

Experimental Evaluation of Functional Splits for 5G Cloud-RANs

Nikos Makris^{†*}, Pavlos Basaras^{†*}, Thanasis Korakis^{†*}, Navid Nikaiein[◇] and Leandros Tassioulas[‡]

[†]Department of Electrical and Computer Engineering, University of Thessaly, Greece

^{*}Centre for Research and Technology Hellas, CERTH, Greece

[◇]Eurecom, Sophia Antipolis, France

[‡]Department of Electrical Engineering, Yale University, New Haven, USA

Email: {nimakris, pabasara, korakis}@uth.gr, navid.nikaiein@eurecom.fr, leandros.tassioulas@yale.edu

Abstract—Centralized RAN processing has been identified as one of the major enablers for 5G mobile network access. By moving the baseband units (BBU) to the Cloud, multiple instances can be instantiated on the fly, serving several Remote Radio Head (RRH) units. The goal is to satisfy the existing demand of particular geographical areas, whereas drastically reducing the overall CAPEX and OPEX costs of the mobile operators. In this work, we present an experimental study of real Cloud-RAN deployments, with respect to different functional splits. We use as a reference architecture the 3GPP LTE stack, and argue about the functional split applicability in contemporary networks. We evaluate Layer 2 functional splits, that can be used for the convergence of multiple heterogeneous wireless technologies in an all-in-one unit. By deploying our approach in a real testbed setup, we extract the backhaul network transfer requirements for the different splits and present our experimental findings, compared with the respective simulation results.

I. INTRODUCTION

The concept of centralizing the Radio Access Networks (RANs) has been identified as one of the keen aspects to be adopted by the 5th Generation (5G) mobile networks. RAN centralization lies in the simple concept of decoupling the base band processing from the actual radio interface, and providing it as a Cloud Computing service. Such Cloud services can enable many real-time operations that are currently not viable in the “full stack” base stations, and can serve as the main enablers for the virtualization of the wireless network, allowing multi-tenancy over the actual same hardware resources.

Cloud-RAN lies on separating the Base Band Units (BBUs) from the actual RF front-end. Existing base station deployments are using a more decentralized approach, where the base station units process a very large part of the networking stack. This fact can pose a big obstacle in the base station syncing, especially for processes like the enhanced Inter Cell Interference Control (eICIC) for LTE-Advanced. Yet, they can all be dealt with when the BBU processes for different RANs are closely located in the Cloud.

Furthermore, the disaggregation of the functions from a single unit to the BBU and a Remote Radio Unit (RRU) can fit extremely the 5G concepts of network virtualization and higher delivered capacity per user in a single geographical area [1]. The RRU can be considered as either a passive element, with the sole purpose to transmit low level data over the air (Remote Radio Head - RRH) or a more intelligent unit, where

part of the processing takes place (e.g. the entire PHY layer or parts of it). Based on the demand and an existing pool of BBUs in the cloud, the serving RRUs covering a single area can be instantiated on the fly. Moreover, since network slicing and virtualization can be handled beyond the RRU, as its sole purpose is only to transmit low level raw data (ideally raw IQ samples for the RRH case), multiple operators can take advantage of the very same physical equipment. Network virtualization can take place as a virtual function in the Cloud, thus enabling multiple tenants (mobile network operators) to take advantage of the same network equipment.

For the implementation of Cloud-RAN architectures high bandwidth connections are needed from the RRU to the BBU. Depending on the point where the split takes place, the transfer requirements of the network may vary; this highly depends on the back/front-hauling technology as well. Employing an IP based scheme can induce delays for the processing and packetization of data, as well as the maximum number of served RRU units [2]. Deduced from these requirements, very stringent delay times need to be met for the realization of the splits. Moreover, interoperability of the Cloud-RAN with the existing network infrastructure is highly desired as well. This essentially means that existing fiber based infrastructure and protocols (e.g. CPRI) or copper links can be exploited for the realization of Cloud-RANs. Copper links have come to the fore due to their high availability, as well as flexibility of protocols that can be executed over a packetized data plane. This is also reflected in the standardization activities of IEEE P1914.3 for enabling Radio over Ethernet communication, mainly for the Cloud-RAN applications.

Towards addressing this emerging research challenge, the Next Generation Fronthaul Interface (NGFI) has identified in [3] the possible splits for the BBU/RRU functions. Both high- and low-level splits have been identified, with the LTE architecture as a reference. We use the following terminology, as it is broadly used in the literature: any functional split inside the PHY layer that has very stringent delay requirements requires a fronthaul (FH) interface, whereas a backhaul (BH) is needed elsewhere.

This article’s main contributions are the following:

- To extract the real-time transfer requirements for a 5G Cloud-RAN in the BH.

- To implement and evaluate different functional splits over the LTE networking stack, complying with NGFI.
- To experimentally evaluate different transport protocols for the aforementioned splits (UDP/TCP/SCTP).

The splits that we evaluate take place at two different points of Layer 2 of LTE stack; 1) PDCP/RLC and 2) MAC/PHY. We employ the open source platform OpenAirInterface [4] for the realization of the splits and evaluate our solutions in a real environment, when using a 1Gbps Ethernet link for our BH.

The rest of the paper is organized as follows. Section II is providing an overview of any previous related work and our motivation. Section III is discussing our choice for the functional splits, as well as the pros and cons of each solution. Section IV is presenting our contributions and experimental setup, whereas in Section V we showcase our experimental findings. Finally, Section VI concludes our work and presents some future directions.

II. MOTIVATION AND RELATED WORK

Cloud-RAN has been identified as one of the key 5G enablers. Next Generation Mobile Networks (NGMN) alliance has pinpointed the advantages, as well as potential interfaces for facilitating the functional splits in [5]. The advantages of employing a centralized processing unit, located in the cloud has been also described in [1]. The authors argue on the Cloud-RAN applicability for 5G schemes, as well as analyze the transfer requirements for the fronthaul (FH), when the functional splits take place at different points of the PHY layer of LTE. Moreover, in [6], the different technologies that are available for realizing the Cloud-RAN architectures are illustrated. Potential splits are identified along with the technologies employed for the data transportation to the Cloud.

Similarly, authors in [7] detail the requirements for the FH network, with respect to low level splits. An analysis of the potential technologies used for FH and BH of 5G networks, based on these specific requirements for PHY layer functional splits is presented in [8].

A study resembling our contributions is presented in [9]. The authors identify high-layer splits for BH, as well as low-layer for FH and extract the transfer requirements for the network. Yet, the work relies on simulation based models for the network setup, while in all of the aforementioned cases the authors do not consider the existing legacy networks as potential technologies for transferring the data to the Cloud.

In order to use existing packet based networks, instead of circuit based fiber connections (e.g. CPRI), the extra delays of packet encapsulation, decapsulation and processing have to be taken into consideration. Authors in [2] and [10] analyze and model these requirements using IP based networks for PHY layer splits.

Yet, experimentally driven results are very scarce regarding the Cloud-RAN modeling. Authors in [11] present a platform where the RRU is composed of all the LTE PHY layer functions, whereas the rest of the eNB processing is taking place as a separate process executed in the Cloud, based on the OpenAirInterface platform. Following up this work,

authors in [12] present through real experiments the delay that is incurred when the BBU is operating inside different virtualization environments (e.g. KVM, LXC, etc.). Similarly, authors in [13] present their own platform for Cloud-RANs.

In this work, we present our contributions to the same Open Source platform for LTE, in which we implement functional splits at two different layers. Our work differentiates from similar former studies in the fact that it is, to the best of our knowledge, the first experimentally driven work on characterizing this type of splits over an IP based BH network. By employing both simulation and real testbed experiments, we evaluate and extract the real time requirements for the operation of such a Cloud-RAN architecture. We use an IP-based BH, and measure the limitations induced by the packetization and processing of different protocols used for transferring the data. We use two approaches for the transportation of data: 1) based on stateless protocols (UDP), for the PHY layer splits, as they are more delay sensitive regarding the scheduling of the transmissions and 2) state-ful protocols (TCP/SCTP) for higher layer splits, as they can operate with more slack delay requirements, if proper buffering of the data is employed.

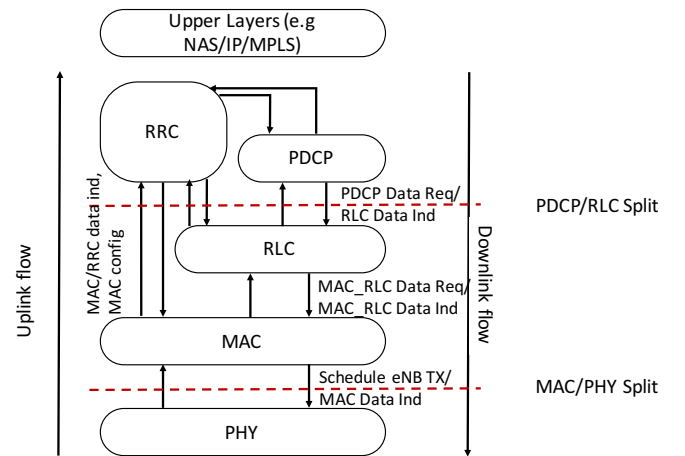


Fig. 1. LTE Reference Architecture and Identified Functional Splits

III. FUNCTIONAL SPLIT ARCHITECTURE

We employ the LTE protocol stack as our reference architecture in order to identify our functional splits and conduct our experiments. In this section we provide a brief overview on the functions of each layer in a bottom-up manner, and the potential of each split when deployed in real systems.

PHY layer is dedicated for the transmission and reception of control and user data over the air. This may include functions such as FEC, encoding/decoding, equalization, FFT and finally the D/A or A/D conversion. Functional splits can be identified at different points of the PHY layer, used mainly for fronthauling the LTE network. MAC layer is endowed with the scheduling processes and allocating resources for the served UEs in the network. Once a stream is scheduled for transmission in a specific subframe in the MAC layer, it is delivered to PHY. RLC is a sublayer used to transfer the higher

layer PDUs to MAC SDUs, by concatenating/segmenting them and reassembling them. PDCP is used as the interface with the IP based networks, used to do packet compression and removing the IP header before giving the packets to lower layers for scheduling their transmission over the air.

Yet, although the splits other than the ones dedicated inside the PHY layer may yield only small performance benefits for Cloud-RAN applications [1], they can be the enablers for novel applications for 5G. The splits that are dealt with in this work are the following:

- **PDCP/RLC split:** Splits over the MAC layer seem to be yielding only small performance benefits for 5G, as they could presumably need more transmissions over the BH link in order to send the same amount of data to the RRU. The data sent are actual IP packets after the PDCP processing, which have not gone through the concatenation process of the RLC. A qualitative disadvantage of such layer splits is that the data sent to the RRU might need significantly more transmissions over the network, as it is of lower size than the ones outputted by the RLC process. Yet, as most of the contemporary networks can transfer packets of up to a specific size (e.g. MTU equal to 1500 bytes), the usage of technologies like Ethernet can be advantageous for the functional splits. Although this might seem a drawback for this type of split, there are clear benefits of using the PDCP as a convergence layer among different technologies [14]; multiple technologies can be coordinated from a single PDCP/IP instance at the base station, enabling seamless mobility experience across several technologies, with a very little overhead for the network operator.
- **MAC/PHY split:** The MAC/PHY split that we examine has been identified as one of the potential splits in [1], [2] and [7]. In this case, the RRU and BBU are synced and operate in a subframe basis. The BBU unit can instruct, based on the output of the MAC scheduling policy, the subframe allocation for each UE. The actual data that needs to be transferred from the BBU to the RRU is equal to the Transport Block Size (TBS), depending on the modulation and the physical resource blocks which are allocated to each specific UE. This split can be beneficial for the real world application of several algorithms and technologies, such as dynamic scheduling of multiple RRUs, spectrum coordination algorithms [15], beamforming coordination [16], etc.

Figure 1 is illustrating the architecture and the splits that we evaluate. We employ the OpenAirInterface platform as our reference implementation of the LTE stack in order to choose the functions which will be split. Regarding the PDCP/RLC split, the splitting function takes place in the following manner: whenever PDCP is receiving a packet, it goes through its normal procedure before being relayed to the next layer. As soon as the packet is processed, it is sent to the RRU where RLC processes it. The resulting stream is placed in a buffer waiting for the MAC protocol to send a request for it. It is

worth to mention here that the existing buffer for handling this type of data in OpenAirInterface needed to be extended for carrying out our experiments.

Regarding the MAC/PHY split, we choose to override the part where the two layers communicate with each other; this is the point where upon the end of the MAC scheduling algorithm, the BBU instructs the RRU at which subframe the data will be transmitted over the air. This means that no buffering of the packets takes place inside the RRU, but are solely handled by the eNB application (BBU). Whenever data needs to be sent over the air, data streams are sent to the RRU along with all the signaling needed to orchestrate the PHY layer, including the subframe scheduled for transmitting, the number of physical resource blocks, the modulation and coding scheme (MCS), the antennas, etc.

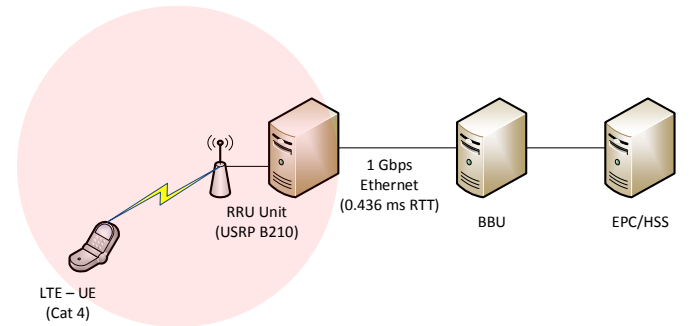
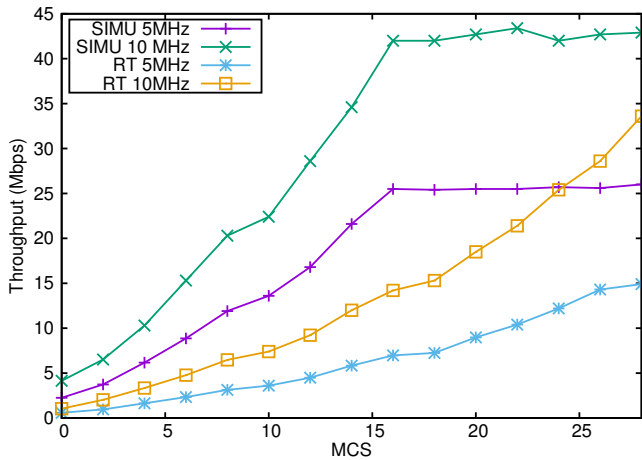


Fig. 2. Experiment setup for the split evaluation

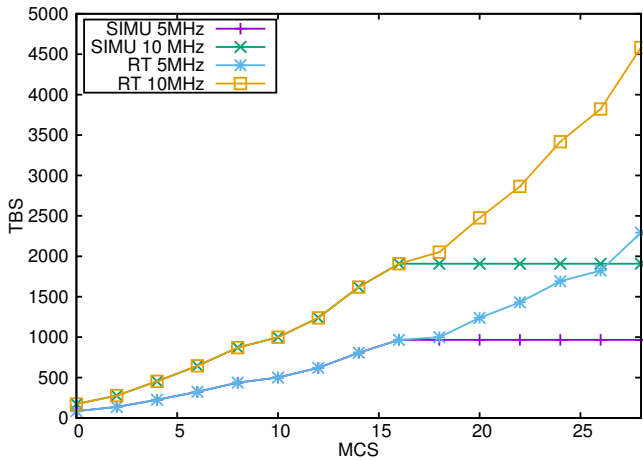
IV. EXPERIMENTAL SETUP

For the evaluation of the chosen splits, we experiment using the NITOS environment. NITOS is a heterogeneous testbed located in the premises of University of Thessaly, in Greece. It offers a very rich experimentation environment with resources spanning from commercial LTE, to WiFi and Software Defined Radio platforms [17]. The topology that we employ is depicted in Figure 2. We split the eNodeB process of OpenAirInterface into two parts, one being executed on a node with a USRP B210 platform, being our RRU, and one on a dedicated testbed node. For the two under study functional splits, different parts of the code are either executed over the BBU or the RRU. A third testbed node is running the EPC and HSS software, while we also employ a fourth node equipped with an LTE Cat. 4 UE.

The splits are configured as follows; we override the default processes for OpenAirInterface and configure one listening server and one client for two different binaries of the code. Each time that a packet is about to be sent as a data request from a higher to a lower layer, it gets packed in a standard message and is sent over the link. We use the *socat* application for redirecting the traffic over the network to the listening server. In case of a data indication message (lower to higher layer direction), a similar process takes place. The parameters of the LTE network setup, as well as the different scenarios that we use are shown in Table I. For all of our cases, the



(a) Throughput achieved per MCS profile



(b) TBS used per MCS profile

Fig. 3. Reference results taken with OpenAirInterface for simulation (SIMU) and real time (RT) operation

testbed nodes are static, and the UE is reporting values of excellent signal quality, with the RSRP ranging from -76 to -83 dBm and reported RSSI values up to -53 dBm. We perform our experiments using two different bandwidth settings, for LTE channels of 5 and 10 MHz.

TABLE I
TESTBED AND SIMULATION PARAMETERS

Network Parameters	Values
LTE mode	FDD Band 7
LTE Frequency	2680 MHz (DL)
No RBs	25, 50
UE	Cat. 4 LTE, Huawei E3272
OAISIM channel emulation	Rayleigh
OAISIM mobility	STATIC
Backhaul RTT	~ 436 msec
Backhaul connection	1Gbps Ethernet
Ethernet MTU size	1400 bytes

Similarly with the real network setup, we use the same testbed nodes in order to run the OpenAirInterface emulation platform [18] (OAISIM). All the functional splits are implemented for both setups, real time and emulator. Regarding OAISIM, we use the setup where the PHY layer is abstracted, meaning that certain functions of the PHY are omitted. This setup is able to yield better results, as the wireless channel is modeled using predefined patterns. For all of our simulation experiments, the multipath model used is Rayleigh, as it is the one that is used by default in OAISIM. The splits are taking place over the same network as happens with the real setup.

As our BH network is an IP based one, we choose to evaluate the performance of different protocols for the splits, depending on the split and real time requirements of the network. Although stateless protocols are the ones that should be adopted for this type of experiments (UDP), we also incorporate TCP and SCTP as our transport solution for the BH for the cases of PDCP/RLC split, which has more loose delay requirements. Our experiments demonstrate that backhauling is viable also with these solutions, although more

capacity for the BH network is needed for achieving similar performance as with the UDP solutions. Regarding the TCP experiments, the congestion control algorithm that we use is *Cubic*, as the rest of the algorithms yielded worst performance results, indicated also in [19]. For the SCTP results, we use 5 parallel streams for each association, and do not use the multi-homing features. We provide experiment results with a resolution of 10 for each measurement. For generating traffic for our measurements, we use the *iperf* traffic generator, set to saturate the wireless link with UDP traffic.

In the following section we present our experimental results, obtained by running the aforementioned functional splits in a real testbed environment as well as with simulation results. The evaluation is broken down in three subsections. Initially we briefly provide some reference measurements from the under study platform without implementing any split. Following, we showcase the experimental results for the PDCP/RLC split, and finally we present our results regarding the MAC/PHY split. Although the splits are applicable for both the Downlink and Uplink data flow, we present measurements for the Downlink channel, as it is the one with the most stringent requirements for transfer. We measure and comment on the total achieved throughput for the LTE UE, for the two under examination functional splits.

V. SYSTEM EVALUATION

A. Reference Measurements

Initially, we present some benchmarking results of the platform that we use for our experiments. In this setup, we use the *vanilla* OpenAirInterface platform, configured as either the LTE emulation platform (OAISIM) or set to operate in real time (RT), running the whole LTE stack in one base station binary application.

In Figure 3 we provide the results regarding the throughput performance achieved per each MCS profile allocated by the eNodeB scheduler (Fig. 3a), as well as the mean TBS used (Fig. 3b). TBS is of paramount importance for the MAC/PHY

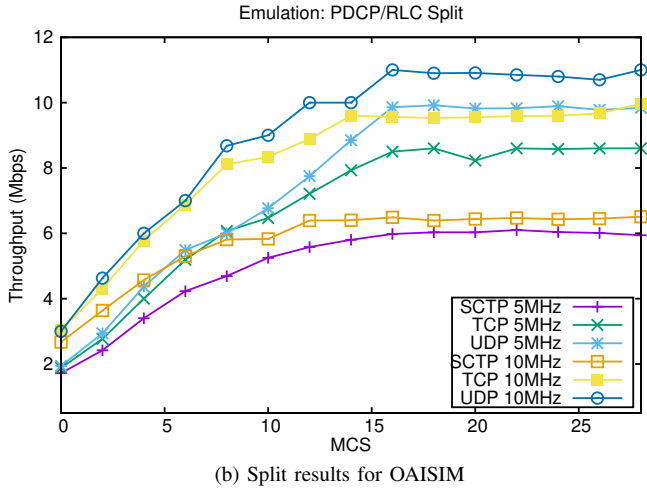
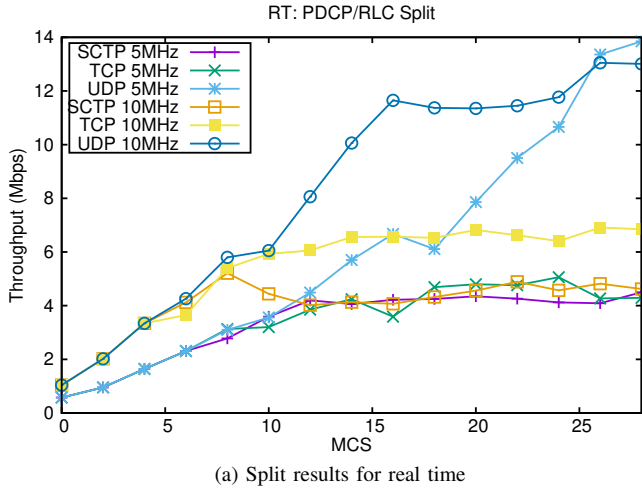


Fig. 4. PDCP/RLC splits when using UDP, TCP and SCTP for backhauling

split, as the output of the MAC processing mandates the transferring of equal sized data to the RRU within the time scheduled for transmission. The bits that are allocated by the LTE scheduler for transmission are the ones that will define the bottleneck in our under investigation backhaul network.

We observe that the OAISIM platform yields the same results for MCS indexes over 16. This happens due to the abstraction flags that are passed to the emulation platform, which omit the execution of certain PHY-layer blocks in favor of better performance. Similarly, the TBS allocated for each transmission follows the same pattern.

B. Evaluation of PDCP/RLC splits

Since the PDCP functions happen at a higher layer, the real time operation can be maintained if proper buffering is used at the RLC level. PDCP is processing every incoming IP packet, and upon the header compression, it delivers it via the BH network to the RRU implementing the LTE protocol below RLC. Whenever MAC layer is finished with the scheduling of its buffered packets, it requests the RLC buffered packets. Based on this fact, real time operation can not be broken if other than stateless protocols are used for the BH. Nevertheless, this fact means that larger memory allocation is needed for enabling such a split. For our experiments, we extended the memory allocation for both the BBU and RRU applications, in order to reassure that the machine does not run out of memory.

Figure 4 is presenting our experimental results when using the real time platform. As we can observe for the real time operation (Figure 4a), and concentrating on the 5MHz transmissions, we see that the worst performing protocol is SCTP. Although SCTP has been introduced as a protocol resolving the head-of-line blocking effect that is present in TCP, its implementations for the Linux kernel are not that mature compared with TCP. For 5MHz, the bottleneck for SCTP when backhauling the LTE data over the 1Gbps Ethernet

link is around 4Mbps. The same bottleneck exists for SCTP backhauling for channels with 10MHz bandwidth.

Regarding TCP experiments, we see that the bottleneck for transferring the 5MHz channels is happening around MCS 18, meaning for TBS sizes over 1000 bits. Similarly, for the 10MHz transmissions, the bottleneck is around MCS 14. Regarding the UDP experiments, for both 5 and 10 MHz transmissions, the bottleneck for the 1Gbps BH link is around 13Mbps. UDP outperforms both SCTP and TCP as it due to its stateless nature, the overhead that is posed on the backhaul only regards the transmission of IP packets, after the PDCP handling and compression to the remote RRU with the RLC layer.

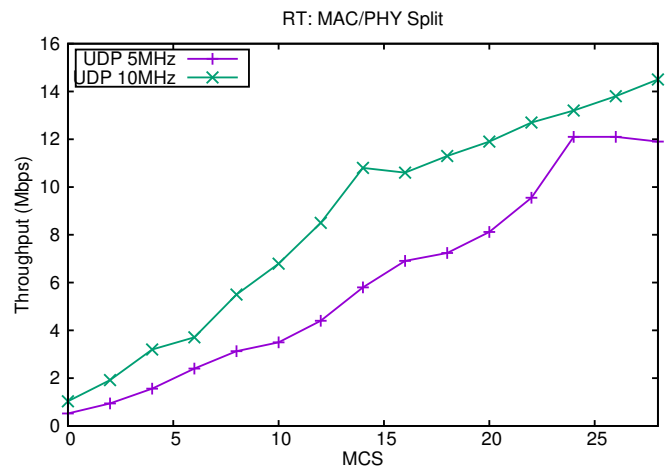


Fig. 5. MAC/PHY real time splits

Regarding the simulation results, a similar pattern as in the real time experiments is witnessed. We observe that the throughput achieved by OAISIM is bounded at approximately 11Mbps for the best case, when using UDP for transferring the data. As illustrated in Figure 3a, for MCS indexes over 16, the data is sent using the same TBS, as several of the PHY

functions are omitted in favor of better performance.

C. Evaluation of MAC/PHY splits

Following the PDCP/RLC splits, we conduct experiments regarding the lower layer split. We present results only for the UDP based data flow, as our first set of experiments denoted that it is the protocol that achieves better performance in such splits. Moreover, the RRU employs a minimal queuing mechanism, so whenever the data is sent over the backhaul to the RRU, they are scheduled for transmission. If they are not sent during the scheduled subframe, they need to be discarded by the RRU. Due to this operation, UDP seems to be the only viable solution for measuring the backhaul network overhead. Apart from the TBS data, information regarding the transmission is also sent, containing the scheduled subframe and physical resource blocks. The split is taking place upon the decision of the scheduler on which subframe the data will be sent (with the subframe duration being 1 msec), the modulation and coding scheme which will be used and the physical resource blocks that will be allocated for each UE.

Figure 5 is illustrating the performance results that are achieved for the MAC/PHY split. Due to the operation of this mechanism, we observe the bottleneck of the backhaul link is around 1500 bits for TBS, when using a 5MHz channel, and around 2000 bits for 10MHz channels. Throughput achieved by the LTE UE is around 14Mbps, whereas in the non-split case it was over 30Mbps. For both cases 5 and 10 MHz we can see that the backhaul network reaches its capacity for MCS indexes over 14. From that point, the achieved throughput is less incremental, compared to the non-split framework.

VI. DISCUSSION AND FUTURE WORK

In this work, we presented experimental results obtained through real experimentation and simulation, about the backhaul performance of two different functional splits over the LTE protocol stack. We investigated the total delivered throughput to an LTE UE, when the eNodeB process is running detached from an RRU, splitting the LTE networking stack at either the PDCP/RLC or MAC/PHY points. Our study concentrates on using existing technologies for the backhaul network of Cloud-RANs, using 1Gbps copper links.

The results that we obtained illustrate that for high layer splits (i.e. PDCP/RLC), the transport protocols can pose performance limitations, but do not break the real-time operation of the base stations. Nevertheless, the results vary and as expected, stateless solutions (e.g. UDP) are found out to be more applicable. Moreover, for lower layer splits, like the MAC/PHY split, where the RRU transmissions are solely scheduled in the Cloud, real time operation mandates the use of high bandwidth solutions, with the least possible overhead.

The proposed PDCP/RLC split can be used as a convergence sublayer among RRUs and BBUs that incorporate more than one heterogeneous wireless technologies. In the future, we foresee to investigate under real-world settings the impact of different functional splits in the low PHY layer, including the splits before the equalization of the signal at the receiver.

ACKNOWLEDGMENTS

The research leading to these results has received funding from the European Union's Seventh Framework Programme under grant agreement no 612050 (FLEX Project).

REFERENCES

- [1] D. Wubben, P. Rost, J. S. Bartelt, V. Savin M. Lalam, M. Gorgoglione, A. Dekorsy, and G. Fettweis. "Benefits and impact of cloud computing on 5G signal processing: Flexible centralization through cloud-RAN". *IEEE signal processing magazine*, 31(6), 2014.
- [2] C. Chang, R. Schiavi, N. Nikaiein, T. Spyropoulos, and C. Bonnet. "Impact of Packetization and Functional Split on C-RAN Fronthaul Performance". In *IEEE International Conference on Communications (ICC)*, 2016.
- [3] China Mobile Research Institute. "White Paper of Next Generation Fronthaul Interface", 2015.
- [4] N. Nikaiein, R. Knopp, F. Kaltenberger, L. Gauthier, C. Bonnet, D. Nussbaum, and R. Ghaddab. "OpenAirInterface 4G: an open LTE network in a PC". In *20th Annual International Conference on Mobile Computing and Networking*. ACM, 2014.
- [5] NGMN Alliance. "Further study on critical C-RAN technologies". *Next Generation Mobile Networks*, 2015.
- [6] A. Checko, H. L. Christiansen, Y. Yan, L. Scolari, G. Kardaras, M. S. Berger, and L. Dittmann. "Cloud RAN for mobile networks - a technology overview". *IEEE Communications surveys & tutorials*, 17(1), 2015.
- [7] J. Bartelt, P. Rost, D. Wubben, J. Lessmann, B. Melis, and G. Fettweis. "Fronthaul and Backhaul Requirements of Flexibly Centralized Radio Access Networks". *IEEE Wireless Communications*, 22(5), 2015.
- [8] J. Bartelt, G. Fettweis, D. Wubben, M. Boldi, and B. Melis. "Heterogeneous Backhaul for Cloud-based Mobile Networks". In *78th Vehicular Technology Conference (VTC Fall)*. IEEE, 2013.
- [9] U. Dötsch, M. Doll, H. Mayer, F. Schaich, J. Segel, and P. Sehier. "Quantitative analysis of split base station processing and determination of advantageous architectures for LTE". *Bell Labs Technical Journal*, 18(1), 2013.
- [10] N. Nikaiein. "Processing radio access network functions in the cloud: Critical issues and modeling". In *Proceedings of the 6th International Workshop on Mobile Cloud Computing and Services*. ACM, 2015.
- [11] N. Nikaiein, R. Knopp, L. Gauthier, E. Schiller, T. Braun, D. Pichon, C. Bonnet, F. Kaltenberger, and D. Nussbaum. "Demo: Closer to Cloud-RAN: RAN as a Service". In *21st Annual International Conference on Mobile Computing and Networking*. ACM, 2015.
- [12] I. Alyafawi, E. Schiller, T. Braun, D. Dimitrova, A. Gomes, and N. Nikaiein. "Critical issues of Centralized and Cloudified LTE-FDD Radio Access Networks". In *IEEE International Conference on Communications (ICC)*. IEEE, 2015.
- [13] S. Bhaumik, S. P. Chandrabose, M. K. Jataprolu, G. Kumar, A. Muralidhar, V. Srinivasan P. Polakos, and T. Woo. "CloudIQ: a framework for processing base stations in a data center". In *Proceedings of the 18th Annual International Conference on Mobile Computing and Networking*. ACM, 2012.
- [14] C. Richards. "Common Radio Resource Control for Cellular Radio and WiFi". US Patent App. 13/707,184.
- [15] J. Bastos, K. M. S. Huq, S. Mumtaz, J. Rodriguez, and C. Verikoukis. "Energy efficiency optimization for downlink OFDMA system in heterogeneous network with QoS constraints". *International Journal of Communication Systems*, 2015.
- [16] Y. Dai, S. Jin, L. Pan, X. Gao, L. Jiang, and M. Lei. "Interference control based on beamforming coordination for heterogeneous network with RRH deployment". *IEEE Systems Journal*, 9(1), 2015.
- [17] N. Makris, C. Zarafetas, S. Kechagias, T. Korakis, I. Sesar, and L. Tassioulas. "Enabling open access to LTE network components; the NITOS testbed paradigm". In *1st Conference on Network Softwarization (NetSoft)*. IEEE, 2015.
- [18] A. Virdis, N. Iardella, G. Stea, and D. Sabella. "Performance analysis of OpenAirInterface system emulation". In *3rd International Conference on Future Internet of Things and Cloud (FiCloud)*, 2015. IEEE, 2015.
- [19] J. Wang, J. Wen, J. Zhang, and Y. Han. "A demonstration of a new TCP Congestion Control Algorithm over LTE and other challenging networks". In *16th Annual International Conference on Mobile Computing and Networking*. ACM, 2010.