



HAL
open science

Assessing the energetical cost of 5G softwarization

Aude Jean-Baptiste, Philippe Owezarski, Pascal Berthou, Isabelle Silvain

► To cite this version:

Aude Jean-Baptiste, Philippe Owezarski, Pascal Berthou, Isabelle Silvain. Assessing the energetical cost of 5G softwarization. International Symposium on Local and Metropolitan Area Networks (LANMAN'2024), IEEE, Jul 2024, Boston (MA), United States. hal-04602752

HAL Id: hal-04602752

<https://laas.hal.science/hal-04602752>

Submitted on 5 Jun 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Assessing the energetical cost of 5G softwarization

1st Aude Jean-Baptiste
LAAS-CNRS, Université de Toulouse, CNRS
Toulouse, France
audejeanbaptist@gmail.com

2nd Philippe Owezarski
LAAS-CNRS, Université de Toulouse, CNRS
Toulouse, France
owe@laas.fr

3rd Pascal Berthou
LAAS-CNRS, Université de Toulouse, UPS
Toulouse, France
berthou@laas.fr

4th Isabelle Silvain
LAAS-CNRS, Université de Toulouse, CNRS
Toulouse, France
isilvain@laas.fr

Abstract—5G is a new key technology for future communication networks. It aims at providing a broad range of new services and capabilities for users as well as facilitating its management for network operators (NetOps). Whereas networks still have monolithic architectures, 5G design takes advantage of softwarization and virtualization of its functionalities, for this purpose. However, 5G is also raising a lot of critics, especially related to its energy consumption. This paper then deals with assessing the energetic cost of the future softwarized 5G facilities. To this aim, an experimental platform has been set-up taking advantage of the software 5G OpenAirInterface (OAI) implementation. This paper shows the complexity of designing energy consumption measurement tools. It then exhibits the level of energy consumption of the main 5G components, pointing out the ones that need to be optimized.

Index Terms—5G, energy consumption, OpenAirInterface, energy measurement

I. INTRODUCTION

5G is aimed at providing extended communication capacities as well as differentiated and guaranteed services thanks to slicing mechanisms. But 5G also has to provide more operational flexibility for the deployment of new users' services, or for facilitating network management by network operators (NetOps). This flexibility is made possible thanks to the network softwarization concept, i.e. using exclusively software components, easily deployable and remotely manageable, without any action on the installed hardware for adding new services or changing the network configuration. 5G takes advantage of the recent trend for softwarizing networks, led by the SDN (Software Defined Network) paradigm, and by the virtualization of their functions (NFV: Network Function Virtualization). Software defined Radio (SDR) techniques are then part of the premiere technologies selected for developing the new services and functionalities to be provided by 5G, and notably OpenAirInterface (OAI), an open-source software platform for 4G and 5G networks. This platform development is still ongoing, but it nowadays provides the most advanced

and successful software platform for developing and experimenting 5G services. Its frequent upgrades moreover represent a sustainability warranty for engineers and researchers that work on the design and development of new high speed cellular communication networks, as 5G.

Having being long involved on 5G research and experimentation, LAAS-CNRS has installed since 2020 an experimental 5G platform based on OAI [1]. Based on two years of experimentations, it appears that the performances of software implementations of 5G modules are limited, despite the use of very powerful USRP and computing servers. Their computing load is always very high. Indeed, with such a full software approach, 5G infers a huge number of operations to be computed by software, whereas they were computed on a dedicated hardware device in previous cellular technologies. This induces an energy consumption level on the computing servers that seems to follow the same evolution as the computing load.

The work presented in this paper then deals with a first attempt for assessing the energy consumption of a software-based implementation of 5G, as it should be for future deployments of 5G facilities by NetOps. The ultimate goal would be to measure the consumption of any single component of the new 5G softwarized architecture, including mandatorily remaining hardware components, notably present on SDR devices. To this end, the work deals with measuring and assessing on our 5G OAI experimental platform the actual energy consumption by both software and hardware components at the different steps of communications. It is then expected to isolate components with high energy consumption costs. The experimental 5G platform consists of USRP devices for implementing SDR functionalities, and computing servers running OAI software modules. Software modules have been installed, debugged and configured during the two last years for providing nowadays a running 5G platform. The development of OAI is indeed not complete, and many modules are still at a beta version development level, and bugs and stability problems still need to be fixed. The experimental platform has also been instrumented with equipment for measuring the electrical energy consumption, both by using physical power meters as well as software tools that take advantage of some

This work got the support of INS2I institute of CNRS as well as from the PEPR 5G program of ANR in the framework of the France 2030 investment plan, specially for NF-MUST and NF-HiSec projects under grants ANR-22-PEFT-0002 and ANR-22-PEFT-0009 respectively

electrical power information collected by processors.

The remainder of the paper is as follows: section II details the experimental OAI platform and energy consumption measurement tools and devices. Section III presents and discusses the results of the obtained measurement with the different experimental scenarios tested. Section IV then presents some related works. Finally, section V concludes the paper by summarizing the results and discussing their quality and certainty, and finishes by presenting future works.

II. DESIGN OF THE 5G EXPERIMENTAL PLATFORM

A. The OAI 5G network deployment

As previously said, LAAS-CNRS has installed since 2020 a 5G OAI experimental platform for testing and validating its researches in the area of future networks. This section gives a short description of this platform. Interested readers can refer to a complete description on [1]. The selection of the equipment for this platform is directly driven by the OAI Software Alliance [2] that develops the OAI 5G version we are using.

The platform currently deployed allows researchers to emulate one gNB (base station) and one UE (User Equipment). It consists of the following equipment (as depicted on 1):

- gNB and CN (Core Network) are installed on a first set consisting of a server and an USRP:
 - The server is a Dell Precision Rack 7920 with two Intel Xeon Gold 6240 processors (18-cores, 2.6-3.9 GHz Turbo, 24.75 MBytes of cache), 64 GBytes of RAM (2933 MHz DDR4 ECC)
 - The USRP is an Ettus X310 connected to the server thanks to a 10Gbit/s Ethernet link. The USRP X310 integrates a KINTEX7-410T FPGA and a UBX-160 extension board (10 MHz - 6 GHz, bandwidth of 160 MHz)
 - The USRP is connected to two wide band log-periodic antennas LP0965 (850 MHz – 6.5 GHz)
- The User Equipment (UE) is installed on a second set consisting of a server and an USRP (UE-SDR):
 - The server is a Dell Precision Rack 7920 with two Intel Xeon Gold 6240 processors (18-cores, 2.6-3.9 GHz Turbo, 24.75 MBytes of cache), 64 GBytes of RAM (2933 MHz DDR4 ECC)
 - The USRP is an Ettus X310 connected to the server thanks to a 10Gbit/s Ethernet link. The USRP X310 integrates a KINTEX7-410T FPGA and a UBX-160 extension board (10 MHz - 6 GHz, bandwidth of 160 MHz)
 - The USRP is connected to two wide band log-periodic antennas LP0965 (850 MHz – 6.5 GHz)

The operating system running on the servers has the following features:

- Ubuntu 20.04 LTS
- Lowlatency kernel




Fig. 1. The LAAS-CNRS 5G experimental platform

- C-states and P-states disabled (except C0 and C1 which allow the cores to run at full speed but also to be able to stop in case of interruption)
- CPU hyperthreading disabled
- Each core configured as Governor Performance
- Cores are manually specialized according to the tasks they execute. It requires to install the i7z package for returning detailed information about the operation states.

The software part takes advantage of the OpenAirInterface 5G facilities in the 4.0.0.0 version:

- Core: <https://gitlab.eurecom.fr/oai/cn5g/oai-cn5g-fed.git>
- RAN: <https://gitlab.eurecom.fr/oai/openairinterface5g.git>

First quality of experience and performance results on this platform are presented in [1]. The first results and feedbacks with this platform in terms of pros and cons are summarized as follows:

1) Cons:

- At this time, the platform is unstable. Many experimental scenarios crash at least for some periods (as it will be exhibited by some experiments described in the evaluation section). It is then difficult to get a constant behavior of the platform for durations of several minutes.
- The communication performances are around 20 MBit/s downlink and 4 MBit/s uplink. The configuration has certainly to be optimized, but such limited performances might also question about the software nature of the platform in face of the servers and USRP computing capacities.

2) Pros:

- The OAI Software Alliance is very active, and software upgrades are regularly available. They contain new functionalities and correct the bugs that users signaled.
- The platform works similarly with 5G smartphones of the market as with SDR UE.
- It is possible to consider many possible platform configurations and experimental scenarios as the platform is fully open.

B. Energy measurement facilities selection and evaluation

Our objective being to measure the energy consumption of any function, process or container of the OAI architecture, we then looked at software tools devoted to this task. These tools read the information about energy consumption from processors. Our servers are equipped with Intel processors providing the RAPL interface (Running Average Power Limit) [3] [4] [5] for that purpose. RAPL provides estimates on energy consumption for:

- The full computer
- The cores of the processors
- The DRAM (in server processors)

There are several tools using the RAPL Intel interface for providing users with energy consumption estimates. Scaphandre [6] is one of them. We selected it among the others just because it is free, and we liked its interface and easiness for recovering the energy consumption data. In addition, it has been evaluated and exhibited very good accuracy in power consumption measurement [7]. We then run some experiments on the OAI 5G platform to confirm the Scaphandre accuracy. For that, we use a physical Fluke Norma 4000 power meter [8] as the ground truth. The Norma 4000 is installed directly on the electrical socket on the wall and is able to measure the energy consumption of the equipment that is plugged on this electrical socket. As an industrial measurement device, the Norma 4000 power meter is certified, i.e. its accuracy is guaranteed, with a known maximum error margin. The error margin being very low with the Norma 4000 (1 microwatt max), it is then used as a ground truth in the following.

The scenario used is very basic. It just consists in making the UE send constant UDP traffic to the gNB at a rate of 5 Mbps (using IPerf). The energy consumption of the gNB process measured by Scaphandre is 500 W. The Norma 4000 at the same time indicates that the full server running the gNB and the CN consumes 280W. The Scaphandre estimation is then completely unrealistic, and confirmed by all measurement led on several scenarios.

This completely unexpected result led us to look for what reason Scaphandre does not provide the same accuracy in its power consumption measurements on our OAI 5G platform as the one reported in related literature. It then appeared that Scaphandre, and more generally all RAPL-based power consumption measurement tools, take advantage of the temporal behavior of the Linux operating system scheduler. But on our OAI 5G platform, the classical asynchronous Linux scheduler is not the one running. Our OAI 5G platform runs the Lowlatency kernel and its real-time scheduler. This real-time scheduler obviously does not report the same temporal behavior as the classical asynchronous one. To the best of our researches, it currently does not exist any RAPL-based power consumption measurement tool that can take advantage of the Lowlatency Linux kernel.

This result limits the scope of energy consumption measurements that can be achieved. It is actually impossible to assess the energy consumption of the software components of the

TABLE I
USRP ENERGY CONSUMPTION

Description	Power consumption
USRP without gNB process	25.62 W
USRP with gNB without traffic	37.15 W (+45%)
USRP with 2 Mbps DL traffic	37.25 W
USRP with 20 Mbps DL traffic	37.59 W

TABLE II
SERVER ENERGY CONSUMPTION

Description	Power consumption
Server without gNB process	274.72 W
Server with gNB without traffic	282.83 W (+2.95%)
Server peak at traffic generation start	289.13 W (+2.23%)

5G architecture. Given that the only trustable equipment for measuring energy consumption is the physical power meter, it is only possible to get some global measurements of the full servers and USRP. Given the architecture of the 5G platform, it remains possible to evaluate the energy consumption due to 5G softwarization on one hand, and to SDR on the other, and this from both the user and network sides. The following section then presents energy consumption assessments on this basis.

III. EXPERIMENTS AND MEASUREMENTS

All measurements presented in this section have been performed using the physical power meter, and then correspond to the full energy consumption of the servers or the USRP. IPerf is used for traffic generation. Generated traffic is a constant bit rate UDP flow with an average throughput that regularly increases for testing 5G from low to higher loads. Experiments have also been done with TCP traffic but are not reported as the conclusions are similar to the ones with UDP. All experiments have been run 100 times. They all exhibit similar results, except for crashes that appear on an irregular basis.

Tables I and II respectively exhibit power consumption for the USRP and servers in the following cases:

- without the gNB process running
- with the gNB running
- with the gNB running and traffic transmission

These first results show, as expected, that the server has, by nature, a higher consumption than the USRP. This is due to the versatile nature of a server that is made of many more diverse components than an USRP whose diversity in possible computings is reduced by nature. Results also show that the gNB process has a significant impact on the USRP with an increase of the power consumption of 45%, whereas it has almost no impact on the server power consumption (+2.95%). On an other side, traffic transmissions only seem to have a limited impact on both the USRP and server power consumption:

- +0.10 W for 2 Mbit/s and +0.34W for 20 Mbit/s compared to the power consumption for 2 Mbit/s on the USRP




Fig. 2. Power consumption of the gNB server when generating downlink traffic from 1 to 20 Mbit/s (1-2-3-4-5-6-7-8-9-10-15-20)

- +6.30 W i.e. +2.23% whatever the throughput of the transmission is, but this increase is only formed by a peak at the beginning of the transmission)

Actually, it appears on Figures 2 and 3 that the server power consumption is not constant during the full transmission duration: the power consumption has a peak at the beginning of the transmission and becomes flat right after. The power consumption then does not increase for the duration of the traffic transmission. The peak however indicates that a virtualized function is called at the beginning of the transmission, and that this function has certainly to be optimized. The function still needs to be identified. This will be done when the platform will be fully instrumented for measuring the power consumption of any software component.

Figures 2 and 3 respectively depict the continuous power consumption of the gNB server when handling downlink and uplink traffic from 1 to 20 and 1 to 30 mbit/s. Traffic is transmitted as short and constant duration bursts. Throughput of bursts on Figure 2 are 1-2-3-4-5-6-7-8-9-10-15-20 Mbit/s. Throughput of bursts on Figure 3 are 5-10-15-20-25-30 Mbit/s. The beginnings of the bursts are indicated by black triangles on the two figures.

Figure 4 shows the same power consumption curve for the USRP for downlink transmission. The figure shows the power consumption for traffic bursts of 1-2-3-4-5-6-7-8-9-10-15-20-25-30 Mbit/s. At the opposite of servers, the power consumption is high for the full duration of the burst, and this power consumption increases with the traffic increase ¹. This represents an unexpected results, as we were initially prone

¹Note the stability issues of the platform already indicated. In the middle of the curve, the platform collapsed and no traffic transmission was possible,

to put the responsibility of most energy consumption to the server. Based on these preliminary results taking advantage of point to point communications between one UE and one gNB, it appears that the USRP experience significant power consumption increases when the gNB is running and when it is forwarding traffic. In addition, the USRP power consumption increases with the traffic throughput. This questions on the ability of USRP SDR of supporting thousands of connections from thousands of users. On the other side, the server shows only very limited impact when the gNB is running and forwarding traffic. Whatever the throughput of the traffic, the server consumption is stable, at the exception of a very short duration peak at the beginning of any transmission.

IV. RELATED WORKS

The recent awareness of the negative impact of digital world on pollution and global warming raised a significant new effort toward greening the information and communication technologies (ICT). Any new ICT is then highly examined from this new point of view. This is the case of 5G that became a casus belli between pro and anti 5G. They all provide figures about the power consumption of a gNB and of all the gNB that need to be deployed on Earth for a full coverage. They however rely on the existing gNB deployed, i.e monolithic ones, and then do not consider the new softwarized gNB architectures. On the other end of the communication, papers presenting power consumption measurements of 5G smartphones exist (smartphone are far more accessible than gNB). This is the case for instance in [9]. To the best of our knowledge, the work for assessing 5G softwarized gNB is still to be done.




Fig. 3. Power consumption of the gNB server when generating uplink traffic from 1 to 30 Mbit/s (5-10-15-20-25-30)

Previously, the problem of assessing power consumption has been addressed for instance for high performance computing (HPC) or cloud platforms [10] and their virtual machines [11] [12] or containers [13] [14]. For this purpose, many software power meters have been designed as PowerAPI [15], Likwid [16], PAPI [17], PowerMon [18], PowerInsight [19], etc. Some of these software power meters have been evaluated and compared as in [20] [21] [22]. However, none of these tools is able to take advantage of the Lowlatency Linux kernel and its real-time scheduler. This opens a new research field to fill this gap.

V. CONCLUSION

Nowadays, a huge effort is made worldwide for designing and deploying 5G as the new standard for cellular communications, with ongoing projects of extension for IoT, non-terrestrial coverage, etc. 5G is nevertheless blamed when considering sustainable development, as 5G opponents are pointing increased energy consumption of 5G compared to 4G. The reason is a supposed increase of the number of 5G gNB for a full terrestrial coverage. However, based on our experience with our experimental 5G platform, its new software architecture implies a heavy computing load, certainly responsible of significant electrical power consumption. This is the evaluation of this second aspect of power consumption that this paper aimed to address.

The first conclusion of this work deals with the limits of power meters. Values returned by software power meters taking advantage of the RAPL function of Intel Processors (as Scaphandre or Intel PCM) have been shown as completely unrealistic when running with a Lowlatency operating system. This completely changed the objectives of this paper that

aimed at measuring the energy consumption of any single component of the new software architecture of 5G networks. Given the lacks of the software power meters, it is then only possible to assess the consumption of full servers or USRP thanks to physical power meters that measure the energy consumption at the level of electrical sockets. The granularity of measurements presented in the paper has then been coarser than initially expected.

Obtained results are nevertheless somewhat interesting and unexpected. The energy consumption of the servers running the software components of 5G is basically high but does not significantly increase when the traffic load is getting higher. Concerning the USRP implementing 5G SDR, the energy consumption is impacted when the traffic increases. This may show that the 5G SDR implementation is not scalable, what was unexpected as it mostly uses hardware components.

These results nevertheless need to be detailed with finer measurements. Developing solutions for measuring power consumption at the level of software components is then a key objective of future works. Two different approaches are currently under consideration. The first one deals with taking advantage of the IDRAC management board of DELL servers. It provides functions giving access to some energy consumption estimates. Despite this solution can only work on DELL server as IDRAC control boards are DELL specific, it will be assessed to check whether the returned values are realistic, expecting that it can at least provide an efficient quick solution for our own LAAS-CNRS 5G experimental platform. The second approach deals with adapting the RAPL-based power consumption estimation mechanism to real-time schedulers.




Fig. 4. Power consumption of the gNB USRP when generating downlink traffic from 1 to 30 Mbit/s (1-2-3-4-5-6-7-8-9-10-15-20-25-30)

REFERENCES

- [1] Q. Douarre, E.-M. Djelloul, P. Berthou, D. Dragomirescu, P. Owezarski, "Design of a 5G experimental platform based on OpenAirInterface", 7th EAI International Conference on Tools for Design, Implementation and Verification of Emerging Information Technologies (TridentCom'2022), Melbourne, Australia, November 23-25, 2022.
- [2] OpenAirInterface Software Alliance, "The general information site," 2024, Accessed: 2024-01-29. [Online]. Available: <https://openairinterface.org/about-us>.
- [3] H. David, E. Gorbatov, U. R. Hanebutte, R. Khanna, and C. Le, "RAPL: memory power estimation and capping," in 2010 ACM/IEEE International Symposium on Low-Power Electronics and Design (ISLPED), Aug. 2010.
- [4] K. Khan, M. Hirki, T. Niemi, J. Nurminen, and Z. Ou, "RAPL in action: experiences in using RAPL for power measurements," ACM Transactions on Modeling and Performance Evaluation of Computing Systems (TOMPECS), vol. 3, Jan. 2018.
- [5] M. Hahnel, B. Dobel, M. Volp, and H. Hartig, "Measuring energy consumption for short code paths using RAPL," ACM SIGMETRICS Performance Evaluation Review, vol. 40, no. 3, pp. 13–17, Jan. 2012.
- [6] B. Petit, "Scaphandre," 2024, Accessed: 2024-01-31. [Online]. Available: <https://hubblo-org.github.io/scaphandre-documentation/>.
- [7] D. Hackenberg, R. Schöne, T. Ilsche, D. Molka, J. Schuchart, R. Geyer, "An energy efficient feature survey of the Intel Haswell processor", IEEE International Parallel and Distributed Processing Symposium Workshop (IPDPSW'2015), 2015.
- [8] FLUKE NORMA 4000 technical description, 2024, Accessed: 2024-01-31. [Online]. Available: <https://www.fluke.com/en-ca/product/electrical-testing/power-quality/norma-4000>.
- [9] D. Xu, A. Zhou, X. Zhang, G. Wang, X. Liu, C. An, Y. Shi, L. Liu, H. Ma, "Understanding Operational 5G: A First Measurement Study on Its Coverage, Performance and Energy Consumption", SIGCOMM '2020: Proceedings of the Annual conference of the ACM Special Interest Group on Data Communication on the applications, technologies, architectures, and protocols for computer communication, July 2020.
- [10] W. Lin, F. Shi, W. Wu, K. Li, G. Wu, and A.-A. Mohammed, "A taxonomy and survey of power models and power modeling for cloud servers," ACM Computing Surveys, vol. 53, no. 5, Oct. 2020.
- [11] M. Colmant, M. Kurpicz, P. Felber, L. Huertas, R. Rouvoy, and A. Sobe, "Process-level power estimation in VM-based systems," in Proceedings of the Tenth European Conference on Computer Systems. Bordeaux France: ACM, Apr. 2015.
- [12] J. Phung, Y. C. Lee, and A. Y. Zomaya, "Lightweight power monitoring framework for virtualized computing environments," IEEE Transactions on Computers, vol. 69, no. 1, Jan. 2020.
- [13] G. Fieni, R. Rouvoy, and L. Seinturier, "SmartWatts: self-calibrating software-defined power meter for containers," in 2020 20th IEEE/ACM International Symposium on Cluster, Cloud and Internet Computing (CCGRID). Melbourne, Australia: IEEE, May 2020.
- [14] S. S. Tadesse, F. Malandrino, and C.-F. Chiasserini, "Energy consumption measurements in Docker," in 2017 IEEE 41st Annual Computer Software and Applications Conference (COMPSAC), vol. 2, Jul. 2017.
- [15] I. SPIRALS, "PowerAPI," 2024, Accessed: 2024-01-31. [Online]. Available: <http://powerapi.org/>.
- [16] J. Treibig, G. Hager, and G. Wellein, "Likwid: a lightweight performance-oriented tool suite for x86 multicore environments," in Proceedings of PSTI2010, the First International Workshop on Parallel Software Tools and Tool Infrastructures, San Diego CA, 2010.
- [17] V. M. Weaver, M. Johnson, K. Kasichayanula, J. Ralph, P. Luszczek, D. Terpstra, and S. Moore, "Measuring energy and power with PAPI," in 2012 41st International Conference on Parallel Processing Workshops. Pittsburgh, PA, USA: IEEE, Sep. 2012.
- [18] D. Bedard, M. Y. Lim, R. Fowler, and A. Porterfield, "PowerMon: fine-grained and integrated power monitoring for commodity computer systems," in Proceedings of the IEEE SoutheastCon 2010 (SoutheastCon), Mar. 2010.
- [19] J. H. Laros, P. Pokorny, and D. DeBonis, "PowerInsight - a commodity power measurement capability," in 2013 International Green Computing Conference Proceedings, Jun. 2013.
- [20] D. Hackenberg, T. Ilsche, R. Schone, D. Molka, M. Schmidt, and W. E. Nagel, "Power measurement techniques on standard compute nodes: a quantitative comparison," in 2013 IEEE International Symposium on Performance Analysis of Systems and Software (ISPASS). Austin, TX, USA: IEEE, Apr. 2013.
- [21] A. Nouredine, R. Rouvoy, and L. Seinturier, "A review of energy measurement approaches," ACM SIGOPS Operating Systems Review, vol. 47, no. 3, pp. 42–49, Nov. 2013.
- [22] M. Jay, V. Ostapenco, L. Lefèvre, D. Trystram, A.-C. Orgerie, et al. "An experimental comparison of software-based power meters: focus on CPU and GPU", CCGrid 2023 - 23rd IEEE/ACM international symposium on cluster, cloud and internet computing, Bangalore, India, May 2023.