

Demonstrating the Energy Consumption of Radio Access Networks in Container Clouds

Venkateswarlu Gudepu[§], Rajashekhar Reddy Tella[§], Carlo Centofanti[•], José Santos[†],
Andrea Marotta[•], Koteswararao Kondepu[§]

[§]Indian Institute of Technology Dharwad, Dharwad, India

[•]University of L'Aquila, DISIM Department, L'Aquila, Italy

[†] Ghent University - imec, IDLab, Department of Information Technology, Gent, Belgium
e-mail: 212011003@iitdh.ac.in, carlo.centofanti1@univaq.it, k.kondepu@iitdh.ac.in

Abstract—The rapid evolution of next-generation mobile networks introduces challenges and opportunities in achieving various use cases with extremely low latency, high data rates, and dense user connectivity. However, these objectives can lead to the higher energy consumption of mobile networks, particularly within the Radio Access Network (RAN), which generally consumes 75% of the mobile networks total energy consumption. The current studies available focus mainly on the energy consumption of the next-generation Core Network (5GC), while this demonstration provides a better understanding of the energy consumption of the RAN components. The demonstration provides energy observability leveraging various open-source software tools to measure and monitor RAN energy consumption trends deployed on the Kubernetes platform — a widely adopted container orchestration platform. Moreover, this demonstration shows energy consumption on different RAN architectures — Monolithic, Disaggregated, and Control Plane and User Plane Separation (CUPS) — utilizing tools such as Kepler and Scaphandre for comprehensive energy measurement.

Index Terms—Radio Access Network (RAN), Energy Consumption, 5G and Beyond (B5G), Kubernetes.

I. INTRODUCTION

The importance of monitoring and reducing energy consumption in mobile networks is a concern operators, as it can reduce the Operating Expense (OPEX) and Carbon Dioxide (CO₂) emissions. As the demand for data-hungry applications, extremely low-latency services, and dense user connectivity continue to increase the energy consumption of Radio Access Network (RAN). The RAN accounts for a significant energy consumption of around 75% from the overall energy consumption within a mobile network, thus it plays a major role in determining the environmental impact of Information and Communication Technologies (ICT) infrastructure [1]. Efficient RAN operations contribute not only to reducing the carbon footprint of ICT networks but also offer economic benefits by optimizing energy usage.

Effectively managing RAN energy consumption necessitates cutting-edge software and hardware tools for measuring and monitoring energy consumption [2]. These tools enable operators not only to track and mitigate carbon emissions but also to comply with regulatory requirements and implement sustainable practices. Beyond these benefits, accurate energy monitoring offers insights crucial for capacity planning, performance optimization, and early issue detection, aligning with

the long-term goals of new Beyond 5G (B5G) use cases. Energy consumption can be measured as power consumption over a period. To reduce energy consumption, a strategic focus on power reduction over shorter intervals emerges as a viable and effective solution.

Recent works [3], [4] have proposed novel approaches to reduce the power consumption in various components of the 5G Core Network (5G CN) and RAN, while this demonstration addresses a critical gap by showcasing an experimental study of energy consumption at the RAN in a containerized environment. This demonstration focuses on different RAN deployments such as *Monolithic*, *Disaggregated*, and *Control Plane and User Plane Separation (CUPS)* architectures while analyzing various open-source energy monitoring tools to understand how RAN energy consumption behaves for different user demands. These insights lay a foundation for informed decisions towards energy-efficient and sustainable network infrastructures. The demonstration includes different RAN flavors while showing its live energy consumption and presenting insights on the impact of diverse RAN parameters on energy consumption. Through this comprehensive study, this demonstration contributes to the design of energy-efficient RAN solutions, crucial for the sustainability of B5G networks in the rapidly evolving ICT landscape.

The main contributions of this demonstration are threefold:

- Visualize the power consumption of different RAN deployment schemes in a containerized environment.
- Assess the accuracy of software-based monitoring tools.
- Discussion about challenges and research opportunities in the field of sustainable cloud-native RAN.

II. SYSTEM ARCHITECTURE

This section outlines the system model, including several components such as the Kubernetes (K8s)-based 5G system architecture, power measuring software tools, and monitoring platforms. The 5G system architecture is composed of three fundamental elements: User Equipment (UE), RAN, and 5G CN, as illustrated in Fig. 1. These components collaborate to assign the functionalities and capabilities of the considered networks. The UE functions as the endpoint device for network connectivity, the RAN facilitates wireless communication between the UE and the network, while the 5G CN manages data

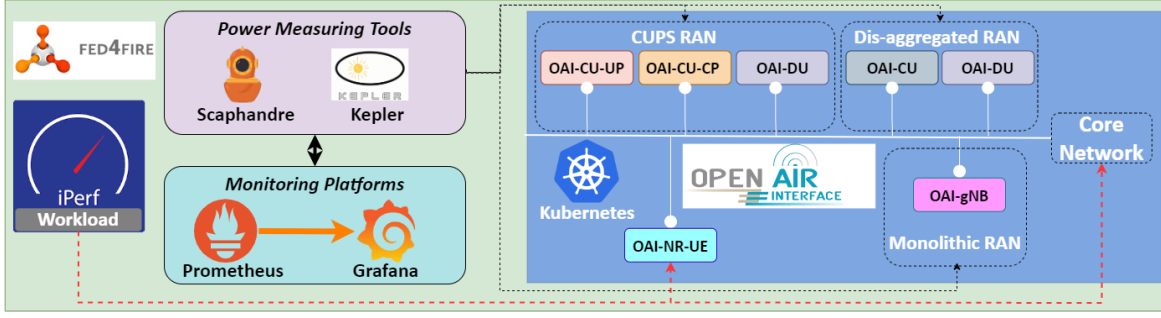


Fig. 1: Overview of the 5G system architecture deployed in Kubernetes (K8s).

routing, network management, and various services. The RAN segment is dynamically evolving, primarily adopting three distinct deployments: (i) Monolithic RAN — characterized by a single gNB; (ii) Disaggregated RAN — the gNB is decoupled into Distributed Units (DUs) and Central Units (CUs); and (iii) CUPS RAN — the CU is further separated into Control Plane (CP) and User Plane (UP). These deployments are designed to address the diverse demands of 5G networks over time.

In the monolithic RAN, all critical RAN functions (e.g., Physical layer, Medium Access Control (MAC) layer, and Radio Link Control (RLC)) are consolidated into a single, integrated unit known as the Next-generation NodeB (gNB), as illustrated in Fig. 1. Although this setup provides advantages such as simplified deployment due to its tightly coupled nature, it brings scalability challenges such as vendor lock-in, performance concerns, and slow development speed. In response to these challenges, the gNB undergoes a process of decoupling into two distinct entities: the DU and the CU, forming the disaggregated RAN. This architectural shift offers notable benefits, including increased flexibility, scalability, and resource optimization. In addition, RAN functions can still be further disaggregated into gNB-CU-Control Plane (gNB-CU-CP) and gNB-CU-User Plane (gNB-CU-UP) to support different use cases. The CP handles the initial attachment procedure between the RAN components, while the UP is responsible for forwarding data packets. The introduction of CUPS RAN brings several advantages, such as improved scalability of UP operations, the ability to tailor user planes, and enhanced system maintenance and upgrades. Moreover, it diminishes operational and maintenance costs for Mobile Network Operators (MNOs) by enabling the hosting of CP and UP nodes at different geographical locations.

Power measurement software tools play a vital role in assessing power consumption in various deployments within the 5G system architecture. These tools offer valuable insights into power consumption of the infrastructure and corresponding CO_2 emissions. The power consumption data can be effectively stored in a dedicated repository, facilitating in-depth analysis through various monitoring platforms. The data repository and monitoring platforms can be deployed in typical container orchestration platforms such as K8s [5], a widely adopted platform by most companies, which seamlessly interface with the end-to-end 5G system architecture.

TABLE I: The Hardware Configuration of the K8s cluster.

Node	Hardware Configuration (CPU and Memory)
Master	Intel(R) Xeon(R) CPU
Worker-{1-4}	E5-2650 v2 @ 2.60GHz, 48 GB RAM

TABLE II: 5G Network Parameters

Description	Value
NR Release - Band - Freq.	3GPP Release 16 - Band 78 - 3.6 GHz
RAN type	5G standalone gNB
CU/DU split	Option 2
Physical Resource Block	106
Radio Channel Bandwidth	40 MHz
Midhaul/Backhaul Capacity	1 Gbps Ethernet
UE	OAI based 5G SA UE

Two software-based power measurement tools have been deployed to monitor the power consumption across the RAN and UE. First, *Scaphandre* [6] measures power consumption using hardware sensors, correlating it with resource utilization and process activities, facilitating comprehensive power analysis. Second, *Kubernetes Efficient Power Level Exporter (Kepler)* [7] utilizes software counters and power sources such as Running Average Power Limit (RAPL), Advanced Configuration and Power Interface (ACPI) and NVIDIA Graphics Processing Unit (GPU) to measure power consumption based on hardware resources within a K8s cluster.

In summary, the power metrics extracted from *Scaphandre* and *Kepler* allow a comprehensive understanding of the power consumption dynamics within the various components of the 5G system.

III. DEMO SETUP AND RESULTS

A K8s cluster with five nodes (one master and four workers) has been deployed within the imec Virtual Wall (VWall) infrastructure [8] at IDLab, Belgium. This cluster serves as the testing ground for evaluating all three RAN deployment schemes: Monolithic, Disaggregated, and CUPS. Tables I and II detail the hardware configuration of the K8s cluster, and the 5G network parameters, respectively. 5G components across the three RAN scenarios are deployed using OpenAirInterface (OAI) [9]. To measure the power consumption across all the RAN scenarios both *Scaphandre* and *Kepler* have been deployed in the K8s cluster. Power metrics can be visualized through *Prometheus* and *Grafana* [10]. *Prometheus* is a widely adopted open-source monitoring tool to scrape the system resource metrics from various exporters, and *Grafana* seamlessly complements *Prometheus*, offering a user-friendly




Fig. 2: Power consumption measurement using *Scaphandre* for different RAN scenarios.

visualization platform that supports the creation of interactive dashboards as shown in Fig. 1 (right side). During the demonstration, different workload patterns will be initiated between UEs and the 5G CN to observe the power consumption of the system in all deployment schemes and its idle state before initiating the workload.

The results obtained for all three scenarios using the *Scaphandre* are as shown in Fig. 2a, 2b, and 2c respectively. It can be observed that during the workload using *iperf3*, power consumption is increasing across all the scenarios due to the processing at RAN functionalities. On the other hand, power monitoring using the *Kepler* tool performs similarly to the *Scaphandre* even though the design of both the tools are different.

IV. CHALLENGES AND OPPORTUNITIES

Reliability of existing measurement tools is one of the primary challenges in energy consumption analysis. Existing tools rely on generic models, leading to deviations from actual energy consumption. These discrepancies limits the accurate assessment of energy efficiency, posing a critical challenge for researchers and practitioners.

The power-aware orchestration of microservices, guided by energy consumption data, presents an opportunity to optimize energy efficiency in workload execution on cluster nodes. Future service orchestration platforms may utilize energy consumption data, spanning from cloud to user, enhancing resource allocation efficiency and influencing decision-making in workload deployment and management.

Energy-efficient hardware and carbon-neutral processing will be a major research topic in the next few years. Existing works mainly address the reduction of power consumption in the cloud infrastructure and its network links. The integration of renewable energy sources (e.g. solar, wind) into data centers could significantly reduce power consumption.

Efficient Carbon Footprint Measurement of containerized applications is crucial. Collaborative research among industry, academia, and non-profit stakeholders is vital to address sustainability challenges in container clouds, paving the way for standardized methodologies and best practices. Open-source communities are key to achieving standardized power consumption measurement in future cloud infrastructures.

In summary, balancing operational efficiency with energy savings will pave the way toward a more sustainable and efficient telecommunication infrastructure.

V. CONCLUSIONS AND FUTURE WORK

The impact of energy consumption on B5G networks, focusing on different RAN deployment scenarios are demonstrated.

By using open-source software-based monitoring tools such as *Scaphandre* and *Kepler*, the RAN power consumption is collected for the multiple scenarios. Our demo highlights the different power consumption of various RAN deployments during idle state and workload initiation, showing differences depending on the applied monitoring tool.

Accurate monitoring of power consumption is vital for efficient energy-aware management in future B5G network infrastructures. We argue that standardizing best practices for power measurement is essential for fostering a sustainable cloud ecosystem. Results show that current power measurement technologies provide slightly different data results, meaning that data collected with different methods is inconsistent. In future work, power-aware scheduling algorithms will be studied toward a sustainable container-based computing environment.

ACKNOWLEDGMENT

This work received funding from DST SERB Startup Research Grant (SRG-2021-001522), the SGNF project (“Reliability Evaluation of Virtualised 5G”). José Santos is funded by the Research Foundation Flanders (FWO), grant number 1299323N.

REFERENCES

- [1] L. M. Larsen, H. L. Christiansen, S. Ruepp *et al.*, “Toward Greener 5G and Beyond Radio Access Networks—A Survey,” *IEEE Open Journal of the Communications Society*, vol. 4, pp. 768–797, 2023.
- [2] D. López-Pérez, A. De Domenico, N. Piovesan, G. Xinli *et al.*, “A survey on 5G radio access network energy efficiency: Massive MIMO, lean carrier design, sleep modes, and machine learning,” *IEEE Communications Surveys & Tutorials*, vol. 24, no. 1, pp. 653–697, 2022.
- [3] M. Sharara, F. Fossati, S. Hoteit *et al.*, “Minimizing energy consumption by joint radio and computing resource allocation in Cloud-RAN,” *Computer Networks*, p. 109870, 2023.
- [4] V. Gudepu, B. Chirumamilla, R. R. Tella *et al.*, “EARNest: Experimental Analysis of RAN Energy with Open-Source Software Tools,” in *16th International Conference on COMMunication Systems NETWORKS (COMSNETS)*, 2024, pp. 1148–1153.
- [5] M. Luksa, *Kubernetes in action*. Simon and Schuster, 2017.
- [6] Scaphandre, “Energy consumption metrology agent,” accessed on 17 January 2024. [Online]. Available: <https://github.com/hubblo-org/scaphandre>.
- [7] M. Amaral, H. Chen, T. Chiba *et al.*, “Kepler: A framework to calculate the energy consumption of containerized applications,” in *IEEE 16th International Conference on Cloud Computing (CLOUD)*. IEEE, 2023, pp. 69–71.
- [8] Virtual Wall, “The Virtual Wall emulation environment,” accessed on 28 May 2023. [Online]. Available: <https://doc.ilabt.imec.be/ilabt/virtualwall/index.html>.
- [9] N. Nikaein, M. K. Marina, S. Manickam *et al.*, “OpenAirInterface: A flexible platform for 5G research,” *ACM SIGCOMM Computer Communication Review*, vol. 44, no. 5, pp. 33–38, 2014.
- [10] M. Holopainen, “Monitoring Container Environment with Prometheus and Grafana,” 2021.