, making it challenging to design a one-size-fits-all style RAN which is capable of simultaneously satisfying each of them.

Hence, the shift in focus in 5G to multifaceted requirements necessitates a flexible RAN that is characterised by its versatility and ability to adapt to different services. Slicing offers a way to introduce this adaptability in 5G through customised virtual sub-networks, and is firmly on 3GPP's radar, with early standardization activities such as specifying requirements and examining potential architectures currently under way. While core slicing has received plentiful attention, the conflicting demands placed on the 5G RAN also motivates the need for RAN slicing. Additionally, advances such as software defined radio (SDR) and cloud-RAN have drastically increased the potential for resource sharing in the RAN, with SDR enabling different transmission schemes to be implemented on the same components and cloud-RAN allowing these components, namely radio heads and baseband processing units, to be shared between different sub-networks.

RAN slicing presents a vastly different set of challenges compared to core slicing. The diversity of requirements demanded of the 5G RAN require physical changes to the manner

in which signals are formed and packed together for transmission, such as variable frame structures and subcarrier spacings. For example, a RAN slice serving automotive services in a high mobility scenario may use a wider subcarrier spacing to combat high Doppler shifts, while a RAN slice serving a latency-sensitive service such as real-time gaming may use fewer symbols in each sub-frame. These lower layer customisations are necessary, but introduce challenges in ensuring isolation between RAN slices.

It has been shown that the use of different frame structures in a system can result in inter-service-band interference [1]. To ensure that transmissions using different numerologies do not interfere with one-another, it is necessary to adopt a combination of guard bands, enhanced spectral filtering or other advanced signal processing techniques. Hence, although RAN slicing increases the ability of the RAN to serve dissimilar services and industries, this flexibility comes at a cost.

Several papers have examined conceptual architectures [2], [3] or proposed radio resource allocation schemes for RAN slicing [4]. Our focus in this paper is instead on how lower layer RAN concerns pertaining to customisation and flexibility affect the development of RAN slicing, particularly in relation to ensuring isolation between RAN slices. Specifically, we direct our attention to the trade-off between adaptability on the one hand, and the cost/overhead of this adaptability on the other hand. The flexibility arises from being able to serve the needs of multiple vertical industries by tailoring the manner in which the signal is transmitted. The cost/overhead arises from the challenges associated with ensuring the coexistence of multiple RAN slices which have been designed to satisfy diverse use-cases.

We begin the paper with an examination of the benefits of various lower-layer design choices in the RAN for serving different use-cases. We then consider four different configurations of the resource grid, with each managing the coexistence of RAN slices with approaches ranging from very static to very dynamic. Focusing on the time-frequency resource grid and how multiple RAN slices can coexist, we examine how these different configurations of the resource grid perform with regard to the above-mentioned adaptability-overhead trade-off. To accomplish this, we provide both a qualitative and quantitative analysis of both sides of the trade-off, and evaluate them for each configuration under consideration.

We note that a network slice represents an end-to-end concept, enveloping the entire network, as defined in 3GPP Technical Specification 28.531. In this paper we are concerned with the RAN, constituting a part of a full network slice, which we refer to as a *RAN slice* throughout the paper.

TABLE I. NUMEROLOGY CONSIDERATIONS FOR SCENARIOS

Channel Measure	Effect	Solution	Relevant Scenario
delay spread	large delay spreads result in inter-symbol interference	use longer symbols and/or cyclic prefixes to absorb symbol tails	extended coverage; high multi-path environments
coherence bandwidth	small coherence bandwidths result in frequency selective fading	use narrow subcarriers to approximate flat fading channels	extended coverage; high multi-path environments
doppler spread	large doppler spreads result in inter-carrier interference	increase subcarrier spacing to separate carrier frequencies	high mobility
coherence time	short coherence times result in outdated feedback and poor channel estimation	use shorter symbols or redesign pilot placements to increase rate of CSI feedback	high mobility

TAILORING RAN SLICES

The benefits afforded by RAN slicing stem from the ability to adapt RAN behaviour to a vertical or service type. The effectiveness of RAN slicing is therefore dependent on the range of options available for adjusting the behaviour of the RAN slice.

Mixed numerologies

The numerology of a multi-carrier system is the set of design parameters which influence how the system performs, such as the subcarrier spacing, the cyclic prefix length, the symbol length, and the number of symbols in a sub-frame. In previous multi-carrier systems such as LTE and WiFi, predefined numerologies were chosen that were suited to the primary service and operating environment being targeted. In contrast, 5G may permit each RAN slice to use a separate numerology that is tailored for a particular vertical.

Subcarrier spacing and symbol length: These two design options are related, and cannot be changed independently. In extended coverage scenarios or high multi-path environments, large delay spreads result in inter-symbol interference, and small coherence bandwidths result in frequency selective fading. The solution is to use longer symbols to absorb symbol tails and narrow subcarriers to approximate flat fading channels. In high mobility scenarios, large Doppler spreads result in inter-carrier interference and short coherence times result in poor channel estimation. The solution is to increase subcarrier spacing to separate carrier frequencies, resulting in shorter symbols. [5] gives a good overview of the design options relating to subcarrier spacing (and symbol length). 5G New Radio (NR) will support scalable subcarrier spacings which are an integer multiple of the LTE spacing of 15 kHz, up to a maximum spacing of 240 kHz as specified in 3GPP Technical Specification 38.211.

Number of symbols per transmission time interval (TTI): 5G NR provides support for shorter TTI times than the fixed length of 1ms in LTE.

Cyclic prefix: 5G NR supports two cyclic prefix lengths: one for regular use, and one to enable extended coverage (at a subcarrier spacing of 60kHz).

Pilot placement: Reference symbols known as pilots are used to estimate the channel for the purpose of one-tap equalisation in OFDM and to allow the UE to provide channel state information to the network to facilitate frequency selective scheduling. If the channel coherence time reduces, such as in a high mobility scenario, the pilot placements are no longer sufficiently close for interpolation to provide accurate channel

estimates for equalisation of OFDM symbols. Similarly, the feedback provided to the network could be outdated when the next round of scheduling is performed. In these cases, pilot symbols should be placed closer together.

Scheduling request period: In the uplink, the scheduling request period can be decreased to reduce the latency. This will increase the control plane overhead and potentially reduce throughput.

Table I displays a number of channel attributes, their effect on performance, how to mitigate any adverse effects through numerology changes, and the scenario it applies to.

Technologies

The use of specialist technologies which are optimal only for particular verticals may be confined to specific RAN slices. For example, massive-MIMO can provide high data rates by scheduling multiple users on the same time-frequency resources. However, up to half of the time in massive-MIMO might be spent on training, making it unsuitable for latency-intolerant and high-mobility devices. Opportunistic off-loading to radio access technologies such as WiFi or millimetre wave is also a possibility for mobile broadband traffic. Different implementations of grant-free multiple access in the uplink may be used in a RAN slice for low latency traffic, or in a RAN slice for machine-type communication (MTC). Non-orthogonal multiple access schemes might also find application in MTC RAN slices. Both grant-free access and NOMA are examined for NR in 3GPP Technical Report 38.812.

Waveforms

Some of the prominent waveforms that were considered for 5G include filter-bank based waveforms which perform filtering on a subcarrier basis such as filter bank multi-carrier (FBMC), and waveforms based on filtering contiguous blocks of subcarriers such as filtered-OFDM (f-OFDM). While these waveforms each proved advantageous in select scenarios, OFDM was retained in 5G NR due to its all-round performance. However, in future networks, we envision that multiple waveforms may be permitted and RAN slices might employ different waveforms depending on their targeted use-cases.

To demonstrate the benefits that employing a different waveform can provide, we adapt some results from our recent paper [6], in which we consider two types of communication. One RAN slice targets regular cellular users (CUEs) communicating with the base station in a synchronous manner. The

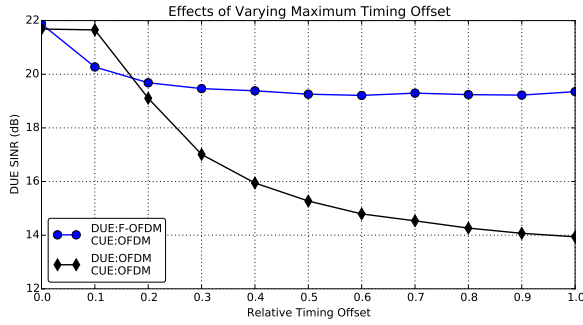


Fig. 1. DUE SINR performance as the maximum permitted timing offset is varied as a fraction of an OFDM symbol time (adapted from our work in [6]).

second RAN slice targets a high-rate clustered machine-type scenario, such as a smart factory, in which directly communicating devices (DUEs) do not achieve time-alignment. As illustrated by Fig. 1, the DUEs achieve greater performance if the second RAN slice employs a waveform such as f-OFDM with improved spectral containment, as there is less leakage interference between devices arising from the misaligned communication.

RAN slice-specific scheduling

RAN slices could be controlled by a single, integrated MAC layer with specific scheduling rules for each RAN slice, allowing RAN slice-specific quality of service differentiation. This is the approach suggested by 3GPP Technical Report 30.801 on New Radio (NR) access technology which proposes that the RAN uses different resource management configurations for different slices.

Alternatively, each RAN slice could have its own instance of the MAC layer and its own scheduler. This idea of slice-specific schedulers is explored in [3], in which a two-tier MAC scheduler consisting of a slice-specific resource manager and a resource mapper is proposed. Each slice possesses its own custom resource manager which allocates virtual resource blocks to users using its own scheduler. These virtual resources are then mapped to physical resources by a system level mapper which considers slice priority.

The first option requires a complex MAC layer that can manage multiple different slice configurations, but incorporates a single scheduler with a view of the entire system. The second option requires more coordination between schedulers, but the schedulers themselves are simpler and offer a more flexible way to customise how scheduling is performed for each RAN slice.

TIME-FREQUENCY RESOURCE HIERARCHY OF THE 5G RAN

Fig. 2 illustrates four potential options for sharing time-frequency resources between RAN slices with different numerologies, which are briefly described below.

- 1) **Fixed contiguous sub-band approach:** Each RAN slice uses a fixed, predefined contiguous sub-band.

- 2) **Variable contiguous sub-band approach:** Each RAN slice uses a contiguous sub-band consisting of a fixed region and a variable region. The variable region may be shared between neighbouring RAN slices according to demand.
- 3) **Sub-band tiling approach:** The resource grid is divided into regular sub-bands of fixed size.
- 4) **Frame tiling approach:** In this approach, tiling occurs at frame granularity in the time domain, and resource block granularity in the frequency domain.

Qualitative and Quantitative Comparison

We employ Monte Carlo simulation to compare the four possible resource grid designs. In the time domain, each block in Fig. 2 represents a frame. As per 5G NR, frames are 10ms in length irrespective of the numerology in use. In the frequency domain, each block represents a resource block consisting of 12 subcarriers, which varies in width depending on the numerology. We consider three RAN slices, each using a different subcarrier spacing (15kHz, 30kHz, 60kHz), and assume the network is fully loaded with all time-frequency resources in use. In each 1 second interval, sub-bands and resource blocks are randomly assigned to each RAN slice in the sub-band and frame tiling approaches, respectively, until each of the three RAN slices has been granted an equal portion of the time-frequency grid. For the sub-band tiling approach, each sub-band is 10 frames in length and 120kHz wide, while for the variable contiguous sub-band approach, the ratio between the fixed and variable part of each RAN slice is 3:1.

1. *Coexistence overhead:* In a multi-service system in which RAN slices utilise different technologies, proactive measures are required to ensure that RAN slices coexist without adversely affecting one another. Of primary concern is that permitting mixed subcarrier spacings and frame structures can result in inter-numerology interference [1], as the side-lobes of the symbols in the frequency domain will no longer overlap in an orthogonal manner. The use of waveforms with enhanced filtering, such as filtered-OFDM, has been considered to overcome this issue [7]. However, even with this enhanced filtering, guard bands may be needed to prevent inter-slice interference. Low complexity inter-service band interference cancellation algorithms which precode information symbols at the transmitter provide another option [8].

We use the general term *coexistence overhead* to capture the effects of leakage interference between RAN slices, resource loss due to inter-slice guard bands and guard times, and the complexity of advanced interference cancellation techniques. The coexistence overhead is dependent on the number (and length) of boundaries between numerologies over a certain time interval. We consider one boundary to be a single border of length 10ms between two neighbouring radio frames employing different numerologies, and use the previously described simulation set-up to determine the average number of boundaries in a 1 second by 20MHz window of the resource grid.

As can be seen in Fig. 3, both the fixed and variable contiguous sub-band approaches minimise the number of boundaries

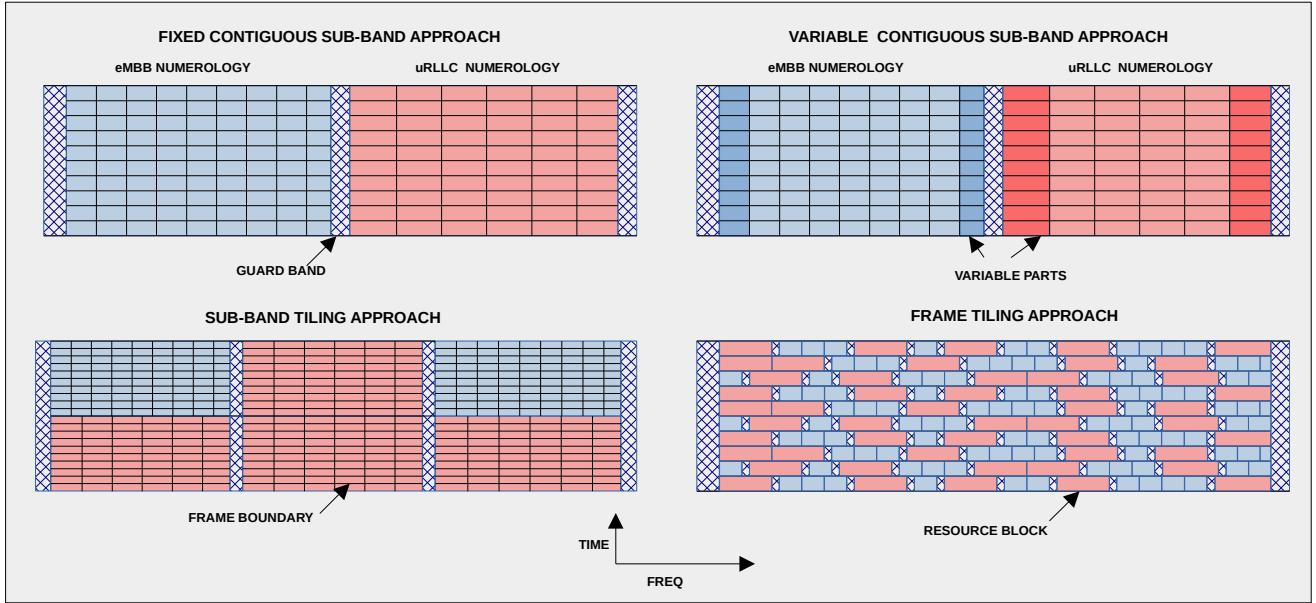


Fig. 2. Four options for sharing time-frequency resources between two services employing different numerologies.




Fig. 3. Comparison of the overhead and level of adaptability associated with each approach.

between different RAN slices; hence the coexistence overhead is low and constant. In the sub-band tiling approach, there are more boundaries between different numerologies, which increases the coexistence overhead to approximately 3.5 times that of the contiguous approaches. This is still relatively low compared to the frame tiling approach, which results in a large number of boundaries between numerologies due to the small granularity at which time-frequency resources can be assigned.

2. *Ability to dynamically adapt to traffic:* RAN slices must be able to adapt to different traffic loads, requiring the active sharing of time-frequency resources between RAN slices. 5G will target many different use cases such as smart healthcare, automated cars, Industry 4.0, and smart city management, and these will likely present different traffic patterns than classic mobile broadband type, smart-phone driven traffic. As a result, the relative split in network load attributed to each RAN slice can reasonably be expected to vary spatially and temporally, and further investigation on the behaviour of different traffic types based on empirical evidence is required.

We estimate the adaptability of each approach by considering how time-frequency resources can be reassigned among RAN slices when one RAN slice is highly loaded and another is lightly loaded. We adapt our simulation model so that, on average, one RAN slice is 50% under-loaded, one RAN slice is fully-loaded, and one RAN slice is 50% over-loaded. For each one second time window, we evaluate the ratio between the total number of requests served and the total number of serviceable requests, which provides us with a measure of utilisation of the time-frequency resources. If each request takes a fixed amount of the time-frequency grid to satisfy, the total number of serviceable requests is equal to the lesser value between the total number of requests submitted to all RAN slices and the total number of requests that can be accommodated by the resource grid.

The average utilisation for a one second window is presented in Fig. 3. The fixed contiguous sub-band approach proves to be the least flexible, with only 84% average utilisation, due to its inability to reassign unused resources from under-loaded RAN slices to over-loaded RAN slices. The variable contiguous sub-band approach offers more flexibility, providing an improved average utilisation of 94%, due to the fact that RAN slices can be rearranged in the resource grid and can borrow resources from their direct neighbours. Similar performance is provided by the sub-band tiling approach, which divides RAN slices into several smaller sub-bands which are distributed throughout the system bandwidth. Finally, the frame tiling approach is able to provide 100% utilisation of resource blocks, ensuring that no RAN slice has unused resources that could be reassigned to another RAN slice.

3. *Control plane considerations:* The regular resource grid consisting of a single numerology in LTE allows time-frequency resources to be conveniently indexed for the purpose

of resource allocation, with the physical downlink control channel (PDCCH) comprising the first few symbols of each sub-frame and spread across the entire channel bandwidth. The irregularity of the resource grid in a multi-service, mixed numerology system requires more complex indexing schemes for referring to portions of the grid. In response to this, a new concept called a bandwidth part (BWP) has been introduced in 5G NR. A BWP is a contiguous set of physical resource blocks with an associated numerology. UEs can be configured with up to four BWPs, with one active at any given time in both the uplink and downlink. The PDCCH for NR is also no longer spread across the entire bandwidth, but is instead localised in a control-resource set (CORESET) within each BWP. Hence, 5G NR allows UEs to be configured to operate in a contiguous bandwidth region using a specified numerology, with the control plane for that region housed internally.

We note that the fixed and variable contiguous sub-band approaches could deploy an LTE-like control plane across the fixed bandwidth allocation of each sub-band. NR could accommodate both of these options and, additionally, the sub-band tiling approach by configuring sub-bands as BWPs with a CORESET defined within each. An NR-style control plane is not an option for the frame tiling approach, as a minimum frequency width for the CORESET prevents the multiplexing of numerologies as BWPs at a frame/resource block granularity. If multiplexing at this granularity is to be considered in future, alternative control-plane designs must be devised.

Trade-off

From the comparison provided in Fig. 3, we can identify a inherent trade-off between adaptability on one hand, and the coexistence overhead on the other hand. At one extreme, the fixed contiguous sub-band approach incurs very little overhead, but is very limited in its adaptability. At the other extreme, the frame tiling approach provides optimal flexibility, but at the cost of a substantial coexistence overhead. The solution appears to be a compromise to both extreme approaches, with both the variable sub-band approach and sub-band tiling approach providing a similar high level of adaptability with a reasonably low overhead.

In short, enabling a highly adaptable system that can adjust to temporally and spatially changing traffic demands for different RAN slices results in an increased number of boundaries between RAN slices, and hence a higher coexistence overhead.

RAN PROFILES

Although RAN slices may be tailored for individual verticals or services, in reality many types of services across different verticals will require similar network behaviour. Hence, we present the idea of RAN profiles: a specific configuration of the air interface that has been selected to meet the demands of a broad category of services. In particular, the RAN profile consists of the more general configuration choices outlined in Section II, relating to lower layer concerns such as subcarrier spacing, TTI length, pilot placement, and scheduling request period. Each of these configuration choices is relevant to an

array of similar services, while individual RAN slices within each profile can further configure their behaviour using some of the more specific configuration options outlined in Section II, such as the choice of scheduler, the use of MIMO, and diversity schemes.

We base RAN profiles on the variable contiguous sub-band approach (Fig. 2, case 2), due to its relatively low overhead and high adaptability (Fig. 3). In the simplest case, the three main 3GPP target areas may be supported in dedicated profiles: enhanced mobile broadband (eMBB), ultra-reliable low latency communication (uRLLC), and massive machine-type communication (mMTC). Basically, eMBB, uRLLC and mMTC would each have their own version of the air interface, using a frame structure suited to their needs. These air interfaces would each have their own resource grid that could then be further sliced using one of the aforementioned tiling approaches (i.e. sub-band or frame based) since the time-frequency resources are homogeneous and hence do not suffer from inter-numerology interference. This allows for further customisation, which can then be applied to the RAN slices through control over a RAN slice-specific scheduler, modulation order, or other parameters.

This idea of supporting multiple configurations of the RAN in 5G was envisioned in the METIS II project, which outlined the need for air interface variants (AIVs) that are optimized for one or more target scenarios/services [9]. A user-plane design framework for a service-tailored 5G RAN is outlined in [10], incorporating the idea of AIVs.

Coexistence overhead

Coexistence must be considered at two levels: coexistence of profiles and coexistence of RAN slices. The coexistence of profiles incurs the same overhead as the variable contiguous sub-band approach, which minimises the number of boundaries between different numerologies, as evident in Fig. 3. Although the coexistence of RAN slices is based on a tiling approach, the resources within a particular profile are homogeneous in nature, resulting in a reasonably low coexistence overhead.

Control plane

Each profile could possess a dedicated control plane which can be implemented similarly to LTE, with the first few symbols of each sub-frame carrying control information. In LTE, the central subcarriers house the primary and secondary synchronization signals and the physical broadcast channel, which is used to broadcast information regarding the configuration of the system. A similar approach could be adopted for a multi-service RAN, and is displayed in Fig. 4.

An NR-style control plane is also an option, with each sub-band implemented as a BWP containing the CORESET for that frequency region. The size of the BWP and location of the CORESET within the BWP can be adjusted in response to the variable sizes of the sub-bands.

Ability to adapt to traffic variations

As depicted in Fig. 5, the profile-based RAN will be a multi-service, multi-carrier system with a 5-tier frequency resource

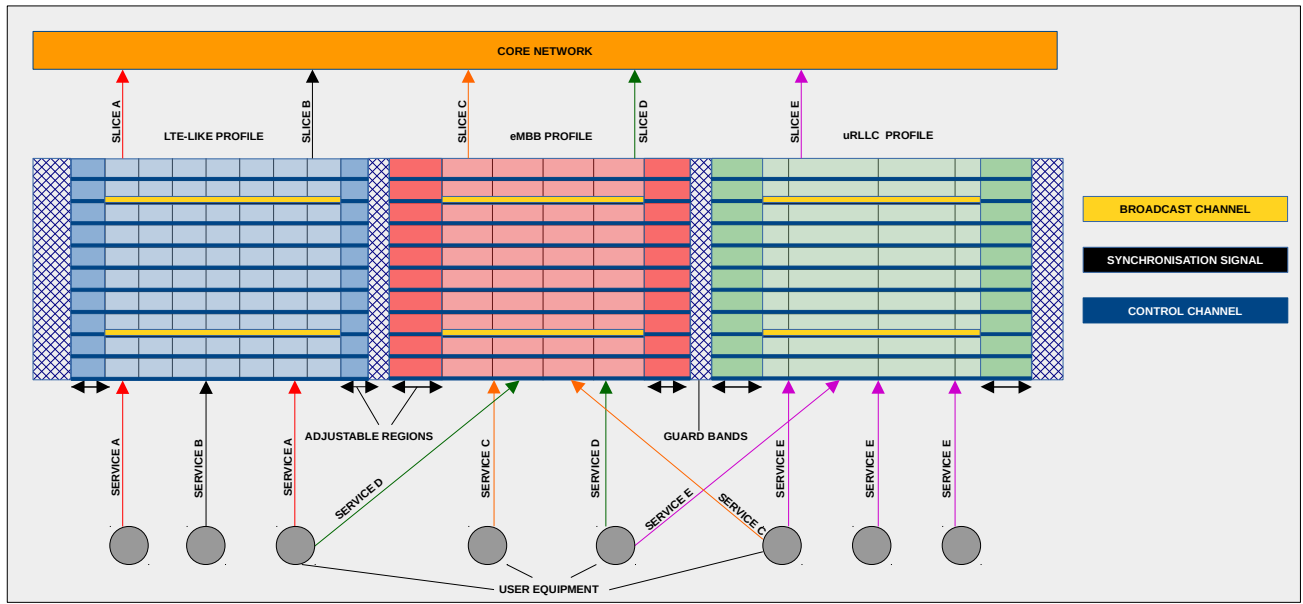


Fig. 4. RAN profiles are based on the variable contiguous sub-band approach, with each profile housing multiple slices.

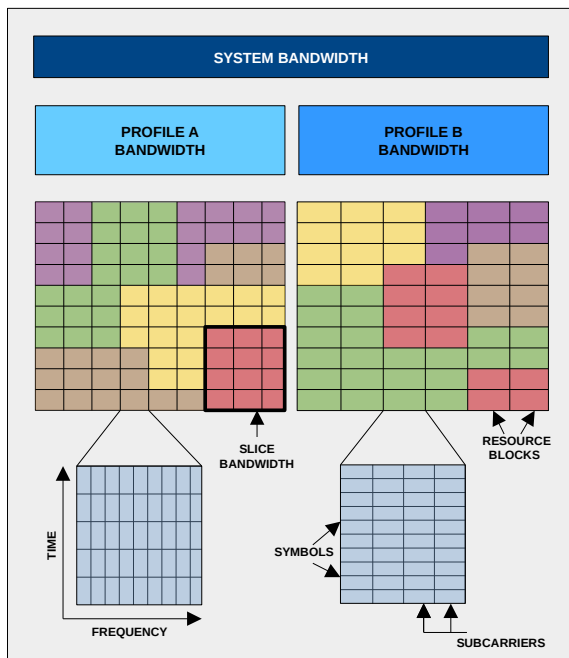


Fig. 5. The 5G RAN may be a multi-service, multi-carrier system with a 5-tier frequency resource structure consisting of system bandwidth, profile bandwidth, RAN slice bandwidth, resource block, and subcarrier.

structure consisting of system bandwidth, profile bandwidth, RAN slice bandwidth, resource block, and subcarrier, listed in descending order of granularity. The system bandwidth is divided into contiguous regions, with each region constituting a different RAN profile. RAN slices are then created dynamically and assigned a number of time-frequency resources

from a RAN profile sub-band. The profile bandwidth can be expanded and contracted within defined limits, while RAN slices can be dynamically allocated in a tiling approach within the profile bandwidth.

Hence, the profile approach offers separate control over the resources granted to a particular service area such as eMBB, and the resources allocated to a particular RAN slice in that service area, such as a video streaming service. In this regard, the profile approach proves to be quite flexible. To reduce signalling overhead and complexity, we suggest that profile bandwidths should be updated less frequently than the minimum scheduling unit for users; for example, every 100 TTIs. Once a profile bandwidth has been assigned to a profile, the profile resources can be allocated to RAN slices up until the next profile bandwidth update.

END-TO-END NETWORK SLICING

In this paper, we focused solely on the lower layer issues associated with RAN slicing. Network slicing, however, is an end-to-end activity, encompassing all layers of the stack and affecting all aspects of the network, including the core, the RAN, and management and orchestration. 3GPP Technical Report 30.801 on NR access technology outlines the key principles for supporting network slicing in the RAN in 5G NR. Importantly, it addresses how a RAN slice interacts with both the UE and core network, specifying network slice selection assistance information (NSSAI) for the purpose of uniquely identifying and selecting slices.

For an overview of higher layer issues associated with slicing we refer the reader to [11], which describes a three-layer technology-agnostic architecture for network slicing. The three-layer architecture addresses important slicing problems such slice definition and slice life-cycle management. The

ORCA project outlines its vision for creating end-to-end slices using separate orchestrators for each network segment (such as the RAN) [12].

In the core network, slicing has received plentiful attention with research on enabling techniques such as software defined networking and network function virtualization at a mature stage. In [13], the authors take an end-to-end view of network slicing in 5G, and survey the state-of-the-art. One of the primary design goals of the 5G core (5GC) is to minimise dependencies between the RAN and the core network, permitting the separate management of core and RAN slicing. In [14], it is envisioned that the mapping among devices, RAN slices and core slices may not simply be 1:1:1, but could also be 1:M:N, with a single RAN slice connecting to multiple core slices (and vice-versa).

CONCLUDING REMARKS

The initial specifications for 5G NR reflect the need for slicing in the next generation of telecommunication systems. The diversity of the services to be supported by different RAN slices will be reflected in the diversity of techniques, technologies and configurations adopted by the RAN slices. This diversity, in both services and RAN slices, must be considered when examining how to share the available time-frequency resources among them. Fundamentally, there is a trade-off between adaptability on one hand, and low coexistence on the other. The concept of RAN profiles can be used to balance this trade-off by creating a distinction between service-types and individual services, and hence providing an ability to schedule resources at different time-scales.

ACKNOWLEDGEMENTS

This publication has emanated from research conducted with the financial support of Science Foundation Ireland (SFI) and is co-funded under the European Regional Development Fund under Grant Number 13/RC/2077.

REFERENCES

- [1] L. Zhang, A. Ijaz, P. Xiao, and R. Tafazolli, "Multi-Service System: An Enabler of Flexible 5G Air Interface," *IEEE Communications Magazine*, vol. 55, no. 10, pp. 152–159, Oct. 2017.
- [2] R. Ferrus, O. Sallent, J. Perez-Romero, and R. Agusti, "On 5G Radio Access Network Slicing: Radio Interface Protocol Features and Configuration," *IEEE Communications Magazine*, vol. 56, no. 5, pp. 184–192, May 2018.
- [3] A. Ksentini and N. Nikaiein, "Toward Enforcing Network Slicing on RAN: Flexibility and Resources Abstraction," *IEEE Communications Magazine*, vol. 55, no. 6, pp. 102–108, 2017.
- [4] O. Sallent, J. Perez-Romero, R. Ferrus, and R. Agusti, "On Radio Access Network Slicing from a Radio Resource Management Perspective," *IEEE Wireless Communications*, vol. 24, no. 5, pp. 166–174, Oct. 2017.
- [5] F. Schaich, T. Wild, and R. Ahmed, "Subcarrier Spacing - How to Make Use of This Degree of Freedom," in *2016 IEEE 83rd Vehicular Technology Conference (VTC Spring)*, May 2016, pp. 1–6.
- [6] C. Sexton, Q. Bodinier, A. Farhang *et al.*, "Enabling Asynchronous Machine-Type D2D Communication Using Multiple Waveforms in 5G," *IEEE Internet of Things Journal*, vol. 5, no. 2, pp. 1307–1322, Apr. 2018.

- [7] J. Abdoli, M. Jia, and J. Ma, "Filtered OFDM: A new waveform for future wireless systems," in *2015 IEEE 16th International Workshop on Signal Processing Advances in Wireless Communications (SPAWC)*, Jun. 2015, pp. 66–70.
- [8] L. Zhang, A. Ijaz, P. Xiao *et al.*, "Subband Filtered Multi-Carrier Systems for Multi-Service Wireless Communications," *IEEE Transactions on Wireless Communications*, vol. 16, no. 3, pp. 1893–1907, Mar. 2017.
- [9] "Draft Overall 5G RAN Design," METIS-II, Tech. Rep. Deliverable D2.2, document ICT-671680, Jun. 2016.
- [10] E. Pateromichelakis, J. Gebert, T. Mach *et al.*, "Service-Tailored User-Plane Design Framework and Architecture Considerations in 5G Radio Access Networks," *IEEE Access*, vol. 5, pp. 17 089–17 105, 2017.
- [11] K. Katsalis, N. Nikaiein, E. Schiller *et al.*, "Network Slices toward 5G Communications: Slicing the LTE Network," *IEEE Communications Magazine*, vol. 55, no. 8, pp. 146–154, Aug. 2017.
- [12] J. F. Santos, J. van de Belt, W. Liu *et al.*, "Orchestrating Next-Generation Services Through End-to-End Network Slicing," *Orca Project*, Oct. 2018. [Online]. Available: https://orca-project.eu/wp-content/uploads/sites/4/2018/10/orchestrating_e2e_network_slices_Final.pdf
- [13] X. Foukas, G. Patounas, A. Elmokashfi, and M. K. Marina, "Network Slicing in 5G: Survey and Challenges," *IEEE Communications Magazine*, vol. 55, no. 5, pp. 94–100, May 2017.
- [14] 5G Americas, "Network Slicing for 5G Networks and Services," *5G Americas White Paper*, Nov. 2016. [Online]. Available: http://www.5gamericas.org/files/3214/7975/0104/5G_Americas_Network_Slicing_11.21_Final.pdf

BIOGRAPHY

Conor Sexton is pursuing a Ph.D. on wireless communications at the CONNECT Telecommunications Research Centre, headquartered in Trinity College Dublin. He received his B.A.I. degree in Computer and Electronic Engineering from Trinity College Dublin in 2015. His research focuses on the system level implications of new PHY technologies, and enabling adaptable 5G networks.

Nicola Marchetti is currently Assistant Professor in Wireless Communications at Trinity College Dublin, Ireland. He leads the Wireless Engineering and Complexity Science lab (WhyCOM) under the Irish Research Centre for Future Networks and Communications (CONNECT). His research interests include Adaptive and Self-Organizing Networks, Complex Systems Science for Communication Networks, PHY Layer, Radio Resource Management.

Luiz A. DaSilva holds the chair of Telecommunications at Trinity College Dublin, where he is the Director of CONNECT, the research centre on Future Networks funded by the Science Foundation Ireland. His research focuses on distributed and adaptive resource management in wireless networks, and in particular radio resource sharing and the application of game theory to wireless networks. Prof DaSilva is a Fellow of Trinity College Dublin, an IEEE Communications Society Distinguished Lecturer and a Fellow of the IEEE.