

Demonstrating a Bayesian Online Learning for Energy-Aware Resource Orchestration in vRANs

Jose A. Ayala-Romero*, Andres Garcia-Saavedra[†], Xavier Costa-Perez^{‡,†}, George Iosifidis[§]

*Trinity College Dublin, Ireland, [†]NEC Laboratories Europe GmbH, Germany

[‡]i2CAT Foundation and ICREA, Spain, [§]Delft University of Technology, Netherlands

Abstract—Radio Access Network Virtualization (vRAN) will spearhead the quest towards supple radio stacks that adapt to heterogeneous infrastructure: from energy-constrained platforms deploying *cells-on-wheels* (e.g., drones) or battery-powered cells to *green edge clouds*. We demonstrate a novel machine learning approach to solve resource orchestration problems in energy-constrained vRANs. Specifically, we demonstrate two algorithms: (i) BP-vRAN, which uses Bayesian online learning to balance performance and energy consumption, and (ii) SBP-vRAN, which augments our Bayesian optimization approach with *safe controls* that maximize performance while respecting hard power constraints. We show that our approaches are *data-efficient*—converge an order of magnitude faster than other machine learning methods—and have provably performance, which is paramount for carrier-grade vRANs. We demonstrate the advantages of our approach in a testbed comprised of fully-fledged LTE stacks and a power meter, and implementing our approach into O-RAN’s non-real-time RAN Intelligent Controller (RIC).

I. INTRODUCTION

Virtualization is considered today the most promising approach for bringing cellular networks up to speed with contemporary service demands. The latest frontier is virtualizing the radio access network (vRAN) by turning the base stations (BSs) into fully-softwarized stacks that can be deployed in diverse platforms such as commodity servers, small embedded devices, or even moving nodes (*cells-on-wheels*). This paradigm shift is expected to offer the much-needed performance flexibility, facilitate network densification, and reduce capital and operating expenses [1]. Hence, not surprisingly, it has spurred numerous industry efforts to build BS stacks [2], fully-open RANs [3], and even launch extensive field tests [4].

However, designing and operating vRANs is far from trivial, since the virtualized BSs differ significantly from their hardware-based counterparts. On the one hand, vBSs are more controllable as one can tune their parameters (transmission power, modulation schemes, etc.) in real time based on the network needs. On the other hand, their softwarization and diverse platforms render less predictable their performance and power consumption. The latter is particularly important both for economic reasons but also because the vBSs, most often, operate under tight energy budgets. Hence, traditional radio control policies run the risk of under-using this new type of BSs, or rendering vRANs economically unsustainable.

Our experimental results in [1] shed light on the relationship between performance (throughput), power consumption, and vBS controls such as the modulation and coding schemes (MCS) and spectrum allocation. For instance, we find that the baseband unit (BBU) consumes power comparable to wireless

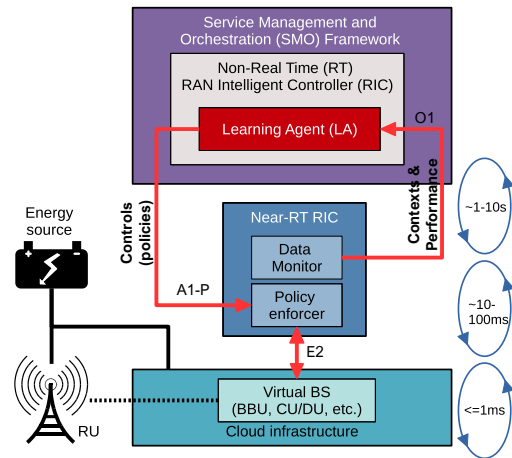


Fig. 1: O-RAN compliant system architecture and workflow.

transmissions, and the vBS power and throughput are affected by the configurations in a non-linear/non-monotonic fashion. These results depend also on the hosting platform.

The take-away message from these extensive measurements is that, *unlike legacy BSs*, virtualized BSs have a complex, poly-parametric, and platform-dependent performance and power consumption profile; and this renders traditional control policies inefficient. Therefore, we propose¹ a powerful machine learning framework that learns on-the-fly the vBS profiles and selects their optimal configurations based on the network needs and power availability. In [1], we formulate two energy-aware vBS control problems and design novel algorithms to solve them: (i) BP-vRAN, which finds the Pareto-optimal trade-off between performance and power consumption; and (ii) SBP-vRAN, which maximizes performance subject to *hard* constraints on the power source. The former allows operators to balance performance and power expenses, while the latter is crucial for vBS running on power-constrained platforms, e.g., power-over-ethernet (PoE) cells.

II. DEMONSTRATION

Our algorithms are founded on Bayesian theory, Gaussian Processes (GPs) and online learning. GPs model the behavior of the vBS in terms of performance and power consumption, using the collected measurements in real time. Accordingly, we use a contextual bandit approach to intelligently *explore* the space of vBS configurations, and *exploit* the best ones for each

¹This proposal demonstrates the algorithms proposed in our paper accepted for presentation at IEEE INFOCOM 2021 [1].

context, namely UL/DL traffic load and SNR patterns. The result is a non-parametric algorithmic framework that makes minimal assumptions about the system, adapts to user needs and network conditions, and *provably* maximizes performance. Furthermore, drawing ideas from *safe* Bayesian optimization, our SBP-vRAN algorithm ensures that the vBS power constraints are not violated even during exploration, hence enabling the vBS deployment on energy-constrained platforms. By design, this framework outperforms other approaches requiring knowledge of the vBS functions [5] or offline data to approximate them [6], and adaptive techniques that do not offer performance guarantees or have strict assumptions [7].

We consider a virtualized Base Station (vBS) comprising a Baseband Unit (BBU), which may correspond to a 4G eNB or 5G gNB² hosted in a cloud platform and attached to a Radio Unit (RU), which are fed by a common and possibly constrained energy source. This type of BSs is relevant for low-cost small cells, Power-over-Ethernet (PoE) cells, and so on. Our goal is to use O-RAN's control architecture to select and adapt radio policies to system dynamics satisfying different energy-driven criteria. Fig. 1 shows the high-level system architecture, which is O-RAN compliant [3]. The Learning Agent (LA) runs online algorithms within the Non-Real-Time (Non-RT) RAN Intelligent Controller (RIC) in the system's orchestrator, and sequentially selects efficient *radio policies* every orchestration period t (in the order of seconds) *given the current context*.

Our testbed is shown in Fig. 2 and comprises a vBS, the user equipment (UE)³, and a digital power meter. Both the vBS and the UE consist of an Ettus Research USRP B210 as RU, srseNB/srsUE (from srsLTE suite [2]) as BBU for both the eNB and UE, and two small factor general-purpose PCs (Intel NUCs with CPU i7-8559U@2.70GHz) deploying each respective BBU and the near-RT RIC of Fig. 1. The vBS and the UE are connected using SMA cables with 20dB attenuators, and we adjust the gain of the RU's RF chains to attain different SNR values. Without loss in generality, we select a 10-MHz band that renders a maximum capacity of roughly 32 and 23 Mbps in DL and UL, respectively. We use the power meter GW-Instek GPM-8213 to measure the power consumption of the BBU and the RU by plugging their power supply cable to a GW-Instek Measuring adapter GPM-001. Finally, we have integrated E2's interface and the ability to enforce control policies *on-the-fly* in srseNB.

We use three auxiliary PCs (not shown in the picture) hosting the non-RT RIC and the network traffic generators. We have finally implemented O1 interface (Fig. 1) using the USB-based power meter SCPI (Standard Commands for Programmable Instruments) interface concerning power consumption measurements and a REST interface for the remainder.

In our demonstration, we verify that both algorithms converge and find the optimal vBS configuration in a variety of scenarios. Using real traffic traces, we show, step-by-step, how our framework explores the configurations, and how it refrains

²5G decouples BBU in 2 logical functions, i.e., a central unit (CU) and a distributed unit (DU). Our scheme controls the DU, or both when co-located.

³We use one UE emulating the load of multiple users




Fig. 2: Experimental vBS and UE testbed.

from violating the power constraints when necessary. We also benchmark our solution with a state-of-the-art Reinforcement Learning (RL) solution. Namely, we implement a Deep Deterministic Policy Gradient (DDPG) algorithm using an actor-critic neural network (NN) architecture, which is presented in [7]. We find that our framework is more data-efficient than such RL approaches which require orders of magnitude more measurements (hence, also more time) to train their NNs.

The details of the design of BP-vRAN and SBP-vRAN are presented in [1]. Our software implementation and the dataset used in our paper are publicly available (see [1]).

III. CONCLUSIONS

We demonstrate an online learning framework to orchestrate vRAN resources in O-RAN's non-RT RAN Intelligent Controllers (RIC). We have two goals: (i) balance performance and power consumption in unconstrained platforms such as data centers; and (ii) maximize performance subject to power constraints vBS, e.g., solar-powered platforms or cells-on-wheels. We have followed a judicious design approach by resorting to Bayesian learning theory. This methodology allowed us to derive two algorithms, BP-vRAN and SBP-vRAN, that achieve the aforementioned goals (i) with theoretical performance guarantees, (ii) with high data-efficiency and convergence speed, and (iii) respecting power constraints even during learning.

ACKNOWLEDGMENTS

This work was supported by the EC through Grant No. 856709 (5Growth) and Grant No. 101017109 (DAEMON); and by SFI through Grant No. SFI 17/CDA/4760.

REFERENCES

- [1] J. A. Ayala, A. Garcia-Saavedra, X. Costa-Perez, and G. Iosifidis, "Bayesian Online Learning for Energy-Aware Resource Orchestration in Virtualized RANs," in *IEEE INFOCOM '21 (to appear)*.
- [2] I. Gomez-Migueluez *et al.*, "srsLTE: an Open-source Platform for LTE Evolution and Experimentation," in *Proc. of ACM WinTech*, 2016.
- [3] O-RAN Alliance, "O-RAN-WG1-O-RAN Architecture Description - v01.00.00." Technical Specification, February 2020.
- [4] "Reimagining the End-To-End Mobile Network in the 5G Era," Cisco, Rakuten, Altiostar. White Paper, 2019.
- [5] D. Bega *et al.*, "CARES: Computation-aware Scheduling in Virtualized Radio Access Networks," *IEEE Trans. on Wireless Communications*, vol. 17, no. 12, pp. 7993–8006, 2018.
- [6] D. Raca *et al.*, "On Leveraging Machine and Deep Learning for Throughput Prediction in Cellular Networks: Design, Performance, and Challenges," *IEEE Communications Magazine*, vol. 58, no. 3, pp. 11–17, 2020.
- [7] J. A. Ayala-Romero *et al.*, "vrAI: A Deep Learning Approach Tailoring Computing and Radio Resources in Virtualized RANs," in *Proc. of ACM MOBICOM*, 2019.