Demo: Service-Oriented Intelligent and Extensible RAN

Robert Schmidt
EURECOM
robert.schmidt@eurecom.fr

Navid Nikaein
EURECOM
navid.nikaein@eurecom.fr

ABSTRACT

Network slicing is considered to be the enabler for a coexistence of a multitude of services on a multi-tenant 5G infrastructure. It is supported through software-defined radio access networking (SD-RAN), bringing programmability to the network in order to enhance performance according to the needs of slice owners. However, SD-RAN so far remained limited to a mere reconfiguration of the base station. In this work, we demonstrate a prototype of a service-oriented RAN on top of the OpenAirInterface and Mosaic5G platforms that brings programmability and extensibility to the RAN with a range of network applications for the purpose of intelligent slicing. We implemented a slice control and management framework, and plug a traffic analysis application that significantly improves the performance of slice users. We observe an improvement of 30 % in application round-trip time with negligible variability for the considered traffic. Further, we demonstrate how to extend control plane functionality from a network store to improve slice performance.

KEYWORDS

5G SD-RAN, Service, Slicing, Intelligence, Functional Extension

ACM Reference Format:
Robert Schmidt and Navid Nikaein. 2020. Demo: Service-Oriented Intelligent and Extensible RAN. In *MobiCom 2020 (MobiCom '20), September 21–25, 2020, London, United Kingdom.* ACM, New York, NY, USA, 3 pages. https://doi.org/10.1145/3372224.3417321

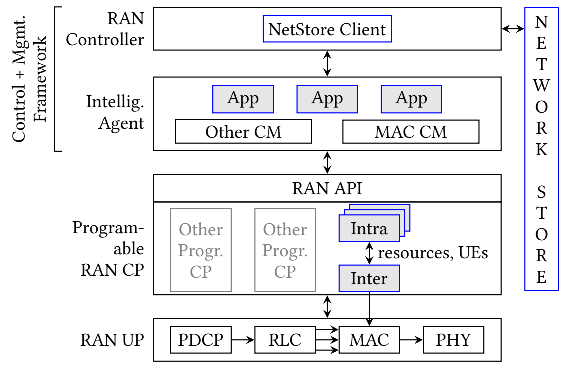

Figure 1: The joint base station-controller platform and a trusted store platform. Blue boxes represent our contributions, and gray boxes mark extensible CP functionality.

1 INTRODUCTION

The 5G mobile communication standard is a paradigm shift beyond the new radio and wider spectrum. It is about the evolution of computing for wireless networks in support of a variety of services and capabilities in diverse usage scenarios, such as low-latency and ultra-reliable communications. As an enabler, the concept of network slicing emerged to overlay a common infrastructure with multiple, logically separated spaces for each service in order to customize the network for those services.

Software-defined radio access networking (SD-RAN) is considered to be a key enabler for RAN slicing. It allows a separation of control and user plane (CP/UP) to program, monitor, control and coordinate multiple network nodes through a central controller. So far, SD-RAN offered programmability through a mere (re-)configuration of the network. In order to enable the service-oriented vision of 5G and beyond, it is paramount to (1) enhance and extend the behavior of the control plane in order to cater for the requirements of each service (similar to the concept of network function virtualization [NFV]) and (2) bring intelligence to the RAN through data-driven control beyond standardized functionality.

The idea of SD-RAN has been explored through FlexRAN [2]. However, FlexRAN does not feature a flexible slice control system, which has only recently been integrated to efficiently manage the base station's radio resources with different slice algorithms [7]. Full slice virtualization has been explored through Orion [1], in which a hypervisor abstracts the RAN on a per-service basis and multiplexes the radio resources of each service onto the common resources. Yet, the customization aspect with respect to slice management and control plane behavior has not been explored. Further, both FlexRAN and Orion did not consider an integration of a data-driven management approach for slice control.

In this demo, we build a prototype on top of the Mosaic5G platform [4] to show how (1) the slice control and management system, together with an intelligent control application, allows to significantly improve the network performance perceived by UEs, and how (2) the control plane behavior can be customized towards service needs through control plane programs that are dynamically fetched from a store.

2 PROTOTYPE OVERVIEW

Our prototype enhances the FlexRAN [2] controller as shown in Figure 1. First, we extend the agent through an application layer that uses the RAN API to gather information about the underlying system and trigger control plane actions through the control modules. This allows arbitrary applications within the agent for platform-independent control (say, of different LTE implementations) instead of FlexRAN's original virtual control functions that

MobiCom '20, September 21–25, 2020, London, United Kingdom

ACM ISBN 978-1-4503-7085-1/20/09…$15.00
https://doi.org/10.1145/3372224.3417321

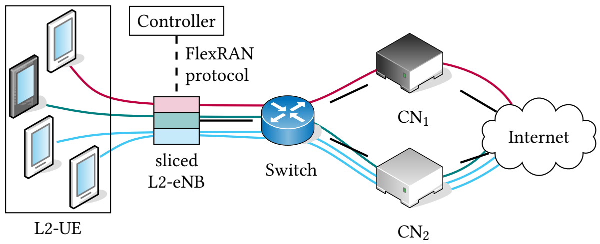

Figure 2: Demo setup using an emulated LTE network (L2-UE, L2-eNB) and two core networks.

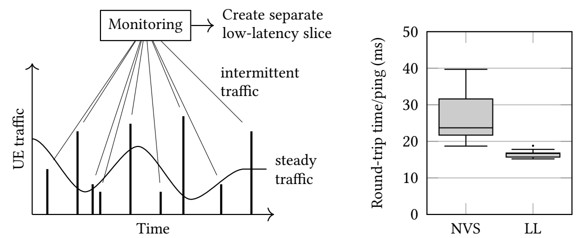

(a) Principle of recognizing traffic patterns of LL users.

(b) RTT of 1 UE in default (NVS) or dedicated slice (LL).

Figure 3: Data-driven slice configuration.

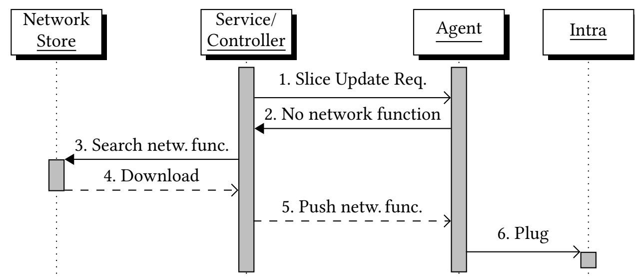

Figure 4: Sequence diagram for CP extension functionality.

are bound to a particular control module. Further, network applications could be plugged and run at the agent to offload the controller from control decisions.

To enable the dynamic slice management within the base station's MAC scheduling functionality, we decoupled the MAC scheduling [7] into (i) an inter-slice scheduler for resource slicing, and (ii) a per-service intra-slice scheduler for user scheduling. Both run within the control plane of the base station. In every slot, the inter-slice scheduler partitions the resources eligible for each slice according to the currently loaded slice scheduling algorithm. It triggers the intra-slice schedulers with a bitmap indicating the allowed RBs, and a list of UEs active in the respective slice, including per-UE information such as channel information.

Both the inter-slice and intra-slice schedulers are exchangeable stateless control plane functions following cloud-native principles. Such a stateless implementation allows the prototype to dynamically and independently manage the life cycle of both the slice and the user scheduler as per service needs by loading platform-/implementation-dependent control functions directly into the control plane of the base station. This enables programmability and extensibility of both the resource slicing and user scheduling problems within the MAC to replace and update the processing behavior.

Finally, the prototype includes a trusted distribution platform providing network applications and control plane functions and allowing to extend and customize the behavior of the base station.

3 DEMO DESCRIPTION

The considered demo scenario is shown in Figure 2. It consists of three hosts: the first contains an LTE network in the form of the "L2 simulator" of the modified OpenAirInterface [5] (unlike the name suggests, it *emulates* an LTE eNB and multiple UEs) as well as the enhanced FlexRAN controller. The eNB and UEs run the full LTE stack except the PHY layer, which allows a larger number of UEs for scalability purposes, and provides a more stable setup. The setup has been verified with commercial UEs connected over radio. Two hosts are used for core networks for end-to-end network slicing.

3.1 Intelligence

New UEs entering the base station need to be associated to a certain slice. Typically, UEs deliver hints on preferred slices, such as the NSSAI in 5G or simply the IMSI or selected PLMN in LTE. To illustrate the concept, we will deploy an agent application that auto-associates UEs to slices based on the IMSI or selected PLMN to offload the controller from control decisions and reduce network load. The association will be based on a prefix-match (e.g., IMSIs starting on a pattern) to choose the correct slice, which assumes that every UE actually sends its IMSI.

The demo will use the NVS slicing algorithm [3] to slice the radio resources for two slices. Due to the specificities of this algorithm, scheduling latencies might be incurred, negatively impacting the round trip time (RTT). We will use a data-driven slicing approach to reduce the RTTs of low latency (LL) users through an application at the agent that analyzes the traffic pattern of every UE, as shown in Figure 3a. If a certain UE has only intermittent, bursty traffic (we use ping with a data size of 200 B and 0.25 s interval time to simulate such traffic), it will be considered "latency-critical" traffic. Such traffic should be handled in a separate LL slice: the application loads a specific LL algorithm [6] and associates the respective UE to their individual LL slice. As can be seen in Figure 3b, the RTT is significantly lowered (the average by 30%), and the variability becomes negligible. As soon as the traffic pattern does not match, the corresponding slice is removed.

3.2 Behavior Extension through Network Store

In the second part of the demonstration, we show how the control plane can be enhanced and extended with new behavior by plugging network functions on-demand from the trusted network store. As shown in Figure 4, upon a slice update, the agent requests a specific network function, e.g., in the demo the intra-slice proportional fair scheduling, that is needed to meet slice owner requirements. The controller therefore checks on the store for a network function, which is then plugged by the agent into the RAN CP. Once the slice is decommissioned, the network function will be disposed.

ACKNOWLEDGMENTS

This work received funding from the European Union's Horizon 2020 Framework Programme under Grant No. 857201 (5G-VICTORI).

REFERENCES

[1] Xenofon Foukas, Mahesh K. Marina, and Kimon Kontovasilis. 2017. Orion: RAN Slicing for a Flexible and Cost-Effective Multi-Service Mobile Network Architecture. In *Proceedings of the 23rd Annual International Conference on Mobile Computing and Networking (MobiCom '17)*. ACM, New York, NY, USA, 127–140. https://doi.org/10.1145/3117811.3117831

[2] Xenofon Foukas, Navid Nikaein, Mohamed M. Kassem, and Kimon Kontovasilis. 2016. FlexRAN: A flexible and programmable platform for software-defined radio access networks. In *CONEXT 2016, 12th International on Conference on Emerging Networking Experiments and Technologies*. Irvine, California, United States. https://doi.org/10.1145/2999572.2999599

[3] R. Kokku, R. Mahindra, H. Zhang, and S. Rangarajan. 2012. NVS: A Substrate for Virtualizing Wireless Resources in Cellular Networks. *IEEE/ACM Transactions on Networking* 20, 5 (Oct 2012), 1333–1346. https://doi.org/10.1109/TNET.2011.2179063

[4] Navid Nikaein, Chia-Yu Chang, and Konstantinos Alexandris. 2018. Mosaic5G: Agile and Flexible Service Platforms for 5G Research. *SIGCOMM Comput. Commun. Rev.* 48, 3 (Sept. 2018), 29–34. https://doi.org/10.1145/3276799.3276803

[5] Navid Nikaein, Mahesh K. Marina, Saravana Manickam, Alex Dawson, Raymond Knopp, and Christian Bonnet. 2014. OpenAirInterface: A Flexible Platform for 5G Research. *SIGCOMM Comput. Commun. Rev.* 44, 5 (Oct. 2014), 33–38. https://doi.org/10.1145/2677046.2677053

[6] Robert Schmidt, Chia-Yu Chang, and Navid Nikaein. 2019. Slice Scheduling with QoS-Guarantee Towards 5G. In *IEEE Global Communications Conference (Globecom) 2019*. https://doi.org/10.1109/GLOBECOM38437.2019.9013258

[7] Robert Schmidt and Navid Nikaein. 2020. Demo: Efficient Multi-Service RAN Slice Management and Orchestration. In *NOMS 2020 - 2020 IEEE/IFIP Network Operations and Management Symposium*. 1–2.