

Integrating Cloud-RAN with Packet Core as VNF using Open Source MANO and OpenAirInterface

Andres F. Ocampo^{*†}, Thomas Dreibholz^{*}, Mah-Rukh Fida^{*}, Ahmed Elmokashfi^{*}, Haakon Bryhni^{*}

^{*} SimulaMet – Simula Metropolitan Center for Digital Engineering

Center for Resilient Networks and Applications

Pilestredet 52, N-0167 Oslo, Norway

[†] OsloMet – Oslo Metropolitan University

Pilestredet 52, N-0167 Oslo, Norway

Email: {andres, dreibh}@simula.no

Abstract—The Cloud-based Radio Access Network (Cloud-RAN) architecture and Network Function Virtualization (NFV) are key enablers to building future mobile networks in a flexible and cost-efficient way. With early deployments of the fifth generation of mobile technologies – 5G – around the world, setting up 4G/5G experimental infrastructures is necessary to optimally design 5G networks. In this demo, we present a custom small-scale 4G/5G testbed based on OpenAirInterface and Open Source MANO. The testbed integrates a Cloud-RAN based on switched Ethernet Xhaul and functional splitting, with an Evolved Packet Core (EPC) deployed as a Virtual Network Function (VNF) in a cloud infrastructure. Using OPEN SOURCE MANO, this demo shows the administration and monitoring of the EPC VNF components. Moreover, as proof of concept, collection and visualization of telemetry will be shown for two smartphones connected to the network through the Cloud-RAN.^{1,2}

Index Terms—Cloud Radio Access Network (Cloud-RAN), Ethernet Xhaul, Functional Splits, Network Function Virtualisation (NFV), Open Source MANO (OSM)

I. INTRODUCTION

The fifth generation of mobile networks – 5G – aims to cater for a wide spectrum of services with diverse requirements. Besides the standard (e.g., 5G New Radio) specifying the Radio Access Technology (RAT), flexible and cost-efficient mobile networks design is needed to realize this vision. For instance, a way to meet the cost requirements for 5G, is to design the Radio Access Network following a centralized cloud-based architecture (Cloud-RAN), leveraging virtualization and shared computing. Cloud-RAN allows to host multiple *Base Band Units* (BBU) on top of the same physical server by instantiating its protocol functions as software application (vBBU). That way, the RAN can be engineered to meet peak network capacity requirements instead of individual cell requirements.

Not only can the vBBUs be instantiated on top of general-purpose servers (GPP), but also the *Evolved Packet Core* (EPC) can be virtualized and instantiated as isolated processes in remote data centers. Such a virtualization of the mobile network functionality allows using NFV technology to manage

and expand network capabilities. In fact, NFV features pave the road for developing network slicing.

Building 4G/5G experimental setups that integrate the so-called 5G key enablers, is essential to test novel solutions and designs for future mobile networks. In this demo, we present a custom small-scale 4G/5G testbed based on the open source LTE implementation from OPENAIRINTERFACE. This testbed integrates a Cloud-RAN architecture with an EPC deployed as a VNF. In the Cloud-RAN part, two vBBUs with different functional splits are multiplexed using an Ethernet-switched mobile transport network. The EPC is instantiated as a VNF by using OPEN SOURCE MANO along with a cloud infrastructure like OPENSTACK. Deploying the EPC as a VNF provides a building block of a *Network Service* (NS). This feature also allows attaching further functionalities into the NS, like e.g., *Mobile Edge Computing* (MEC) services.

II. BACKGROUND

A. The Cloud-RAN architecture

Despite its appeal as a key enabler for 5G, the Cloud-RAN architecture faces critical latency constraints and capacity requirements. Deploying vBBUs in a centralized server far away from antennas imposes stringent latency constraints to the transport network (referred to as Fronthaul network) connecting the antennas and the vBBU. Although a fiber solution meets the requirements for the Fronthaul, deploying dedicated fibers per cell is extremely expensive. One way to overcome this problem is to use a switched Ethernet-based Fronthaul. Switched Ethernet Fronthaul can enable multiplexing of several vBBUs into the same network [1]. However, Ethernet-switched networks bring latency, latency variation, and time-frequency synchronization issues that can be harmful for mobile services [2].

To ease the strict bandwidth and latency requirements on the Fronthaul, the 3GPP proposed a set of functional splits of the BBU processing functions [3]. Processing part of the functions in a *Distributed Unit* (DU) close to the antennas, the central *Compute Unit* (CU) processes the remainder functions. The more functions instantiated in the DU, the more delay would be afforded in the Fronthaul network. Introduction of both the functional split units of DU and CU

¹The authors would like to thank Ann Edith Wulff Armitstead for her friendly support.

²The authors would like to thank the anonymous reviewers for their helpful comments.

redefines the mobile transport network segments and their requirements in terms of latency and capacity: the Fronthaul is the segment between the Remote Radio Head (i.e., which performs analog/digital conversion of the signal) and the DU, which transports time-domain quadrature samples; the Midhaul is the segment connecting the DU and the CU, with data rate and requirements depending upon the chosen functional split; and the Backhaul, which connects the Cloud-RAN with the Core Network. These transport segments are referred to as mobile crosshaul (Xhaul).

B. VNF based EPC

We are developing the SIMULAMET EPC [4]–[6] VNF, which is an open source VNF for an EPC based on OPENAIRINTERFACE and OPEN SOURCE MANO³ (OSM). OSM is an open source *Management and Orchestration* (MANO) framework for NFV. NFV with OSM can be split into three parts [7, Chapter 1]:

- 1) *NFV Infrastructure* (NFVI) denotes hosting virtual machines and containers, as well as connecting them by *Virtual Links* (VLs). OSM is independent of an underlying hosting mechanism like e.g. OPENSTACK [8], [9].
- 2) Collection of VNFs and their interconnection and composition into a NS. NSs can further be composed and shared to form network slices.
- 3) The MANO controls the life-cycle of VNFs, NSs and network slices, including configuration and monitoring.

Realising and deploying a VNF as part of a NS is explained in detail in [4], [10]. In short, a VNF is created in form of a *VNF Descriptor* (VNFD). The VNFD contains a definition of its *Virtual Deployment Units* (VDU), each corresponding to its own *Virtual Machine* (VM). *Connection Points* (CP) define interfaces. CPs of a VDU are network interfaces in a VM instance. CPs can be connected by *Virtual Links* (VL). VDUs of the VNF are connected by using internal CPs. This VL is referenced by its VL Descriptor (VLD), which can be a name like e.g. “S11”. CPs of VDUs can also be connected to external CPs of the VNF. Finally, A *NS Descriptor* (NSD) connects VNFs with VLs. Furthermore, it has the possibility to also define external NS CPs, e.g. to attach them to physical networks in the underlying NFVI.

III. TESTBED SETUP

We consider the network scenario depicted in Figure 1. This network scenario is composed of two fully virtualized LTE-BBUs, called vBBU1 and vBBU2, respectively. vBBU1 adopts the functional split 7.1, which performs low physical layer functions (L1) locally at the DU. The CU instantiates the remainder L1 functions, as well as layer two (L2) and layer three (L3) functions. On the other hand, vBBU2 adopts the functional split 2. In this case, while the DU instantiates both L1 and L2 functions, the CU instantiates L3 functions. Furthermore, we assume that both of these vBBUs belong

to the same network operator, meaning that both vBBUs are under the domain of the same EPC.

The EPC which controls both the vBBU1 and the vBBU2 in the Cloud-RAN, as depicted in Figure 1, is deployed using the SIMULAMET EPC VNF. The SimulaMet EPC VNF is instantiated through a NFVI in form of OPENSTACK. OSM is the orchestration framework, which uses JUJU Charms for configuration (i.e. there is one Charm container assigned for the management of each of the VDU VMs).

A particular property of the SIMULAMET EPC VNF is that the software for each EPC component (i.e. HSS, MME, SPGW-C and SPGW-U), during instantiation, is built from scratch using the component sources of given Git repositories and commits. This enables the usage of tailor-made EPCs, making the usage of customised software versions, e.g. to include own bug fixes or experimental extensions, easy. All of the SIMULAMET EPC VNFD sources, together with the example NSD, build tool-chain and test scripts, operating system image build script as well as documentation is available as open source in the form of a public GIT repository under <https://github.com/simula/5gvinni-oai-ns> [4], [5].

IV. MAIN FEATURES, COMPONENTS AND OPERATIONS

This demo shows the operation of a mobile system which integrates a Cloud-RAN based on functional split and switched Ethernet Xhaul, with an EPC deployed as a VNF using OSM. As proof of concept, two phones will be connected to the network through a different vBBU each. The main features and components implemented in our testbed are described below:

1) *RRH*: Both RRHs are deployed using the ETTUS B210 *Universal Software Radio Peripheral* (USRP) platform. These RRHs involves one antenna port in a *Single Input Single Output* (SISO) antenna array configuration.

2) *DU*: Both DUs use INTEL NUC7i7BNB microcomputers. These systems are equipped with two Intel Core i7-7567U processors, 32 GiB of memory, and an integrated Intel *Network Interface Card* (NIC) I219-V at 1 Gbit/s. Through this NIC, the DU connects its hosted vBBU to the Midhaul network. In addition, the host OS for each of these DUs is Ubuntu 16.04 with low-latency Linux kernel version 4.19.58, i.e., real-time (RT) kernel. On the other hand, each vBBU is deployed on top of Linux Containers (LXC) at the DU, adopting the Linux RT-kernel libraries for real-time guarantees.

3) *CU*: The CU is a GPP equipped with twelve INTEL i7-8750H processors at 2.20 GHz, and 32 GiB system memory. While the host OS is Ubuntu 18.04 with low-latency Linux kernel version 5.3.28, i.e., RT kernel, both of the vBBUs are deployed on top of LXC. To provide real-time guarantees, these LXCs adopt the Linux RT-kernel libraries. Moreover, the CU uses the SUPERMICRO AOC-SG-i2 Gigabit Ethernet adapter, equipped with two INTEL 82575 Gigabit Ethernet ports (e.g., *eno1* and *eno2*), as physical NIC. These two ports implement the SR-IOV mechanism for NIC sharing among LXCs. More specifically, two *Virtual Functions* (e.g., VF1 and VF2) are created on top of the *Physical Function* (PF)

³OPEN SOURCE MANO: <https://osm.etsi.org>.

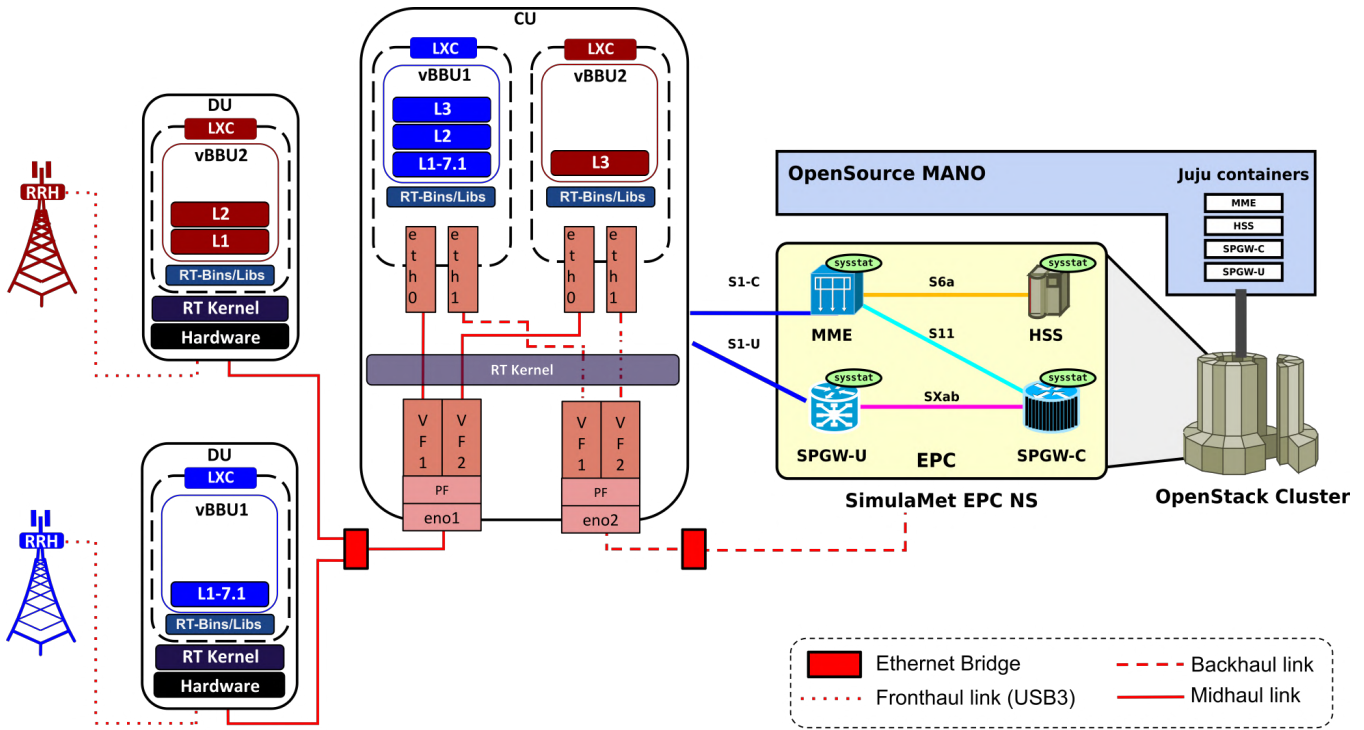


Fig. 1. Testbed setup: a CloudRAN with two functionally split vBBUs, switched Ethernet Xhaul, and a SIMULAMET EPC with OPENSTACK as NFVI

of each port. These VFs allow to pass-through the LXC to the Midhaul and to the Backhaul, respectively. As shown in Figure 1, each LXC hosting a vBBU deploys two virtual interfaces. While the interface eth0 connects the vBBU to eno1 – Midhaul network, the interfaces eth1 connects the vBBU to eno2 – Backhaul network.

4) *Xhaul*: For the Fronthaul link connecting the RRH with the DU, the B210 provides a fast SuperSpeed USB 3.0 connectivity with bus-power at 5.0 Gbit/s. Both, the Midhaul network and the Backhaul network, use the JUNIPER EX4200 Ethernet switch, with physical interfaces at 1 Gbit/s.

5) *vBBUs*: Based on their functional split, the vBBUs deploy the eNodeB-LTE implementation from OPENAIRINTERFACE. Particularly, in the case of vBBU1, OPENAIRINTERFACE implements split 7.1 using the NGFI-IF4p5 interface specification [11]. In the case of vBBU2, OPENAIRINTERFACE implements split 2 using the F1 Application Protocol (F1AP) [12], [13]. The wireless transmission setup for both vBBUs, consists of 25 Physical Resource Blocks (PRB), which provides 5 MHz bandwidth in band 7.

6) *UE*: A single user equipment (UE) is connected to each vBBU. In both the cases, ONEPLUS-5 phones are used as UEs.

7) *The EPC*: For our demo, we are going to instantiate the EPC by using SIMULAMET EPC VNF (as described in Subsection II-B). The EPC will be parametrised with the settings for our mobile network parameters, i.e. realm, *Mobile Country Code* (MCC), *Mobile Network Code* (MNC), *Operator Code* (OP), *Subscriber Key* (K), as well as the users' *International Mobile Subscriber Identities* (IMSI)

and *Mobile Station International Subscriber Directory Numbers* (MSISDN).

V. LIVE DEMONSTRATION

With the live demonstration, we show how the different components in the Cloud-RAN architecture can be integrated into the same network infrastructure. Moreover, this demonstration highlights the benefits of this deployment as an experimental facility for testing and experimenting, which is relevant to test and put in place the best solutions to design future mobile networks. In the live demonstration, the features described in the following subsections will be shown:

A. Administration and monitoring of the VNF parts with Open Source MANO

From the perspective of OSM, we would like to highlight its NS and VNF management features. Particularly, we would also like to show to the audience how we solved the challenges of managing the complex SIMULAMET EPC VNF [4], [5], with some details on the build procedures, configuration options, and the management of the day-0/day-1 configuration of its instances in OSM.

OSM furthermore provides telemetry [6] for the instantiated NSs and VNFs, based on the collection of standard metrics (CPU utilisation, network statistics, etc.) from the underlying NFVI (in our case: OPENSTACK with CEILOMETER [8], [9]) and the collection of customised metrics with JUJU directly from the VNF instances [6]. The collected metrics are stored in and can be queried from PROMETHEUS [5]. Furthermore,

it is possible to display the metrics with various visualization features by using GRAFANA.

B. End-to-end network performance

For the Cloud-RAN architecture, a main challenge is to deal with the impact on the end-to-end performance when multiple vBBUs share both, network and compute resources. In this specific case, two vBBUs with different functional splits share network resources in the Xhaul network, where their mobile traffic is aggregated using a switched Ethernet Xhaul. Moreover, both vBBUs share compute resources (e.g., NIC, memory, CPU time) at the CU. To demonstrate the feasibility of this architecture, we perform the following experiments during the live demonstration:

1) *Demonstrating traffic aggregation:* To show the feasibility of aggregating mobile traffic in the Xhaul network, we stream a pre-recorded video from a video server deployed at the same edge computer hosting the VNF-EPC. The two phones are served by a different vBBU and both are connected to the video streaming. Performance metrics are shown at the video server.

2) *Measuring the RTT:* This test consists of sending ICMP packets from the UE to a server in the Internet, e.g. using the HiPERCONTRACER [14] tool. Measured at the UE, RTT statistics from those packets provide insight into the end-to-end network latency. This test allows comparing the end-to-end network latency obtained by each vBBU using different functional splits.

3) *Measuring end-to-end network performance:* Using the NETPERFMETER [15]–[17] tools, we benchmark server-UE (client) data transmissions in terms of throughput (i.e., received bitrate), jitter, and packet loss. For the network scenario illustrated in Figure 1, the edge computer hosting the VNF-EPC runs the IPERF3 or NETPERFMETER server, whereas the UE runs the IPERF3 or NETPERFMETER client. The test consists of sending downstream UDP flows from the server to the UE at different bitrates.

VI. CONCLUSIONS

Several studies have been conducted demonstrating the benefits of Cloud-RAN and NFV as key enablers for future mobile networks. However, most research has been focused on developing the individual components within this architecture. Integrating those components into the same network infrastructure is needed to design and build future mobile networks. In this paper, we present a small-scale 4G/5G testbed, which integrates the following elements into the same mobile network:

- 1) Cloud-RAN where vBBUs are deployed using different functional splits;
- 2) switched Ethernet Xhaul allowing the aggregation of vBBUs into the same mobile transport network;
- 3) an EPC deployed as a VNF using OSM.

The deployment presented in this paper constitutes an experimental facility, which aims to help on finding and place best solutions for future mobile networks. To show the

benefits of this deployment, the live demonstration performs experiments to measure end-to-end network performance, as well as demonstrates the management of telemetry data from the EPC.

REFERENCES

- [1] N. J. Gomes, P. Sehier, H. Thomas, P. Chanclou, B. Li, D. Munch, P. Assimakopoulos, S. Dixit, and V. Jungnickel, “Boosting 5G through Ethernet: How evolved Fronthaul can take Next-Generation Mobile to the Next Level,” *IEEE Vehicular Technology Magazine*, vol. 13, no. 1, pp. 74–84, 2018.
- [2] N. J. Gomes, P. Chanclou, P. Turnbull, A. Magee, and V. Jungnickel, “Fronthaul Evolution: from CPRI to Ethernet,” *Optical Fiber Technology*, vol. 26, pp. 50–58, 2015.
- [3] L. Larsen, A. Checko, and H. L. Christiansen, “A Survey of the Functional Splits proposed for 5G Mobile Crosshaul Networks,” *IEEE Communications Surveys & Tutorials*, vol. 21, no. 1, pp. 146–172, 2018.
- [4] T. Dreiholz, “Flexible 4G/5G Testbed Setup for Mobile Edge Computing using OpenAirInterface and Open Source MANO,” in *Proceedings of the 2nd International Workshop on Recent Advances for Multi-Clouds and Mobile Edge Computing (M2EC) in conjunction with the 34th International Conference on Advanced Information Networking and Applications (AINA)*, Apr. 2020, pp. 1143–1153.
- [5] —, “A 4G/5G Packet Core as VNF with Open Source MANO and OpenAirInterface,” in *Proceedings of the 28th IEEE International Conference on Software, Telecommunications and Computer Networks (SoftCOM)*, Hvar, Dalmacija/Croatia, Sep. 2020.
- [6] M. Xie, J. S. Pujol-Roig, F. Michelinakis, T. Dreiholz, C. Guerrero, A. G. Sánchez, W. Y. Poe, Y. Wang, and A. M. Elmokashfi, “AI-Driven Closed-Loop Service Assurance with Service Exposures,” in *Proceedings of the 29th IEEE European Conference on Networks and Communications (EuCNC)*, Dubrovnik, Dubrovnik-Neretva/Croatia, Jun. 2020.
- [7] A. Reid, A. González, A. E. Armengol, G. G. de Blas, M. Xie, P. Grønsund, P. Willis, P. Eardley, and F.-J. R. Salguero, “OSM Scope, Functionality, Operation and Integration Guidelines,” ETSI, White Paper, Jun. 2019. [Online]. Available: https://osm.etsi.org/images/OSM_EUAG_White_Paper_OSM_Scope_and_Functionality.pdf
- [8] OpenStack, *OpenStack Architecture Design Guide*, Dec. 2019.
- [9] —, *OpenStack Operations Guide*, Dec. 2019.
- [10] G. Lavado, “OSM VNF Onboarding Guidelines,” ETSI, White Paper, Jun. 2019.
- [11] R. Knopp, “Overview of Functional Splits in OAI,” Feb. 2020. [Online]. Available: https://www.openairinterface.org/docs/workshop/5_OAI_Workshop_20180620/KNOPP_OAI-functional-splits.pdf
- [12] 3GPP, “5G; NG-RAN; F1 Application Protocol (F1AP) (3GPP TS 38.473 version 15.2.1 Release 15),” May 2020. [Online]. Available: https://www.etsi.org/deliver/etsi_ts/138400_138499/138473/15.02.01_60/ts_138473v150201p.pdf
- [13] OAI, “F1 Interface,” Jul. 2020. [Online]. Available: <https://gitlab.eurecom.fr/oai/openairinterface5g/wikis/f1-interface>
- [14] T. Dreiholz, “HiPerConTracer - A Versatile Tool for IP Connectivity Tracing in Multi-Path Setups,” in *Proceedings of the 28th IEEE International Conference on Software, Telecommunications and Computer Networks (SoftCOM)*, Hvar, Dalmacija/Croatia, Sep. 2020.
- [15] —, “Evaluation and Optimisation of Multi-Path Transport using the Stream Control Transmission Protocol,” Habilitation Treatise, University of Duisburg-Essen, Faculty of Economics, Institute for Computer Science and Business Information Systems, Mar. 2012.
- [16] —, “NetPerfMeter: A Network Performance Metering Tool,” *Multi-path TCP Blog*, Sep. 2015.
- [17] T. Dreiholz, M. Becke, H. Adhari, and E. P. Rathgeb, “Evaluation of A New Multipath Congestion Control Scheme using the NetPerfMeter Tool-Chain,” in *Proceedings of the 19th IEEE International Conference on Software, Telecommunications and Computer Networks (SoftCOM)*, Hvar, Dalmacija/Croatia, Sep. 2011, pp. 1–6.