Demo: BeGREEN Intelligence Plane for AI-driven Energy Efficient O-RAN management

Miguel Catalan-Cid, David Reiss *i2CAT Foundation, Spain* {miguel.catalan, david.reiss}@i2cat.net German Castellanos Accelleran, Belgium german.castellanos@accelleran.com Joss Armstrong Ericsson, Ireland joss.armstrong@ericsson.com

Abstract—Cellular networks management is being enhanced by O-RAN architecture and AI/ML solutions, enabling automated intelligent control loops for RAN optimization across various use cases. Ensuring energy sustainability is crucial to minimizing the impact of mobile networks on global energy consumption. This demo showcases the BeGREEN Intelligence Plane, an AI-driven solution for energy-efficient management of O-RAN networks. The presented workflow focuses on controlling the operational status of emulated cells, highlighting the integration of key components such as the AI Engine and the optimizations achieved through rApps and xApps.

Index Terms-O-RAN; Energy Efficiency; AI/ML; Demo;

I. INTRODUCTION

As emphasized in the IMT-2030 report by ITU [1], a key objective for 6G is to minimize network-wide energy consumption. To tackle this challenge, the O-RAN Alliance is developing control mechanisms to manage energy-saving features in multi-vendor environments. The incorporation of Artificial Intelligence and Machine Learning (AI/ML) will be essential to analyze historical data, adapt proactively to changing network conditions, and drive automated decision-making for energy optimization.

To address these challenges, the SNS BeGREEN project¹ proposes an Intelligence Plane, which works as a crossdomain management entity, integrating control and monitoring functions across RAN, Core and Edge domains, and fostering the creation of advanced ML models hosted in its AI Engine component. This paper presents the implementation of the integrated BeGREEN architecture demonstrating one of the targeted use cases, which aims at automating the operational state of network cells, dynamically adapting to forecasted traffic demands.

II. ARCHITECTURE

The demo showcases the three key components of the BeGREEN Intelligence Plane, as illustrated in Figure 1. The Non-Real-Time RAN Intelligent Controller (**Non-RT RIC**), hosts rApps for RAN optimization, utilizing AI/ML models

¹https://www.sns-begreen.com/



Fig. 1. Demo architecture including the BeGREEN Intelligence Plane

trained on data collected from the emulated E2 nodes. These models are deployed in the **AI Engine**, which leverages the MLRun² and Nuclio³ frameworks to enable serverless model inference through AI Engine Assist (AIA) rApps. AIA rApps decouple ML model management from control loops, acting as a proxy between ML models and control rApps. During inference, AIA rApps function as data producers, exposing model outputs via the Data Management and Exposure (DME) service of the R1 interface. In BeGREEN, this interface is implemented using the Information Coordinator Service (ICS) from the O-RAN Software Community [2], a data subscription platform that streamlines interactions between data producers and consumers. [3].

The AI Engine also provides **Energy Score and Energy Rating** functions to identify network areas or components requiring energy-saving policies and to assess the benefits of performed optimizations [3]. The Energy Score reports an absolute measure of energy efficiency (bits/Joule) based on data volume and energy consumption, while the Energy Rating monitors relative performance by comparing historical

This work is supported by the Smart Networks and Services Joint Undertaking (SNS JU) under the European Union's Horizon Europe research and innovation programme under Grant Agreement No 101097083, BeGREEN project. Views expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or SNS-JU. Neither the European Union nor the granting authority can be held responsible for them.

²https://www.mlrun.org/

³https://nuclio.io/

data from equivalent network entities, such as cells of the same vendor within the same area. Based on these insights, control rApps determine appropriate A1 energy-saving policies [4] and communicate them to the Near-RT RIC over the A1 interface. Both control and AIA rApps depend on RAN node data collected and exposed through a Key Performance Measurement (KPM) producer rApp, which connects to the Near-RT RIC via the O1 interface.

The **Near-RT RIC** handles fast telemetry and control operations for the RAN. It uses an Energy Savings xApp to manage RAN configurations and implement energy-saving policies received from the Non-RT RIC. Integrated into Accelleran's dRAXTM platform⁴, it features a Telemetry Collector that gathers data from both O-RAN and non-O-RAN compliant interfaces. This telemetry is processed into 3GPP-compliant messages, distributed via the dRAX data bus, exposed through the O1 interface, and preprocessed to be visualized in a Grafana dashboard. Additionally, the Near-RT RIC enforces E2-based RAN control and monitors RAN KPIs by interfacing with the emulated RAN environment.

III. DEMO DESCRIPTION

This demo showcases the complete workflow involving the components⁵ described in Section II and illustrated in Figure 1, demonstrating AI-driven management of network cells to improve energy efficiency while maintaining traffic performance. The scenario involves multiple 5G SA cells, with a subset performing as capacity layer cells that can be switched off during low-demand periods. Stationary and mobile UEs are introduced into the scenario to generate downlink traffic. The emulated RAN is implemented using Viavi's TeraVM AI RSG⁶. This tool facilitates the evaluation of rApps and xApps by emulating O-RAN compliant deployments. It also generates 3GPP-compliant KPMs, enabling the collection of data for AI model training and testing.

Three components are essential to realize the intelligent automated control loop in this scenario, as described below.

Cell Load-Energy ML Model: Supports control rApp decisions on cell operational status. Specifically, it includes a traffic load predictor trained to forecast the expected load for each cell. Managed by its associated AIA rApp, it requires several KPMs as input, including the cell identifier, current load demand, and the number of connected UEs. These inputs are periodically obtained through a subscription to the KPM producer rApp.

Cell control rApp: Implements the control logic to generate A1 energy-saving policies. Once deployed in the non-RT RIC k8s cluster, it creates subscriptions to the AIA and KPM producer rApps to gather the required inputs for decision-making. Based on the predicted load and the Energy Score/Rating of the cells, it identifies candidate cells to be switched on or off. The generated A1 policies adjust the percentage of energy



Fig. 2. Demo visualization

consumption gradually, as defined by O-RAN [4], until the cells are fully switched on or off. The Energy Score/Rating is then used as feedback to assess the efficiency of the applied policies.

Energy-saving xApp: Based on the received energy-saving policies, it determines the operational status of the cells and enforces control via the E2 interface. When policies include multiple cells, the xApp must select the required action for each cell to meet the policy target. Additionally, when a cell is switched off, the required handovers are managed by the Smart Handover xApp integrated into the near-RT RIC framework, with a conflict avoidance mechanism ensuring smooth network operation.

The demo provides insights into utilization of the DME services within the R1 interface, highlighting how data is exchanged among the involved rApps. This is illustrated through the non-RT RIC Swagger interface, depicted in Figure 2a. Additionally, Viavi's TeraVM AIA RSG dashboard (Figure 2b) presents the scenario details, including the status of cells and UEs. Finally, the near-RT RIC Grafana dashboard (Figure 2c) displays the received A1 policies and the evolution of KPMs during the demo runtime. Results show energy savings of up to 40% without degrading traffic allocation.

REFERENCES

- International Telecommunication Union Radiocommunication Sector (ITU-R). Future technology trends of terrestrial international mobile telecommunications systems towards 2030 and beyond. Report ITU-R M.2516-0, International Telecommunication Union, November 2022.
- [2] O-RAN Software Community. OSC Information Coordination Service. https://docs.o-ran-sc.org/projects/ o-ran-sc-nonrtric-plt-informationcoordinatorservice/en/latest/overview. html. Accessed: 2024-02-06.
- [3] Miguel Catalán-Cid (Ed.) and BeGREEN Consortium. D4.2: Initial Evaluation of BeGREEN O-RAN Intelligence Plane, and AI/ML Algorithms for NFV User-Plane and Edge Service Control Energy Efficiency Optimization. Technical report, BeGREEN Project, September 2024.
- [4] O-RAN A1 Interface: Type Definitions 9.0. Technical Specification R004, O-RAN Alliance, WG2, October 2024. O-RAN.WG2.A1TD-v09.00.

⁴https://accelleran.com/ran-intelligent-controller/

⁵The required components are hosted in different premises and interconnected through a VPN.

⁶https://www.viavisolutions.com/en-us/products/teravm-ai-rsg