Architecture and Methodology for Green MEC Services Using Programmable Data Planes in 5G and Beyond Networks

Jorge Andrés Brito
Departamento de Ingeniería de Sistemas Telemáticos, ETSI de Telecomunicación
Universidad Politécnica de Madrid
Madrid, Spain
jorge.brito@upm.es

José Ignacio Moreno
Departamento de Ingeniería de Sistemas Telemáticos, ETSI de Telecomunicación
Universidad Politécnica de Madrid
Madrid, Spain
joseignacio.moreno@upm.es

Luis M. Contreras
Telefónica Innovación Digital
Madrid, Spain
luismiguel.contrerasmurillo@telefonica.com

Marta Blanco Caamaño
Telefónica Innovación Digital
Madrid, Spain
marta.blancocaamano@telefonica.com

***Abstract*— In the era of 5G/6G developments and Multi-access Edge Computing (MEC) adoption, the telecommunications sector faces the dual challenge of enhancing network capabilities while mitigating environmental impact. This paper introduces an innovative architecture and methodology that leverages programmable data planes to address these concerns. Employing the P4 programming language and QUIC protocol, we propose dynamic load-balancing schemes for switches and servers, aiming to minimize energy consumption and carbon footprint without compromising network performance. Our methodology aims to deliver potential energy savings and enhanced reliability, providing a scalable solution for environmentally sustainable 5G/6G future network deployments.**

Keywords—Programmable data planes, green computing, energy consumption, 5G/6G, MEC, P4, service availability.

I. Introduction

The push towards 5G/6G and edge computing (for instance, based on MEC) aims to build ultra-reliable, high-bandwidth, and low-latency communication networks with cloud-native services. These attributes enable a wide range of applications from autonomous vehicles (V2X) to virtual/augmented reality (VR/AR) and 3D gaming, as well as industrial IoT [1]. However, the deployment of these networks demands substantial infrastructure required to support advanced computing capabilities, which in turn scales-up energy consumption and contributes significantly to the carbon footprint of telecommunications operators [2,3]. This introduces a critical need for innovative solutions that not only embrace the capabilities of 5G/6G and MEC but also address the environmental implications of their widespread adoption.

One promising approach for mitigating the energy demands of the aforementioned infrastructure lies in the strategic application of Software-Defined Networking (SDN) principles by offloading 5G/6G network functions to programmable data plane devices [4]. These devices have the capability of enhancing system performance (e.g., higher throughput, lower latency) and support advance traffic management by customizing packet processing [5]. Thus, the data plane can be programmed to dynamically adjust network device usage based on the detected traffic flow. During periods of low traffic, certain components can be powered down or operated in a low-power state, significantly reducing energy consumption This strategy is particularly relevant in the context of the Access Gateway Function (AGF), a 5G network function (NF) that bridges fixed network users with the 5G/6G core [6]. Furthermore, by hosting MEC services, the AGF can benefit from the proximity to end-users so enhancing service delivery while optimizing energy consumption.

This paper introduces a novel architecture and methodology dedicated to a dynamic power management scheme of the network infrastructure. This scheme aims to achieve a dual objective: Aggregate network traffic across the fewest possible switches in accordance with traffic volume demand and distribute the incoming workload among servers based on their existing energy resources. To achieve these goals, this framework proposes P4 as the programming language for the data plane [7] and adopts QUIC as a transport protocol, offering enhanced adaptability to network variations [8]. Additionally, we present an impact analysis to explore the tradeoffs between achieving energy savings and maintaining service availability.

II. Architecture

The Third Generation Partnership Project (3GPP) introduces a Service-Based Architecture (SBA) for 5G [9], complemented by the Broadband Forum's (BBF) 5G Wireless Wireline Convergence (WWC) framework for integrating fixed and mobile services (Fixed Mobile Convergence, FMC) [10]. A fundamental component within this framework is the AGF, which serves as a bridge linking the FN-RG with the 5G Core (5GC).

In addition to managing user traffic and executing subscriber services like authentication, authorization, and accounting (AAA) [11], the AGF can extend its capabilities to interface with MEC services. These services can include content delivery, data caching, real-time analytics, and support for industrial Internet of Things (IoT) applications, leveraging the AGF's strategic position near end-user devices. Fig. 1 illustrates the AGF-MEC integration in a 5G system.

The architecture proposed in this study delves into a specific setup where the AGF accesses MEC services via a spine-leaf network topology, as illustrated in Fig. 2. This chosen topology, typically employed in data centers, streamlines the separation of control and data planes while optimizing the distribution of network traffic [12]. It thereby

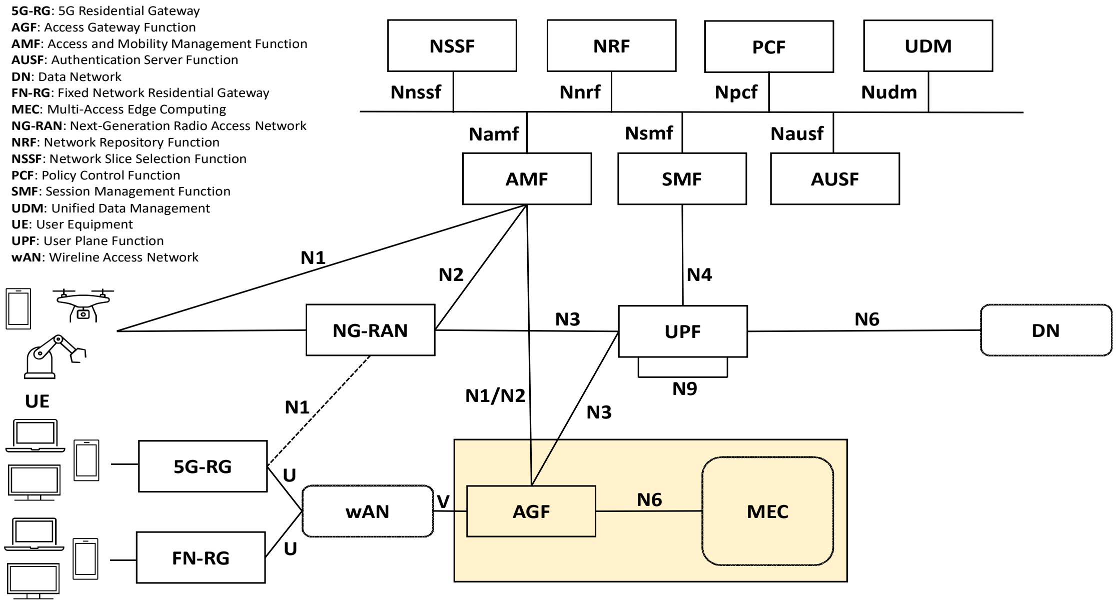

Fig. 1. AGF-MEC integration in a SBA 5G architecture

creates an ideal environment for implementing sustainable load distribution strategies aimed at enhancing energy efficiency directly within the data plane. It is worth mentioning that in this study, the MEC architecture is presented as a generic computing platform. Furthermore, Fig. 3 further exposes the hierarchical structure of this setup which include the following type of switches:

- Server switches: nodes that link the MEC servers with the broader network.
- Aggregation switches: nodes that manage the traffic flows between MEC hosted services and the AGF.
- AGF switch: node that direct the traffic to and from the AGF, ensuring seamless connectivity.

It can also be noted that the communication between the SDN controller and the programmable switches is done through gRPC-based southbound APIs (e.g., P4Runtime and gNMI) [13], [14].

III. Methodology

The proposed methodology details the design of two P4-based load balancing schemes that need minimum involvement of the control plane. Using data plane programmable devices, we establish a unified AGF-MEC management platform. This integration significantly reduces latency and mitigates network congestion when compared to traditional software-based solutions. Our procedure takes a similar approach of the studies in [15] and [16] but employs QUIC as a transport protocol due to its ability to identify flows efficiently independently of network changes.

The first load balancing scheme will consolidate traffic in the least number of aggregation switches handling flows between AGF and MEC, and then deactivating the ones that are not being used for reducing overall device consumption. The second load balancing scheme will use service switches to forward workload to MEC servers that have available energy resources. This represents a double approach procedure to optimize energy consumption in a vital part of a telecommunications network. The operational details are described in the following sub-sections:

A. Load balancing for switches

As it can be seen in Fig. 4, this procedure is composed of 3 stages: Register Initialization, Traffic Volume Estimation and Dynamic Active Switch Adjustment.

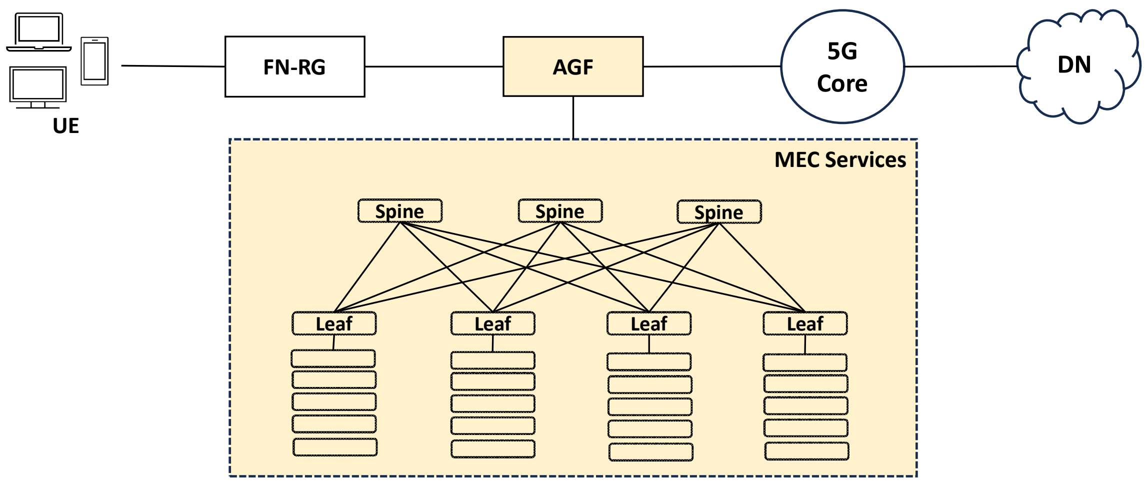

Fig. 2 Spine-leaf topology for MEC services

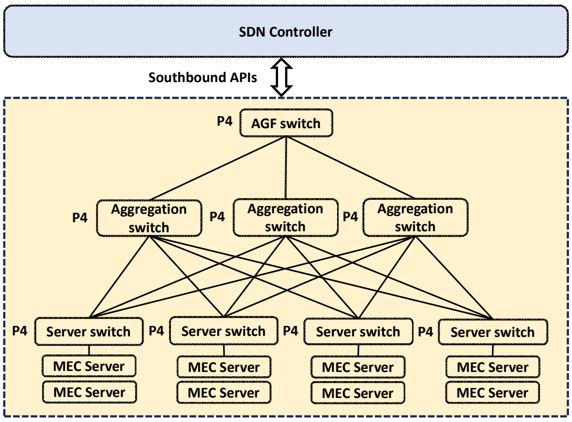

Fig. 3 Spine-leaf topology for MEC services

1) Register initialization: The process starts with the definition and initialization of P4-registers that mark the beginning of a time interval for traffic measurement (R1), estimate the incoming traffic volume (R2), and track the number of active aggregation switches (R3). These registers are initialized and updated directly in the data plane. Additional registers for setting the switch type (R4), defining the time interval length (R5) and the threshold traffic (R6), are also required to be incorporated to fulfill the network's specific needs and topology, underlining the process's adaptability and scalability. This last group of registers are initialized from the control plane but will be updated in the data plane.

2) Traffic volume estimation: Unlike traditional TCP traffic, QUIC encrypts almost all packet information, including the payload and most header fields. The operation leverages QUIC's visible headers for packet identification and load balancing through ECMP routing. Specifically, the hashing function relies in QUIC's Connection ID header (one of the few unencrypted fields), alongside source and destination IP addresses. The hash width is provided by R3, i.e., the number of aggregation switches. This estimation is necessary for understanding traffic patterns and making informed decisions on the allocation of traffic. The traffic volume, represented by the R2 register, is updated for each incoming packet in the time interval R1, relying on packet counts and sizes as proxies for the actual load due to QUIC's encryption.

3) Dynamic active switch adjustment: The core of the process lies in its capability to dynamically adjust the number of active aggregation switches in response to incoming traffic conditions. This is achieved by monitoring the traffic register and comparing its value against the predefined threshold (R6) at the end of each time interval, as determined by R5. Depending on the outcome, the R3 register is updated to either increase or decrease the number of switches engaged in traffic forwarding. Finally, the information relative to the number of active switches needed, stored in R3, is sent to a power management feature in the control plane. This module will deactivate the switches that are not processing any traffic thereby optimizing hardware energy resource utilization.

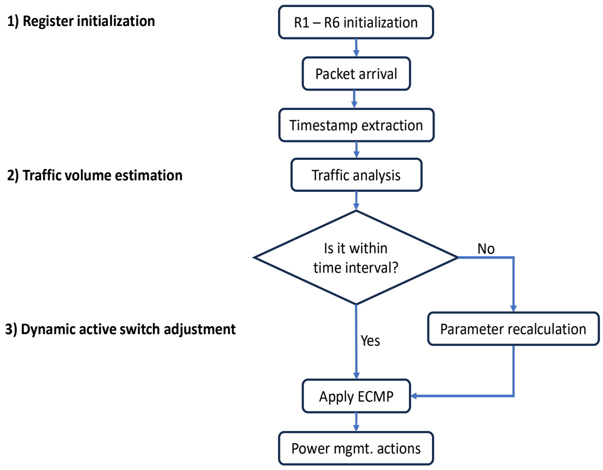

Fig. 4 Load balancing for switches flowchart

B. *Load balancing for servers*

Three phases are identified for this procedure: State Initialization, Info-packet Handling and Client Request Processing. The details are depicted in Fig. 5.

1) State initialization: Registers and match-action tables are configured within the programmable switches to manage server resources and client connections. Each server's green resource availability index is stored in a dedicated register array, initially set to zero (R7). A virtual IP address (VIP) is assigned for processing incoming client packets. A host information table is also initialized. This table details actions for packet forwarding based on the server ID, including MAC header rewriting and destination IP address adjustments.

2) Info-packet handling: Servers periodically send outgoing info-packets to the programmable switches. These packets contain a resource availability index, encoded within the payload of a QUIC frame, using a predefined format for ease identification and parsing by the switch. Upon receiving an info-packet, the switch parses it to extract the server ID (utilizing QUIC's source connection ID for initial packets) and the resource availability index. This information is stored in the server data registers and used for subsequent server selection decisions.

3) Client request processing: Client requests incoming to servers are directed using VIPs. New connections are established through QUIC's initial handshake process, which incorporates a unique, unencrypted connection ID field. This ID aids in selecting an appropriate server based on current energy resource availability and connection requirements, ensuring an efficient allocation of workloads to servers with sufficient resources. To maintain connection affinity, the switch maps the initial connection ID to a server ID. This mapping is dynamically updated to reflect any changes in connection IDs, ensuring consistent routing of packets to the correct server. This method replaces the TCP timestamp-based approach with a mechanism that is inherently supported by QUIC's connection management features. For incoming client requests, the system evaluates the resource availability index of each server. If a server with sufficient

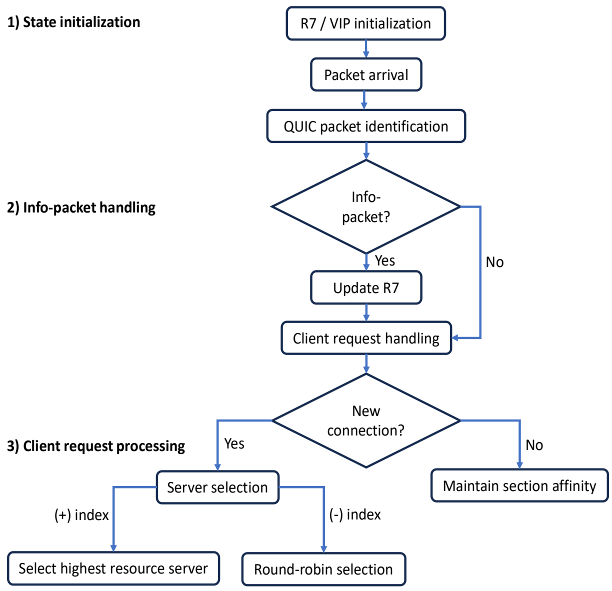

Fig. 5 Load balancing for servers flowchart

resources is available, it is selected to handle the request. Otherwise, a round-robin method is employed. The switch modifies the client packet by replacing the VIP with the server's actual IP address, facilitating direct communication between the client and the selected server.

IV. Analysis of Impact on Availability

When deciding on power down part of the infrastructure, the service availability can be compromised. Here we will analyze the tradeoff between availability and energy savings. It is important to clarify that this analysis focuses exclusively on the energy savings associated with the networking components of the architecture. Therefore, considerations regarding energy savings from cloud computing are outside the scope of this study.

A. Calculation of availability

Let’s assume the scenario depicted in Fig. 6 with a MEC node hosting number of servers providing value added services, such as video streaming, gaming, etc. The MEC node has two Top-of-the-Rack (ToR) switches, namely Leaf-1 and Leaf-2 providing connectivity to the servers running applications, which are interconnected by means of two additional switches, Spine-1 and Spine-2. These servers can host standalone functions, microservices, etc. This MEC structure is connected to the AGF User Plane, so providing connectivity towards the end user and other network segments. Note that the AGF Control Plane can be even an application running on one of those servers.

Different redundancy schemes could be considered. For instance, if the servers are only connected to the ToR of its rack, then different instances of the same application should run on different servers of each rack. An alternative could be that each server is connected to the two ToRs, so that the application can be deployed in one single server (with the redundancy being provided by another server in a different location / AGF).

In principle, this infrastructure should be powered up even during low traffic periods, for instance, by night. The extreme case would be the one in which just one single server is injecting traffic to the residential users. Even for that case, in a conventional scenario, all the spine and leaf nodes run and consume energy.

Let us assume that situation of minimal traffic in which a single server is delivering traffic. We can consider for example a video streaming application serving content to an area served by the AGF. In that minimal traffic situation, we can assume that just one single 10 GE port is used in the server for the purpose of delivering content.

Assuming, for simplification, a unique coding scheme for the on-demand content, with an average bit rate of 5 Mbps per video stream [17], this implies that up to 2,000 residential users can be served in such minimal traffic scenario. If we consider an average subscription rate of 8 €/month for video streaming services, the revenue on risk for service unavailability is of the order of 16,000 € in the extreme case that the full subscription rate per customer is counted as penalty.

Let us now compare the availability values in the situation of switching off the networking infrastructure (that is, spine and leaf nodes) for a minimal traffic situation. In this case we assume that the only active server is connected to its corresponding ToR (i.e., leaf switch), which at the same time is connected to one single active spine node providing connectivity to the AGF User Plane.

In order to calculate the availability of such scenario, we reference the Mean Time Between Failures (MTBF) and Mean Time To Repair (MTTR) metrics as reported in [18] These calculations are based on the performance of an Extreme SLX9150 switch used in both leaf and spine configurations. The relationship between the aforementioned parameters and device availability is given by:

$$Availability = \frac{MTBF}{(MTBF + MTTR)} \quad (1)$$

As it can be noted, this expression links reliability metrics to the availability of a network device.

The MTBF for the SLX9150 switch is 450,938 hrs. The calculations in [18] assume an MTTR value of 8 hrs., which will be dependent of the type of facility where the equipment is deployed and/or the availability of spares. Here we are considering a MEC environment, thus geographically distributed, where could be the case of not storing locally an additional unit for fast substitution of failed equipment. Thus 8 hrs. are assumed also here for MTTR in conventional operation scenario. This deals to an availability figure per switch of 99.99823%.

With these values in mind, the availability of the scenario depicted in Fig. 6 is as follows:

- *Situation in which all leaf and spine nodes are powered up.* The availability of the scenario (*A*) is determined by the availability of the alternative paths between AGF User Plane and the server. This means that the connectivity will be in place if either S1-L1, or S1-L2, or S2-L1 or S2-L2 paths are available. The availability in this case can be calculated as:

$$A = 1 - \{(1 - A_{S1-L1}) \times (1 - A_{S1-L2}) \times (1 - A_{S2-L1}) \times (1 - A_{S2-L2})\} \quad (2)$$

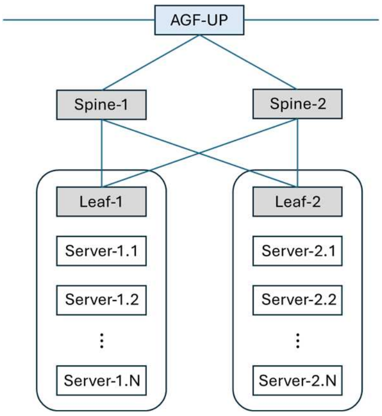

Fig. 6 AGF-MEC scenario under analysis

The resulting figure of availability is ≈ 100%.

- *Situation in which one spine and one leaf switches are powered down during low traffic demand periods.* Taken as extreme case the minimal traffic scenario described before, the availability is determined just by the availability of a single path, e.g., S1-L1. The availability in this case can be calculated as:

$$A_{S1-L1} = A_{S1} \times A_{L1} \quad (3)$$

With the value of availability for the SLX9150 switch, this leads to an availability of 99.99645%.

The corresponding unavailability, calculated as 1 - 99.99645%, translates to an estimated annual downtime of ~ 18 min. It is important to note that this unavailability figure strongly depends on the MTTR value used in the analysis. Unlike typical scenarios where repair time is the primary factor, in this specific case, the recovery time is determined by the time required to detect the failure and the triggering of power-up procedures for switches at the MEC facility.

In this context, the deployment of a P4-based infrastructure within the MEC facility becomes relevant for mitigating the impacts of failures. According to findings presented in [20], P4-based hardware switches can be effectively employed to rapidly detect link failures, achieving detection times of less than one millisecond. This capability would allow to react fast in an event of node or link failure, triggering the process of powering up the dormant leaf and/or spine switches and recovering connectivity. If we assume that the detection plus the powering up of the spine and leaf switches is less than 15 min. (the usual reboot time for this kind of devices), the immediate detection of failure in P4 infrastructure reduces the expectation of unavailability to ~ 1 min. This significantly minimizes the risk of revenue loss due to service unavailability while reducing energy consumption.

As it is evident, a clear tradeoff exists between system availability and energy consumption. To further evaluate such tradeoff, we will next analyze the potential savings in energy consumption that can be achieved.

B. *Calculation of energy savings*

Let us now focus on the potential energy savings. Revisiting the SLX9150 switch, which serves as our reference model, its specifications indicate an energy consumption of 104 W in typical operational mode [21]. With this energy consumption value, we calculate the total energy consumption for a conventional scenario where all leaf and spine switches are in working mode. The total consumption per month is given by:

$$E_{all} = E_{SLX9140} \times 4 \times 24 \times 30 = 299{,}52 \text{ kW} \quad (4)$$

where the unitary consumption of a switch is multiplied by the 4 units in the MEC facility, running 24 hrs. for 30 days.

On the other hand, we can consider the strategy of switching tow of the SLX9150 switches (one spine and one leaf) during low demand traffic situations. The typical pattern of traffic in an operational network has diurnal shape, with clear peaks and valleys in specific hours of the day. We can assume that during certain periods all the demand could be served from a single rack in the MEC facility, then strictly needing just a spine and a leaf node for that purpose. Let us assume that such period is 8 hrs. per day for our calculations. This would imply that the energy savings per month at a single MEC facility, specifically from networking switches can be quantified as follows:

$$E_{saved} = E_{SLX9140} \times 2 \times 8 \times 30 = 49{,}92 \text{ kW} \quad (5)$$

where the individual consumption of a switch is multiplied by the 2 units in the MEC facility, the ones switched-off during 8 hrs. along 30 days.

Extending such saving to the total footprint of edge computing sites can represent an important amount of energy. These MEC facilities could be deployed in PoPs concentrating on main distribution areas (at region/province level) or in central offices in such a country. The number of MEC facilities could account from 100 to 1,000 in those cases, for a mid-size country. This is compatible with the number of aggregation (the former) and pre-aggregation (the latter) sites considered in reference networks like in [22].

V. Conclusion

The proposed architecture and methodology for green MEC services using programmable data planes offer a promising path toward achieving energy efficiency in 5G and beyond networks. By adaptively managing traffic flows and device usage, our approach aims to achieve substantial energy savings and a reduced carbon footprint. The employment of P4 for data plane programming and QUIC for transport protocols highlights the adaptability and efficiency of this solution in handling dynamic network conditions. Furthermore, we conducted an availability impact analysis to explore the balance between enhancing energy savings and maintaining service availability. Future work will extend to the practical implementation and evaluation of our framework through hardware deployments, such as P4-based hardware switches and SmartNICs, within a controlled testbed environment. This step will allow a full deployment analysis of the benefits and potential of our framework in green network engineering.

Acknowledgment

This research is supported by the Universidad Politécnica de Madrid Project "SDGine for Healthy People and Cities". This project has received funding from the European Union's Horizon 2020 Research and Innovation Programme under the Marie Sklodowska-Curie Grant Agreement No. 945139 and from Telefónica I+D. This work is also partially supported by the HORIZON-JU-SNS2022 Research and Innovation Action project 6Green (Grant Agreement No. 101096925) and the

UNICO-5G I+D Program project B5GEMINI-INFRA (Grant Agreement No. TSI-063000-2021-81).

REFERENCES

[1] MEC in 5G Networks. ETSI White Paper No. 28. June 2018. Available online: https://www.etsi.org/images/files/ETSIWhitePapers/etsi_wp28_mec_in_5G_FINAL.pdf

[2] Vertiv and 451 Research Survey on 5G energy consumption at MWC19. Available online: https://www.vertiv.com/en-emea/about/news-and-insights/news-releases/2019/mwc19-vertiv-and-451-research-survey-reveals-more-than-90-percent-of-operators-fear-increasing-energy-costs-for-5g-and-edge/

[3] Bill Shock: Orange, China Telecom Fret About 5G Energy Costs. Available online: https://www.lightreading.com/5g/bill-shock-orange-china-telecom-fret-about-5g-energy-costs#close-modal

[4] J.A. Brito, J.I. Moreno, L.M. Contreras, M. Alvarez-Campana, M. Blanco Caamaño, "Programmable Data Plane Applications in 5G and Beyond Architectures: A Systematic Review," in Sensors 2023, 23, 6955.

[5] E. F. Kfoury, J. Crichigno and E. Bou-Harb, "An Exhaustive Survey on P4 Programmable Data Plane Switches: Taxonomy, Applications, Challenges, and Future Trends," in IEEE Access, vol. 9, pp. 87094-87155, 2021.

[6] What is a 5G Access Gateway Function (AGF)? Available online: https://www.metaswitch.com/knowledge-center/reference/what-is-a-5g-access-gateway-function-agf

[7] P. Bosshart, D. Daly, G. Gibb, M. Izzard, N. McKeown, J. Rexford, C. Schlesinger, D. Talayco, A. Vahdat, G. Varghese et al., "P4: Programming protocol-independent packet processors," ACM SIGCOMM Computer Communication Review, vol. 44, no. 3, pp. 87–95, 2014.

[8] A. Langley et al., "The QUIC transport protocol: Design and Internetscale deployment" in Proc. Conf. ACM Special Interest Group Data Commun., pp. 183196, 2017.

[9] System Architecture for the 5G System, Document TS 23.501, V17.5.0, 3GPP. Available online: https://www.etsi.org/deliver/etsi_ts/123500_123599/123501/17.05.00_60/ts_123501v170500p.pdf

[10] BBF, "TR-470: 5G Wireless Wireline Convergence Architecture," BroadBand Forum, Tech. Rep., 2022.

[11] BBF, "TR-456: AGF functional requirements, Issue: 02," BroadBand Forum, Tech. Rep., 2022.

[12] L. Peterson, C. Cascone, B. O'Connor, T. Vachuska, and B. Davie. "Software-Defined Networks: A Systems Approach," Systems Approach LLC. Chapter 2: Use Cases. 2022. Available online: https://sdn.systemsapproach.org/uses.html

[13] The P4.org API Working Group, "P4Runtime Specification version 1.3.0," Available online: https://p4.org/p4-spec/p4runtime/main/P4Runtime-Spec.html

[14] L. Peterson, C. Cascone, B. O'Connor, T. Vachuska, and B. Davie. "Software-Defined Networks: A Systems Approach," Systems Approach LLC. Chapter 5: Switch OS. 2022. Available online: https://sdn.systemsapproach.org/stratum.html

[15] G. Grigoryan and M. Kwon, "Towards Greener Data Centers via Programmable Data Plane," 2023 IEEE 24th International Conference on High Performance Switching and Routing (HPSR), Albuquerque, NM, USA, pp. 62-67, 2023.

[16] G. Grigoryan, Y. Liu, and M. Kwon, "iLoad: In-network load balancing with programmable data plane," in Proceedings of the 15th International Conference on emerging Networking EXperiments and Technologies, pp. 17–19, 2019.

[17] L.M. Contreras, A. Solano, F. Cano, J. Folgueira, "Analysis of Network Function Sharing in CDN-as-a-Service Slicing Scenario", International Journal of Network Management, Vol. 33, Issue 4, pp. 1055-7148, 2023.

[18] The Extreme Networks Federal Data Center Design Series. Volume 3: Evaluation of Reliability and Availability of Network Services in the Data Center Infrastructure. Available online: https://www.extremenetworks.com/resources/solution-brief/the-extreme-networks-federal-data-center-design-series-volume-3

[19] M. Alonso, S. Coll, J.M. Martínez, V. Santonja, and P. López, "Power consumption management in fat-tree interconnection networks," Parallel Comput. 48, C, pp. 59–80, 2015.

[20] D. Franco, M. Higuero, E. Ollora Zaballa, J. Unzilla, and E. Jacob, "Quantitative measurement of link failure reaction time for devices with P4-programmable data planes," Telecommun. Syst. 85, 2, 2023.

[21] ExtremeSwitching SLX 9150 Technical Specifications, February 2020. Available online: https://documentation.extremenetworks.com/slxos/hw/20.1.1/slx9150_specs/slx9150-techspecs.pdf

[22] B. Naudts, M. Kind, F.-J. Westphal, S. Verbrugge, D. Colle, M. Pickavet, "Techno-economic analysis of software defined networking as architecture for the virtualization of a mobile network", European Workshop on Software Defined Networking, 2012.