

Learning-Based Orchestration for Dynamic Functional Split and Resource Allocation in vRANs

Fahri Wisnu Murti[†], Samad Ali[†], George Iosifidis*, Matti Latva-aho[†]

[†]Centre for Wireless Communications, University of Oulu, Finland

*Delft University of Technology, Netherlands

Abstract—One of the key benefits of virtualized radio access networks (vRANs) is network management flexibility. However, this versatility raises previously-unseen network management challenges. In this paper, a learning-based zero-touch vRAN orchestration framework (LOFV) is proposed to jointly select the functional splits and allocate the virtualized resources to minimize the long-term management cost. First, testbed measurements of the behaviour between the users' demand and the virtualized resource utilization are collected using a centralized RAN system. The collected data reveals that there are non-linear and non-monotonic relationships between demand and resource utilization. Then, a comprehensive cost model is proposed that takes resource overprovisioning, declined demand, instantiation and reconfiguration into account. Moreover, the proposed cost model also captures different routing and computing costs for each split. Motivated by our measurement insights and cost model, LOFV is developed using a model-free reinforcement learning paradigm. The proposed solution is constructed from a combination of deep Q-learning and a regression-based neural network that maps the network state and users' demand into split and resource control decisions. Our numerical evaluations show that LOFV can offer cost savings by up to 69% of the optimal static policy and 45% of the optimal fully dynamic policy.

I. INTRODUCTION

Virtualizing the radio access network (vRAN) is one of the most promising technologies for accommodating the increased service demands with diverse requirements at a reasonable cost [1]. Its latest development enables the base station (BS) functions to turn into virtualized components that can be executed across cloud platforms. This paradigm brings unprecedented flexibility to RAN operations, mitigates vendor lock-in, offers fast deployment and potentially reduces operational expenses [2]. Therefore, it is not surprising that many standardization bodies adopt virtualization for Next-Generation RANs such as Open RANs [3] and 5G+ RANs [4].

In vRANs, the virtualized BS (vBS) functions can be disaggregated then hosted at virtualized distributed units (vDUs) and central units (vCUs) through *functional split*. The network operators can flexibly deploy their vBS functions over vCUs and vDUs based on their resource availability and the network load, facilitating cost-efficient and high-performance RANs operation [5]. However, this flexibility also carries non-trivial decisions on splitting the functions and allocating the virtualized computing resources¹ to implement these functions. Each split induces different fronthaul/backhaul (xHaul) load and virtualized resources. The suitability of each split also

depends on the network properties (xHaul capacity, computing capacity, etc.) and might change abruptly over time due to the varying user needs. Therefore, it is not only important to design the splits and resource allocation in vRANs carefully, but also to update these decisions and reconfigure the system in order to adapt to varying conditions (resource availability and user needs). Otherwise, we risk inducing high operating expenditures and performance degradation.

On the other hand, orchestrating the dynamic split selection and virtualized resource allocation is also a challenging problem as the decisions take place before the actual users' traffic is observed. Hence, there is a risk for resource *overprovisioning* (e.g., the allocated resources are higher than the actual usages resulting in waste resources) and *underprovisioning* (e.g., insufficient allocated resources resulting in declined users' demand). Meanwhile, reconfiguring the splits and virtualized resources at runtime can induce additional costs, potentially disrupts network operations during the migration of the virtual machines (VMs) [6], and *is therefore not always beneficial*. Conversely, deploying a static policy can not unleash the potential flexibility of the vRAN system. Therefore, it is necessary to manage and reconfigure the splits and virtualized resources in an intelligent manner.

The authors of [6] have experimentally analyzed the migration activity of VMs in vRANs and have successfully validated the possibility of deploying an adaptive functional split practically, albeit have not discussed designing the cost-efficient split reconfiguration policies. Recent work in [7] has proposed a framework for the split that dynamically adapts its configuration at runtime to maximize the users throughput. Similar works have proposed flexible split selection to minimize the inter-cell interference and fronthaul utilization [8] and the network cost [5]. However, these works assume complete knowledge models of the split performance and resource utilization. We argue that such approaches can be inaccurate in practice as vRANs softwarization is deployed together with other workloads in the diverse cloud platform, which are hardly precise in predicting its resources and behaviour.

Using Machine Learning (ML) techniques for tackling resource allocation problems in wireless networks is becoming increasingly popular [9]. The authors in [10] have proposed a learning framework that successfully manages the interplay between computing and radio resources. It models the problem as a contextual bandit, then utilizes an actor-critic neural network structure and a classifier to map contexts into resource control decisions. The follow-up work [11] has studied an energy-aware resource orchestration that uses Bayesian online learning to balance performance and energy

This research has been supported by the Academy of Finland, 6G Flagship program under Grant 346208.

¹It is common in Network Function Virtualization (NFV) that Virtualized Network Functions (VNFs) needs CPU, memory, I/O, and GPU for their virtualized computing resources. The operators typically use these parameters to calculate their billing units to charge the amount of monetary cost.

consumption. The authors in [12] have proposed an ML-based predictor that learns to share the unutilized CPU resources with the other workloads such as video analytics. Recent works have brought the importance of ML-based optimization [13] and dynamic selection [14] of functional splits, albeit not discussing virtualized resource management. Although the mentioned works have addressed complex vRANs problems, they still did not answer how to design a framework that intelligently decides the splits and allocates the virtualized resources. Moreover, they also did not consider the effects of resource reconfiguration in their vRAN problems.

Contributions. We firstly use a vBS prototype implementing the srsRAN platform [15] to collect measurements regarding the behaviour between the users' demand and the virtualized resource utilization (details in Sec. IV). Our findings suggest that this behaviour varies with the demand and the platform resources and, importantly, is *non-linear* and *non-monotonic*; hence it is hard to model the underlying system precisely. Inspired from network slicing [16], we also propose a new cost model accounting for resource overprovisioning, declined service demand, reconfiguration and instantiation, representing the virtualized resource management in vRANs. Besides, depending on the splits, the transferred load between vCU, vDU and the radio unit (RU) induces a different cost for reserving the xHaul link bandwidth.

Our goal is to develop a learning-based zero-touch orchestration framework (LOFV) that intelligently selects the splits and allocates the virtualized resources to minimize the long-term management cost while serving the users' demand. We model the vRAN operation as a time-slotted system, where each slot has an arbitrary incoming users' demand and network state conditions. At every time stage, LOFV decides whether to preserve the previous network settings or reconfigure them by reselecting the splits and reallocating virtualized resources. We formulate this sequential decision-making problem as Markov Decision Process (MDP). In our solution, LOFV is tailored from a model-free reinforcement learning paradigm that does not make particular assumptions about the underlying system and state transition. It is constructed from two functional blocks based on neural network structures: *i*) the functional split orchestration and *ii*) virtualized resource orchestration. The functional split orchestration is a function that maps the input state into the selected network setting and deployed split. It is constructed from a deep Q-network (DQN) and target network [17], and utilizes Q-learning [18] as the learning step. The virtualized resource orchestration is a regression-based neural network that maps the input of split selection and users' demand into resources at vDU and vCU. This function utilizes α -OMC loss function [19] that considers the penalty fee incurred from prediction errors of virtualized resources.

We conduct a battery of tests using our collected measurements from container-based virtualization of srsRAN [15]. We evaluate our vBS measurements, learning convergence of LOFV and long-term accumulated cost during the online stages. Our evaluations show that the cost-saving of LOFV can be as high as 69% of the optimal static policy (STAO)²

²STAO is the optimal static policy that selects the single best split and resource allocation based on the peak traffic. Consequently, STAO only incurs overprovisioning and xHaul costs. We use STAO to normalize all of our monetary cost evaluations.

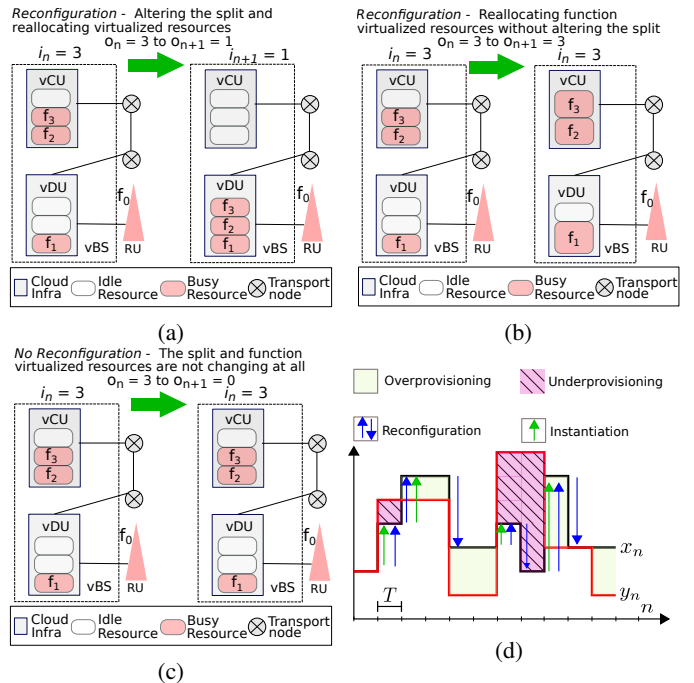


Fig. 1: **Virtualized Resource Management.** Reconfiguration activity occurs when the operators decide to: (a) alter the deployed split and reallocate the virtualized resources; (b) only reallocate the resources but keep the deployed split. Otherwise, (c) the operators may not want to reconfigure their resources and the deployed split at all. (d) The virtualized resource behaviour.

and 45% of the optimal fully dynamic policy (DYNO)³.

The rest of this paper is organized as follows. The background, model and trade-offs are presented in Sec. II. The problem formulation and LOFV framework are discussed in Sec. III. Our detailed experiments and results are in Sec. IV and the conclusion is in Sec. V.

II. MODEL AND TRADE-OFFS

A. Background

Our model refers to the O-RAN compliant system model [3]. We consider a vBS comprising a vCU and vDU connected to RU, corresponds to 4G eNodeB or 5G gNodeB. The vDU is typically hosted at a far-edge cloud and vCU is at an edge cloud. Let suppose f_0 is a function that encapsulates RF. Then, we denote f_1, f_2 and f_3 for the respective functions of Layer 1 (PHY), Layer 2 (MAC, RLC) and Layer 3 (PDCP, RRC, GTP). These functions can be deployed at the vDU and vCU (except f_0) following a chain: $f_0 \rightarrow f_1 \rightarrow f_2 \rightarrow f_3$. We consider four split options that have been well standardized [4], [20] and experimentally validated as a prototype [6]. **Split 1 (S1):** All functions are at vDU except f_0 is at RU (a fully distributed-RAN). **Split 2 (S2):** f_3 is deployed at vCU, but f_1 and f_2 are at vDU. **Split 3 (S3):** f_2 and the higher layer are at vCU, while f_1 is at vDU. **Split 4 (S4):** All functions are at vCU except f_0 (a fully centralized vRAN). Hence, we define the respective split $i \in \mathcal{I} = \{1, 2, 3, 4\}$. LOFV is to be executed from the Learning Agent (LA) inside Non-Real-Time (Non-RT) RAN Intelligent Controller (RIC).

³It knows an oracle of users' demand and resource availability. It uses a fully dynamic policy by always reconfiguring the vRAN settings at every time stage to obtain the current optimal, i.e., the current best split and resources.

B. Virtualized Resource Management

Let λ_n (Mbps) denote the incoming traffic demand from multiple users at time stage n which has an interval of duration T . At the beginning of each interval, the operators need to decide network configuration setting $o_n \in \mathcal{O} = 0 \cup \mathcal{I}$, select split $i \in \mathcal{I}$, and allocate function virtualized resources at the vDU $x_n \in \mathbb{R}$ and the vCU $\hat{x}_n \in \mathbb{R}$. The option $\mathbf{0}$ ($o_n = 0$) is dedicated for *no reconfiguration*⁴ (Fig. 1c). Then, a *reconfiguration*⁵ activity ($o_n \neq 0$) is any change in the deployed split and virtualized resources (Figs. 1a and 1b). In this sense, deciding o_n will directly determine split- i as: $i_n := o_n \iff o_n \neq 0$ and $i_n := i_{n-1} \iff o_n = 0$. Then, the allocated resources x_n and \hat{x}_n are determined through a mapping function $\omega: (\lambda_n, o_n) \mapsto (x_n, \hat{x}_n)$ (details in Sec. III).

In practice, x_n and \hat{x}_n may differ to the actual resource utilization. We hence define $y_n \in \mathbb{R}$ and $\hat{y}_n \in \mathbb{R}$ as the actual resource utilization at vDU and vCU while deploying split $i_n \in \mathcal{I}$ and serving λ_n . The difference between the allocated resources and actual resource utilization is the prediction error resulting in overprovisioning or underprovisioning cost. If the operators do not want to reconfigure their resources, they may suffer from these costs, i.e., miss the opportunity to share unused resources for other workloads or even have declined demand due to insufficient capacity. Otherwise, the operators can decide to resize their resources by $\beta_n := |x_n - x_{n-1}|$ and $\hat{\beta}_n := |\hat{x}_n - \hat{x}_{n-1}|$ at vDU and vCU, respectively. Motivated by network slicing management [16], we propose a new cost model that can capture virtualized behaviour in vRANs. This model is illustrated in Fig. 1d and described as follows.

1) *Overprovisioning*: The operators have to pay more expensive cost for unused resources if the allocated resources are higher than the actual utilization. This cost is defined for resource overprovisioning, which is denoted:

$$f_o(\max(0, x_n - y_n) + \max(0, \hat{x}_n - \hat{y}_n)), \quad (1)$$

where $f_o(\cdot)$ is a function cost for overprovisioning. We assume this function is proportional with the input, e.g. $f_o(z) := \kappa_o z$, where κ_o is the estimated fee for one unit capacity (\$/unit).

2) *Declined service demand*: The underprovisioning and violating the split requirements can cause a disrupted or declined demand resulting in a penalty for service level agreement (SLA) violation and initializes a monetary compensation. We define this cost as:

$$f_d(\mathbb{1}_{<y_n}(x_n) \vee \mathbb{1}_{<\hat{y}_n}(\hat{x}_n) \vee \mathbb{1}_{\neq 0}(C_n)), \quad (2)$$

where $\mathbb{1}(\cdot)$ is an indicator function that takes value 1 if the condition is satisfied, otherwise 0. $C_n = (c_{np})_{\forall p}$ is a function of constraint dissatisfaction that captures the penalization for violating each p constraints in vRANs at time n . The indicator functions in (2) are activated when the allocated resources x_n and \hat{x}_n do not meet the respective incurred resource utilization y_n and \hat{y}_n . It is also activated when there exists constraint violations. In this case, the operator has to pay a penalty with

⁴The last network settings are preserved; consequently, the selected split and virtualized resources at the current sequence same as the previous one: $o_n = 0 \iff (i_n = i_{n-1} \wedge x_n = x_{n-1} \wedge \hat{x}_n = \hat{x}_{n-1})$.

⁵The operators may decide to alter the deployed split, but it will change the incurred resources at vDU and vCU. Therefore, the operators have to reallocate their virtualized resources accordingly. The operators may also decide only to resize the virtualized resources without altering the split. In this case, $o_n \neq 0 \iff (i_n \neq i_{n-1} \vee x_n \neq x_{n-1} \vee \hat{x}_n \neq \hat{x}_{n-1})$.

a value κ_d (\$) for any SLA violation. It complies with the monetary fee penalty used in network slicing [19].

3) *Instantiation & Reconfiguration*: The operators may decide to reconfigure the split and virtualized resources following the users demand to reduce resource overprovisioning and avoid declined demand. However, there are overhead costs as:

$$f_i(\beta_n \mathbb{1}_{>x_{n-1}}(x_n) + \hat{\beta}_n \mathbb{1}_{>\hat{x}_{n-1}}(\hat{x}_n)) + f_r((\beta_n + \hat{\beta}_n) \mathbb{1}_{\neq 0}(o_n)). \quad (3)$$

Instantiating and reconfiguring VMs have capital expenses in NFV [16]. The first term in (3) captures the amount of instantiating additional resources at vDU and vCU. The second term then captures the cost initiated from any reconfiguration activities ($o_n \neq 0$). For instance, resizing the VMs' resources initiates a price of management delay [21] for load balancing setups and migrating the resources, which may interrupt the user sessions. We found that resizing a VM instance in CSC cPouta cloud (<https://www.csc.fi/>) induces delay around 23 seconds. Modern software architecture such as Kubernetes also needs several seconds for executing new pods [16], [21]. The seamless migration for altering the split deployment also requires creating new vBS functions while preserving the old migrated functions active [6]. Thus, the reconfiguration cost is affected by the migrated resources measured from the difference between the current and the previous virtualized resources. In our evaluation, we assume that both costs $f_i(\cdot)$ and $f_r(\cdot)$ are proportional to the input, e.g., $f_i(z) := \kappa_i z$ and $f_r(z) := \kappa_r z$, where κ_i (\$/unit) is the parameter that captures the estimated cost for instantiation and κ_r (\$/unit) is for reconfiguration.

xHaul cost. O-RAN has proposed an open interface between vCUs, vDUs and RUs [3]. The operators can enjoy benefits such as computational and performance gains from a more centralized function [22], but it incurs a higher transferred data load (δ) [5]. S1 and S2 generate λ , S3 incurs $1.02\lambda + 1.5$, while S4 transfers 2500 (Mbps) of data load [23]. We define the xHaul cost as: $f_h(\delta_i) := \kappa_h \delta_i$, where δ_i is the data load for selecting split i and κ_h is the estimated fee for reserving a bandwidth (Mbps) of xHaul link.

C. Trade-offs and Problem Statement

Trade-offs. (i) Centralizing more functions gains a lower computational cost, but it has tighter constraint requirements and requires a higher xHaul load. (ii) STAO only incurs resource overprovisioning, but the amount of underutilized resources can be large. (iii) DYNO can reduce the overprovisioning and insufficient allocated resources; however, it needs additional costs for reconfiguration and instantiation. (iv) A standard loss function may not capture prediction error in virtualized resource management as the penalty fee is different for each behaviour. (v) The relationship between the users' demand and the incurred virtualized resources is non-linear and non-monotonic. (vi) A sequential decision for orchestrating the split selection and the virtualized resources at vCU and vDU is intricate problem, particularly without making assumptions of the underlying system and state transitions.

Problem statement. Given the above trade-offs, users' demand, and network state, what is the most suitable split and allocated resources at each time stage to minimize the long-term cost? Next, we present how we formulate the above sequential decision-making problem as an MDP, then discuss how LOFV solves the problem.

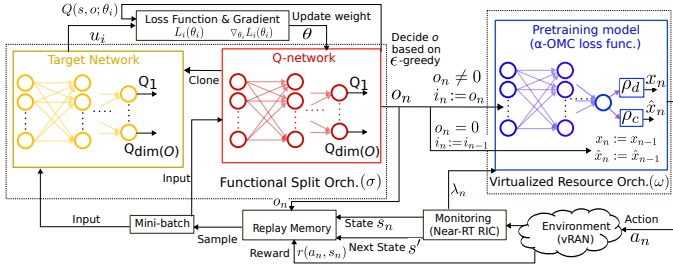


Fig. 2: **LOFV Architecture.** LOFV works following a model-free reinforcement learning and comprising of two functional blocks.

III. PROBLEM AND LEARNING FRAMEWORK

We formulate our problem in Sec. II as an MDP, which is specified by a tuple $\{\mathcal{S}, \mathcal{A}, \mathcal{P}(s_{n+1}|s_n, a_n), r_n, \gamma\}$. LOFV follows sequential decision-making based on a model-free reinforcement learning paradigm to solve the problem. At every time stage n , our agent observes a *state* from state spaces $s_n \in \mathcal{S}$ drawn from the environment, takes an action a_n , and expects to receive a *reward* signal $r(s_n, a_n)$ as a feedback, where $\gamma \in (0, 1]$ is the discount factor. The state may not be stationary as the network load and conditions are changing over time with the sequence state arrival $(s_n)_{n \in \mathcal{N}}$. Then, $\mathcal{P}(s_{n+1}|s_n, a_n)$ is the state transition probability that maps a state-action pair at time step n into the distribution of next states. LOFV operation relies on two functional blocks: *i) functional split orchestration* (σ) and *ii) virtualized resource orchestration* (ω). Fig. 2 illustrates how LOFV operates.

State. Our state comprises: *(i)* the incoming traffic demand at time stage n , λ_n ; *(ii)* the average traffic over T period, $\bar{\lambda}_n := \frac{1}{T} \sum_{t=1}^T \lambda_{nt}$; *(iii)* the variance traffic for T period, $\tilde{\lambda}_n := \sum_{t=1}^T (\lambda_{nt} - \bar{\lambda})^2 / T$; *(iv)* the previous allocated resources at vDU x_{n-1} and *(v)* vCU \hat{x}_{n-1} ; and *(vi)* the previous deployed split i_{n-1} . It provides *time dynamic* of our variable interests: *(i)* the input traffic demand that affects the split and resources; *(ii-iii)* the characteristics and variation of the users demand that helps to decide doing reconfiguration or not; and network conditions regarding *(iv)* the availability of resources at vDU and *(v)* vCU and *(vi)* the last deployed split. Therefore, our agent receives an input state as a tuple $s_n = \{\lambda_n, \bar{\lambda}_n, \tilde{\lambda}_n, x_{n-1}, \hat{x}_{n-1}, i_{n-1}\} \in \mathcal{S} \subseteq \mathbb{R}^6$.

Action. Our action space consists all pairs of $o \in \mathcal{O}$, $i \in \mathcal{I}$, $x \in \mathcal{X}$ and $\hat{x} \in \hat{\mathcal{X}}$. Thus, we define the action to be taken by our agent at the time stage n as a tuple of $a_n = \{o_n, i_n, x_n, \hat{x}_n\} \in \mathcal{A}$, where $\mathcal{A} = \mathcal{O} \times \mathcal{I} \times \mathcal{X} \times \hat{\mathcal{X}}$ is our action spaces.

Reward. Our objective is to minimize the long-term accumulated management cost over the time. Thereby, we define our reward function at time stage n as:

$$r(a_n, s_n) := -J(a_n, s_n) \quad (4)$$

where $J(\cdot) := (f_o(\cdot) + f_d(\cdot) + f_i(\cdot) + f_r(\cdot) + f_h(\cdot))$ is the management cost defined in Sec. II. The long-term accumulated reward starting at time step n is $G_n := \sum_{\tau=0}^{\infty} \gamma^\tau r_{n+\tau}$. Our goal can be redefined by maximizing the expected long-term accumulated reward as $\mathbb{E}[G_n] = \mathbb{E}[\sum_{\tau=0}^{\infty} \gamma^\tau r_{n+\tau}]$. Therefore, our agent aims to learn the optimal policy that maximizes the expected long-term accumulated reward as: $\pi_* := \arg \max \mathbb{E}_\pi[\sum_{\tau=0}^{\infty} \gamma^\tau r_{n+\tau} | \pi]$, where the policy π is a function that maps from state to action $\pi(s): \mathcal{S} \mapsto \mathcal{A}$. To solve this maximization problem, we decompose our agent into two functional blocks and summarize its operation in Algorithm 1.

Algorithm 1: LOFV operation

- 1 **Initialize:** Replay memory \mathcal{D} with capacity D , Function ω_φ with pretraining weights φ , Q-network Q_θ with weight θ , Target Network \hat{Q}_θ with weights $\hat{\theta} \leftarrow \theta$.
- 2 **for** $e = 1, \dots, E$ **do**
- 3 **Initialize:** $s_1 = \{\lambda_1, 0, 0, x_{\max}, \hat{x}_{\max}, 1\}$.
- 4 **for** $n = 1, \dots, N$ **do**
- 5 Select random o_n with probability ϵ , otherwise $o_n := \max_o Q^*(s_n, o; \theta)$.
- 6 Determine $i_n := o_n \iff o_n \neq 0$ and $i_n := i_{n-1} \iff o_n = 0$
- 7 Allocate resources $(x_n, \hat{x}_n) \leftarrow \omega(\lambda_n, o_n)$
- 8 Execute $a_n = \{o_n, i_n, x_n, \hat{x}_n\}$ and observe reward r_n
- 9 Set $s_{n+1} := s_n$
- 10 Store transion $\{s_n, o_n, r_n, s_{n+1}\}$
- 11 Sample random minibatch of transion $\{s_j, o_j, r_j, s_{j+1}\}$ from \mathcal{D} .
- 12 Set TD target $u_j := r_j + \gamma \max_{o'} \hat{Q}(s_{j+1}, o'; \hat{\theta})$
- 13 Perform a gradient descent on (5) with Adam
- 14 $\hat{Q}_\theta \leftarrow Q_\theta$ for every C steps.
- 15 **end**
- 16 **end**

A. Functional Split Orchestration (σ)

In this block, we design a function σ that maps the input state to the selected configuration setting o and split i , thus the long-term accumulated reward is maximized. The reward function r is affected by the configuration o , split i and virtualized resources at vDU x and vCU \hat{x} . Given o_n , we can determine i_n directly while allocate x_n and \hat{x}_n through a deterministic orchestrator ω (the virtualized resource orchestration). Consequently, in this block, we can treat ω as a part of the environment. We can redefine our action for this block with $o \in \mathcal{O}$ instead of using $a \in \mathcal{A}$ which has a high dimensional action space. Then, the goal is to learn an optimal function that maximizes the long-term accumulated reward as: $\sigma_* := \arg \max_\sigma \sum_{\tau=0}^{\infty} \gamma^\tau r_{n+\tau}$. This problem is then has a high-dimensional state space but discrete action spaces. Hence, the function σ can be designed using deep Q-learning [17]. It utilizes a neural network to approximate the action-value function (Q-function) while the learning step is based on Q-learning [18].

We define the optimal action-value function as the maximum expected reward that follows some strategies based on a function σ , after observing some sequences s , then taking actions o as: $Q^*(s, o) := \max_\sigma \mathbb{E}[G_n | s_n = s, o_n = o]$. If the optimal value $Q^*(s', o')$ of the sequence at the next time step s' for all possible action o' is known, we can define the optimal function σ_* , which is to select the action o' that maximizes the expected value of $r + \gamma Q^*(s', o')$ as: $Q^*(s, o) := \mathbb{E}_{s' \sim \mathcal{E}}[r + \gamma \max_{o'} Q^*(s', o') | s', o']$. Using *value iteration*, the optimal action-value function can converge to the optimality when $i \rightarrow \infty$ [18]; however, this method is impractical [17]. Therefore, we use a function approximator to estimate the action-value function parameterized by weight

θ as $Q(s, o; \theta) \approx Q(s, o)$. A neural network can be applied to approximate our action-value function (Q-network). The Q-network is trained by minimizing the sequence of a loss function that changes at iteration i :

$$L_i(\theta_i) := \mathbb{E}_{s, o \sim p(\cdot)} [(u_i - Q(s, o; \theta_i))^2], \quad (5)$$

where $u_i := \mathbb{E}_{s' \sim \mathcal{S}} [r + \gamma \max_{o'} Q(s', o'; \theta_{i-1} | s, o)]$ is the Temporal-Difference (TD) target and $p(s, a)$ is the *behaviour distributions* (probability distribution over s and a). Note that the parameters from the previous iteration θ_{i-1} is frozen while we are optimizing $L_i(\theta_i)$. Instead of directly calculating the full expectation of the gradient $\nabla_{\theta_i} L_i(\theta_i)$, the loss function in (5) can be efficiently calculated using a stochastic gradient descent method, e.g., Adam optimizer [24]. Then, our Q-network comprises of an input layer receiving $(\lambda_n, \bar{\lambda}_n, \hat{\lambda}_n, x_{n-1}, \hat{x}_{n-1}, i_{n-1})$, output layers activated by a linear function with a size of $\dim(O)$, and the hidden layers constructed from a fully connected linear layer activated by ReLU function. This approach is a model-free and off-policy [17]. It also learns about the greedy strategy $o := \max_o Q(s, o; \theta)$, while follows the behaviour distribution for exploration. In this case, we can tailor an ϵ -greedy strategy by setting the probability of selecting a random action with ϵ ; otherwise, we can select a greedy action by probability $1 - \epsilon$. LOFV can directly determine split- i from o_n by $i_n := o_n \iff o_n \neq 0$ and $i_{n-1} := i_{n-1} \iff o_n = 0$.

B. Virtualized Resource Orchestration (ω)

The virtualized resource orchestration is a function ω that maps an input pair of selected configuration o (an output of a function $\sigma : s \mapsto o$) and the traffic demand λ into the allocated virtualized resources at vDU x and vCU \hat{x} . Thus, we define $\omega : (o, \lambda) \mapsto (x, \hat{x})$. This function has been simply defined as an affine function for vRAN network design [23]. However, our measurement shows that the relations between users' demand and virtualized resource utilization is non-linear and non-monotonic (Sec. IV), which is hard to predict exactly. Therefore, we leverage a data-driven model that learns to approximate a function ω . Our design is constructed from a simple neural network ω_φ parameterized by weight φ . The neural network consists of a single-neuron input layer receiving λ , a single-neuron output layer activated by a linear function and the hidden layers activated by a ReLU function. The output of neural network ω_φ is then scaled with parameters ρ_d (vDU) and ρ_c (vCU) where each value depends on the split selection. The value of both parameters are obtained using a linear regression from [22, Fig. 6a]. Then, we leverage α -OMC [19] loss function to train ω_φ as a regressor, where y (vDU) and \hat{y} (vCU) are the ground truth. The α -OMC is able to capture the prediction error with penalty fee incurred from resource overprovisioning and underprovisioning. As ω_φ is a deterministic orchestrator for σ , we can train separately and use it as a pretraining model. If $o_n = 0$, our agent does not reallocate the resources, hence $x_n := x_{n-1}$ and $\hat{x}_n := \hat{x}_{n-1}$.

IV. RESULTS AND DISCUSSION

A. Experimental Setup

Our testbed consists of the vBS, virtualized evolved packet core (vEPC) and virtualized user equipment (vUE), where each entity is built using an open-source srsRAN (e.g., srsENB, srsEPC, srsUE) [15] and virtualized using a containerized system from Docker. The radio interface of vBS and vUE are

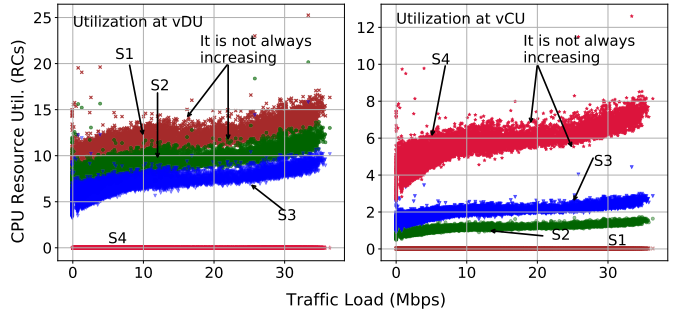


Fig. 3: The relationship between the users' traffic demand to the incurred vBS resources at vDU (left) and vCU (right).

emulated via radio samples in the Linux shared memory and inter-process communication. We deploy all of the virtualized entities in a general-purpose PC (AMD Ryzen 7 PRO 4750U) and use this computing specification as a *reference core* (RC)⁶. We can control, set a capacity limit, and isolate each container resource through `-cpus` in Docker. The interval of time stage t is a second and n is a minute with default total duration per episode is 2 hours. Our default penalty coefficients are $\kappa_o=1$, $\kappa_d=2$, $\kappa_i=\kappa_r=0.5$ and $\kappa_h=0.0005$. The traffic demand follows a Poisson-generated user datagram protocol. The computing consumption of f_1 , f_2 and f_3 yield 65%, 15% and 20% of total vBS computing usage, respectively [23]. The vCU consumes a twice of a lower computational processing than vDU (see [22, Fig. 6a] with ≈ 10 BSs). Thus, we have scaled parameters for deploying split- i at vDU with $\rho_i^d = \{1, 0.8, 0.65, 0\}$ and at vCU with $\rho_i^c = \{0, 0.1, 0.175, 0.5\}$. The hyperparameters of our neural networks are set after exploratory assessments. The neural network ω_φ is constructed from an input, an output and three hidden layers with the sizes of 128, 64 and 16. We use Adam optimizer [24] with learning rate is set to 5×10^{-5} , mini-batch with the size of 128 and α -OMC [19] loss function, then train ω_φ with 200 epochs. In the functional split orchestration, our Q-network consists of an input, outputs with the size of $|\mathcal{O}|$ and three hidden layers with the size of 512 each. The target network is updated every 10 time stage n . The batch size is set with 256 and the replay buffer has a capacity of 10^6 . Our exploration and exploitation strategy is based on ϵ -greedy, where we set $\epsilon_{\max} = 0.95$ at the beginning of episode, then it exponentially decays to $\epsilon_{\min} = 0.02$. We use Adam optimizer [24] with learning rate is set to 0.0003 and (5) for the loss function, then train our model over 5000 episodes.

B. Measurements and Evaluation Results

Our first evaluation studies the relationship between the users' demand and the virtualized resource utilization of vBS for selecting various splits. Fig. 3 illustrates that the resource utilization at vDU and vCU for each split does not always increase with the growth of traffic load. Instead, it is non-linear and non-monotonic. Although the traffic load reaches a peak (35 Mbps), the virtualized resources are still highly underutilized with less than 25 RCs and 13 RCs of used resources at vDU and vCU, when LOFV is not implemented.

Fig. 4 illustrates the convergence behaviour of LOFV during the training process. We increase the penalty coefficients of

⁶We focus on the CPU resources as the most affected parameters; however, it is not limited and can be extended to other virtualized computing resources. We have 16 CPU threads in total, which translates to 100 RCs.

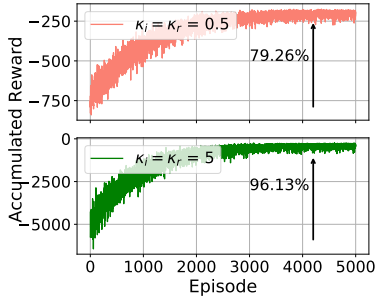


Fig. 4: Convergence performance of LOFV during training.

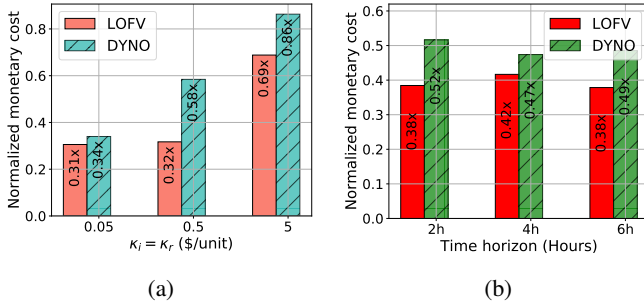


Fig. 5: Performance of LOFV during the online stages with (a) various penalty coefficients and (b) over different time-horizon settings. All of the monetary costs are normalized toward STAO.

reconfiguration and instantiation from $\kappa_i = \kappa_r = 0.5$ (default) to $\kappa_i = \kappa_r = 5$. At the beginning of learning, LOFV has a higher probability from using a random policy for exploration, resulting in a high variance and expensive costs for both cases. Also, a higher penalty coefficient ($\kappa_r = 5$) induces a considerably more expensive cost (almost eight times) than a lower penalty ($\kappa_r = 0.5$) when choosing random actions. After some episodes, LOFV successfully learns the optimal policy and converges to the best policy that the agent can learn, and improves around 96.13% ($\kappa_r = 5$) and 72.26% ($\kappa_r = 0.5$).

We analyze the effect of penalty coefficients (reconfiguration & instantiation) and the duration of time-horizon that LOFV can achieve during the online stages. Fig. 5a depicts that LOFV has the lowest long-term cost across all coefficient settings compared to STAO and DYNO. It shows that LOFV achieves the highest cost-saving (69%) compared to STAO when the penalty fees for reconfiguration and instantiation are low ($\kappa_r = 0.05$). Although having a higher penalty coefficient ($\kappa_r = 5$), LOFV still outperforms STAO with 31% of cost-saving and has a cheaper cost compared to DYNO across all coefficients. LOFV manages to have 9% of cost-saving at a low penalty coefficient ($\kappa_r = 0.05$) and gains considerably cost-saving as high as 45% of DYNO.

Fig. 5b shows the results of altering the duration of time-horizon (up to 6 hours) during the online stages that LOFV can achieve. It shows that the duration of the time horizon does not significantly degrade the performance of LOFV, where LOFV can achieve a stable cost-saving compared to our two baselines. LOFV saves around 58% (4h) to 62% (2h & 6h) across all time-horizon settings compared to STAO. Similarly, LOFV also gains cost-saving to 27% (2h) of DYNO.

V. CONCLUSION

In this paper, we have proposed LOFV that jointly selects the splits and allocates the virtualized resources to minimize

the long-term management cost. Aside from the xHaul cost, we have also proposed a new cost model accounting for resource overprovisioning, declined service demand, re-configuration and instantiation, representing the NFV behaviour in vRANs. A model-free reinforcement learning paradigm has been leveraged to solve this sequential decision-making problem, constructed from deep Q-learning and a regression-based neural network. We also have performed a battery of tests to evaluate the effectiveness of LOFV using testbed measurements of a containerized RAN system. The results have shown that LOFV successfully learns the split selection and resource allocation and saves the cost by up to 69% of STAO and 45% of DYNO.

REFERENCES

- [1] "5G immersive service opportunities with Edge Cloud and Cloud RAN (White Paper)," Nokia, Tech. Rep., 2019.
- [2] "Open and Virtualized - The Future Radio Access Network," NEC, Tech. Rep., 01 2020.
- [3] "O-RAN-WG1-O-RAN Architecture Description v01.00.00," O-RAN Alliance, Technical Specification (TS), 02 2020.
- [4] 3GPP, "Architecture description (Release 16)," 3rd Generation Partnership Project (3GPP), Technical Specification Group Radio Access Network (NG-RAN) 38.401, 03 2020, version 16.1.0.
- [5] F. W. Murti *et al.*, "Optimal Deployment Framework for Multi-Cloud Virtualized Radio Access Networks," *IEEE Transactions on Wireless Communications*, 2020.
- [6] A. M. Alba *et al.*, "An adaptive functional split in 5G networks," in *IEEE INFOCOM Workshop*, 2019.
- [7] A. M. Alba, S. Janardhanan, and W. Kellerer, "Enabling dynamically centralized RAN architectures in 5G and beyond," *IEEE Transactions on Network and Service Management*, 2021.
- [8] D. Harutyunyan and R. Riggio, "Flex5G: Flexible Functional Split in 5G Networks," *IEEE Transactions on Network and Service Management*, 2018.
- [9] S. Ali *et al.*, "6G White Paper on Machine Learning in Wireless Communication Networks," 2020.
- [10] J. A. Ayala-Romero *et al.*, "vrAIN: Deep Learning based Orchestration for Computing and Radio Resources in vRANs," *IEEE Transactions on Mobile Computing*, 2020.
- [11] J. A. Ayala-Romero *et al.*, "Bayesian Online Learning for Energy-Aware Resource Orchestration in Virtualized RANs," in *Proc. of IEEE INFOCOM*, 2021.
- [12] X. Foukas and B. Radunovic, "Concordia: Teaching the 5G VRAN to Share Compute," in *Proc. of ACM SIGCOMM*. ACM, 2021.
- [13] F. W. Murti, S. Ali, and M. Latva-aho, "Deep Reinforcement Based Optimization of Function Splitting in Virtualized Radio Access Networks," in *IEEE ICC Workshops*, 2021.
- [14] T. Pamuklu *et al.*, "Reinforcement Learning Based Dynamic Function Splitting in Disaggregated Green Open RANs," in *Proc. of IEEE ICC 2021*, 2021.
- [15] I. Gomez-Miguel *et al.*, "SrsLTE: An Open-Source Platform for LTE Evolution and Experimentation," in *Proc. of ACM WinTech*, 2016.
- [16] D. Bega *et al.*, "AZTEC: Anticipatory Capacity Allocation for Zero-Touch Network Slicing," in *Proc. of IEEE INFOCOM*, 2020.
- [17] V. Mnih *et al.*, "Human-level control through deep reinforcement learning," *Nature*, vol. 518, 2015.
- [18] R. S. Sutton and A. G. Barto, *Introduction to Reinforcement Learning*, 1st ed. Cambridge, MA, USA: MIT Press, 1998.
- [19] D. Bega *et al.*, "DeepCog: Optimizing Resource Provisioning in Network Slicing With AI-Based Capacity Forecasting," *IEEE Journal on Selected Areas in Communications*, 2020.
- [20] S. C. Forum, "R6.0. small cell virtualization functional splits and use cases, document 159.07.02," Tech. Rep. Release 7, 2016, version 14.0.0.
- [21] 5G-CORAL, "Refined design of 5G-CORAL orchestration and control system and future directions," Public Deliverable D3.2, May 2019.
- [22] P. Rost *et al.*, "The complexity-rate tradeoff of centralized radio access networks," *IEEE Transactions on Wireless Communications*, 2015.
- [23] A. Garcia-Saavedra *et al.*, "Joint Optimization of Edge Computing Architectures and Radio Access Networks," *IEEE Journal on Selected Areas in Communications*, vol. 36, no. 11, 2018.
- [24] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," in *Proc. of ICLR*, 2015.