

Energy Monitoring and Management in 5G Integrated Fronthaul and Backhaul

Osamah Ibrahim Abdullaziz¹, Marco Capitani², Claudio Ettore Casetti³, Carla Fabiana Chiasserini³,
ShahzooB Bilal Chundrigar¹, Giada Landi², Xi Li⁴, Francesca Moscatelli², Kei Sakaguchi⁵, Samer T. Talat¹

¹ITRI, ²Nextworks, ³Politecnico di Torino, ⁴NEC, ⁵FhG-HHI

Abstract— Energy efficiency is likely to be the litmus test for the sustainability of upcoming 5G networks. Before the new generation of cellular networks are ready to roll out, their architecture designers are motivated to leverage the SDN technology for the sake of its offered flexibility, scalability, and programmability to achieve the 5G KPI of 10 times lower energy consumption. In this paper, we present Proofs-of-Concept of Energy Management and Monitoring Applications (EMMAs) in the context of three challenging, realistic case studies, along with a SDN/NFV-based MANO architecture to manage converged fronthaul/backhaul 5G transport networks.

Keywords—Fronthaul/Backhaul; energy management and monitoring; SDN/NFV

I. INTRODUCTION

In the past, actions aimed at improving hardware efficiency have led to high energy savings at device and infrastructure levels in mobile communication. To address the expected densification brought about by the upcoming 5G networks [1], gradual hardware improvements are no longer sufficient, especially energy efficiency is set to become a crucial feature. The reduction of the energy footprint of 5G networks, while maintaining the expected Quality of Service (QoS) for each Mobile Network Operator (MNO) or for end users, requires novel control-plane solutions that leverage the flexibility of Software-Defined Network (SDN) concepts.

For this purpose, Energy Management and Monitoring Applications (EMMAs) have been designed in the 5G-Crosshaul project [2] to monitor the energy parameters of the fronthaul and backhaul network elements, to estimate their energy consumption and trigger reactions. Conveniently deployed along with a standard ETSI MANO, EMMAs can also collect information about several network aspects like traffic routing paths, traffic load levels, user throughput, number of active sessions, radio coverage, interference of radio resources, and equipment activation intervals. Such data can be used to compute a virtual infrastructure energy budget for subsequent analysis and reactions using machine learning and optimization techniques.

The fundamental EMMA design addresses the optimal schedule of power operational states and levels of power consumption of nodes, jointly performing load balancing and frequency bandwidth assignment, in highly heterogeneous SDN domains.

EMMAs are based on algorithms that provide a heuristic solution to an optimization problem for energy-efficient flow routing in the integrated backhaul/fronthaul network [3]. Furthermore, although it is outside the scope of this paper, the re-allocation of virtual functions across backhaul and fronthaul can be done as part of the optimization actions. Virtual network functions can thus be mode to less power-consuming or less-loaded servers, reducing the overall energy demand in the network. Arguably, such actions can be instrumental in meeting the target KPI of 10 times lower energy consumption.

The literature shows several efforts toward more widespread energy saving in SDNs. In general, all authors agree that finding minimum-power network subsets in an SDN is an NP-hard problem, therefore all resort to some form of heuristics, as we did. This approach is exemplified by [4], where hybrid and partially deployed SDNs were considered and the heuristic is based on a spanning-tree approach. In [5], the authors solve the energy-saving optimization problem through a heuristic algorithm that minimizes the in-band control traffic by properly placing controllers. Greedy algorithms for energy-efficient routing in SDN using link utilization and packet delay as constraints are also defined in [6]. However, all the above works evaluate the performance of their algorithms in simple idealized settings through either analysis, simulation or limited emulation, while in this paper we strive to present real implementations of the EMMAs through Proof-of-Concepts (PoC) for realistic use cases. Additionally, our EMMA applications are thoroughly integrated in an SDN ecosystem that includes monitoring and management capabilities for single and multi-tenant cases, as explained below.

II. THE EMMA ARCHITECTURE

The monitoring and management of power consumption in 5G-Crosshaul infrastructures operate over different kinds of network technology domains, as exemplified in Figure 1. In particular, a monitoring layer, developed on top of an SDN controller, collects, aggregates and elaborates energy-related measurements for network domains such as: (1) networks composed of software-based switches named Crosshaul Packet Forwarding Elements (XPFEs), (2) mmWave links, and (3) analogue Radio over Fibre (RoF) technologies. Importantly, energy consumption information can be collected not only for

network paths, but also for virtual network slices, network services and tenants.

Energy management is then implemented above the monitoring application to determine an optimal resource allocation for both network connections and cloud-based services. Specifically, as depicted in Figure 1, energy-based optimization is achieved through several modules, each implementing a different task: routing (and re-routing) of traffic flows, Virtual Network Functions (VNFs) placement, and regulation of network node power states (including their On/Off switching) depending on the network resource demand. We remark that such optimization is either performed upon on-demand instantiation or automatically triggered by the monitoring application when re-planning is needed. EMMAs indeed include management interfaces that can put in place, through the southbound interface, the decisions made by the aforementioned modules, within each domain. This is achieved through signaling across the XCI, i.e., the Crosshaul Control Infrastructure, composed of a hierarchy of network and cloud controllers, together with orchestration and management entities.

In the next sections, we present PoC prototypes of the EMMAs, applied to the three domains shown in Figure 1: a software switch network domain, a mmWave domain, and an analogue RoF domain for ground-to-train radio access. In particular, we focus on the energy management module for the automatic power regulation of nodes and their On/Off switching.

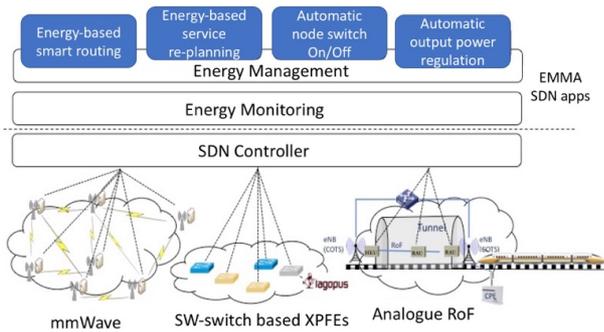


Figure 1: EMMA architecture: EMMAs operate over different domains and effectively monitors energy consumption for system optimization

III. EMMA FOR SOFTWARE SWITCHES NETWORKS

The PoC prototype of the EMMA for the energy optimization of software switch networks has been developed on the SDN/NFV-based integrated backhaul/fronthaul transport network, defined by the 5G-Crosshaul project [7]. The implementation is based on the OpenDaylight Beryllium framework¹ and includes features for (i) power consumption monitoring of both physical and virtual infrastructures; and (ii) optimization of the network device power states. The prototype is now deployed in the 5TONIC² laboratory in Madrid.

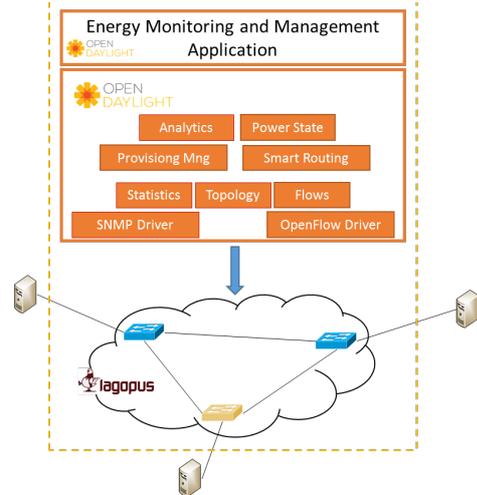


Figure 2: Software architecture of the EMMA PoC

The EMMA PoC software architecture is illustrated in Figure 2. The data plane is based on Lagopus³ software switches controlled via the OpenFlow protocol. The switches implement an SNMP (Simple Network Management Protocol) agent based on the Energy Management (EMAN) MIB [8] for the monitoring of power consumption and dynamic regulation of power states. Each power state is associated with an increasing static value of power consumption (i.e., independent of the current traffic load of the device) and it is characterized by some forwarding capacity. For example, in *sleeping mode*, which guarantees the minimum power consumption, only management messages can be handled, while in *medium active mode* also data plane traffic is supported but with some restrictions in terms of bandwidth or delay. In *full active mode*, the nodes operate at full speed but with maximum static power consumption. Furthermore, in our model we assume that the total power consumed by a node also includes an additional variable component that linearly depends on the real-time traffic load.

The OpenDaylight SDN controller implements the southbound plugins for both OpenFlow and SNMP protocols, thus enabling the programmability and monitoring of the devices at the EMMA application level. These low-level protocols are abstracted to the application through the north-bound APIs exposed by the controller services for flow configuration, topology management and statistics collection. On top of them, new EMMA software modules implement the logic to compute the real-time power consumption for the whole physical infrastructure and for the virtual instances running on top, as well as to establish connection services while adapting the power states of the data plane devices to the current traffic load.

The *Analytics* module elaborates traffic statistics as well as data about power states configuration and consumption in physical nodes, to identify the variable component of the infrastructure power consumption and its mapping on the active

¹ <https://www.opendaylight.org/odlbe>

² <https://www.5tonic.org/>

³ <http://www.lagopus.org/>

connections and tenants (see Figure 3). The OpenDaylight DLUX web interface has been extended to show the related real-time graphs for the EMMA application.

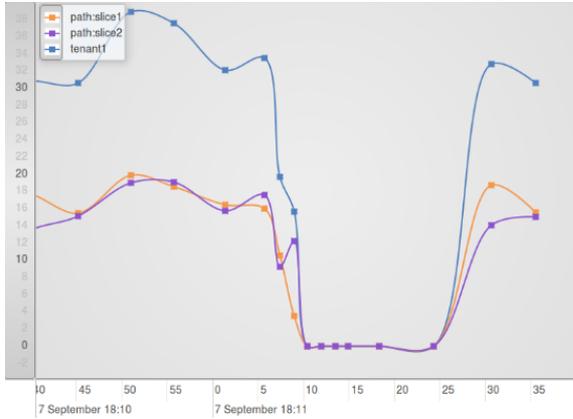


Figure 3: Screenshot of power consumption monitoring per service and per tenant over time

The *Provisioning Manager* handles the setup of new network connections, as well as rerouting of existing traffic flows, in the data plane using Open-Flow, based on the paths computed by the energy-based *Smart Routing* service. These algorithms consider the current power states of the nodes, together with the current load of links and traffic demand for requested paths. The switches are usually switched off or kept in sleeping mode to minimize the power consumption of the whole physical infrastructure, and they are moved to the active power states only when forwarding traffic for active flow connections. We refer the interested reader to [3] for a detailed description of such algorithms. The *Power State Manager* configures the power states via SNMP and it is automatically triggered by the Provisioning Manager, which selects the suitable states depending on traffic requirements as an integrated step of the connection setup procedure.

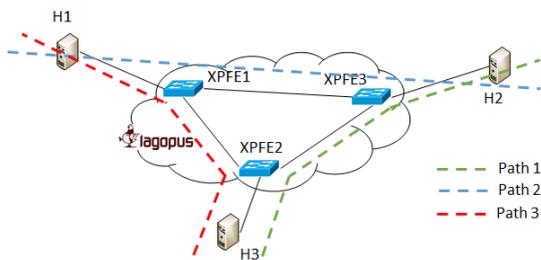


Figure 4: Test scenario and instantiated paths for EMMA experiments on XPFEs domain

The prototype has been functionally validated in the 5TONIC testbed using a simplified deployment with 3 XPFEs (Figure 4), with the objective to evaluate the impact of the automated power state management on the service provisioning time.

During the test, path 1, 2 and 3 are progressively instantiated, as detailed in Table 1. Path 1 setup requires activating XPFE2 and XPFE3, while Path 2 setup requires the additional activation of XPFE1. On the other hand, when Path 3 is established all the XPFEs are already active and no further changes in the power state of XPFE are needed.

Table 1: Paths installed for EMMA tests on XPFEs domain

Path	Source	Destination	Traversed nodes
Path 1	H2	H3	XPFE 2 – XPFE 3
Path 2	H1	H2	XPFE 1 – XPFE 3
Path 3	H1	H3	XPFE 1 – XPFE 2

Figure 5 compares the average power consumption in W with a different number of established paths when EMMA is adopted (blue) and using a shortest path first algorithm without EMMA, i.e., always keeping all nodes in full active mode (red). In this deployment, the EMMA approach yields benefits just when zero or one path is established (29% and 10% of power saving respectively), since keeping some nodes in sleeping mode reduces the power consumption of the global network. From the second path on, all the nodes are already in full active mode and no further gain can be achieved.

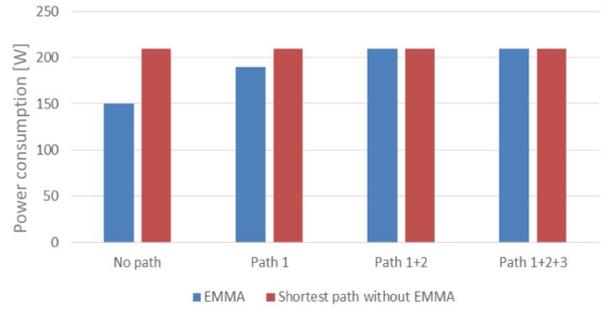


Figure 5: Power consumption saving: comparison between EMMA and shortest path without EMMA

Figure 6 shows the setup time (in seconds) for the three paths (values computed averaging from 10 executions), together with the time required to change the power state of the XPFEs. Setup time for Path 1 and Path 2, which includes the activation of the XPFEs, is much longer than the time required to establish Path 3, where all the XPFEs are already active. The configuration of the power states of devices clearly introduces the largest delay. While the internal procedures (i.e., provisioning coordination and path computation) takes just few milliseconds, the actions that require an interaction between controller and devices are in the order of hundreds of milliseconds for the configuration of OpenFlow flow rules and few seconds for the configuration of the power state. The configuration of the power state constitutes indeed the main difference in the provisioning procedure, if compared with the traditional connection setup where all the nodes are always maintained active, and it adds a further component of the end-to-end provisioning time.

However, since the SNMP interaction between controller and devices is implemented through asynchronous and parallel

messages, this delay does not increase linearly with the number of nodes to be switched on but varies in a limited range. This means that in wider networks, where EMMA can reach its best performance due to the high number of nodes that can be maintained in sleeping mode, this delay can be kept in a reasonable range independently on the increasing number of XPFEs to be activated.

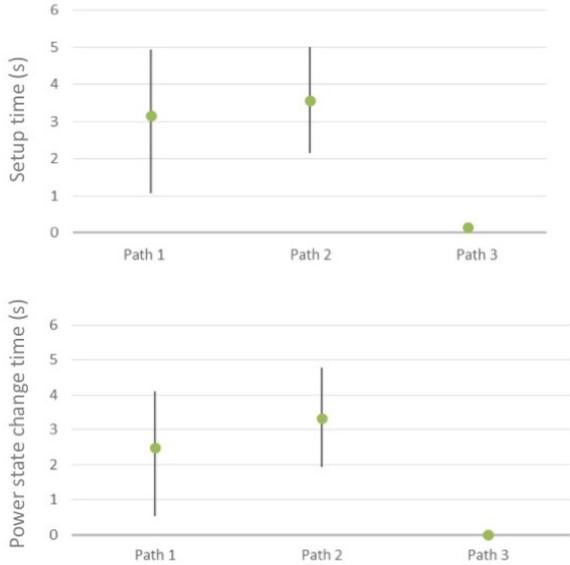


Figure 6: Average, max and min values of setup time (top) and time for power state change (bottom) for three paths

IV. EMMA FOR MMWAVE MESH NETWORKS

In the dense urban scenario, one of the important scenarios in 5G, network densification is necessary to accommodate the high traffic volume generated not only by smartphones/tablets but also by augmented reality information such as from sensors and wireless-connected cameras. Typical environments are shopping malls, airports, open squares, street canyons, etc., where mobile users tend to gather and move as large, dynamic crowds while expecting to keep their connectivity to the Internet.

In this section, we propose another EMMA algorithm to provide efficient deployment and management of a mmWave meshed network for such densely-deployed access networks, as the one depicted in Figure 7. In this scenario, mmWave nodes are overlaid on an LTE macro cell to play a role of both relay (i.e., XPFE) and (mmWave) access with three or four sectors for both access and backhaul/fronthaul [9]. The LTE macro base station also plays a role of mmWave gateway in the cell to accommodate time-variant and spatially non-uniform traffic by forming a mmWave meshed network dynamically. The prominent goal of EMMA in this scenario is to reduce the energy consumption of mmWave meshed networks by switching off as many mmWave nodes as possible in an area with low traffic demand [10]. The EMMA controller is located in the LTE macro BS and it is in charge of signaling the switch

On/Off control messages over the XCI. Signaling messages are sent over LTE as out-of-band control plane messages. Thus, EMMA can control the On/Off state of mmWave nodes and set up physical paths between them. In the following, we will evaluate the performance of our proposed EMMA algorithm for the mmWave meshed network scenarios by dynamically changing the density of users.

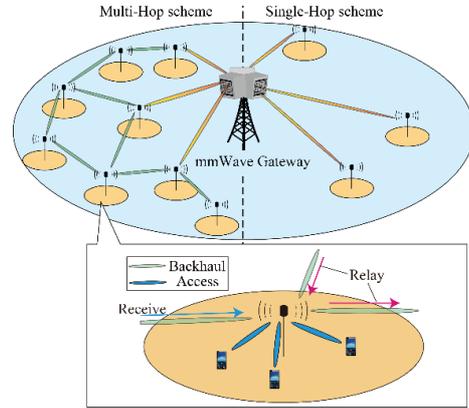


Figure 7: mmWave meshed network overlaid on a LTE macro cell.

The evaluation has been conducted with a system-level simulation, whose parameters are given in Table 2. In total, 90 mmWave nodes are deployed uniformly on a LTE macro cell with a radius of 250 m, and 5000 users are placed over the service area. To spatially reproduce non-uniform user distribution, the user location is assumed to follow a 2-D Gaussian distribution with mean μ and standard deviation σ . The mean μ corresponds to a hotspot location, where user density is higher, while σ captures the non-uniformity of user distribution.

Table 2: Simulation Parameters

Parameter	LTE macro	mmWave node
# of nodes	1	90
Frequency	2.1 GHz	60 GHz
Bandwidth	10 MHz	2 x 2.16 GHz
Antenna gain	17 dBi	26 dBi
Tx power	46 dBm	10 dBm

Figure 8 shows an example of the On/Off state of the mmWave nodes (marked as APs) and the physical paths established by the EMMA algorithm in the case of $\sigma = 100$ m with $\mu = (200, 0)$ m. In this example, the mmWave nodes located far from the hotspot are switched off, and mmWave backhaul resources are concentrated on the area of the hotspot to accommodate all traffic from/to users.

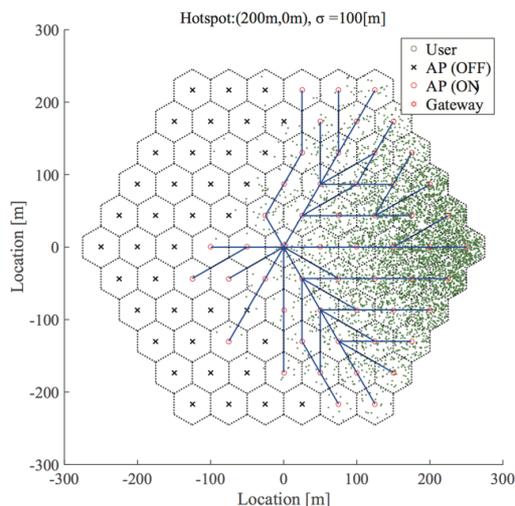


Figure 8: Map showing the On/Off state of mmWave nodes when $\sigma = 100m$.

The overall energy consumption of the mmWave meshed network is evaluated in Figure 9 by changing σ , to analyze the impact of the non-uniformity of user distribution. The smaller the σ , the less uniform the spatial distribution of users.

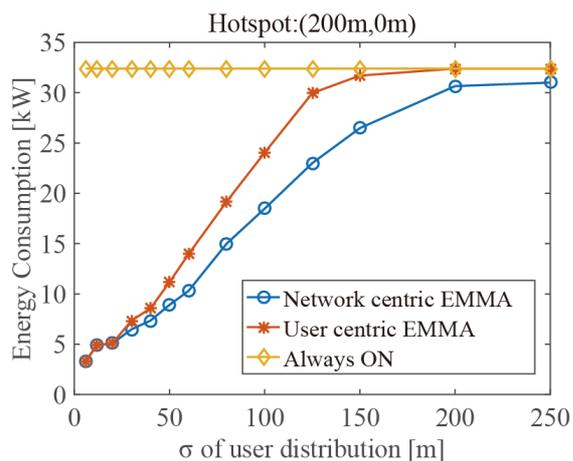


Figure 9: Energy consumption of network against σ .

Three approaches for mmWave node activation are compared. The first solution is our baseline scheme, denoted by “Always On”, namely without EMMA algorithm. The second and third ones are our proposed schemes: the “User centric EMMA” where mmWave nodes are turned on based on the location of users, and the “Network centric EMMA” where mmWave nodes are turned on to lower the overall energy consumption with minimum degradation of user experience. From the plot, it is obvious that EMMA has a large margin to save energy. Compared to the “always on” case, the EMMA algorithm can save nearly 60% of energy consumption when $\sigma = 50$ m. Besides, both user and network centric EMMA can save more energy when the distribution of users is less uniform.

It is to be noted that the network-centric EMMA is the most effective solution, by switching off more mmWave nodes in areas with low traffic load and forcing the corresponding users to associate with other mmWave nodes or LTE macro BS.

V. EMMA FOR HIGH-SPEED TRAINS

In high-speed train deployment, RoF is integrated as the fronthaul technology for ground-to-train radio access to overcome doppler effects. RoF nodes are deployed inside tunnels to provide constant coverage. Currently, the deployed nodes are active all the time regardless of the presence of the trains within the tunnel. This leads to a waste of energy, thus our goal is to develop a software-defined energy-efficient RoF management system in an attempt to move toward greener communication. Therefore, in this section, we propose an EMMA for the high-speed train scenario.

In this scenario, EMMA controls the power state of RoF whenever the presence of a high-speed train in close proximity of the nodes is reported. The application signals to the XCI that idle RoF nodes should be switched off when unused. The goal is to minimize the energy footprint of the deployed distributed RoF nodes in the crosshaul network without degrading the QoS of ground-to-train communication.

EMMA is composed of the following three modules:

- *Context Information Module*, in charge of gathering the context information related to train mobility to determine the current location of the train. It collects the information from the eNBs and stores the information in the database;
- *Statistics Module*, in charge of storage and retrieval of context information. It allows the XCI to periodically retrieve new records from the database and updates the XCI about the current location of the train;
- *Management Module*, responsible for the actual control of RoF nodes. It decides to switch on or off the RoF nodes via SNMP protocol based on the received context information.

An example scenario is described in Figure 10, which shows that the RoF nodes are connected to the eNB B and C.

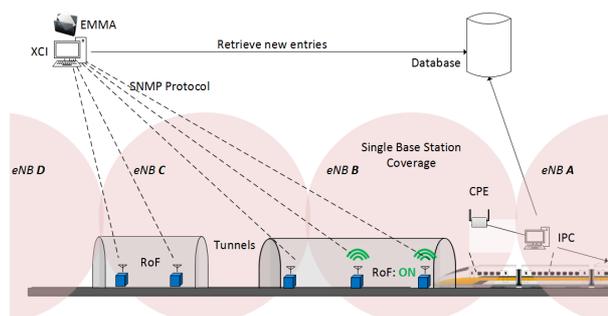


Figure 10: Scenario of EMMA-specific High-speed Train

As the moving train enters the coverage of eNB B, the context information log (such as the physical cell id) is pushed to the IPC (Industrial Personal Computer) server, which is installed on the train via CPE (Customer Premises Equipment). The IPC then extracts the relevant information and posts it to the database. The Database then notifies the EMMA application sitting on top of

XCI upon the reception of new entries. After the retrieval of records, the EMMA application decides based on the mapping table (RoF nodes mapped to Physical Cell ID of eNBs), if it needs to switch On or Off the connected RoF nodes via SNMP protocol and the specific eNB it has to associate to.

In our scenario, EMMA switches on the connected nodes as the moving train enters the coverage area and switches them off when it leaves. OpenDaylight is used for the XCI control plane which exposes two northbound APIs: snmp-get and snmp-set. Specifically, once the train enters the tunnel in the coverage area of eNB B, RoF nodes connected to eNB B are switched on via the SNMP protocol. When the train hands over from eNB B to eNB C, the RoF nodes connected to both eNBs are switched on without degrading the QoS of ground-to-train communication.

Considering R_E east-bound and R_W west-bound trains per day, and assuming that the serving time for each train by a single eNB is T_S seconds, the fraction of the idle time τ_{idle} of a RoF node in a day period can be expressed as:

$$\tau_{idle} = 1 - \frac{(R_E + R_W) \cdot T_S}{86400}$$

For instance, when $R_E=90$, $R_W=83$, and $T_S=50$, then each RoF node is idle for nearly 90% of the time. This behavior matches the design where RoF can be turned on only when there are trains to serve. The energy consumption for RoF nodes can be reduced to 10% compared to the current usage.

In our experimental setup, the offline mobility context information is used by EMMA to emulate a real-time scenario. In such a scenario, the high-speed train operates daily between 6 AM and 11 PM and RoF nodes are switched on constantly, regardless of operational time of the train. Figure 11 illustrates the comparison of energy consumption of RoF nodes with and without EMMA. The x-axis represents the time of day by hour increments. The y-axis represents the energy consumption per hour in percentage. With the EMMA-integrated solution, RoF nodes will be switched on only to serve the high-speed train when it is approaching, thus saving significant energy.

VI. CONCLUSIONS

In this paper, we have presented proof-of-concept prototypes of Energy Management and Monitoring Applications, applied to three realistic use cases: a software switch network, a mmWave mesh network, and an analogue RoF domain for ground-to-train radio access. We have shown the versatility of our solutions and their potential to curb the energy consumption in upcoming 5G networks, without compromising the user experience

ACKNOWLEDGMENT

This work has been supported by the H2020 project “5G-Crosshaul: The 5G Integrated fronthaul/backhaul” (671598).

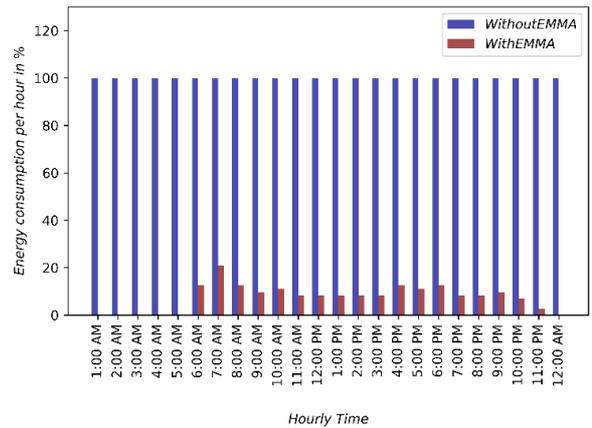


Figure 11: Comparison of energy consumption without (blue) and with (red) EMMA

REFERENCES

- [1] N. Bhushan, J. Li, D. Malladi, R. Gilmore, D. Brenner, A. Damjanovic, R. Sukhvasi, C. Patel, and S. Geirhofer, “Network densification: the dominant theme for wireless evolution into 5g,” *IEEE Communications Magazine*, vol. 52, no. 2, pp. 82–89, February 2014.
- [2] X. Li et al., “Innovations Through 5G-Crosshaul Applications,” *EUCNC 2017*.
- [3] S. S. Tadesse; C. Casetti; C.-F. Chiasserini; G. Landi, “Energy-efficient Traffic Allocation in SDN-based Backhaul Networks: Theory and Implementation”, *IEEE CCNC 2017*.
- [4] H. Wang, Y. Li, D. Jin, P. Hui, and J. Wu, “Saving energy in partially deployed Software Defined Networks,” *IEEE Trans. Comput.*, vol. 65, no. 5, pp. 1578–1592, 2016.
- [5] A. Fernández-Fernández; C. Cervelló-Pastor; L. Ochoa-Aday, “Energy-Aware Routing Algorithm in Software-Defined Networks”, *IEEE LCN 2016*.
- [6] R. Wang, Z. Jiang, S. Gao, W. Yang, Y. Xia, and M. Zhu, “Energy-aware routing algorithms in Software-Defined Networks,” *IEEE WoWMoM’14*, June 2014.
- [7] 5G-Crosshaul project: <http://xhaul.eu>
- [8] Chandramouli, et al., “Monitoring and Control MIB for Power and Energy”, *IETF RFC 7460*, March 2015.
- [9] K. Sakaguchi, G.K. Tran, H. Shimodaira, S. Nanba, T. Sakurai, K. Takinami, I. Siaud, E.C. Strinati, A. Capone, I. Karls, R. Arefi, and T. Haustein, “Millimeter-wave evolution for 5G cellular networks,” *IEICE Trans. Commun.*, vol. E98-B, pp. 388-402, Mar. 2015
- [10] H. Ogawa, G.K. Tran, K. Sakaguchi, T. Haustein, “Traffic adaptive formation algorithm for mmWave meshed backhaul networks,” accepted to *IEEE ICC 2017*, WDN-5G WS, May 2017.