

vrAIIn: Deep Learning based Orchestration for Computing and Radio Resources in vRANs

AYALA-ROMERO Jose, GARCIA-SAAVEDRA Andres, COSTA PEREZ Xavier

Abstract

We present vrAIIn, a resource orchestrator for vRANs based on deep reinforcement learning. First, we use an auto-encoder to project high-dimensional context data into a latent representation. Then, we use a deep deterministic policy gradient (DDPG) algorithm based on an actor-critic neural network structure and a classifier to map contexts into resource control decisions. Our results show that: (i) vrAIIn provides savings in computing capacity of up to 30% over CPU-agnostic methods; (ii) it improves the probability of meeting QoS targets by 25% over static policies; (iii) upon computing capacity under-provisioning, vrAIIn improves throughput by 25% over state-of-the-art schemes; and (iv) it performs close to an optimal offline oracle.



O-RAN, vRAN, deep reinforcement learning, resource control, CAPEX/OPEX reduction

1. Introduction

Radio Access Network virtualization (vRAN) is well-recognized as a key technology to accommodate the ever-increasing demand for mobile services at an affordable cost for mobile operators. vRAN centralizes software-defined radio access point (RAP)* stacks into computing infrastructure in a cloud location — typically at the edge, where CPU resources may be scarce. **Fig. 1** illustrates a set of vRAPs sharing a common pool of CPUs to perform radio processing tasks such as signal modulation and encoding (light gray arrows). This provides several advantages, such as resource pooling (via centralization), simpler update roll-ups (via softwareization) and cheaper management and control (via commoditization).

It is thus not surprising that vRAN has attracted the attention of academia and industry. O-RAN or Rakuten Mobile's vRAN — led by key operators (such as AT&T, Verizon or China Mobile), manufacturers (such as Intel, Cisco or NEC) and research leaders (such as Stanford

University) — are examples of publicly disseminated initiatives towards fully programmable, virtualized and open RAN solutions based on general-purpose processing platforms and decoupled base band units (BBUs) and remote radio units (RRUs).

Despite the above, the gains attainable today by vRAN are far from optimal, and this hinders its deployment at scale. In particular, computing resources are inefficiently pooled since most implementations over-dimension computational capacity to cope with peak demands in real-time workloads.

Dynamic resource allocation in vRAN is an inherently hard problem:

- (1) The computational behavior of vRAPs depends on many factors, including the radio channel conditions or users' load demand, that may not be controllable. More specifically, there is a strong dependency with the context (such as data bit-rate load and signal-to-noise-ratio (SNR) patterns), the RAP configuration (e.g., bandwidth, MIMO

* The literature uses different names to refer to different radio stacks, such as base station (BS), eNodeB (eNB), gNodeB (gNB), access point (AP), etc. We will use RAP consistently to generalize the concept. made to the content, wording, positioning, priority, and emphasis of the constituent elements.

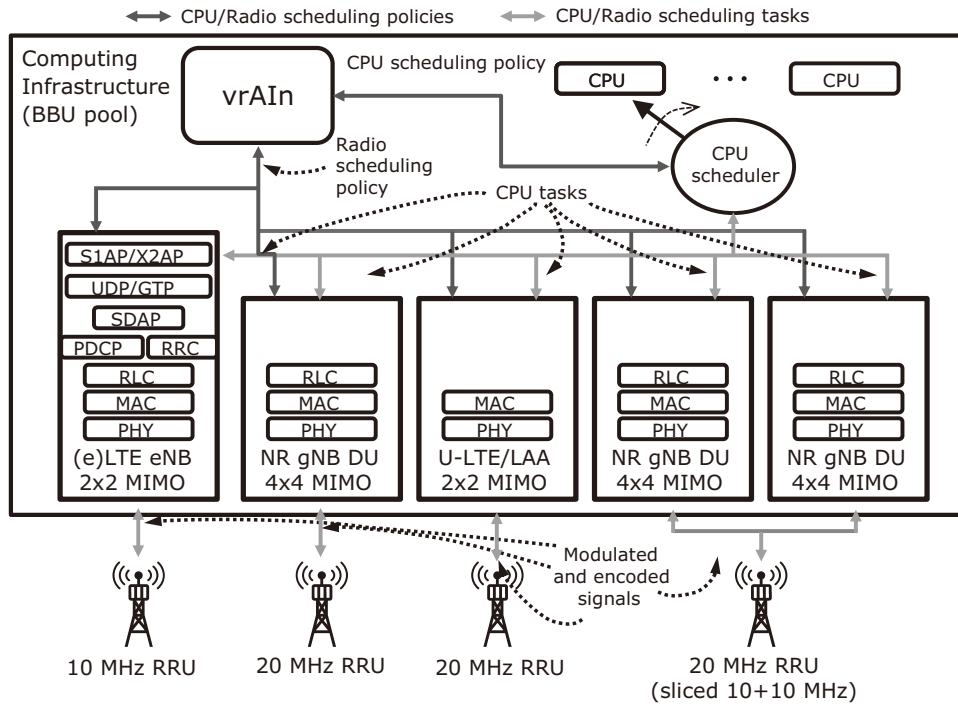


Fig. 1 Overview of vRAN using vrAIn.

setting, etc.) and on the infrastructure pooling computing resources.

- (2) Upon shortage of computing capacity (e.g., with nodes temporarily overloaded due to orchestration decisions) CPU control decisions and radio control decisions (such as scheduling and modulation and coding scheme (MCS) selection) are coupled.

We present vrAIn, an artificial intelligence-powered (AI) vRAN orchestrator that governs the allocation of computing and radio resources (dark gray arrows in Fig. 1).

2. vrAIn design

The vRAN landscape powered by vrAIn (Fig. 1) consists of a feedback control loop where:

- (1) Contextual information (SNR and data load patterns) is collected and encoded.
- (2) An orchestrator that maps contexts into computing and radio scheduling policies.
- (3) A reward signal assesses the decisions taken and fine-tunes the orchestrator accordingly.

We hence formulate our resource control problem as a contextual bandit (CB) problem, a sequential decision-making problem where, at every time stage $n \in \mathbb{N}$, an agent observes a context or feature vector drawn from an arbitrary feature space $\mathbf{x}^{(n)}$, chooses an action $\mathbf{a}^{(n)}$ and receives a reward signal $r(\mathbf{x}^{(n)}, \mathbf{a}^{(n)})$ as feedback.

Context Space: At each stage, T context samples are collected. Each sample consists of the buffer size, the mean SNR, and the variance SNR, measured for all users across all vRAPs.

Action Space: Our action space comprises all pairs of CPU and radio scheduling policies. $c_i^{(n)} \in \mathcal{C}$ and $m_i^{(n)} \in \mathcal{M}$ denote, respectively, the maximum computing time share (CPU scheduling policy) and the maximum MCS (radio scheduling policy) allowed to vRAP i in stage n . We also let $c_0^{(n)}$ denote the amount of CPU resource left unallocated (to save costs). Thus, a resource allocation action on vRAP i consists of a pair $a_i := \{c_i, m_i\}$ and a system action $\mathbf{a} = (a_i)_{\forall i \in \mathcal{P}} \in \mathcal{A} := \{(c_i \in \mathcal{C}, m_i \in \mathcal{M})\}_{\forall i \in \mathcal{P}}$.

Reward: The objective in the design of vrAIn is twofold:

- (1) When the CPU capacity is sufficient, the goal is to minimize the operation cost (in terms of CPU usage) as long as vRAPs meet the desired performance.
- (2) When there is a deficit of computing capacity to meet such performance target, the aim is to avoid decoding errors that lead to resource wastage, thereby maximizing throughput and minimizing delay.

To meet this objective, we design the reward function as follows. Let q_{i,x_i,a_i} be the (random) variable capturing the aggregate buffer occupancy across all users of vRAP i




Fig. 2 vrAIIn Resource Orchestrator.

given context x_i and action a_i at any given slot. As a quality-of-service (QoS) criterion, we set a target buffer size Q_i for each vRAP. Note that this criterion is closely related to the latency experienced by end-users (low buffer occupancy yields small latency) and throughput (a high throughput keeps buffer occupancy low). Thus, by setting Q_i , an operator can choose the desired QoS, which can be used to, e.g., provide differentiation across network slices. We let $J_i(x_i, a_i) := \text{Prob}[q_{i,x_i,a_i} < Q_i]$ be the probability that q_{i,x_i,a_i} is below the target per vRAP i and define reward as:

$$r(\mathbf{x}, \mathbf{a}) := \sum_{i \in \mathcal{P}} J_i(x_i, a_i) - M\varepsilon_i - \lambda c_i$$

where ε_i is the decoding error probability of vRAP i (which can be measured locally), and M and λ are parameters that determine the weight of decoding errors and the trade-off between resource usage and performance, respectively. We set M to a large value to avoid decoding errors due to low allocations and λ to a small value to meet QoS requirements (while minimizing the use of compute resources).

To solve our CB problem, we design the system depicted in **Fig. 2**. First, we encode each context $\mathbf{x}^{(n)}$ into a latent representation using Sparse Autoencoders (SAE), which yield a lower-dimensional vector $\mathbf{y}^{(n)}$. Then, we decouple both action policies into two orchestrators: CPU orchestrator and Radio orchestrator. We build the radio orchestrator as a simple classifier ν that, given an encoded context $\mathbf{y}^{(n)}$, provides the maximum MCS that can be processed reliably given computing resource policy $\mathbf{c}^{(n)}$. Conversely, we design our CPU orchestrator with a deep deterministic policy gradient (DDPG) algorithm using a model-free actor-critic neural network structure as shown in the figure.

3. Evaluation

To assess the performance attained by vrAIIn in a real-world scenario, we simulate vrAIIn over a production RAN deployment in a European city, with 197 access points distributed as shown in **Fig. 3** (a). As we can observe from the figure, there is a higher density of RAPs




Fig. 3 (a) Deployment of an operational RAN in Romania. Gray and black dots represent, respectively, radio sites and backhaul aggregation nodes. (b) Traffic load pattern over a period of 24h of a regular weekday.

Table Simulation parameters.

vRAP bandwidth	10 MHz
Subframe duration	1 ms
Transmission power	46 dBm
Antenna pattern	$A_n(\varphi) = -\min(12 \frac{\varphi^2}{\varphi_{3dB}}, A_m)$ $\varphi_{3dB} = 70^\circ$ $A_m = 25$ dB
Antenna gains	14 dBm
Path loss	$128.1 + 37.6 \cdot \log_{10}(R[\text{Km}])$, R: vRAP to UE distance
Shadow fading	Lognormal distribution 10 dB standard deviation
Thermal noise	-176 dBm
Number of vRAPs	197
Mean Nr. UEs per vRAP	10
UE mobility	20% @ 100 km/h, 80% @ 3 Km/h

in the center (a big city) and the RAN is sparser by the outskirts (covering mostly highways and small commuter suburbs). In order to leverage our training data, and without loss in generality, we assume all RAPs are SISO 10-MHz LTE vRAPs with the same behavior as our LTE vRAPs analyzed in Section 1.

Our custom-built simulator follows the 3GPP guidelines for LTE performance evaluation and its parameters are detailed in **Table**. The Signal-to-interference-plus-noise-ratio (SINR) perceived by the UEs is obtained by aggregating the interference of all active RAPs. For a given SINR, we compute the CQI of this UE, and then the maximum allowed MCS associated with this CQI according to 3GPP specification. Further, we implement a random mobility model for the UEs, ensuring a minimum distance between UE and RAP of 35 m, as recommended by 3GPP. Finally, due to the difficulty to capture the computing behavior of our vRAPs in a tractable model for simulation, we have implemented a deep neural network (DNN), trained using our dataset,

to determine when decoding errors occur due to lack of computing resources.

We simulate one regular week using synthetic traffic patterns, which emulate the behavior of real RANs at scale. To simplify the analysis, in the following we focus on a 24-hour period during weekdays — with the traffic profile (relative to the capacity of the system) shown in Fig 3 (b). We further assume that the aggregate computing capacity of the whole vRAN is dimensioned to the minimum amount of CPU resources required such that no violations of the encoding deadlines due to CPU deficit occur during the load peak of the day when not using vrAI, which we refer to as “100% provisioning”. We also study the behavior of vrAI when the system is under-provisioned to 70% and 85% of that computing capacity, which enables capital cost savings.

3.1 Provisioning computing capacity for the peak

Let us first focus on the evolution of throughput performance, buffer states and cost savings when the computing capacity is provisioned to accommodate the

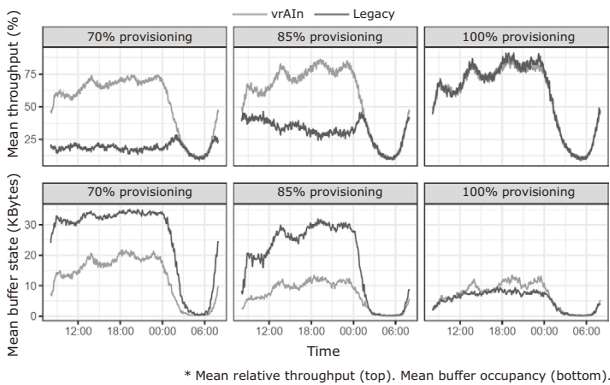


Fig. 4 Performance evolution over time for a computing capacity dimensioned for the peak load (“100% provisioning”) and for 70% and 85% of that computing capacity (“70% provisioning” and “85% provisioning”, respectively).




Fig. 5 vrAI’s savings relative to a static computing-agnostic approach provisioned with sufficient CPU resources for the peak demand.

peak load (right-most plots in Fig. 4 and Fig. 5 labeled as “100% provisioning”). From Fig. 4 we observe that vrAI achieves roughly the same throughput performance and slightly higher buffer sizes (up to 5%) than “Legacy”. This is explained because vrAI trades off this slightly higher delay for substantial OPEX savings. This is shown in Fig. 5 (right-most plot, labeled as “100% provisioning”), where vrAI achieves between 10% and 20% of OPEX savings. Note however that this difference in delay vanishes when vrAI is configured to favor performance over OPEX savings (results omitted here for the sake of space).

3.2 Under-provisioning of computing capacity

We now analyze the case where we impose under-provisioning to obtain aggressive CAPEX gains by reducing the availability of computing capacity to 85% and 70% relative to the dimensioning strategy discussed before, and labeled as “85% provisioning” and “70% provisioning”, respectively, in Fig. 4 and Fig. 5. The evolution of throughput and buffer states in Fig. 4 for these two scenarios (left-most and middle plot, respectively) make it evident how vrAI enables aggressive CAPEX savings while retaining high performance gains when compared to legacy computing-agnostic approaches. Specifically, vrAI provides up to and throughput gains over our legacy scheme when, respectively, the computing capacity is under-provisioned to 85% and 70% of the peak load, in average, and during the time of peak demand (between 18:00 and 00:00). The reason lies on the fact that vrAI’s ability to optimize jointly radio and computing scheduling policies allows vrAI to better accommodate the demand along the time domain, trading off delay for near-zero violations of the encoding deadlines due to a deficit of computing capacity. In contrast, legacy’s agnostic behavior with respect to the availability of computing capacity yields substantial throughput loss due to a high rate of violations of the encoding deadlines during instantaneous peak demands. This produces a cascade effect causing large amounts of time wasted in re-transmissions and rendering even larger perceived latency: up to and higher buffer occupancy over vrAI for 85% and 70% of computing provisioning, respectively, in average, and in the same period of peak demand.

A final remark is that vrAI adapts, i.e., maximizes performance, to constrained computing environments. This can be observed by comparing the performance indicators of both “70% provisioning” and “85% provisioning” to those shown for “100% provisioning”. In particular, vrAI has no loss in throughput performance for “85% provisioning” and only up to 10% throughput loss

for “70% provisioning”, in average, and during the same period of peak demand discussed earlier. This contrasts with the substantial throughput loss attained by our legacy approach: up to 65% and 80% of throughput loss in those same conditions. In terms of latency, vrAI only suffers from 1% increased in mean buffer occupancy in the case of “85% provisioning” relative to the buffer occupancy for “100% provisioning” (in contrast to the 282% increase of the legacy approach), and 73% increased in mean buffer occupancy with 75% under-provisioning (in contrast to the 337% increase of the legacy approach).

-
- * Intel and the Intel logo are trademarks of Intel Corporation or its subsidiaries.
 - * Cisco, Cisco Systems, and the Cisco logo are trademarks or registered trademarks of Cisco Systems, Inc. in the United States and other countries.
 - * The names O-RAN ALLIANCE, O-RAN and their logo are trademarks or registered marks of O-RAN ALLIANCE e.V.
 - * Rakuten Mobile is a trademark of Rakuten Group, Inc. or its affiliates.
 - * All other company names and product names that appear in this paper are trademarks or registered trademarks of their respective companies.

Authors' Profiles

AYALA-ROMERO Jose

Senior Researcher
6G Networks
NEC Laboratories Europe

GARCIA-SAAVEDRA Andres

Principal Research Scientist
6G Networks
NEC Laboratories Europe

COSTA PEREZ Xavier

Manager
6G Networks
NEC Laboratories Europe

Information about the NEC Technical Journal

Thank you for reading the paper.

If you are interested in the NEC Technical Journal, you can also read other papers on our website.

Link to NEC Technical Journal website

Japanese

English

Vol.17 No.1 Special Issue on Open Network Technologies

– Network Technologies and Advanced Solutions at the Heart of an Open and Green Society

Remarks for Special Issue on Open Network Technologies
NEC's Technological Developments and Solutions for Open Networks

Papers for Special Issue

Open RAN and Supporting Virtualization Technologies

Innovations Brought by Open RAN
Reducing Energy Consumption in Mobile Networks
Self-configuring Smart Surfaces
Nuberu: Reliable RAN Virtualization in Shared Platforms
vrAIn: Deep Learning based Orchestration for Computing and Radio Resources in vRANs

Wireless Technologies for 5G/Beyond 5G

NEC's Energy Efficient Technologies Development for 5G and Beyond Base Stations toward Green Society
Millimeter-wave Beamforming IC and Antenna Modules with Bi-directional Transceiver Architecture
Radio-over-Fiber Systems with 1-bit Outphasing Modulation for 5G/6G Indoor Wireless Communication
28 GHz Multi-User Massive Distributed-MIMO with Spatial Division Multiplexing
28 GHz Over-the-Air Measurements Using an OTFS Multi-User Distributed MIMO System
Comprehensive Digital Predistortion for improving Nonlinear Affection and Transceivers Calibration to Maximize Spatial Multiplexing Performance in Massive MIMO with Sub6 GHz Band Active Antenna System
Black-Box Doherty Amplifier Design Method Without using Transistor Models
39 GHz 256 Element Hybrid Beam-forming Massive MIMO for 8 Multi-users Multiplexing

Initiatives in Open APN (Open Optical/All Optical)

NEC's Approach to APN Realization — Towards the Creation of Open Optical Networks
NEC's Approach to APN Realization — Features of APN Devices (WX Series)
NEC's Approach to APN Realization — Field Trials
Wavelength Conversion Technology Using Laser Sources with Silicon Photonics for All Photonics Network
Optical Device Technology Supporting NEC Open Networks — Optical Transmission Technology for 800G and Beyond

Initiatives in Core & Value Networks

Technologies Supporting Data Plane Control for a Carbon-Neutral Society
NEC's Network Slicing Supports People's Lives in the 5G Era
Application-Aware ICT Control Technology to Support DX Promotion with Active Use of Beyond 5G, IoT, and AI
Using Public Cloud for 5G Core Networks for Telecom Operators

Enhancing Network Services through Initiatives in Network Automation and Security

NEC's Approach to Full Automation of Network Operations in OSS
Autonomous Network Operation Based on User Requirements and Security Response Initiatives
Enhancing Information and Communications Networks Safety through Security Transparency Assurance Technology
Enhancing Supply Chain Management for Network Equipment and Its Operation

Network Utilization Solutions and Supporting Technologies

Positioning Solutions for Communication Service Providers
The Key to Unlocking the Full Potential of 5G with the Traffic Management Solution (TMS)
Introducing the UNIVERGE RV1200, All-in-one Integrated Compact Base Station, and Managed Services for Private 5G
Vertical Services Leveraging Private 5G to Support Industrial DX
Integrated Solution Combining Private 5G and LAN/RAN

Global 5G xHaul Transport Solutions

xHaul Solution Suite for Advanced Transport Networks
xHaul Transformation Services
xHaul Transport Automation Solutions
Fixed Wireless Transport Technologies in the 5G and Beyond 5G Eras
SDN/Automation for Beyond 5G
OAM Mode-Multiplexing Transmission System for High-Efficiency and High-Capacity Wireless Transmission

Toward Beyond 5G/6G

NEC's Vision and Initiatives towards the Beyond 5G Era

NEC Information

2022 C&C Prize Ceremony



Vol.17 No.1
September 2023

Special Issue TOP