



Design, implementation and experimental validation of a 5G energy-aware reconfigurable hotspot

Oriol Font-Bach^{*,a}, Nikolaos Bartzoudis^a, Marco Miozzo^a, Carlos Donato^{b,c}, Pavel Harbanau^a, Manuel Requena-Esteso^a, David López-Bueno^a, Pablo Serrano^c, Josep Mangués-Bafalluy^a, Miquel Payaró^a

^a Centre Tecnològic de Telecomunicacions de Catalunya (CTTC)-CERCA, Parc Mediterrani de la Tecnologia (PMT), Av. Carl Friedrich Gauss 7, Barcelona, 08860 Castelldefels, Spain

^b IMDEA Networks Institute, Av. Mar Mediterráneo 22, Leganés 28918, Spain

^c Department of Telematics Engineering, Universidad Carlos III de Madrid, Av. de la Universidad 30, Leganés 28911, Spain

ARTICLE INFO

Keywords:

HW-SW partitioning
5G networks
5G dynamic hotspots
Energy-aware design
Reconfigurability
Real-time prototype
Experimental validation

ABSTRACT

Flexibility and energy efficiency are considered two principal requirements of future fifth generation (5G) systems. From an architectural point of view, centralized processing and a dense deployment of small cells will play a vital role in enabling the efficient and dynamic operation of 5G networks. In this context, reconfigurable hotspots will provide on-demand services and adapt their operation in accordance to traffic requirements, constituting a vital element of the heterogeneous 5G network infrastructure. In this paper we present a reconfigurable hotspot which is able to flexibly distribute its underlying communication functions across the network, as well as to adapt various parameters affecting the generation of the transmitted signal. The reconfiguration of the hotspot focuses on minimizing its energy footprint, while accounting for the current operative requirements. A real-time hotspot prototype has been developed to facilitate the realistic evaluation of the energy saving gains of the proposed scheme. The development flexibly combines software (SW) and hardware (HW) accelerated (HWA) functions in order to enable the agile reconfiguration of the hotspot. Actual power consumption measurements are presented for various relevant 5G networking scenarios and hotspot configurations. This thorough characterization of the energy footprint of the different subsystems of the prototype allows to map reconfiguration strategies to different use cases. Finally, the energy-aware design and implementation of the hotspot prototype is widely detailed in an effort to underline its importance to the provision of the flexibility and energy efficiency to future 5G systems.

1. Introduction

Fifth generation (5G) networks will need to tackle the massive growth of the data traffic (i.e., a seven-fold increase is expected in 2021 [1]), while accommodating a wide range of new applications (i.e., industry verticals) and providing ubiquitous communication services to mobile users across a heterogeneous radio access network (RAN) infrastructure. In this context, centralized processing and very dense deployments are envisioned as key enablers of future 5G systems. Additionally, there is an imperative necessity to design 5G networks to be sustainable by bounding their energy consumption in spite of their increased service provision. Because of all the previous, flexibility and reconfigurability need to be secured from conception and at different

levels, starting from the underlying hardware (HW) elements and arriving up to the adaptive management of the network operation.

In the described environment, flexibly distributing the communication functions across different network elements and adopting enhanced cloud computing schemes allows to transversely optimize the use of infrastructure resources. Yet, selecting the optimal functional split is a complicated task which requires carefully balancing the performance, delay and energy efficiency requirements [2]. From an architectural point of view, this reconfigurability needs to be enabled by intelligently combining a cloudified multi-layer network, embracing both cloud RAN (C-RAN) and multi-access edge computing (MEC) paradigms, with dynamically operated small cells.

Ultradense small cell deployments are thus not only meant to extend

* Corresponding author.

E-mail addresses: oriol.font@cttc.cat (O. Font-Bach), nikolaos.bartzoudis@cttc.cat (N. Bartzoudis), marco.miozzo@cttc.cat (M. Miozzo), carlos.donato@imdea.org (C. Donato), pavel.harbanau@cttc.cat (P. Harbanau), manuel.requena@cttc.cat (M. Requena-Esteso), david.lopez@cttc.cat (D. López-Bueno), pablo@it.uc3m.es (P. Serrano), josep.mangués@cttc.cat (J. Mangués-Bafalluy), miquel.payaro@cttc.cat (M. Payaró).

<https://doi.org/10.1016/j.comcom.2018.06.008>

Received 25 August 2017; Received in revised form 1 June 2018; Accepted 14 June 2018

Available online 19 June 2018

0140-3664/ © 2018 Elsevier B.V. All rights reserved.

the macro cell coverage in dense urban environments, but are fundamental to provide a dynamic 5G network architecture. Besides the standard static small cell installations, dynamic hotspots are also required to provide on-demand services in localized spaces and time periods (e.g., due to a large increase of capacity requirements). In such heterogeneous environment, optimizing the energy footprint is key to ensure the sustainability and economic viability of small cells [3] (e.g., temporary deployments might be battery-powered). Adapting the transmission strategy is also fundamental from a practical point of view, since small cells need to coexist with other RAN infrastructure elements (and deal with the associated problems, e.g., interferences, limited spectrum resources) and to efficiently serve varying traffic and quality of service (QoS) requirements. Consequently, beyond providing a rigid macro-cellular traffic offloading scheme, small cells are principal facilitators of the intelligent distribution of communication functions with respect to the dynamically varying demands of the mobile users.

This ongoing quest toward flexibility and agile reconfiguration is one of the main drivers of the current efforts to softwarize 5G networks. The pairing of software (SW) defined networking (SDN) and network function virtualization (NFV) techniques with programmable HW will naturally play a critical role in the implementation of 5G systems and dynamic hotspots in particular, as it will be shown in this paper.

1.1. Related work

In the ongoing activities to define and standardize the cornerstone 5G technologies, there is a common goal to attain a flexible, scalable and energy efficient network architecture. Indeed, energy efficient wireless networks have been a subject of intense research for over a decade, as illustrated by numerous relevant survey papers on the topic, e.g., [4–8]. More recently, the flexible distribution of communication functions across the 5G network started attracting more attention among the proposed solutions. The study of these solutions starts from theoretical and numerical analyses covering both their fundamental challenges [9,10] and energy efficiency benefits [11,12].

In a narrower scope, there is a growing interest in the literature to investigate small cell deployments which are aimed at addressing the capacity requirements in dense indoor and outdoor urban environments. In particular, many research efforts are focusing on proposing methods to efficiently reduce the amount of energy drained by these network installations. The great majority of such analyses are based on computer simulations. A relevant example of the latter is found in [13] where a deployment of nomadic nodes is proposed (i.e., vehicle mounted battery-powered relays) in order to provide temporary demand-driven service provisioning to mobile users. The authors analyse the energy saving gains obtained by dynamically switching off those nomadic cells with no traffic demand. Similarly, the authors in [14] analyse a database-aided scheme that enables the use of deep sleep modes (i.e., completely switch off the transceiver HW) in clustered small cell deployments (e.g., train stations, shopping malls). Energy-saving gains up to 30% are observed in the simulated scenarios.

The centralization of resources is also a topic of elevated interest in the recent literature, with special attention to C-RAN architectures and communication function splits. In this context the first prototyping efforts have been also identified, mostly based on pure SW implementations and not accounting for energy-related aspects. Notably, the authors in [15] present a long term evolution (LTE)-based testbed to analyse a medium access control (MAC)-PHY split featuring an Ethernet fronthaul (i.e., favoring the reuse of the existing packet based infrastructure). The analysis is conducted from the standpoint of incurred latencies. A closely related contribution is found in [16] where an LTE-based prototype for flexible C-RAN deployments introduces an implementation of the next generation fronthaul interface (NGFI), which aims at redefining the functional split between the baseband units (BBUs) and the remote radio heads (RRHs). In this respect the authors test two different PHY-layer splits from a functional point of view, with

the second one executing part of the digital signal processing (DSP) functions on the RRH. A low-latency time division duplex (TDD)-based physical (PHY) air interface is proposed and modeled in [17] to be used in 5G dense deployments. Focusing on the user equipment (UE), the authors argue that the battery lifetime can be largely extended by using the proposed scheme instead of LTE.

Numerous studies in the field of heterogeneous C-RAN networking (i.e., combination of C-RAN and small cells) can also be found in the literature. In many cases, the authors of these studies propose an energy-aware management of the network that consists in switching off those parts which are not required given the current operative requirements. For instance, in [18] a flexible C-RAN prototype for small cells based on commercial WiMAX BBUs is described. The authors coarsely evaluate the energy savings obtained as a function of the number of BBUs that are deactivated during low activity periods. In contrast, the work detailed in [19] presents experimental energy measurements from a similar traffic-aware C-RAN prototype based on standard server (i.e., computing) nodes. Another relevant development is found in [20], where an LTE-based C-RAN testbed featuring few HWA functions is described. Whereas the authors do not provide power consumption measurements, they are making a qualitative energy saving analysis.

Finally, the use of field programmable gate arrays (FPGA) devices to implement HW accelerated (HWA) DSP functions in a SW defined radio (SDR) context has been continuously growing with the increased capacity of modern programmable devices [21,22]. A relevant effort is described in [23], where the authors introduce an FPGA-based prototype that combines SW and HWA functions to test different synchronization procedures in an Ethernet-based fronthaul. Moreover the utilization of advanced digital design techniques to reduce the energy consumption of FPGA implementations is a well-established topic [24–26]. Recently, the emergence of FPGA-based system-on-chip (SoC) devices featuring an unprecedented combination of performance and flexibility, has introduced the means to efficiently implement SDN systems [27,28] and opened the door to the dynamic distribution of SW and HWA functions. However, to this date few contributions could be found in the literature, which mainly cover partial FPGA implementations, such as the common public radio interface (CPRI)-based C-RAN systems presented in [29,30]. Furthermore, a fixed function split is showcased in [31], where the authors describe a simple C-RAN system featuring dynamic offloading of MAC functions onto the Cloud and analyse the energy gains obtained by optimizing the number of active processors.

1.2. Motivation and contribution

The aim of this work is to show that the use of reconfigurable small cells can effectively contribute to reduce the energy consumption of 5G networks. To that end, we propose the deployment of dynamic hotspots that can be reconfigured in order to optimize the utilization of available HW resources. Specifically, a hotspot reconfiguration scheme that distributes communication functions across the network (e.g., C-RAN), and that also adapts various aspects affecting the generation of the transmitted signal is envisioned. In this context, the reconfiguration of the hotspot should account for the current operative requirements, while simultaneously attempting to minimize its energy footprint.

Another important objective of the paper is to complement the contributions found in the literature by presenting a complete and realistic validation of the proposed system, supported by actual power consumption measurements. In that respect, our work revolves around the development of a fully reconfigurable real-time hotspot prototype that combines HWA and SW functions in a flexible manner. In addition, the energy-aware design and implementation of the underlying functions is widely detailed. Furthermore, the energy saving gains are evaluated at a subsystem level which provides a fine-grained characterization of the benefits resulting from the hotspot reconfiguration.




Fig. 1. Considered scenario: dynamic 5G hotspots.

All the previous finally allows us to determine which of the evaluated adaptation schemes could better serve the goals of the dynamic hotspot under different operating scenarios.

2. System description

In this section, the use case that is considered in the paper will be briefly described underlining the relevance of energy-awareness in 5G networks. It must be noted that, given the experimental flavour of the presented work and the absence of a standardized and fully defined 5G air interface, the 4G LTE is being used throughout the paper. As it will be detailed later in the text, the main ideas and reported concepts regarding the implementation of different function splits will remain relevant with independence from the new radio (NR) waveform that will be finally adopted.

2.1. Considered scenario

This paper takes as a use case a dense urban scenario where, in a limited space and time, the network needs to handle a sudden increase in the number of user connections (e.g., crowded venues, shopping malls or large public buildings among others). Hence, the nearby macro basestations, or evolved node Bs (eNBs), cannot adequately serve the compound coverage, number of subscribers and overall capacity requirements. In this situation, data traffic is offloaded onto a dynamic hotspot (see Fig. 1). That is, a number of strategically deployed small cells provide the otherwise lacking network resources. In order to ensure an optimized utilization of the infrastructure with a reduced energy footprint, these small cells need to adapt their operation to the instantaneous performance requirements. A flexible split of their underlying 5G communication stack could be applied to leverage traffic dynamics and optimize the utilization of resources across the network, addressing in this manner the optimization of given key performance indicators (KPIs; e.g., latency, capacity). Here we propose to use the energy efficiency KPI as the driving factor of the presented small cell reconfiguration. We believe it is particularly relevant considering that the baseband processing consumes a significant fragment of the total energy budget of the small cell (i.e., nearly a 50% in a femtocell [32]). Moreover, if we consider battery powered dynamic hotspots (e.g., temporary deployments to increase the network capacity in reduced spaces during crowded events, such as in a large music festival), their energy-aware reconfiguration would play a vital role in ensuring their availability (i.e., by reducing the consumed energy whenever possible).

In this paper we present the design and experimental validation of an energy-aware hotspot. In more detail, the presented work revolves around the dynamic reconfiguration of the home eNBs (HeNBs; i.e., basestations of the small cell) in order to minimize their related energy footprint, while fulfilling the QoS requirements of the connected users. Towards that end, the HeNB flexibly partitions its communication functions across different nodes in 5G networks, either at stack-level, which is the focus of this paper, or at algorithm-level, in order to adapt the most energy-efficient configuration. On top of that, the HeNB is also

able to reconfigure its primary wireless communication parameters (WCPs) to optimize the balance between consumed energy and system performance. Relevant examples of WCPs are: scaling the downlink (DL) signal bandwidth (BW), limiting the maximum allowable modulation and coding scheme (MCS) index and/or resource block (RB) group (RBG) allocation, and constraining the transmit power.

As described in [33,34], there are many different function partitioning possibilities. However, our work considers three specially relevant network configurations (NETCFGs), as shown in Fig. 2, each one presenting a different degree of communication function offloading:

- i. The first considered partition has all L1 functions executed locally at the HeNB, whereas all higher layer processing is offloaded onto the Cloud. From this point onwards, this particular function split will be referred to as *NETCFG1*.
- ii. As in the previous case, the PHY-layer is executed locally at the HeNB. Nevertheless, in this second configuration, or *NETCFG2*, a MEC-like approach is adopted. That is, the remaining protocol stack functions will be virtualized as a specialized type of application in a server located in the vicinity of the small cell (i.e., MEC-like node).
- iii. *NETCFG3* considers the typical C-RAN setup, where the HeNB will act as a RRH, while all protocol stack functions are placed onto the Cloud.

The presented scenario assumes that high-speed communication links are available where required, in order to let the small cell communicate with the Cloud and/or MEC-like nodes. As it can be observed in Fig. 2, depending on the adopted NETCFG, the communication links between the different 5G nodes present quite diverse latency and data-rate requirements [35]. In more detail, the interconnection of the HeNB to the Cloud or to the MEC-like node has stringent traffic requirements, which can be fulfilled using an ideal transport channel (i.e., 250 μ s of one-way latency and 2.5 Gbps for supporting a standardized LTE 2×2 20 MHz scheme) and the CPRI specification.¹ The interconnection of the MAC with the PHY-layer (i.e., L2-L1 interface) poses more relaxed needs, that can be satisfied with sub-ideal transport channels (i.e., 6 ms of one-way latency and a capacity of 150 Mbps). Finally, the interconnection to the evolved packet core (EPC; i.e., S1 traffic) is the less demanding and can be satisfied with a non-ideal transport channel (i.e., up to 30 ms of one-way latency and variable BW).

It is important to underline that, while the main focus of the proposed HeNB reconfigurations is set on the energy efficiency, this flexibility could also be exploited to satisfy a wide range of KPIs in different operative 5G scenarios (e.g., latency, availability or performance among others).

2.2. System architecture

A real-time reconfigurable hotspot prototype has been developed with the objective to facilitate a realistic yet simple validation of the proposed concept described in the previous section. To make this possible, the prototype flexibly combines HWA and SW communication functions, by using standard computers and HW components. Regarding the implementation of HWA blocks, the target platform is an FPGA-based SoC device: The latter embeds an integrated processing system (PS) and programmable logic (PL) on a single die, providing likewise high flexibility (i.e., run-time reconfigurability) and computational capacity (i.e., massive parallelism).

Focusing on the goal of obtaining an empirical assessment of the energy savings attainable by reconfiguring the system, the prototype

¹ When looking forward to NR it is relevant to underline that the latest CPRI specification (7.0), with rates up to 24 Gbps, is able to accommodate significantly larger BWs than the 20 MHz considered in our scenario (e.g., an LTE-based 2×2 200 MHz scheme requires 19.2 Gbps). Furthermore, the first specification of the CPRI evolution for 5G (eCPRI 1.0) has been recently released [36].

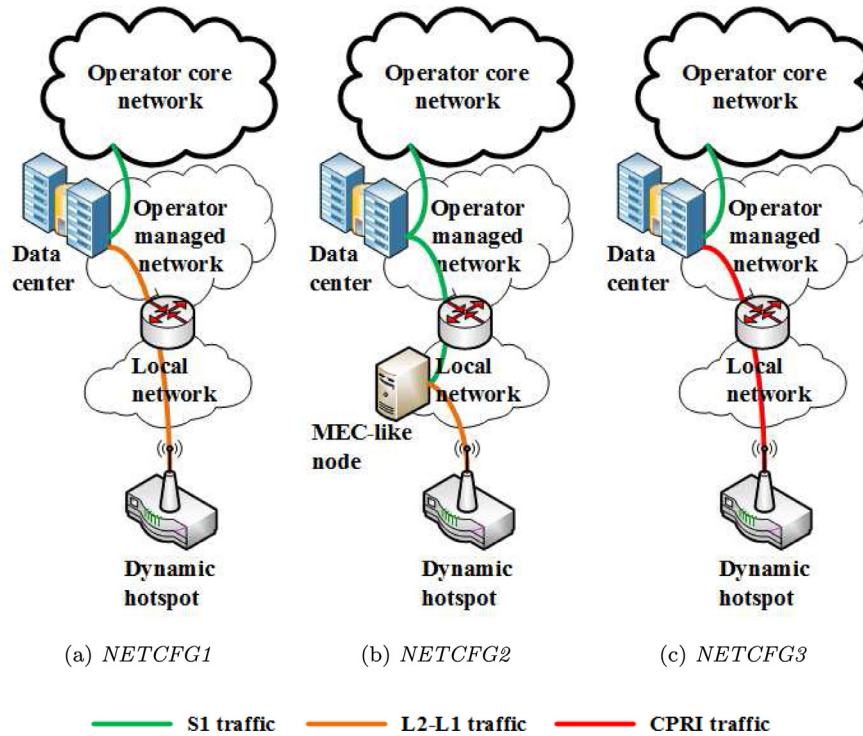


Fig. 2. Considered NETCFGs and their related traffic loads.

has been kept as simple as possible. Hence, in the presented work the implemented system includes a single LTE-based HeNB, which uses different HW elements to execute its functions, depending on the selected NETCFG, and a single UE.² Moreover the required EPC functionalities are also included, jointly with an emulated Internet, to complete the system implementation.

3. Energy-aware design and system implementation

The principal objective of the development presented herein is to allow the realistic evaluation of the gains that could be provided by applying energy-aware reconfigurations at the dynamic hotspot. For this purpose, a real-time prototype has been designed and implemented, where the communication functions can be moved among different processing nodes in order to improve the energy efficiency of the system, accounting for the actual network conditions as well. The prototype combines a fully SW-based LTE emulator, implementing L2 and above stack functions, with a HWA DL PHY-layer based on a custom HW description language (HDL) realization.³ This enables both the real-time operation and run-time reconfiguration of the hotspot. Moreover, the implementation complies with the most relevant features described in the LTE standard and targets commercial off-the-shelf HW elements. In the following, the design and implementation fundamentals of the dynamic hotspot prototype are briefly discussed, setting the focus on the HeNB. The full technical details can be found in [Appendix A](#).

3.1. L2 and upper-layers

All L2 and above layer functionalities have been implemented in the SW domain by extending the existing features of the LTE-EPC network simulator (LENA) [37]. LENA is an open-source implementation of the

² Given that the current HeNB implementation supports multi-user transmissions, adding more UEs to the prototype only requires replicating the related HW setup.

³ The uplink implementation is kept in the SW domain, making use of the native functionalities of the LTE emulator.

LTE and EPC standards, originally born to enable the simulation of realistic scenarios relying on the widespread network simulator ns-3 [38]. It possesses two main characteristics that result crucial to the development of the dynamic reconfigurable hotspot. First, given that ns-3 is a full stack simulator, all upper-layer functionalities are accurately implemented (which is not the case for link and system level simulators). Second, the core design of LENA is based on the Small Cell Forum MAC scheduler application programming interface (API). This enables its fast integration to realistic developments (i.e., by considering commercial product requirements), as well as modelling their constraints in the scheduler design. On top of that, the LENA module also implements the EPC network elements and their protocol stacks, including the serving gateway (SGW), the packet data network gateway (PGW) and the mobility management entity (MME). A full internet protocol stack is also incorporated. It provides cornerstone networking functions, including the transmission control protocol (TCP), the user datagram protocol (UDP) and the internet protocol (IP). With all those combined features, LENA allows to accurately simulate the performance of end-to-end services in LTE-based network configurations, considering both the fronthaul and backhaul, as well as all communication stack protocols ([Fig. 3](#)).

The original purpose of LENA was to be used as an advanced SW simulator. Consequently, LENA was developed as single process that encapsulates all network elements for a given simulation scenario. Similarly, all interactions between those network elements were constrained to the ns-3 simulation-space (i.e., no interaction to third party SW or external network elements was enabled; e.g., video applications). Moreover, neither real-time computing nor HW interaction support was meant to be originally provided. All these limitations have been addressed in the current work by adequately extending the original LENA code. Whereas the full details of such extension are provided in [Appendix A.1](#), a brief description follows.

Certain advanced ns-3 features were exploited to facilitate time-constrained interactions with external SW and HW elements (e.g., FPGA-based PHY-layer). Moreover, the original LTE interfaces have been updated to work with real IP packets, that can be sent through the network and reinjected to the simulator. These modifications allow to




Fig. 3. LTE-EPC data plane protocol stack.

realistically emulate 5G scenarios considering different SDN/NFV configurations [39], and facilitate MEC-based deployments (e.g., in the line of *NETCFG2*).

Several modifications have also been introduced in LENA to enable the distributed execution of the communication stack among different RAN processing nodes. To start with, the EPC protocols have been extended to facilitate the emulation of their functions on different network elements. Furthermore, all processes implementing the RAN have also been separated. Specifically, the functionalities corresponding to the (H)eNB and UE are now completely splitted from one another, as shown in Fig. 4. As a result, all underlying SW communication functions can be executed in a completely flexible and distributed manner, covering a wide range of function-split cases.

The specifications defined by the Small Cell Forum for the scheduler API have been exploited in order to facilitate a virtualized small cell architecture [33]. As a result, the RAN functions of LENA are now fully virtualizable, enhancing accordingly the splitting capabilities of the whole LTE protocol stack. A new L2–L1 interface has also been implemented to facilitate the dynamic reconfiguration of the HWA L1, based on the exchange of time-stamped packets. This design relaxes the requirements both in terms of latency and BW in C-RAN and PHY splits as defined in [33], such as the one used in *NETCFG1*. Similarly, the MAC-split helps attaining an increased flexibility when moving the L2 and EPC network elements. Finally, specific HW requirements have been integrated in the scheduler.

3.2. L2–L1 interface

The L2–L1 interface constitutes a key element to facilitate the function split by enabling the agile communication of the HWA and SW blocks. In this regard, the main objective of this interface is to provide a

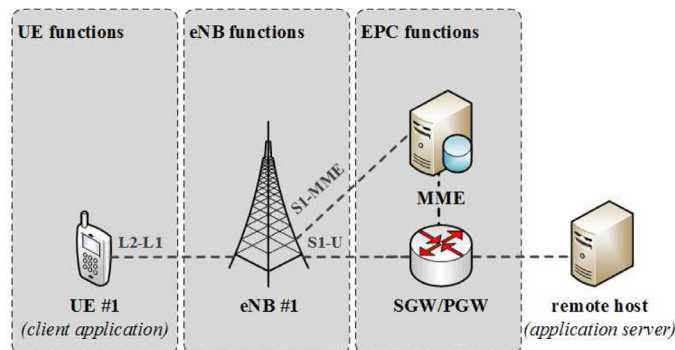


Fig. 4. Architecture of the LTE-EPC network.

reliable communication means with strict latency requirements. Moreover, its inner procedures are transparent to the HWA and SW blocks, simplifying likewise their design. This means that the communication of the partitioned functions is abstracted from the DSP design and only presents a minimal impact on the internal structure of each block (i.e., by conforming with essential interfacing specifications). This allows optimizing each stage independently and promotes the modularity of the overall system design.

The real-time L2–L1 interfacing SW is executed in the PS of the target SoC. A customized distribution of the Linux operating system able to serve real-time tasks has been utilized. More specifically, a fully preemptive kernel is used to satisfy the stringent latency requirements of real-time applications. By this way, application-level and kernel-level jitters are reduced enabling likewise a reliable communication between the partitioned L2 process and the HWA L1 with a deterministic behavior. A series of SW techniques were employed toward that end, including real-time scheduling policies, assignation of real-time priority to critical tasks, memory locking mechanisms and pre-faulting stacks [40].

From an architectural point of view, the L2–L1 interface is connected to the 5G node hosting the L2 process of LENA on the one end, and to the HWA implementation of the L1 residing in the same chip on the other end. The interaction with L2 and above layers is based on the exchange of messages on a subframe basis. Given the strict latency requirements of the real-time L2–L1 communication, the transfer of information between the SW processes relies on a UDP connection. In more detail, the messages are assembled onto L2–L1 interfacing frames, which use a custom and flexible format that efficiently adapts its contents according to the current system configuration (e.g., the use of different WCPs greatly affects the amount of information to be exchanged between L2 and L1). As it can be observed in Fig. 5, the interfacing frames are comprising a number of control and data messages that need to be passed from LENA to the HWA L1. The DL control information (DCI) is of key importance, since it defines the contents of each radio frame (i.e., a DCI is generated for each subframe, according to the Small Cell Forum API). For each transport block (TB) directed to




Fig. 5. Simplified view of the L2–L1 interfacing frame (with an example DL-DCI message).

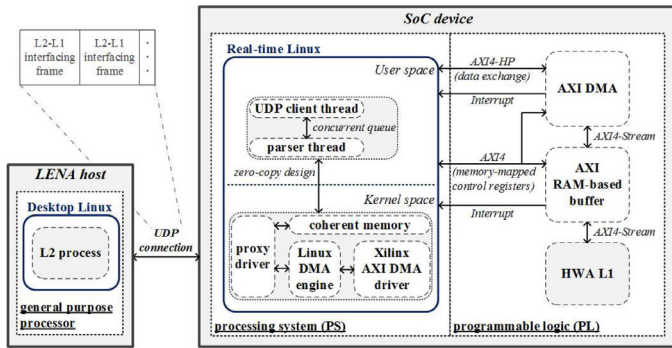


Fig. 6. General overview of the implemented L2-L1 interface.

a given UE, the following parameters are specified: size (i.e., amount of user-data bits), MCS index and bitmap of allocated RBs (i.e., bit-mask indicating which set of physical DL resource elements (REs) are dedicated to transmit the TB). Other control plane messages that are supported include the system information block (SIB) or the master information block (MIB). The implemented interface is also enabling the time-constrained exchange of the data plane. Given that the hybrid automatic repeat request (HARQ) mechanism has not been implemented, cyclic redundancy check (CRC) codes are used as a basic error detection mechanism in order to avoid passing corrupted packets to the upper layers. The first task of the L2-L1 interfacing SW is thus to parse the received L2-L1 interfacing frames in order to recover the control information (i.e., MCS index, TB size and RB allocation bitmap) and user data (i.e., TB) required by L1.

As for the interaction with the HWA L1, the PS and PL communicate through an embedded dedicated high-speed interface which is based on a proprietary bus specification known as advanced extensible interface (AXI).

Fig. 6 presents an overview of the custom SW architecture implementing the L2-L1 interface. As it can be observed, the application is divided in two main components. The user-space process comprises two threads that share a concurrent queue. The first thread is responsible for establishing the network connection with the L2 host and receiving the UDP packets generated in the communication. The second thread

implements the parsing functions, in order to interpret the received information and produce the control information that will be forwarded to the HWA L1. A detailed technical description of the L2-L1 interface is provided in Appendix A.2.

3.3. Energy-aware HWA L1

All required DL L1 features have been implemented as real-time FPGA-based HWA functions by using advanced digital design techniques. A low-level optimized register transfer level (RTL) architecture was designed, focusing on two major goals: i) to minimize the utilized logic resources and its related energy-consumption, and ii) to enable a flexible on-the-fly reconfiguration of its operation. This was achieved by combining a highly resource-efficient RTL design, which reused Xilinx intellectual property cores (IP-cores; i.e., reusable logic blocks) to implement the most complex DSP functions (e.g., inverse FFT, channel coding), and dedicated control units to ensure the energy-efficient yet very flexible operation of the logic. In this respect, the operation of the HWA L1 can be adapted according to the requirements provided by the SW L2 or due to the adoption of a new NETCFG (including DL BW adaptations), in a subframe basis. A general overview of the energy-aware RTL architecture implementing the HWA L1 of the HeNB is depicted in Fig. 7.

From a functional point of view, the L2 dictates the configuration to be utilized by the PHY-layer according to the instantaneous operative requirements of the HeNB. Among others, this includes the specific allocation of frequency resources to each UE attached to the HeNB (i.e., RBGs) and its related QoS constraints (i.e., channel coding parameters). All the necessary control information to reconfigure the HWA L1 is provided from the SW domain (PS) through L2-L1 interfacing-frames (recall Fig. 5). The main control unit of the PHY-layer parses this information in order to determine whether there is data allocated onto the available RBGs or not, and generates the corresponding user-data (i.e., DL shared channel, DLSCH; management of the turbo encoding stage) and control channels contents (i.e., physical DL control channel, PDCCH/physical broadcast channel, PBCH; management of the convolutional encoding stage). Additionally, this complex state machine also manages the errors in the L2-L1 communication: in case of missing or wrongly decoded control information, the HeNB will interrupt its

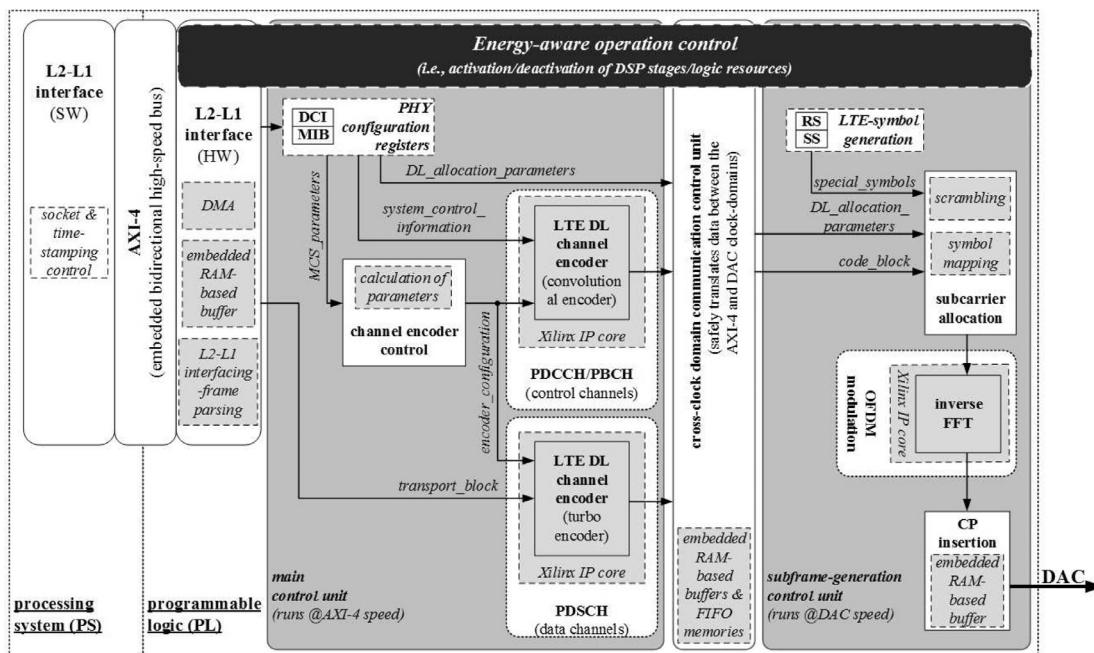


Fig. 7. RTL architecture of the HWA L1.

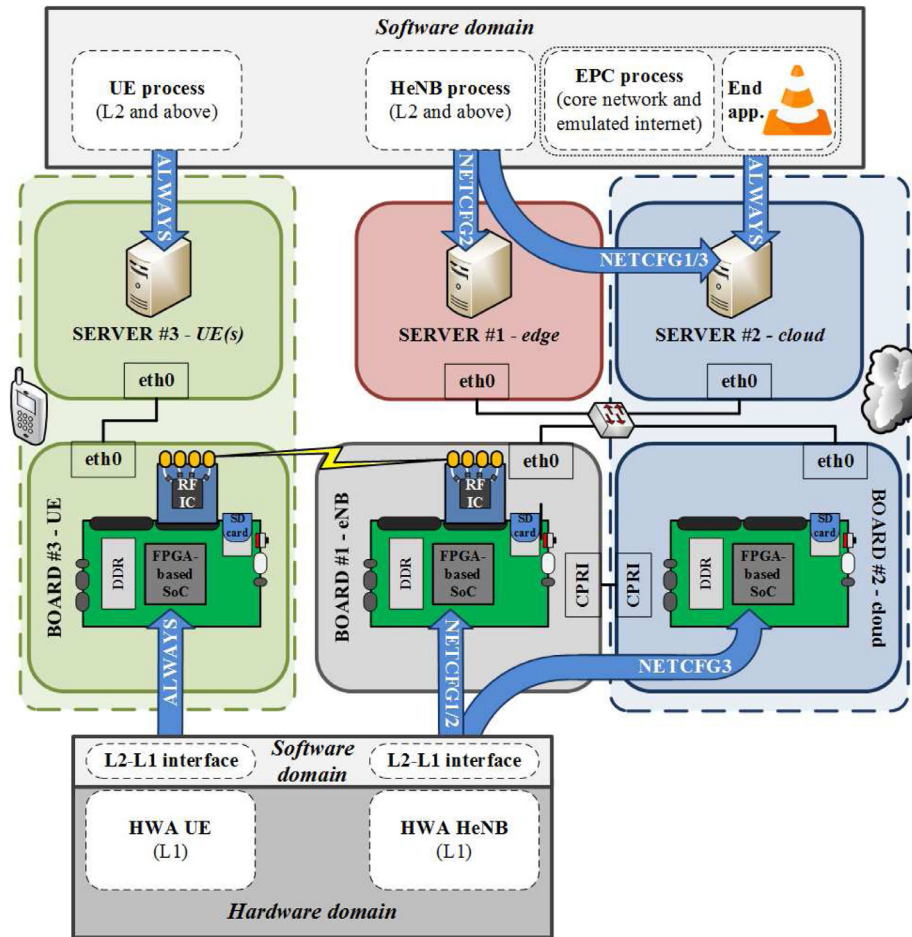


Fig. 8. Overview of the HW setup implementing the hotspot prototype, including the different supported communication function splits.

transmission until valid data is received from L2 to generate the following frame (i.e., starting from subframe 0). This situation will be signaled to the embedded memory buffer within the HWA part of the L2–L1 interface which, in turn, will generate an interrupt to alert the PS. This ensures that the exception will be correctly handled at all communication levels. A second control unit is in charge of allocating the required contents to each RE in the DL signal (i.e., generation of the frequency-subframe), driving the operation of the inverse FFT and CP insertion stages, and, finally, providing further management of L2–L1 communication errors. In more detail, in case the required DLSCCH control is not available when needed, a request will be made to the principal state machine to interrupt the HeNB transmission.

Besides reconfiguring and driving the operation of the different stages comprising the PHY-layer, the flexible control of the logic provided by the control units forms the basis of the energy-aware design, when intelligently combined with standard clock-gating techniques. By this way, it is minimized the amount of logic which is actively utilized at each moment. This design-time decision contributes to the overall reduction of the energy budget of the HeNB: minimizing the switching activity of the implemented gates (e.g., flip-flops), actively reduces the energy consumed by the digital circuit (experimental studies show dynamic power savings up to 30% for the complete design, and up to 90% at a block/IP-core level [41,42]). This is especially relevant for those logic elements residing in the fast clock domain, considering that the dynamic power consumption of any FPGA design has a linear dependency on the clock frequency [43]. Our energy-aware RTL design exploits this trait by minimizing the amount of active logic at any instant. Similarly, the energy-saving efforts are propagated across the entire RTL hierarchy, through low-level optimizations within each

designed DSP block. A fully detailed description of the HWA L1 RTL design is provided in Appendix A.3.

4. Dynamic hotspot prototype

A schematic of the basic HW setup of the energy-aware dynamic hotspot prototype is presented in Fig. 8. Its flexible real-time operation supports the typical performance requirements described by the LTE standard, as well as the reconfiguration of the system to adopt different NETCFGs and/or wireless communication parameters (e.g., DL BW). The following section details the different HW components that are hosting the presented HWA and SW functions. Even though the focus of the described work resides in the adaptive HeNB prototype, its UE counterpart functions were also developed and included in the HW demonstrator in order to enable the end-to-end operation of the system under test. In this regard, the design and implementation details of the receiver subsystems are out of the scope of this work.

4.1. Hardware components

4.1.1. EXTREME[®] testbed

The EXTREME[®] Testbed is used to host the different SW processes of LENA (i.e. HeNB, UE and EPC). The EXTREME[®] Testbed [44] is an experimental framework for testing wireless access and backhaul/fronthaul architectures featuring generic purpose server pools (e.g., SDN control, NFVs such as vEPC), cellular and other wireless equipment, ns-3 emulation/simulation, and tools for fast prototyping and evaluation. Here, we only describe the components directly related to the hotspot prototype. The core of the EXTREME[®] Testbed lies on two

Table 1
FPGA-resource utilization of the implemented HWA L1 entities.

Utilization	HWA eNB functions (%)	HWA UE functions (%)
LUT	22.30	34.52
LUTRAM	2.38	13.38
FF	9.57	26.88
BRAM	8.53	61.65
DSP	2.67	60.78
IO	28.73	19.61
BUFG	9.38	21.88
MCMC	0	37.50

central management servers, which act as interface between the final users and the SDN/NFV experimentation services. A series of reconfigurable multi-purpose servers and high-performance laptops can be customized and used as network elements for experimentation purposes. In this paper, Supermicro servers have been the preferred option to execute the distributed LENA processes. These servers are equipped with two Intel Xeon E5-2640v4 processors (20 cores/40 threads running at 2.4 GHz), 64 GB of RAM and 6 Gigabit Ethernet (GigE) ports. Finally, the Cisco C6513 switch router is used to manage the GigE connections and is dynamically reconfigured according to the adopted scenario to build the required fronthaul and backhaul networks.

4.1.2. FPGA-based SoC and RF boards

The Xilinx ZC706 board is used to host all HWA functions, as well as the related L2–L1 SW interfaces (for both HeNB and UE sides). In more detail, the board features the Zynq XC7Z045 all-programmable SoC, which integrates a dual-core ARM Cortex-A9 central processing unit (CPU; clocked at 667 MHz) on the PS side. These are paired with internal memory resources, a dedicated high-speed AXI-based bus and DMA interfaces with the PL, as well as with a set of standard input/output interfaces and peripherals (including Ethernet). The resource utilization metrics of the implemented HWA L1 functions can be observed in Table 1. It should be noted, that these metrics do not account for the HDL firmware utilized by the ZC706 board (i.e., they only account for the designed DSP functionality).

The Analog Devices (AD) AD-FMCOMMS3 board was used as the radio frequency (RF) front-end. It is connected to the ZC706 board through a FPGA Mezzanine Card (FMC) interface and includes the AD9361 RF integrated chip (RFIC). Toward that end, a Linux kernel space application residing in the PS of the Zynq SoC, allows to fully tune and program the AD9361 RFIC. Optionally, power amplifiers, RF band filters and antennas are enabling over-the-air communications.

4.2. Testbed setup and measurement devices

A different number of ZC706 (and AD-FMCOMMS3) boards is used depending on the specific NETCFG that is adopted, as it can be observed in Fig. 8. When a C-RAN architecture is adopted, two FPGA-based platforms are required on the HeNB side (i.e., one for the RRH and another one for the HWA L1), interfaced through a CPRI link (i.e., using coaxial cables as our ideal transport channel), and a third one is used to implement the HWA L1 of the UE. On the contrary, when the L1 functions are executed locally at the HeNB (NETCFG1/2), then only two ZC706 boards are used (i.e., one for the HWA HeNB and another for its UE counterpart). In all cases, GigE links are utilized to provide the required interconnections between the DL signal and the EXTREME® testbed (i.e., sub-ideal and non-ideal transport channels).

In order to accurately assess the energy-saving benefits of the proposed reconfiguration of both WCPs and function splits, specialized HW has also been utilized in order to obtain experimental measurements at different key subsystem elements of the hotspot prototype. Namely the energy consumption of the HeNB has been analysed at the SoC baseband




Fig. 9. HW setup utilized to measure the energy consumption of the RF IC.

processor (i.e., HWA L1), at the RF transceiver IC and at the GigE interfaces. Additionally, the CPU load resulting from the different LENA configurations has also been monitored in order to (indirectly) assess its effect on the energy consumption of the Supermicro servers. In more detail, the servers provide an intelligent platform management interface (IPMI), which enables monitoring the usage of CPU and memory resources, as well as the overall server consumption. Nevertheless, given the very low granularity of this embedded energy measurement HW⁴, the CPU loads resulting from different system configurations have been captured with the objective to complete the presented analysis.

As for the HWA L1, the Xilinx ZC706 board hosts a power system based on the Texas Instruments (TI) UCD90120A power supply sequencer. The latter integrates a 12-bit ADC enabling to monitor up to 12 power-supply voltage lines. Moreover, a power management bus (PMBus) compliant controller is also included. By using a proprietary universal serial bus (USB)-based cable and the TI Fusion Digital Power Designer graphical user interface (GUI), the voltage and current utilized by the baseband SoC can be measured at run-time. In our case, the VCCINT rail has been monitored for the power consumption measurements, taking into account that it is the one powering both the internal logic of the PL and also the PS.

Regarding the RF IC, a custom measurement setup was implemented as it can be observed in Fig. 9. In more detail, two adhoc measurement circuits were specially designed to work with the 3.3 V power rail of the AD9361 chip. A shielded connector block (SCB), based on the National Instruments (NI) SCB-68A device, allows to interconnect the adhoc measurement circuits to a data acquisition (DAQ) card. Specifically, the NI PCI-6289 multifunction DAQ device is used to quantify the measurements. Toward that end, the DAQ is hosted in a general purpose computer through the peripheral component interconnect express (PCIe) bus. The DAQAcquire SW application tool [DAQ] is then in charge of demultiplexing the measurements of the DAQ card and dumping the measured values onto a file, enabling their posterior post-processing. A similar setup was employed to obtain energy measurements at the GigE network interface card (NIC) of the server hosting the HeNB LENA processes.

5. Experimental results and discussion

This section presents the experimental evaluation of the energy-aware hotspot prototype under different function splits and WCP configurations. It must be noted that the aim of the presented results is to prove that important power-saving benefits can be obtained by applying the proposed HeNB reconfigurations. In that sense, we are aware that the power consumption measurements detailed hereafter are strongly dependent on the specific HW elements comprising the prototype. Consequently, the focus of our evaluation is laid on the trade-offs and energy savings observed when comparing the power consumption of the prototype under different system configurations. From our past implementation experiences, we believe that similar results should be obtained (i.e., same order of magnitude) when using different HW setups and ICs.

⁴ The IPMI measurement solution provided the compound power consumption of the Supermicro server (i.e., including not only the CPU, but the hard disk, fans and remaining HW components).

5.1. Methodology

A measurement campaign has been carried out with the objective to characterize the energy savings obtained when employing a reconfigurable hotspot. In that regard, a set of operating scenarios have been defined by modifying the applied NETCFG and the values of the WCPs. Namely, different DL BW configurations, RBG allocation loads, MCS indexes and RF transmit power settings have been considered. For each given scenario a series of power measurements were then performed. In order to procure statistically significant data the time resolution of the DAQ card responsible for gathering the energy measurements was fixed at $1 \mu\text{s}$ (i.e., 1 MHz sampling frequency). Moreover, the data has been captured in uninterrupted sequences of 30 seconds, with several repetitions per experiment. Similar values have been also considered for the measurement solution of the baseband, in spite of its inferior specifications with respect to the time resolution of the samples when compared to the DAQ card.

In all experiments the HeNB has been configured and operating according to the defined scenario, and with the primary objective of setting the focus on a given reconfiguration parameter on each experiment. That is, only a single parameter has been modified at a time, whereas the remaining ones were kept fixed. This has allowed to observe the result of adapting that isolated parameter (e.g., DL BW) with respect to the power consumption of the HeNB. Considering the utilized measurement setup, 30 million samples have been obtained for each experimental iteration. Nevertheless, only selected sets of data have been used in order to bound the cost of the post-processing, without compromising the validity of the analysis. Finally, the evaluation of the results has been performed offline using custom scripts based on R, an open-source SW environment for statistical computing [45]. The obtained curves represent the cumulative distribution function (CDF) of the energy consumed by the HeNB under each implemented scenario.

5.2. Observed energy consumption results

The power consumption of the hotspot prototype is analysed for the different measured subsystems, according to the results obtained after post-processing the captured data for the different experiments.

5.2.1. Ethernet

Several tests have been conducted in order to characterize the power consumption of the GigE NICs of the servers⁵ hosting LENA. Each experiment utilized a different hotspot configuration. For instance, Fig. 10 depicts the consumption observed at the server hosting the L2 and above functionalities (HeNB side) for different DL BW and MCS index settings. As it is presented in the figure, there are no variations in the consumed power due to changes in the traffic load (i.e., MCS index, RBG allocation).

While this behavior was expected for legacy Ethernet devices [46], it is also observed when NICs are equipped with modern power saving features, such as those defined by the energy efficient Ethernet (EEE; which is the case of the Supermicro servers). This is because energy savings are linked to low and bursty data activity, which supports the use of aggregation techniques (i.e., coalescing) [47]. However, this is not the case for the dynamic hotspot traffic. First, the traffic load of the CPRI link is not low in those NETCFGs where L2 and above functions are offloaded onto the Cloud. Moreover, coalescing can introduce jitter, which might result in a disrupted system performance. Therefore, the contribution of the network interfacing HW elements to the energy budget of the dynamic hotspot can be hardly reduced and its energy-aware reconfigurations should target other subsystems.




Fig. 10. Power consumption observed at the GigE NIC for different HeNB configurations.

5.2.2. SW L2 (CPU)

Independently of the particular WCP settings being adopted (i.e., DL BW, MCS index and RBG allocation), the CPU utilization observed for the HeNB process (i.e., L2 and upper-layers) of LENA is always in the range between 25% and 30%, as shown in Fig. 11. This is due to the massive computation capacity of the Supermicro servers (i.e., a single UE is attached to the HeNB; e.g., resulting in the simplest scheduling), as well as to the low measurement granularity provided by the monitoring interface. Accordingly, from the point of view of the PS, it seems that the only strategy to save energy in the HeNB would be to offload the processing of the L2 and above functions to remote or neighboring edge servers, which is actually the case in all considered NETCFGs. In this respect, the energy consumption of a PS can be analysed based on its activity. In more detail, an important fraction of its consumption is related to the usage of its CPU(s) and memories, including their related cooling systems. Regarding the processors, they present an elevated baseline power consumption when in idle state (i.e., the idle power consumption represents an important fraction of its peak consumption when at full load). This consumption increases with every added workload (i.e., process). On top of that, the drained energy depends on its operating frequency [48]. Moreover, the more accesses that these processes require to the memory system, the more energy it will be utilized. Based on that, it can be argued that minimizing the workload of a processor at any given time, will help reducing the overall consumed system energy. Even more, if the workload is sufficiently low (i.e., low number of attached UEs), a reduced clock frequency could be used in the CPU, helping likewise to reduce its energy footprint [49]. This latter fact is empirically verified in the evaluation of the power consumed by the HWA L1.

5.2.3. HWA L1 (FPGA-based SoC)

In the analysis of the power consumption observed at the Zynq XC7Z045 SoC it must be first noted that *NETCFG1* and *NETCFG2* are indistinguishable (i.e., both feature a locally executed HWA L1). Consequently only the power measurements for *NETCFG2* and *NETCFG3* are presented. In detail, the first three experiments assume a locally executed HWA L1, whereas the fourth examines the benefits of moving the DSP computation to the Cloud.

The first experiment analyses the variation of the power consumption at the baseband processor as a function of the utilized DL BW configuration. In this respect, a fixed MCS index of 24 (i.e., highest modulation order) and a fully allocated DLSC (i.e., user-data in all available RBGs) is combined with the four considered DL BW values

⁵ A similar behavior is expected on the local HeNB HW (e.g., Xilinx board), but it was less cumbersome measuring the NIC of a standard server.




Fig. 11. CPU load resulting from the LENA HeNB process when using a 10 MHz DL BW configuration (MCS index 25).

(i.e., 1.4, 5, 10 and 20 MHz). As it can be seen in Fig. 12 downscaling the signal BW can provide important power savings, which scale up to 44% when changing from 20 MHz to 1.4 MHz (e.g., end of the venue, when most attendees leave the hotspot coverage area). The experiments have shown that these gains remain stable over time (Fig. 12b). The reason behind these findings is that the energy-aware RTL design is able to minimize the amount of active logic for each applied DL BW configuration (i.e., selecting a lower BW results in a lower circuit switching activity). Moreover, downscaling the BW also implies that the active logic operates at a lower clock frequency (i.e., further reducing the circuit switching activity).

The impact that different RBG allocations have on the energy consumption of the SoC has been evaluated on the second experiment for two different DL BW settings and MCS indexes. Three different RBG loads are considered, ranging from a very low utilization of the available DL-SCH resources to a fully allocated case. Fig. 13a shows a fixed baseline MCS configuration with index 7 (i.e., lowest modulation order) for the 10 and 20 MHz DL BW cases. Similarly, Fig. 13b adopts a fixed MCS index of 24. In both cases minimal variations in the energy consumption of the HWA L1 are reported (i.e., around 2% in the best case). In this case, thus, the reduction of the circuit switching activity resulting from reconfiguring the hotspot is marginal.

A third experiment investigates the relation between the MCS index and the power consumption observed at the HWA L1. This is done for the 5 MHz and 10 MHz DL BW configurations and with only 2 RBGs allocated. Then, three different MCS index values have been used,

namely 7, 13 and 24 (i.e., ranging from the lowest to the highest modulation order). As it can be observed in Fig. 14, nearly identical results with the previous experiment are reported in this case, with the MCS index having little impact on the energy being drained by the baseband processor.

After analysing the influence of reconfiguring the WCPs on the energy budget of the HWA L1, the fourth experiment is focusing on the dynamic split of communication functions. Specifically, the energy saving gains obtained by moving from NETCFG1/2 to NETCFG3 have been evaluated (i.e., adoption of a C-RAN scheme, where the reconfigurable hotspot acts as a RRH). As reported in Fig. 15, the consumed power was considerably reduced (i.e., up to 46.38%) independently of the utilized DL BW configuration. The reason behind these benefits is that the clock-gated design can be fully exploited in this case. Hence, a considerable amount of energy-hungry DSP logic can be put to an idle state, minimizing likewise the switching activity of the digital circuit; larger gains are obtained for larger BWs, because of the larger amount of logic operating at a faster clock frequency that can be deactivated.

From the point of view of the baseband processor, as it was expected, most energy saving benefits come from reducing the activity of the circuit. Thus, dynamically adapting the WCPs is not always the most efficient way to reduce the energy. More specifically, we have found that the reconfigurations affecting the MCS index or RBG allocation settings should focus on satisfying the frequently changing QoS requirements. In that case, the energy-aware RTL design is only capable of intermittently preventing the unnecessary operation of some parts of the system, which results in a minimal reduction of the power drained by the PL. On the other end, reducing the operating frequency of the circuit results in a notable reduction of the consumed power, exactly as it is expected for the CPU case. Thus, downscaling the DL BW has been proved as an effective means to minimize the energy consumption of the HeNB during those periods where a reduced performance can serve the needs of the attached users. Similarly, distributing the underlying DSP functions across the network can be also exploited to effectively minimize the energy consumption of the HeNB.

5.2.4. RF IC

The power measurements presented for the AD9361 RFIC apply to all three considered NETCFGs (i.e., the RF stage is always active and operates locally, independently of the underlying distribution of functions). As in the HWA L1 case, different settings for the MCS index, DL BW and RBG allocation load have been considered. Additionally, variations in the output power of the RFIC have also been evaluated.

The weight of the DL BW configuration in the power drained by the RF stage is investigated in the first experiment presented here. Fig. 16a shows the results for all four signal BW settings, when using a fixed MCS index of 7 (i.e., QPSK modulation), a complete allocation of the




Fig. 12. Impact of the DL BW configuration on the consumption of the baseband SoC.

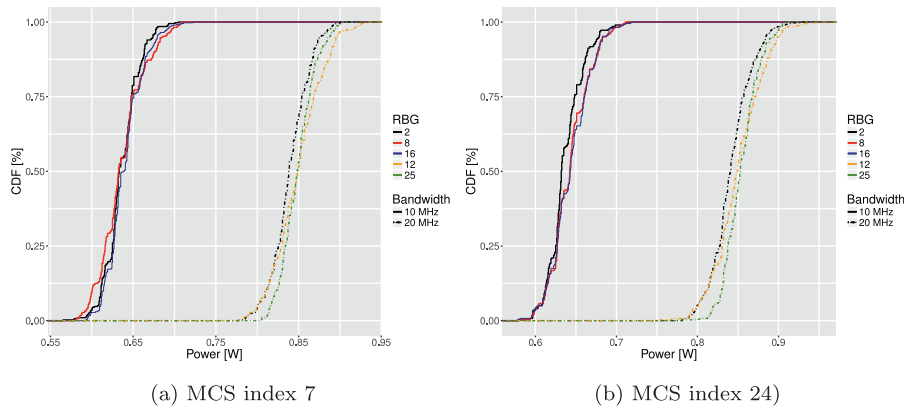


Fig. 13. Impact of the RBG allocation on the energy consumption of the baseband SoC.




Fig. 14. Impact of the MCS index on the power drained by the baseband SoC.

available RBG resources and an RF output power of -19 dBm. Similarly, the output power of the RFIC was attenuated to -39 dBm in Fig. 16b. Analogous and elevated savings are observed in both cases, which grow above 34% when downscaling the DL BW from 20 MHz to 1.4 MHz.

In the second experiment conducted at the RF transceiver IC, the effect of attenuating the RF output power on the power consumption is investigated. As in the previous case, a fixed MCS index of 7 and a fully

allocated DSLCH has been considered. Moreover, three different RF output power settings have been utilized (i.e., -19 , -29 and -39 dBm). In Fig. 17a, the resulting energy consumption for the 1.4 MHz DL BW case is shown, where modest energy savings can be observed (around 6%). Higher gains (i.e., up to 20%) were measured in the 5 MHz BW setting as it can be seen in Fig. 17b.

A third experiment analyses the impact of allocating a different number of RBGs on the power consumption of the RFIC. A fixed MCS index of 7 and output power of -19 dBm was used, for two different DL BW settings (i.e., 10 MHz and 20 MHz). As it can be observed in Fig. 18 the variation of the power consumption resulting from varying the RBG allocation is negligible. The exact same behavior has been found when adapting the MCS index WCP.

As a summary, adapting the MCS index or the RBG allocation does not help in reducing the energy footprint of the RFIC. On the contrary, energy saving gains can be obtained by attenuating the RF output power in those use cases where the link quality allows it. Furthermore, adopting the most efficient DL BW configuration given the actual hotspot requirements is the optimum reconfiguration in power saving terms.

6. Summary and conclusion

Flexibility and reconfigurability are two principal requirements for future 5G systems. Centralized processing in a multi-layered Cloud architecture (i.e., C-RAN, MEC) and the massive deployment of small cells will play a central role toward providing a dynamic network management. Additionally, energy-awareness is essential to ensure the sustainability and economical viability of this heterogeneous network infrastructure. Dynamic hotspots are an important piece of this puzzle, extending the macro network coverage in a flexible and efficient




Fig. 15. Energy savings reported by adopting NETCFG3.




Fig. 16. Variation of the power consumption at the RFIC as a result of changing the DL BW.

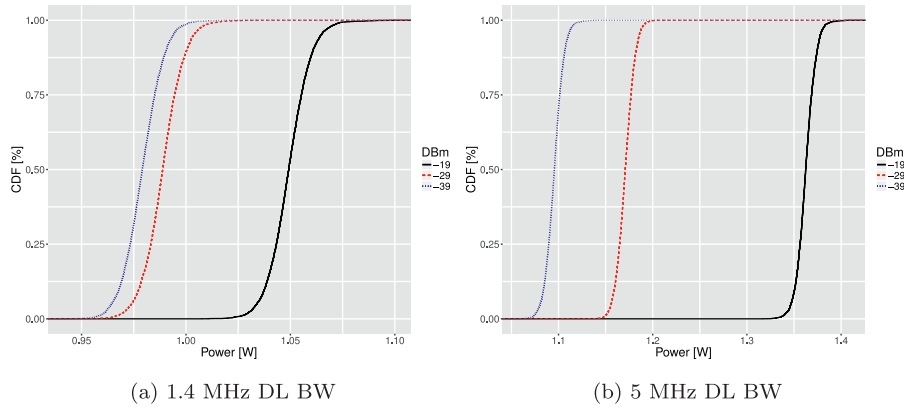


Fig. 17. Variation of the power drained by the RFIC as a result of attenuating the RF output power.




Fig. 18. Energy consumed by the RFIC under different RBG allocation cases.

manner. To that end, both the distribution of communication functions across the network, and the adaptation of the transmission strategy, need to be optimized accounting for the current operative requirements and the minimization of the consumed energy.

This paper has focused on the energy saving benefits that yield from the dynamic reconfiguration of the hotspot. The presented work has a strong applied component and revolves around the development of a real-time hotspot prototype. This has enabled the realistic validation of the proposed concepts. Moreover, actual power consumption measurements of the different subsystems comprising the HeNB have been presented. We believe, thus, that in the paper it has been experimentally demonstrated that important energy saving gains can be obtained with the energy-aware reconfiguration of the HeNB. Part of these gains

are enabled at design time and need to be intelligently exploited taking into account the operative needs of the dynamic hotspot at each moment. In this respect, the energy-aware design of the underlying SW and HWA functions has been detailed in an effort to underline its importance in providing the required flexibility and energy efficiency in future 5G systems.

The analysis of the power measurements has also allowed to select the most interesting reconfiguration strategies in the considered scenario. In more detail, both distributing the communication functions and/or downscaling the utilized signal BW can help to greatly reduce the power consumption of the HeNB (i.e., savings up to 50% at a subsystem level). Hence, the energy-aware reconfiguration of the hotspot might help optimizing its energy footprint during those periods where

the system presents low performance requirements or is in need to save energy (e.g., battery powered HeNB). It has also been observed that adapting the MCS index or RBG load results in minimal energy savings and, thus, should not be used as a main trigger to reconfigure the system (from a power saving perspective). Nevertheless, balancing their configuration according to the requested performance might help to fully minimize the energy consumption of the HeNB in extreme situations (e.g., in a low-battery state). Finally, it is also worth mentioning, that on top of the energy-related objectives, the dynamic reconfiguration of the NETCFG or WCPs adopted by the dynamic hotspot can also help optimizing other KPIs of interest for may 5G use cases (e.g., latency, reliability, throughput).

We genuinely believe that the results reported herein will remain

relevant, and with a similar impact in terms of the presented figures, when different ICs will be employed to serve the 5G NR air interface. Hence, the presented prototype can be seen as a flexible platform that can be extended in many ways to fulfil future reconfigurable network requirements.

Acknowledgments

This work was supported by the European Commission in the framework of the H2020-ICT-2014-2 project Flex5Gware (Grant agreement no. 671563). The work of CTTC was also partially supported by the Generalitat de Catalunya (2017 SGR 891) and by the Spanish Government under project TEC2014-58341-C4-4-R.

Appendix A. Technical details of the system design and implementation

In an effort to ease the readability of the paper, most of the technical aspects regarding the development of the HeNB have been stripped out of the main text. Nevertheless, considering the relevance of the design and implementation decisions to the fulfilment of the flexibility and energy efficiency requirements of the HeNB, this appendix provides its fully detailed description.

A1. L2 and upper-layers

The real-time interaction of LENA with external HW elements (e.g., FPGA-based PHY-layer) has been facilitated by exploiting the ns-3 emulation functions. Specifically, the *RealTime* scheduler has been used to synchronize the events generated by the different network models with the wall clock of the local host (i.e., instead of the virtual simulated clock that would be normally used). As a result, all packet traffic is generated and processed in real-time, according to the timing requirements of the HWA L1 implementation, as well as those related to the currently adopted NETCFG. Moreover, the original LTE interfaces have been updated to work with real IP packets, that can be sent through the network and reinjected to the simulator. In more detail, the *File Descriptor NetDevice* ns-3 functionality has been used to read/write traffic from/to a physical device in the host machine, permitting likewise the implemented LTE protocol stack to send/receive packets/frames to/from external HW components, third party SW and/or other ns-3 processes by means of standard UDP-IP packets. This extension considers the L2–L1, S1–U and S1–MME interfaces.

With respect to those modifications based on the scheduler API, a custom header has been added encapsulating the standard UDP-IP packets to facilitate the interchange of protocol-related information among distributed nodes. Following the principles of the MAC-split defined in the virtual small cells paradigm, the API has been extended to generate time-stamped header frames. These frames contain the relevant primitives that need to be exchanged with the HWA L1 (i.e., as a serialized bit sequence), enabling likewise its reconfiguration through the L2–L1 interface. The changes applied to the EPC enable a number of primitives defined by the LTE standard, including the procedure to attach new users, create new connections and transmit data and control packets over them.

Concerning the scheduler functionalities, the baseline round robin implementation has been updated to include specific constraints originated at the HWA implementation of the downlink (DL) L1. In more detail, natively LENA does not consider any limitation regarding the number of UEs that can be simultaneously allocated in the DL at any given moment, whereas in reality the physical resources that can be dedicated to convey the DL control information (DCI) is limited and depends on the adopted WCPs (i.e., each DL BW configuration uses a different number of REs). In the same way, the size of the allocated TBs takes into account the specifications of the channel encoder implementation provided by the FPGA-based PHY-layer.

A2. L2–L1 interface

The communication of the L2–L1 interface and the HWA L1 is controlled by an AXI-based direct memory access (DMA) core. The DMA block is connected to a high-performance and high-bandwidth port, AXI4-HP, on the PS side. On the PL end, the DMA is connected to an AXI4-Stream interfacing block, which provides a simplified low-latency master-slave communication mechanism. By using this specific interfacing architecture, the PS-PL communication supports an exchange of data up to 6400 Mbits/s.⁶

Regarding the two components shown in Fig. 6, the user-space parsing process will produce a set of 32-bit words that will be forwarded to the HWA L1 through the AXI bus. Consequently, this second thread also manages the required intermediate buffering. An effective implementation of the concurrent queue, combined with the use of modern C++ programming features (e.g. move semantics), enables to minimize the time required by the PS to complete this task and guarantees that no incoming data will be queued for prolonged periods, even if the current system configuration requires an elevated data-rate.

A custom driver executed in the kernel-space efficiently manages the interaction with the different HW elements. It is in charge of configuring the DMA core and HWA L1 according to the selected hotspot configuration, by programming a set of memory-mapped registers. Moreover, the driver is also responsible for setting up the interrupt service routines that allow the HWA L1 to request new L2–L1 interfacing frames to the PS. The real-time communication has been made feasible by implementing a random access memory (RAM)-based buffer that acts as a jitter-absorbing first in first out (FIFO) memory in the PL side; a small number of L2–L1 interfacing frames will be initially stored internally at the FPGA, accounting for the network delays. The memory controller associated to this embedded buffer implements a simple control mechanism which is in charge of generating interrupts to the PS when necessary (i.e., by controlling the FIFO contents) and forwarding the stored L2–L1 interfacing frames to the HWA PHY-layer (i.e., basic control of the underlying memory read and write operations). Similarly, the kernel driver uses a ring-buffer memory element mapped onto the user-space, where the associated thread pushes the 32-bit words generated by the parser. A zero-copy design has been implemented to avoid

⁶ This data-rate calculation accounts for the 32-bit AXI-4 bus and the specific operating frequencies of the embedded PS (i.e., 666.7 MHz) and AXI logic within the PL (i.e., 200 MHz).

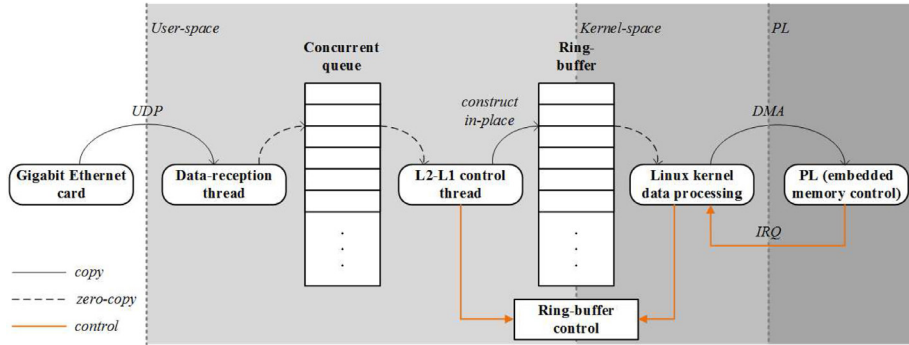


Fig. A19. PS to PL communication and data processing flow.

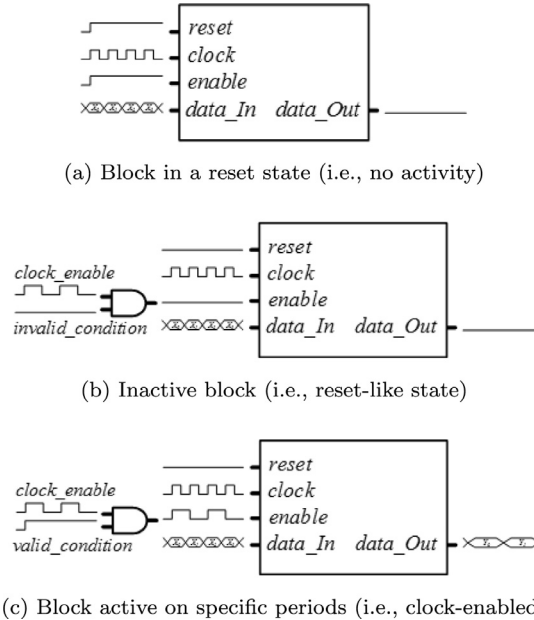


Fig. A20. Control the activity of the DSP stages through clock-gating techniques in order to save energy.

degrading the performance of the PS during the interactions between the user-space and kernel-space code. Like this, unnecessary data copying and context switching operations are avoided. Additionally the interrupt service routine is also used to configure the DMA and start a new transaction when required, by transferring the data that was last fetched from the ring-buffer descriptor to the PL buffer. Finally, when the transaction is completed the driver updates the ring-buffer pointers accordingly. A detailed diagram of the communication and data processing flow taking place inside the L2–L1 interface is provided in Fig. A.19.

A3. Energy-aware HWA L1

In order to deal with the bit-intensive computation resulting from the adaptive real-time operation of the HWA L1, the RTL architecture is articulated around two differentiated clock domains (see Fig. 7). Apparently, the first one is directly related to the sampling frequency of the digital-to-analog conversion (DAC) circuitry (i.e., derived from the principal LTE sampling frequency, $\frac{1}{N} \cdot 30.72$ MHz, with $N = [1, \dots, 16]$). The second clock domain needs to be aligned with the PS from which data is received through the dedicated high-speed AXI-4 streaming interface. Considering the specifications of the channel coding stage, the most convenient is to use a clock with a frequency, M , at least six times the one derived from the DAC (i.e., $M \geq 6 \cdot \frac{1}{N} \cdot 30.72$ MHz). That is accounting for the worst possible case, when the MCS index ≥ 17 , which requires to work with groups of six bits (i.e., 64-QAM) for each allocated RE. In that case, since we do consider a 20 MHz DL BW, six bits will then need to be sequentially fed to the channel encoder IP-core during each 30.72 MHz clock cycle (i.e., $N = 1$ and $M \geq 184.32$ MHz). Using a fixed 200 MHz⁷ clock, thus, allows processing one RE per slow-clock cycle independently of the current system configuration (i.e., DL BW and MCS). Furthermore, it accelerates the computation of the channel encoding operations resolving likewise all the associated latency issues.

As detailed, beyond reconfiguring and driving the operation of the PHY-layer, the two control units are also in charge of controlling all synchronous communications among them. This demands to efficiently exploit the limited time budget while accounting for the characteristics of cross-clock domain data sharing, for which complex memory structures have been built. More importantly, the control units provide a flexible control of the logic which is combined with clock-gating techniques to attain an energy-aware design. In more detail, each (synchronous) component/process of the digital circuit will be active on the rising edge of the clock directing its operation, only when the related clock-enable signal is high. On top of

⁷ This specific frequency value has been selected by considering the specifications of the PS clock subsystem of the target SoC device.

that, the clock-enable signal is combined with other control signals which account for the current configuration of the HeNB (e.g., selected DL BW, allocated RBGs or MCS values), in order to minimize the amount of logic which is actively utilized at each moment, keeping the rest in a reset-like state, as shown in Fig. A.20. At a system level, all those DSP components related to the turbo encoding stage will remain completely inactive during the generation of those subframes where no user data is being transmitted. That is, only those stages related to the parsing and encoding of the PDCCH or PBCH, as well as the ones generating the synchronization or reference signals, will present activity during the generation of each subframe comprising the transmitted signal (i.e., there are certain REs that are always occupied even when no DL user data is allocated at all).

An illustrative example of the low-level optimizations implemented across the RTL hierarchy is found in the management of the internal storage: whereas it is dimensioned to support the highest user data capacity requirements defined by the LTE standard (i.e., all RBGs allocated to a single user, maximum MCS value, 20 MHz DL BW), only the minimum required number of registers/memory elements are actively used; the others will remain in an idle/reset state. A similar approach has been used for the FIFOs, and related control-logic, that are utilized to pass the encoded user data, generated in the AXI-4 clock domain, to the symbol mapping stage residing in the DAC clock domain. In that case, the logic with the faster operating frequency is only active each $\frac{M}{6}$ clock cycles, where M depends on the MCS value (being $M=6$ the worst case, for $MCS \geq 17$).

An added yet substantial benefit of using clock-gating techniques is that the use of the frequency synthesizer primitives found in the clock management (CM) tiles of the FPGA devices is avoided. CM conveniently allows deriving the required clock signals from those received by the FPGA (i.e., from an external source), providing highly precise phase and jitter specifications, but also incurring in a nonnegligible increase in the consumed energy. According to estimations that we obtained by using the *Xilinx power analyzer* (XPA) SW tool, the increase in the consumed energy derived from the use of CMs can be as high as 33.44%. Specifically, we compared a simplified version of the clock-enabled HeNB design presented above (i.e., not interfaced with the L2, but using a fixed DL allocation) with a second implementation that was modified to exploit the features provided by the CM tiles (i.e., different clock signals were generated internally instead of using clock-enable signals). Precise energy consumption estimates for each implementation were finally obtained by loading the placed and routed (PAR) designs to the XPA, jointly with the related post-PAR signal-toggle activity files (i.e., realistic modelling of the timing delays of the implemented circuit).

Appendix B. Acronyms

5G	fifth generation (wireless communications systems)
AD	Analog Devices
ADC	analog-to-digital converter
API	application programming interface
AXI	advanced extensible interface
BBU	baseband unit
BW	bandwidth
C-RAN	cloud RAN
CDF	cumulative distribution function
CM	clock management
CPRI	common public radio interface
CPU	central processing unit
CRC	cyclic redundancy check
DAC	digital-to-analog conversion
DAQ	data acquisition
DCI	DL control information
DL	downlink
DLSCH	DL shared channel
DMA	direct memory access
DSP	digital signal processing
eCPRI	CPRI evolution (for 5G)
EEE	energy efficient Ethernet
eNB	evolved node B
EPC	evolved packet core
FIFO	first in first out
FFT	fast Fourier transform
FMC	FPGA Mezzanine Card
FPGA	field programmable gate array
GigE	Gigabit Ethernet
GUI	graphical user interface
HARQ	hybrid automatic repeat request
HDL	HW description language
HeNB	home eNB
HW	hardware
HWA	HW accelerated
IP	internet protocol
IP-core	intellectual property core
IPMI	intelligent platform management interface
KPI	key performance indicator
LENA	LTE-EPC network simulator
LTE	long term evolution
MAC	medium access control (layer)

MCS	modulation and coding scheme
MEC	multi-access edge computing
MIB	master information block
MME	mobility management entity
NETCFG	network configuration
NFV	network function virtualization
NGFI	next generation fronthaul interface
NI	National Instruments
NIC	network interface card
NR	new radio
PAR	placed and routed
PCIe	peripheral component interconnect express
PGW	packet data network gateway
PHY	physical (layer)
PL	programmable logic
PMB	power management bus
PS	processing system
QoS	quality of service
RAM	random access memory
RAMB	RAM block
RAN	radio access network
RB	resource block
RBG	RB group
RE	resource element
RF	radio frequency
RFIC	RF integrated chip
RRH	remote radio head
RTL	register transfer level
SCB	shielded connector block
SDN	SW defined networking
SDR	SW defined radio
SGW	serving gateway
SIB	system information block
SoC	system-on-chip
SW	software
TB	transport block
TCP	transport control protocol
TDD	time division duplex
TI	Texas Instruments
UDP	user datagram protocol
UE	user equipment
UL	uplink
USB	universal serial bus
WCP	wireless communication parameter
XPA	Xilinx power analyzer

References

- [1] Cisco Visual Networking Index: Global Mobile Data Traffic Forecast Update, 2016–2021 Technical Report, 2017.
- [2] P. Rost, C.J. Bernardos, A.D. Domenico, M.D. Girolamo, M. Lalam, A. Maeder, D. Sabella, D. Wübben, Cloud technologies for flexible 5G radio access networks, *IEEE Commun. Mag.* 52 (5) (2014) 68–76, <http://dx.doi.org/10.1109/MCOM.2014.6898939>.
- [3] R.L.G. Cavalcante, S. Stanczak, M. Schubert, A. Eisenblaetter, U. Tuerke, Toward energy-efficient 5G wireless communications technologies: tools for decoupling the scaling of networks from the growth of operating power, *IEEE Signal Process. Mag.* 31 (6) (2014) 24–34, <http://dx.doi.org/10.1109/MSP.2014.2335093>.
- [4] S. Buzzi, C.L. I, T.E. Klein, H.V. Poor, A.Z. C. Yang, A survey of energy-efficient techniques for 5G networks and challenges ahead, *IEEE J. Sel. Areas Commun.* 34 (4) (2016) 697–709, <http://dx.doi.org/10.1109/JSAC.2016.2550338>.
- [5] R. Mahapatra, Y. Nijssure, G. Kaddoum, N.U. Hassan, C. Yuen, Energy efficiency tradeoff mechanism towards wireless green communication: a survey, *IEEE Commun. Surv. Tutor.* 18 (1) (2016) 686–705, <http://dx.doi.org/10.1109/COMST.2015.2490540>.
- [6] K. Davaslioglu, E. Ayanoglu, Quantifying potential energy efficiency gain in green cellular wireless networks, *IEEE Commun. Surv. Tutor.* 16 (4) (2014) 2065–2091, <http://dx.doi.org/10.1109/COMST.2014.2322951>.
- [7] X. Chen, J. Wu, Y. Cai, H. Zhang, T. Chen, Energy-Efficiency oriented traffic offloading in wireless networks: a brief survey and a learning approach for heterogeneous cellular networks, *IEEE J. Sel. Areas Commun.* 33 (4) (2015) 627–640, <http://dx.doi.org/10.1109/JSAC.2015.2393496>.
- [8] D.A. Temesgene, J.N. nez Martínez, P. Dini, Softwareization and optimization for sustainable future mobile networks: a survey, *IEEE Access* 5 (2017) 25421–25436, <http://dx.doi.org/10.1109/ACCESS.2017.2771938>.
- [9] H. Zhang, N. Liu, X. Chu, K. Long, A.H. Aghvami, V.C.M. Leung, Network slicing based 5G and future mobile networks: mobility, resource management, and challenges, *IEEE Commun. Mag.* 55 (8) (2017) 138–145, <http://dx.doi.org/10.1109/MCOM.2017.1600940>.
- [10] R. Trivisonno, X. An, Q. Wei, Network slicing for 5G systems: a review from an architecture and standardization perspective, *Proceedings of the 2017 IEEE Conference on Standards for Communications and Networking (CSCN)*, (2017), pp. 36–41, <http://dx.doi.org/10.1109/CSCN.2017.8088595>.
- [11] D. Sabella, A. de Domenico, E. Katranaras, M.A. Imran, M. di Girolamo, U. Salim, M. Lalam, K. Samdanis, A. Maeder, Energy efficiency benefits of RAN-as-a-service concept for a cloud-Based 5G mobile network infrastructure, *IEEE Access* 2 (2014) 2169–3536, <http://dx.doi.org/10.1109/ACCESS.2014.2381215>.
- [12] K. Zhang, Y. Mao, S. Leng, Q. Zhao, L. Li, X. Peng, L. Pan, S. Maharjan, Y. Zhang, Energy-Efficient offloading for mobile edge computing in 5G heterogeneous networks, *IEEE Access* 4 (2016) 5896–5907, <http://dx.doi.org/10.1109/ACCESS.2016.2597169>.
- [13] O. Bulakci, Z. Ren, C. Zhou, J. Eichinger, P. Fertl, D. Gosalvez-Serrano, S. Stanczak, Towards flexible network deployment in 5G: nomadic node enhancement to

- heterogeneous networks, Proceedings of the 2015 IEEE International Conference on Communication Workshop (ICCWC), (2015), pp. 2572–2577, <http://dx.doi.org/10.1109/ICCWC.2015.7247565>.
- [14] E. Ternon, P. Agyapong, L. Hu, A. Dekorsy, Energy savings in heterogeneous networks with clustered small cell deployments, Proceedings of the 2014 11th International Symposium on Wireless Communication Systems (ISWCS), (2014), pp. 126–130, <http://dx.doi.org/10.1109/ISWCS.2014.6933333>.
- [15] G. Mountaser, M.L. Rosas, T. Mahmoodi, M. Dohler, On the feasibility of MAC and PHY Split in Cloud RAN, Proceedings of the 2017 IEEE Wireless Communications and Networking Conference (WCNC), (2017), pp. 1–6, <http://dx.doi.org/10.1109/WCNC.2017.7925770>.
- [16] S.S. Kumar, R. Knopp, N. Nikaiein, D. Mishra, B.R. Tamma, A.A. Franklin, K. Kuchi, R. Gupta, FLEXCRAN: cloud radio access network prototype using OpenAirInterface, Proceedings of the 2017 9th International Conference on Communication Systems and Networks (COMSNETS), (2017), pp. 421–422, <http://dx.doi.org/10.1109/COMSNETS.2017.7945423>.
- [17] E. Lähtekangas, K. Pajukoski, J. Vihriälä, G. Berardinelli, M. Lauridsen, E. Tiirola, P. Mogensen, Achieving low latency and energy consumption by 5G TDD mode optimization, Proceedings of the 2014 IEEE International Conference on Communications Workshops (ICC), (2014), pp. 1–6, <http://dx.doi.org/10.1109/ICCWC.2014.6881163>.
- [18] K. Sundaresan, M.Y. Arslan, S. Singh, S. Rangarajan, S.V. Krishnamurthy, Fluidnet: a flexible cloud-based radio access network for small cells, IEEE/ACM Trans. Netw. 24 (2) (2016) 915–928, <http://dx.doi.org/10.1109/TNET.2015.2419979>.
- [19] N. Saxena, A. Roy, H. Kim, Traffic-aware cloud RAN: a key for green 5G networks, IEEE J. Sel. Areas Commun. 34 (4) (2016) 1010–1021, <http://dx.doi.org/10.1109/JSAC.2016.2549438>.
- [20] C. Bluemm, Y. Zhang, P. Alvarez, M. Ruffini, L.A. DaSilva, Dynamic energy savings in Cloud-RAN: an experimental assessment and implementation, Proceedings of the 2017 IEEE International Conference on Communications Workshops (ICC Workshops), (2017), pp. 791–796, <http://dx.doi.org/10.1109/ICCWC.2017.7962755>.
- [21] A. Sadek, H. Mostafa, A. Nassar, Y. Ismail, Towards the implementation of multi-band multi-standard software-defined radio using dynamic partial reconfiguration, Int. J. Commun. Syst. 2017 (2017) 1099–1131, <http://dx.doi.org/10.1002/dac.3342>.
- [22] M. Braun, J. Pendium, M. Ettus, RFNoC: RF network-on-chip, Proceedings of the GNU Radio Conference 1, 1 (2016). <https://pubs.gnuradio.org/index.php/grcon/article/view/3>.
- [23] J. Paulo, I. Freire, I. Sousa, C. Lu, M. Berg, I. Almeida, A. Klautau, FPGA-based testbed for synchronization on Ethernet fronthaul with phase noise measurements, Proceedings of the 2016 1st International Symposium on Instrumentation Systems, Circuits and Transducers (INSCIT), (2016), pp. 132–136, <http://dx.doi.org/10.1109/INSCIT.2016.7598202>.
- [24] A.F. Beldachi, J.L. Nunez-Yanez, Accurate power control and monitoring in ZYNQ boards, Proceedings of the 2014 24th International Conference on Field Programmable Logic and Applications (FPL), (2014), pp. 1–4, <http://dx.doi.org/10.1109/FPL.2014.6927415>.
- [25] M. Shafique, L. Bauer, J. Henkel, Adaptive energy management for dynamically reconfigurable processors, IEEE Trans. Comput.-Aided Des. Integr. Circuits Syst. 33 (1) (2014) 50–63, <http://dx.doi.org/10.1109/TCAD.2013.2282265>.
- [26] C. Ravishankar, J.H. Anderson, A. Kennings, FPGA power reduction by guarded evaluation considering logic architecture, IEEE Trans. Comput.-Aided Des. Integr. Circuits Syst. 31 (9) (2012) 1305–1318, <http://dx.doi.org/10.1109/TCAD.2012.2192478>.
- [27] R. Marlow, C. Dobson, P. Athanas, An enhanced and embedded GNU radio flow, Proceedings of the 2014 24th International Conference on Field Programmable Logic and Applications (FPL), (2014), pp. 1–4, <http://dx.doi.org/10.1109/FPL.2014.6927427>.
- [28] J. van de Belt, P.D. Sutton, L.E. Doyle, Accelerating software radio: Iris on the Zynq SoC, Proceedings of the 2013 IFIP/IEEE 21st International Conference on Very Large Scale Integration (VLSI-SoC), (2013), pp. 294–295, <http://dx.doi.org/10.1109/VLSI-SoC.2013.6673295>.
- [29] H. Zeng, X. Liu, S. Megeed, N. Chand, F. Effenberger, Real-time demonstration of CPRI-compatible efficient mobile fronthaul using FPGA, J. Lightwave Technol. 35 (6) (2017) 1241–1247, <http://dx.doi.org/10.1109/JLT.2017.2660484>.
- [30] D. Riscado, J. Santos, D. Dinis, G. Anjos, D. Belo, N.B. Carvalho, A.S.R. Oliveira, A flexible research testbed for C-RAN, Proceedings of the 2015 Euromicro Conference on Digital System Design, (2015), pp. 131–138, <http://dx.doi.org/10.1109/DSD.2015.72>.
- [31] B. Guan, X. Huang, G. Wu, C. Chan, M. Udayan, C. Neelam, A pooling prototype for the LTE MAC layer based on a GPP platform, Proceedings of the 2015 IEEE Global Communications Conference (GLOBECOM), (2015), pp. 1–7, <http://dx.doi.org/10.1109/GLOCOM.2015.7417473>.
- [32] D. Zeller, M. Olsson, O. Blume, A. Fehske, D. Ferling, W. Tomaselli, I. Gódor, A. Galis, A. Gavras, Sustainable Wireless Broadband Access to the Future Internet – The EARTH Project, Springer, Berlin, Heidelberg, pp. 249–271. doi:10.1007/978-3-642-38082-2_21.
- [33] Small Cell Virtualization Functional Splits and Use Cases, Small Cell Forum Release 5.1 (159.05.1.01), Technical Report, 2015.
- [34] 3GPP, Technical Specification Group Radio Access Network, Study on New Radio Access Technology, Radio Access Architecture and Interfaces (Release 14), 3GPP TR 38.801 V1.0.0, Technical Report, (2016).
- [35] Virtualization for small cells: Overview, Small Cell Forum White Paper (106.06.01), Technical Report, 2015.
- [36] Common Public Radio Interface eCPRI Interface Specification V1.0 (2017–08–22), Technical Report, 2017.
- [37] N. Baldo, M. Miozzo, M. Requena-Esteso, J. Nin-Guerrero, An open source product-oriented LTE network simulator based on Ns-3, Proceedings of the 14th ACM International Conference on Modeling, Analysis and Simulation of Wireless and Mobile Systems (MSWiM), (2011), pp. 293–298, <http://dx.doi.org/10.1145/2068897.2068948>.
- [38] The Network Simulator – ns-3. Available on-line: <http://www.nsnam.org>. [Accessed: June- 2018].
- [39] R. Martínez, A. Mayoral, M. Requena-Esteso, N. Baldo, R. Vilalta, R. Casellas, M. Miozzo, J.M. R. Muñoz, Application of SDN-based orchestration for the automated deployment of fixed and mobile convergent services in future 5G networks, Proceedings of the 1st International Workshop on Elastic Networks Design and Optimisation (ELASTICNETS), (2016), <http://dx.doi.org/10.5281/zenodo.438938>.
- [40] The Linux Foundation, The Real Time Linux Collaborative Project. Available on-line: <https://wiki.linuxfoundation.org/realtime/start>. [Accessed: June- 2018].
- [41] Reducing Switching Power with Intelligent Clock Gating, Xilinx White Paper (WP370), Technical Report, 2013.
- [42] H. Blasinski, F. Amiel, T. Ea, Impact of different power reduction techniques at architectural level on modern FPGAs, Proceedings of the 1st IEEE Latin American Symposium on Circuits and Systems (LASCAS), (2010).
- [43] P. Jamieson, W. Luk, S.J.E. Wilton, G.A. Constantinides, A flexible research testbed for C-RAN, Proceedings of the 2009 19th International Conference on Field Programmable Logic and Applications (FPL), (2009), pp. 324–327, <http://dx.doi.org/10.1109/FPT.2009.5377675>.
- [44] EXTREME Testbed. Available on-line: http://networks.cttc.cat/mobile-networks/extreme_testbed/. [Accessed: June-2018].
- [45] The R Project. Available on-line: <https://www.r-project.org/>. [Accessed: June-2018].
- [46] J.A. Aroca, A. Chatzipapas, A.F. Anta, V. Mancuso, A measurement-based characterization of the energy consumption in data center servers, IEEE J. Sel. Areas Commun. 33 (12) (2015) 2863–2877, <http://dx.doi.org/10.1109/JSAC.2015.2481198>.
- [47] A. Chatzipapas, V. Mancuso, An m/g/1 model for gigabit energy efficient ethernet links with coalescing and real-trace-based evaluation, IEEE/ACM Trans. Netw. 24 (5) (2016) 2663–2675, <http://dx.doi.org/10.1109/TNET.2015.2477090>.
- [48] J. Phung, Y.C. Lee, A.Y. Zomaya, Application-agnostic power monitoring in virtualized environments, Proceedings of the 2017 17th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing (CCGRID), (2017), pp. 335–344, <http://dx.doi.org/10.1109/CCGRID.2017.100>.
- [49] S. Leibson, Reduce SOC energy consumption through processor ISA extension, Proceedings of the 2007 International Symposium on System-on-Chip (SoC), (2010), pp. 1–4, <http://dx.doi.org/10.1109/ISSOC.2007.4427427>.