

Constrained Deep Reinforcement Based Functional Split Optimization in Virtualized RANs

Fahri Wisnu Murti¹, Samad Ali¹, and Matti Latva-aho¹, *Senior Member, IEEE*

Abstract—In virtualized radio access network (vRAN), the base station (BS) functions are decomposed into virtualized components that can be hosted at the centralized unit or distributed units through functional splits. Such flexibility has many benefits; however, it also requires solving the problem of finding the optimal splits of functions of the BSs in such a way that minimizes the total network cost. The underlying vRAN system is complex and precise modelling of it is not trivial. Formulating the functional split problem to minimize the cost results in a combinatorial problem that is provably NP-hard, and solving it is computationally expensive. In this paper, a constrained deep reinforcement learning (RL) approach is proposed to solve the problem with minimal assumptions about the underlying system. Since in deep RL, the action selection is the outcome of inference of a neural network, it can be done in real-time while training to update the neural networks can be done in the background. However, since the problem is combinatorial, even for a small number of functions, the action space of the RL problem becomes large. Therefore, to deal with such a large action space, a chain rule-based stochastic policy is exploited in which a long short-term memory (LSTM) network-based sequence-to-sequence model is applied to estimate the policy that is selecting the functional split actions. However, the utilized policy is still limited to an unconstrained problem, and each split decision is bounded by vRAN's constraint requirements. Hence, a constrained policy gradient method is leveraged to train and guide the policy toward constraint satisfaction. Further, a search strategy by greedy decoding or temperature sampling is utilized to improve the optimality performance at the test time. Simulations are performed to evaluate the performance of the proposed solution using synthetic and real network datasets. Our numerical results show that the proposed RL solution architecture successfully learns to make optimal functional split decisions with the accuracy of the solution is up to 0.05% of the optimality gap. Moreover, our solution can achieve considerable cost savings compared to C-RAN or D-RAN systems and a faster computational time than the optimal baseline.

Index Terms—Radio access networks (RANs), network virtualization, Cloud RAN (C-RAN), deep reinforcement learning (RL), optimization.

Manuscript received 18 May 2021; revised 15 December 2021 and 15 February 2022; accepted 27 May 2022. Date of publication 9 June 2022; date of current version 11 November 2022. This work was supported by the Academy of Finland, 6G Flagship Program, under Grant 346208. An earlier version of this paper was presented at IEEE ICC 2021 Workshop [DOI: 10.1109/ICC-Workshops50388.2021.9473703]. The associate editor coordinating the review of this article and approving it for publication was T. Q. S. Quek. (Corresponding author: Fahri Wisnu Murti.)

The authors are with the Centre for Wireless Communications, University of Oulu, 90570 Oulu, Finland (e-mail: fahri.murti@oulu.fi; samad.ali@oulu.fi; matti.latva-aho@oulu.fi).

Color versions of one or more figures in this article are available at <https://doi.org/10.1109/TWC.2022.3179811>.

Digital Object Identifier 10.1109/TWC.2022.3179811

I. INTRODUCTION

THE increase in mobile data traffic of emerging applications with diverse requirements has driven the efforts to re-design the radio access networks (RANs). Cloud/Centralized-RAN (C-RAN) has become a favourable solution to enable the low-cost deployment and high-performance systems by pooling the baseband functions of the base station (BS) to a central server which is also known as Cloud/Central unit (CU). This idea is motivated by the growth of network densification, enabled by the concept of *softwarization*, to offer cost-efficient solutions and high-performance network operations through centralized control [2]. However, a fully centralized RAN is typically not implementable for many reasons [3]. For instance, it requires a low-latency and high-capacity fronthaul, which is frequently absent in current RANs and prohibitively expensive to develop from scratch. This challenge motivates the transition from rigid C-RAN designs to flexible architectures, where only a subset of BS functions is centralized at the CU, and the other functions are hosted at the distributed units (DUs) and radio units (RUs).¹ Further, the term virtualized RAN (vRAN) is coined to describe these architectures [4].

In vRANs, the BS functions (except RF functions) can be decomposed into virtualized components and executed on commodity hardware across a geo-distributed edge cloud system² [5]. Then, the operators can uniquely select the *functional splits* suited to their needs by deciding which functions will be centralized at the CU and which will be kept at the DUs. This paradigm brings flexibility to the RAN operations, potentially offers a cost-saving, and accommodates diverse use cases and applications in 5G+ systems [6], [7]. However, selecting the functional splits of all the BSs is challenging. Each split has a different delay requirement, initiates a different computing load to the CU and DUs, and induces a different data flow. The initial design of vRAN fronthaul using point-to-point Common Public Radio Interface (CPRI) is also suggested to be updated with the new Crosshaul/xHaul architecture based on an open interface and packet-switch (shared) network such as the enhanced CPRI [6]–[8]. As a result, in addition to sharing the same computing nodes, each BS has to share the same network links, which leads to complex interdependence

¹RUs are the radio hardware units to run the RF functions.

²Each CU and DU is to run as virtualized software, e.g., virtualized CU (vCU) and virtualized DU (vDU). A DU is typically executed at the far-edge server (co-located or close to the RU), while a CU is at the edge server (a more centralized server).

between split decisions. Therefore, it is important to optimize the splits carefully to ensure that the deployment is beneficial; otherwise, it can lead to high operating costs and performance degradation.

On the other hand, optimizing the functional splits produces a high degree of complexity. In addition to the mentioned challenges, unlike legacy RANs, the behaviour of vRAN system performance such as computing utilization [9] and power consumption [10] is highly non-trivial. This non-triviality is also reinforced by vRAN deployment over the same platform with other workloads such as video analytics [11]. As a result, it is complex and difficult to model the underlying system precisely. Meanwhile, traditional mathematical optimization approaches rely on complete knowledge of the system behaviour to define the models and solve the problems; and this can be unfeasible in practice. These challenges motivate us to shift to machine learning (ML)-based approaches, which can be best to address our vRAN splitting problem amidst minimal assumptions about the underlying system.

A. Related Works

1) *Optimization-Based Approaches*: 3GPP [12], [13] and a seminal white paper [14] have defined the detail vRAN split specifications. Although the authors in [6] have discussed the gains and requirements for each split, there are still few works on the optimization issues. Energy consumption for various splits has been evaluated in [15]. The authors in [16] have studied optimizing the function centralization of vRANs over xHaul. Follow-up works, [17] and [18] offered an optimal solution of minimizing total cost for integration vRANs with Mobile Edge Computing (MEC). Then, the work in [19] has proposed an optimized multi-cloud vRAN framework with balancing its centralization [3]. The authors in [20] have proposed the PlaceRAN framework to minimize the computing resources while maximizing the radio function aggregations using the IBM CPLEX solver. These mentioned works [3], [15]–[20] have addressed various optimization problems in vRANs. However, their frameworks require assuming complete models of the underlying system to define their problem structures and solve the problems. We argue that such strong assumptions can be inaccurate as the underlying vRAN system is complex and difficult to model precisely. Additionally, those frameworks need heavy mathematical solutions with exponential complexity and slow execution time, which typically are unsuitable for large networks and online execution. The above problems are also often complex combinatorial and difficult to solve optimally. Therefore, we opt out to use optimization-based approaches to formulate and solve our vRAN splitting problem.

2) *ML-Based Approaches*: ML techniques recently have been spurred to address complex optimization and control problems in wireless networks [21], [22]. The authors in [9] have proposed a learning framework that successfully solves a contextual bandit problem of dynamic computing and radio resource controls in vRANs using a deep reinforcement learning (RL) paradigm. Further, they leveraged Bayesian learning for energy efficient-based resource orchestrator in [10].

ML-based predictor also has been developed in [11] that learns to share the CPU resources between a vRAN workflow and other workflows in the same server. Although the authors in [9]–[11] have shown the non-triviality of vRAN performance and the importance of learning based-framework to manage the vRAN resources, they still did not discuss how to design a framework that learns to optimize the functional splits.

Recent work in [23] has studied user-centric slicing and split optimization problems using a deep learning method. The authors modelled their problem as supervised learning, which relies on high-quality labelled datasets (e.g., optimal labels) to assess the quality of the decisions. Once trained, the model can be used quickly in an online manner, offering a real-time solution for assigning the split for each user slice. However, in vRANs, obtaining such high-quality labels is expensive. To construct the labelled datasets, we still need complete knowledge of the system performance to model the problem mathematically and solve multiple instances of the problem. The work in [24] has addressed the functional split problem for green Open RANs using Q-learning and SARSA; however, they assumed each DU/CU as an independent agent that focuses on its own utility. We argue that every split decision in vRANs is interdependent as the BSs share the same network links and computing nodes with limited capacity. Besides, [23] focused on the split assignment for the users, and [24] studied the effectiveness of energy sources, but we aim for a different goal.

B. Methodology & Contributions

Our goal is to develop a zero-touch optimization framework that optimizes the functional splits of the BSs to minimize the total network cost in the vRAN system. First, we formulate and present the functional split problem mathematically to provide a better understanding of its objective and constraints. Our formulation yields a combinatorial and NP-hard problem. Therefore, it is computationally expensive to solve optimally, especially for large-scale networks and real-time execution. Moreover, solving such a problem often relies on the assumptions of the underlying system to define and model the problem structure (e.g., mathematical optimization). However, in practice, the behaviour of vRAN performance and resources is highly non-trivial, which is complex and can be unfeasible to model the system precisely.

Motivated by the above challenges, we formulate the functional split problem with constrained neural combinatorial RL. We use neural networks to approximate the policy that maps the state observations to the actions. Then, the idea is to estimate the neural network model's parameters iteratively by taking instances from the problem spaces using a constrained deep RL paradigm. For every interaction with the environment (vRAN system), we expect to receive a reward (the induced total network cost) and penalty (constraint violation) as feedback signals and the output returned by the neural network to learn and improve the model. This paradigm considers the vRAN system as a black-box environment, making minimal assumptions about the underlying system. It also does not need

the optimal labelled datasets, which are highly expensive to obtain in vRANs.

Further, we propose a novel constrained deep reinforcement-based functional split optimization framework (CDRS) to solve the problem. Due to the combinatorial nature of the problem finding the optimal splits, the action space of our RL problem becomes enormously large. Therefore, we develop CDRS using a chain rule-based stochastic policy [25] in which policy network architecture using a long short-term memory (LSTM) network-based sequence-to-sequence model is applied to estimate the policy [26], [27]. However, this policy is still limited to an unconstrained problem, which is not directly applicable to our vRAN splitting problem. Therefore, we leverage a constrained policy gradient method to train and guide the policy toward constraint satisfaction. Then, CDRS can be tailored into CDRS-Fixed and CDRS-Ada. CDRS-Fixed uses a fixed penalty coefficient [28], [29] while CDRS-Ada updates the penalty coefficient adaptively [30], [31]. A self-competing baseline is also utilized with an auxiliary network to improve the policy further. Once the model is trained, finding the solution for the problem is computationally efficient as it only requires a forward pass through the trained neural network. Therefore, we provide a search strategy to improve the optimality performance at the test time. Following the search strategy, CDRS can be further tailored into CDRS-Fixed-G, CDRS-Ada-G, CDRS-Fixed-T and CDRS-Ada-T. CDRS-Fixed-G and CDRS-Ada-G utilize greedy decoding while CDRS-Fixed-T and CDRS-Ada-T use temperature sampling.

CDRS is evaluated in terms of training behaviour, optimality performance, the impact of altering the traffic load and routing cost and the computational time. The evaluations are performed using a synthetic network generated by the Waxman algorithm [32] and a real network dataset [33]. The used system parameters are from a measurement-based 3GPP-compliant system model. To assess the effectiveness of our approach, we compare it to the optimal value obtained from a Python-MIP solver.³ Following our evaluations, CDRS successfully learns the optimal functional splits and solves the problem with 0.05% of the optimality gap⁴ (e.g., CDRS-Fixed-T). Our results also show that CDRS is the most cost-effective compared to two extreme cases: fully C-RAN and D-RAN. All of our CDRS settings are also faster than the MIP solver, where CDRS-Ada-G can attain as high as 22.82 times faster. Our contributions can be summarized:

- We formulate the vRAN split problem to constrained neural combinatorial RL, which takes minimal assumptions about the underlying system and does not require the optimal labelled datasets to solve the problem. We also consider the interdependence between split decisions capturing the network links and computing resources sharing among the BSs.

³We use a solution obtained from a mixed-integer programming solver (<https://www.python-mip.com/>) as an optimal baseline comparison. It offers an exact solution through a well-known method, Branch-&-Cut algorithm.

⁴We use the term *optimality gap* to define our solution's error compared to the optimal value obtained from the MIP solver.

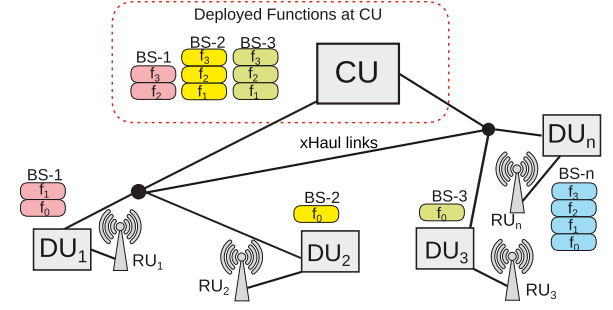


Fig. 1. vRAN over integrated fronthaul/midhaul (xHaul). It has many degrees of design freedom by possibly hosting BS functions at the CU or DUs.

- We propose CDRS as a novel solution framework. CDRS adopts a chain rule-based stochastic policy to deal with the interdependence between split decisions and the combinatorially large discrete action space of the problem in which an LSTM networks-based sequence-to-sequence model is applied to estimate the policy. We utilize a constrained policy gradient method with a self-competing baseline to train and guide the policy toward constraint satisfaction. Following the penalty coefficient and search strategy settings, CDRS can be tailored into CDRS-Fixed-G, CDRS-Fixed-T, CDRS-Ada-G and CDRS-Ada-T.
- We conduct a series of evaluations using synthetic and real network datasets. We investigate the training behaviour, the accuracy of the solution, the impact of routing cost and traffic load and the computational time.

To the best of our knowledge, this work is first to optimize the functional splits of the BSs to minimize the total network cost in the vRAN system using a constrained deep RL paradigm, which takes minimal assumptions about the underlying system and adopts a chain rule-based stochastic policy to deal with the large action space and interdependence between decisions.

The rest of this paper is organized as follows. The background and system model of vRAN are presented in Section II. The functional split problem is formalized mathematically in Section III. Our proposed framework is described in Section IV. Our simulation and experiment results are discussed in Section V. Finally, our work is concluded in Section VI.

II. SYSTEM MODEL

A. Background

In C-RAN, all BS functions are centralized at the Base Band Unit (BBU) except RF layers at the RU. In vRANs, the BBU is decoupled into the CU and DU [13]. Hence, functions of a BS can be deployed at the CU, DU and RU. Fig 1 illustrates that a CU is typically executed at a bigger and more centralized server (e.g., edge server), while a DU is at a smaller server (e.g., far-edge server) and located near (or co-located) with an RU.

Our model refers to the standardization of 3GPP [12], [13] and seminal white paper [14], where each split has a different performance gain [3], [6]. Although 3GPP has

TABLE I

DATA AND DELAY REQUIREMENTS OF vRAN SPLIT WHEN THE TRAFFIC LOAD IS λ MBPS [14], [18]

	Flow (Mbps)	Delay Req. (ms)
Split 0 (S0)	λ	30
Split 1 (S1)	λ	30
Split 2 (S2)	$1.02\lambda + 1.5$	2
Split 3 (S3)	2500	0.25

defined eight options for the splits, several are still hardly implemented. Therefore, we consider four splits that have been experimentally validated in a prototype [34], [35]. **Split 0:** All functions are at the DU, except the RF layers at the RU. It is a typical D-RAN setup. **Split 1 (PDCP-RLC):** RRC, PDCP, and upper layers are hosted at the CU, while RLC, MAC, and PHY are at the DU. This split enables a separate user plane and control plane with centralized RRC. **Split 2 (MAC-PHY):** MAC and upper layers are at the CU, while PHY is at the DU. It allows improvement for CoMP by centralized HARQ. **Split 3 (PHY-RF):** All functions are at the CU, except RF layers. It is a fully centralized version of vRANs. It gains power-saving and improved joint reception CoMP with uplink PHY level combining. Going from Split 0 to Split 3, more functions are hosted at the CU. In addition to increasing network performance, a higher centralization level can lead to more computing cost savings [3]. However, centralizing more functions increases the data load to be transferred to the CU, going from λ in Split 1 to 2.5 Gbps in Split 3 for each BS, and has a stricter delay requirement. Table I summarizes vRAN split options and their requirements.⁵

B. RAN

We model a vRAN architecture with a graph $G = (\mathcal{I}, \mathcal{E})$ where \mathcal{I} has a subsets \mathcal{N} of the $N = |\mathcal{N}|$ DUs, \mathcal{L} of the $L = |\mathcal{L}|$ routers and a CU (index 0). Each node is connected through a link of (i, j) with a set \mathcal{E} of links and has capacity c_{ij} (Mbps) each. The DU- n is connected to $\{0\}$ with a single path (e.g., shortest path) p_{n0} ; hence, we define $r_{p_{n0}}$ as the amount of data flow (Mbps) to be transferred and routed through a path $p_{n0} := \{(n, i_1), \dots, (i_k, 0) : (i, j) \in \mathcal{E}\}$. The BS functions are deployed in servers using virtual machines (VMs).⁶ Each server has a processing capacity, i.e., H_n for DU- n and H_0 for CU. Naturally, a central server has a higher computing performance and capacity, hence $H_0 \geq H_n$. We define ρ_o^c and ρ_o^d as the incurred computational load (cycle/Mb/s) in results of deploying the split configuration $o \in \{0, 1, 2, 3\}$ at each CU and DU, respectively.

C. Demand & Cost

We focus on the uplink transmission where $\lambda_n \geq 0$ (Mbps) is the aggregate data flow of DU- n to serve the users traffic; hence, there are N different flows in the network. We denote

⁵The requirements are tailored from [14], [18] by following settings: 1 user per TTI, 20MHz channel bandwidth, 1 carrier component, UE IP MTU 1500 bytes, 2×2 MIMO.

⁶Each BS function can operate as a virtual network function (VNF), and the VNFs can be executed on top of a single VM or multiple VMs.

$\alpha = (\alpha_n, n \in \mathcal{N})$ and $\beta = (\beta_n, n \in \mathcal{N})$ as the VM instantiation cost (monetary units) and the computing cost (monetary units/cycle) at the DUs, respectively, while α_0 and β_0 are the respective cost for the CU. We also have a routing cost $\zeta_{p_{n0}}$ (monetary units/Mbps) for each path p_{n0} . This cost arises from the network links being leased from third parties or maintaining the links.

D. Problem Statement

We have four choices of the splits for each BS in vRANs. What is the best-deployed split for each BS that minimizes the total network cost? The decision leads to interesting problems. Each split generates a different DU-CU data flow and has a distinct delay requirement. Executing more functions at the CU is more efficient in computing cost; however, it produces a higher load for xHaul links. The BSs share the same capacitated servers and network links, where each split decision is interdependent. Moreover, the behaviour of the vRAN system (e.g., resources, performance) is complex and highly non-trivial, which makes complete assumptions of the model can be unfeasible or inaccurate. The goal is to design a framework to solve this problem by taking minimal assumptions about the model of the system.

III. FORMALIZATION OF vRAN SPLIT PROBLEM

The BS functions can be deployed at the DUs or CU depending on the splits, as seen in Table I. Each split must respect to the *chain of functions* $f_0 \rightarrow f_1 \rightarrow f_2 \rightarrow f_3$.⁷ Thus, we define $x_{on} \in \{0, 1\}$ as the decision for deploying split $o \in \{0, 1, 2, 3\}$ at DU- n . For instance, $x_{0n} = 1$ is for deploying f_0, f_1, f_2, f_3 (Split 0); $x_{1n} = 1$ for f_0, f_1, f_2 (Split 1); $x_{2n} = 1$ for f_0, f_1 (Split 2); or $x_{3n} = 1$ for f_0 (Split 3) at DU- n . We only deploy a single split configuration for each BS. Therefore, a set of eligible splits is:

$$\mathcal{X} = \left\{ \mathbf{x}_n \in \{0, 1\} \mid \sum_{o=0}^3 x_{on} = 1, \quad \forall n \in \mathcal{N} \right\}, \quad (1)$$

where $\mathbf{x}_n = (x_{on}, \forall o)$ and $\mathbf{x} = (\mathbf{x}_n, \forall n)$. The BS functions, f_1, f_2 and f_3 , are deployed using VMs at each server. We have computational processing at the CU and DU- n that must respect to its capacity as:

$$\sum_{n \in \mathcal{N}} \lambda_n \sum_{o=0}^3 x_{on} \rho_o^c \leq H_0, \quad (2)$$

$$\lambda_n \sum_{o=0}^3 x_{on} \rho_o^d \leq H_n, \quad \forall n \in \mathcal{N}. \quad (3)$$

A. Data Flow & Delay

Let define $r_{p_{n0}}$ (Mbps) as the amount of data flow (Mbps) to be transferred through a path p_{n0} . Hence, the flow must respect capacity of each link:

$$\sum_{n \in \mathcal{N}} r_{p_{n0}} I_{p_{n0}}^{ij} \leq c_{ij}, \quad \forall (i, j) \in \mathcal{E}, \quad (4)$$

⁷ f_0 is a function that encapsulates RF layers. Then, f_1, f_2 and f_3 are the functions for Layer 1 (PHY), Layer 2 (MAC, RLC) and Layer 3 (PDCP, RRC and the upper layers), respectively.

where $I_{p_{n0}}^{ij} \in \{0, 1\}$ indicating whether the link (i, j) is used by path p_{n0} . Assuming a single path (e.g., shortest path), the amount of data flow depending on each split configuration is [17]:

$$r_{p_{n0}} = \lambda_n(x_{0n} + x_{1n}) + x_{2n}(1.02\lambda_n + 1.5) + 2500 x_{3n}. \quad (5)$$

We let $d_{p_{n0}}$ denote the incurred delay for routing through path p_{n0} from DU- n to the CU. Each split has to satisfy the respective delay requirement (Table I):

$$x_{on}d_{p_{n0}} \leq d_o^{\max}, \quad \forall o, \forall n \in \mathcal{N}. \quad (6)$$

B. Objective Function

We aim to minimize the total network cost consisting of the computational costs at the DUs and CU and the routing cost.⁸ The needs of computing cost for each BS- n at DU- n is:

$$V_n(\mathbf{x}_n) = \alpha_n + \beta_n \lambda_n \sum_{o=0}^3 \rho_o^d x_{on}. \quad (7)$$

We also have the required computing cost of BS- n at the CU:

$$V_{n0}(\mathbf{x}_n) = \sum_{o=0}^3 x_{on}(\alpha_o + \lambda_n \beta_o \rho_o^c). \quad (8)$$

The first terms in (7) and (8) represent the required instantiating cost at each DU and CU for BS- n . The last terms in (7) and (8) are the required data processing cost by each DU and CU to serve BS- n load. Next, we have the cost to route the data flow from DU- n to the CU:

$$U_{n0}(\mathbf{x}_n) = \zeta_{p_{n0}} r_n(\mathbf{x}). \quad (9)$$

Finally, we have the total vRAN cost as:

$$J(\mathbf{x}) = \sum_{n \in \mathcal{N}} (V_n(\mathbf{x}_n) + U_{n0}(\mathbf{x}_n) + V_{n0}(\mathbf{x}_n)), \quad (10)$$

which leads to the following problem:

$$\mathbb{P} : \text{minimize } J(\mathbf{x}), \quad \text{s.t. (2) - (6).}$$

\mathbb{P} is a combinatorial problem to decide the function placement \mathbf{x} for all the BSs and serve the traffic load λ with DU-CU path p_{n0} for each BS- n in the network graph $G = (\mathcal{I}, \mathcal{E})$. Next, we discuss the complexity of \mathbb{P} .

C. Complexity Analysis

The complexity of \mathbb{P} can be identified from the polynomial reduction of *multiple-choice multidimensional knapsack problem (MMKP)*.

⁸In this case, we follow the linear objective cost function similar to the previous studies [17], [18]. However, our solution approach does not restrict only to the linear objective. Our approach relies on the scalar reward and penalty as feedback; hence, it can also be tailored to a non-linear objective.

1) MMKP: Let suppose there are N items with values v_1, v_2, \dots, v_N . We also have r_1, r_2, \dots, r_N correspond to the required resources to pick the items. In the 0-1 knapsack problem (KP), the aim is to pick the items $x_i \in \{0, 1\}, \forall i$ that maximize the total value $\sum_{i=1}^N x_i v_i$, subject to constraint $\sum_{i=1}^N x_i r_i \leq R$. This is a well-known NP hard problem and there is a pseudo-polynomial algorithm using a dynamic programming concept that has complexity $\mathcal{O}(NR)$ [36]. MMKP is a variant of 0-1 KP where there are M groups of items, e.g., group i has l_i items. Each item has a specific value v_{ij} corresponds to j -th item of i -th group and needs K resources. Hence, each item in a group has a resource vector $\mathbf{r}_{ij} = (r_{ij1}, \dots, r_{ijK})$ and $\mathbf{R} = (R_1, \dots, R_K)$ is the resource bound of the knapsack. The aim is to exactly pick one item from each group, e.g., $\sum_{j=1}^{l_i} x_{ij} = 1, x_{ij} \in \{0, 1\}$ that maximizes the total value: $\sum_{i=1}^M \sum_{j=1}^{l_i} x_{ij} v_{ij}$, subject to the resource constraint: $\sum_{i=1}^M \sum_{j=1}^{l_i} x_{ij} r_{ijk} \leq R_k, k = 1, \dots, K$.

Finding an exact solution for MMKP is also NP-hard [36]. It is also worth noting that the search space for solution in MMKP is smaller than other KP variants; hence, exact solution is not implementable in many practical problems as there is more limitation of picking items from a group in MMKP instance [36]. Next, We prove that \mathbb{P} is harder than MMKP.

Theorem 1: *MMKP can be reduced to \mathbb{P} in polynomial time, e.g., $\text{MMKP} \leq_P \mathbb{P}$.*

Proof: Let suppose we have unlimited link capacity, no routing cost and no delay requirement. Hence, all paths of the DU-CU pair are eligible and (4)-(6) are always satisfied. This problem then can be mapped to MMKP by setting: 1) M groups to N BSs, 2) each i -th group with l_i items to each BS- n with $|o| = 4$ of split options, 3) j -th item of i -th group to the split o_n of BS- n , 4) r_{ij} to the incurred computing loads, e.g., $\lambda_n \rho_o^c$ and $\lambda_n \rho_o^d$, and 5) the knapsack constraints to computing constraints H_n and H_0 . The value v_{ij} of item- j in group- i also can be mapped with the costs (e.g., computing and routing) of deploying split- o of BS- n , where the MMKP is a maximization problem and \mathbb{P} is a minimization problem. We can see the reduction is of polynomial time: we select the functional split for every BS correspond to that we activate an item that we pick to a knapsack in each group. Therefore, we can conclude that if we can solve \mathbb{P} in polynomial time we also can solve any MMKP problem.

IV. CONSTRAINED DEEP REINFORCEMENT BASED FUNCTIONAL SPLIT OPTIMIZATION FRAMEWORK

We leverage a constrained deep RL paradigm to solve our vRAN splitting problem by treating the vRAN system as a black-box environment, which makes minimal assumptions about the underlying system. Consequently, our RL agent does not need to know the information about the formulation in (1)-(10) to decide the splits. Our agent relies on the scalar reward and penalization returned from the environment to assess the quality of the solutions. At each episode, our agent observes a *state* of incoming a sequence of all BS functions drawn from the *environment* of vRANs, takes an *action* to decide the splits for all the BSs, and expects to receive feedback signals of the *reward* (total network cost) and

penalization (for violating the constraints). Our state comprises of a sequence information of BS functions: $\mathcal{F} = \{\mathcal{F}_n\}_{n=1}^N$, where \mathcal{F}_n is a set of BS- n functions. Given the input state, our agent assigns $\mathcal{O} = \{o_n \in \{0, 1, 2, 3\}, \forall n \in \mathcal{N}\}$ as a set of selected splits for all the BSs, which decides the placement of BS functions at the CU or DUs. Our objective is to minimize the total network cost while enforcing the constraint requirements. Given the selected splits, our agent expects to receive scalar values from the environment consisting of: *i*) $J(\mathcal{O}|\mathcal{F})$, the total induced cost and *ii*) $\xi(\mathcal{O}|\mathcal{F})$, the weighted sum of penalization. Further, we consider a particular RL algorithm using one-step constrained policy optimization and neural network architecture, where the interactions are narrowed to a single time step at every episode, and our agent learns iteratively over episodes.

Our goal is to design a stochastic policy $\pi_\theta(\mathcal{O}|\mathcal{F})$ parameterized by a neural network with weights θ to predict the splits for all the BSs to minimize the total cost while satisfying constraint requirements. However, we have the N BSs that need to deploy the splits together, where each has four possible split options. Each split decision is also interdependent as the BSs share the same network links and computing servers. Consequently, our problem has a combinatorially large discrete action space with a total of 4^N possible actions. Such a curse dimensionality in high dimensional spaces can be avoided by modelling complicated joint probability distributions using the chain rule decomposition. Therefore, we design our policy based on a chain rule by factorizing the output probability, parameterized by a neural network with weights θ as:

$$\pi_\theta(\mathcal{O}|\mathcal{F}) = \prod_{n=1}^N \pi_\theta(o_n | o_{(<n)}, \mathcal{F}_n). \quad (11)$$

This policy strategy assigns a higher probability to the splits for having a lower cost and vice versa for every BS, which also can be represented by individual softmax modules (e.g., at the output layer). Motivated by [26], [27] that uses neural networks to estimate the same factorization of our stochastic policy for machine translation, we design our policy network using an encoder-decoder sequence-to-sequence model based on LSTM networks. Our policy network architecture, which also utilizes an attention mechanism, captures the dependency and correlation between split decisions. This architecture allows our policy to read input information from all BS functions, then maps them into split selections for all the BSs. In the training, we use a batch of B i.i.d samples on the stochastic policy to select the splits and generate several pretraining models. In the test, we perform an inference through a search strategy by greedy decoding or temperature sampling.

A. Policy Network Architecture

Our policy network infers a strategy to deploy the splits for all the BSs, given a sequence information of BS functions as an input $\mathcal{F} = \{\mathcal{F}_1, \dots, \mathcal{F}_N\}$. It is constructed from an encoder decoder sequence-to-sequence model with an attention mechanism based on LSTM networks [26], [27]. We also consider a batch training by drawing a batch of B i.i.d samples with different sequence order to encourage the exploration further.

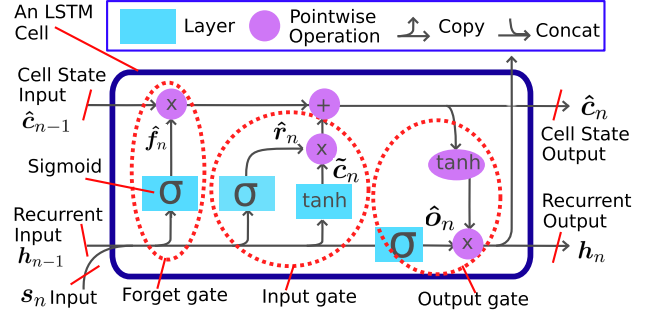


Fig. 2. A generic architecture of an LSTM cell.

1) *LSTM Structure*: We leverage LSTM networks, a particular RNN architecture [37], to construct our sequence-to-sequence model that maps the input BS functions into split decisions for all the BSs. An LSTM cell has three main structures comprising of: (i) a forget gate that receives the cell state input and learns how long should memorize or forget from the past; (ii) an input gate that aggregates the current input and the output of past steps, then feeds them to the activation function; and (iii) an output gate that provides the LSTM output from the combination of current cell state and the output of input gate. The relationship of these blocks can be expressed as:

$$\hat{f}_n = \sigma(W_f[h_{n-1}^T, s_n^T]^T + b_f), \quad (12)$$

$$\hat{r}_n = \sigma(W_r[h_{n-1}^T, s_n^T]^T + b_r), \quad (13)$$

$$\tilde{c}_n = \tanh(W_c[h_{n-1}^T, s_n^T]^T + b_c), \quad (14)$$

$$\hat{c}_n = \hat{f}_n * \hat{c}_{n-1} + \hat{r}_n * \tilde{c}_n, \quad (15)$$

$$\hat{o}_n = \sigma(W_o[h_{n-1}^T, s_n^T]^T + b_o), \quad (16)$$

$$h_n = \hat{o}_n * \tanh(\hat{c}_n), \quad (17)$$

where function $\sigma(x) \triangleq \frac{1}{1+\exp(-x)}$ is the sigmoid function and symbol $*$ is element-wise multiplication. The weight and bias matrices for the respective forget, input and output gates of the LSTM cell are represented by W_f, W_r, W_c, W_o and b_f, b_r, b_c, b_o . Multiple LSTM layers can be further stacked one on top of another (a stacked LSTM) to create a deeper model, which may obtain more accurate prediction. Each LSTM cell reads an input of embedding vector representation $s_n \in [-1, 1]^E$ translated from each input \mathcal{F}_n , where E is the embedding size. The structure of an LSTM cell is illustrated in Fig 2 and utilized to construct our sequence-to-sequence model.

2) *Policy Network*: Our policy network is built from an encoder-decoder sequence-to-sequence model based on LSTM networks. One main drawback of vanilla sequence model is generally unable to learn accurately long sequence. Therefore, the vanilla model may not be able learn our problem with large number of BSs. An attention mechanism comes to address this issue as it considers all the hidden state from all input sequences. The encoder read the entire input sequence to a fixed-length vector. The decoder decides the deployed split of each BS at each step from an output function based on its own previous state combined with an attention over the

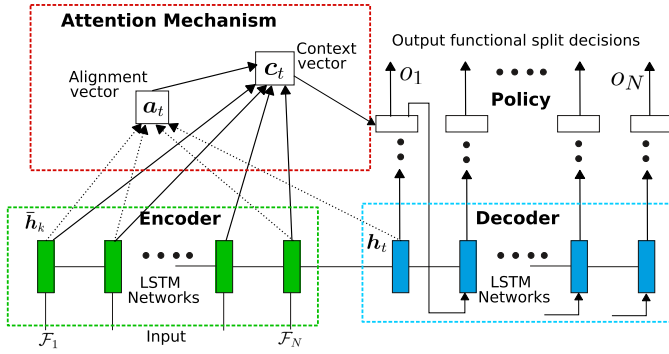


Fig. 3. **Policy Network.** CDRS utilizes a neural network architecture to approximate the stochastic policy over the solution. It is constructed by an encoder-decoder sequence-to-sequence model with attention mechanism based on LSTM networks.

encoder hidden states [27]. The decoder network hidden state is defined with a function: $h_t = f(h_{t-1}, \bar{h}_{t-1}, c_t)$, where c_t and \bar{h}_t are the context vector and the source hidden state at time step t . Our model derives the context vector c_t that captures relevant source information that helps to predict the splits. The main idea is to use an attention mechanism, where the context vector c_t takes consideration of all the hidden states of the encoder and the alignment vector a_t :

$$c_t = \sum_{k \in \mathcal{N}} a_{tk} \bar{h}_k. \quad (18)$$

Note that the alignment vector has an equal size to the number of steps in the source side, which can be calculated by comparing the current target hidden state of decoder h_t with each source hidden state \bar{h}_k as:

$$a_{tk} = \frac{\exp(\text{score}(h_t, \bar{h}_k))}{\sum_{k'=1}^N \exp(\text{score}(h_t, \bar{h}_{k'}))} \quad (19)$$

This alignment model gives a score a_{tk} which describes how well the pair of input at position k and the output at position t . The alignment score is parameterized by a feed-forward network where the network is trained jointly with the other models [27]. The score function is defined by a non-linear activation function following Bahdanau's additive style:

$$\text{score}(h_t, \bar{h}_k) = v_a^\top (\tanh(w_1 h_t + w_2 \bar{h}_k)), \quad (20)$$

where $v_a^\top \in \mathbb{R}^n$, $w_1 \in \mathbb{R}^{n \times n}$ and $w_2 \in \mathbb{R}^{n \times n}$ are defined as the weight matrices to be learned in the alignment model, and n is the size of hidden layers. The overall architecture of our policy network is illustrated in Fig. 3.

B. Constrained Policy Gradient With Baseline

We train the above neural network model using a constrained policy gradient method with a self competing baseline. We define the objective of \mathbb{P} as an expected reward that is obtained for every vector of weights θ . Hence, the expected cost J in associated with the selected split o_n given BS- n functions is denoted as:

$$J^\pi(\theta | \mathcal{F}_n) = \mathbb{E}_{o_n \sim \pi(\cdot | \mathcal{F}_n)} [J(o_n)], \quad (21)$$

and we have the expected of total cost from all BSs:

$$J^\pi(\theta) = \mathbb{E}_{o_n \sim \mathcal{O}} [J(\theta | \mathcal{O})]. \quad (22)$$

The vRAN system has constraints of delay requirement and computational and link capacity. Therefore, our original problem turns to a primal problem as:

$$\mathbb{P}_{1P} : \min_{\pi \in \Pi} J^\pi(\theta); \quad \text{s.t. } J_{C_i}^\pi(\theta) \leq 0, \forall i,$$

where we define $J_C^\pi(\theta) = (J_{C_i}^\pi(\theta), \forall i)$ as a function of constraint dissatisfaction to capture the penalization that the environment returns for violating each i constraint requirement, e.g., computing, link, delay. In this problem, we consider parametrized stochastic policy using a neural network. In order to ensure the convergence of our policy to constraint satisfaction, we follow [30] and make assumptions:

Assumption 1: J^π is bounded for all policies $\pi \in \Pi$.

Assumption 2: Each local minima of $J_C^\pi(\theta)$ is a feasible solution.

Assumption 2 describes that any local minima π_θ satisfies all constraints, e.g., $J_{C_i}^\pi(\theta) \leq 0, \forall i$. It is the minimal requirement that guarantees the convergence of a gradient algorithm to a feasible solution. The stricter assumptions, e.g., convexity, may guarantee the optimal solution.

Next, we reformulate \mathbb{P}_{1P} to unconstrained problem with Lagrange relaxation method [38]. The penalty signal is also included aside from the original objective for infeasibility, which leads to a sub-optimality for infeasible solutions. Given \mathbb{P}_{1P} , we have the dual function:

$$\begin{aligned} g(\mu) &= \min_{\theta} J_L^\pi(\mu, \theta) = \min_{\theta} J^\pi(\theta) + \sum_i \mu_i J_{C_i}^\pi(\theta) \\ &= \min_{\theta} J^\pi(\theta) + J_\zeta^\pi(\xi), \end{aligned} \quad (23)$$

where $\mu = (\mu_i, \forall i)$, $J_L^\pi(\mu, \theta)$ and $J_\zeta^\pi(\xi)$ are the penalty coefficients (Lagrange multipliers), Lagrange objective function and the expected penalization, respectively. Then, we define the dual problem:

$$\mathbb{P}_{1D} : \max_{\mu} g(\mu).$$

\mathbb{P}_{1D} aims to find a local optima or a saddle point $(\theta(\mu^*), \mu^*)$, which is a feasible solution. The feasible solution is a solution that satisfies: $J_{C_i}^\pi(\theta) \leq 0, \forall i$. To compute the weights θ that optimize the objective, we use Monte-Carlo policy gradient and stochastic gradient descent by the following update:

$$\theta_{k+1} = \theta_k - \eta_a(k) \nabla_{\theta} J_L^\pi(\mu, \theta), \quad (24)$$

where $\eta_a(k)$ is the step-size. The gradient $\nabla_{\theta} J_L^\pi(\mu, \theta)$ with regards to weights θ can be calculated using a log-likelihood method as:

$$\nabla_{\theta} J_L^\pi(\theta) = \mathbb{E}_{\mathcal{O} \sim \pi_{\theta}(\cdot | \mathcal{F})} [L(\mathcal{O} | \mathcal{F}) \nabla_{\theta} \log \pi_{\theta}(\mathcal{O} | \mathcal{F})]. \quad (25)$$

$L(\mathcal{O} | \mathcal{F})$ represents the total cost with penalization obtained from: $L(\mathcal{O} | \mathcal{F}) = J(\mathcal{O} | \mathcal{F}) + \xi(\mathcal{O} | \mathcal{F})$, where $J(\mathcal{O} | \mathcal{F})$ is the total network cost in each iteration and $\xi(\mathcal{O} | \mathcal{F}) = \mu C(\mathcal{O} | \mathcal{F})$ is the weighted sum of constraint dissatisfaction of $C(\mathcal{O} | \mathcal{F})$.

The penalty coefficient μ is set manually [29], [39] for CDRS-Fixed within a range $[0, \mu_{\max}]$.⁹ In this case, the selection of μ can be set following intuition approach in [29] (Appendix C), i.e., agent will not pay attention to penalty if $\mu = 0$, and it will only converge to penalization if $\mu = \infty$. Hence, selecting the appropriate penalty coefficient is important in this case. Otherwise, we can follow a less intuitive approach by adaptively updating the penalty coefficient (CDRS-Ada). CDRS-Ada is updated based on the primal-dual optimization (PDO) method inspired from [31]. Hence, we update the penalty coefficient in the ascent direction as:

$$\mu_{k+1} = \mu_k + \eta_d(k) \nabla_{\mu} J_L^{\pi}(\mu, \theta) \quad (26)$$

$$= \mu_k + \eta_d(k) (J_C^{\pi}(\theta))_+, \quad (27)$$

where $\eta_d(k)$ is the step-size (Dual) and $\nabla_{\mu} J_L^{\pi}(\mu, \theta) = \mathbb{E}_{\mathcal{O} \sim \pi_{\theta}(\cdot|\mathcal{F})} [C(\mathcal{O}|\mathcal{F})]$ is the gradient with respect to μ . The penalty coefficient μ_k is updated for every k -th iteration and will converge to a fixed value once the constraints are satisfied [30], [31]. Then, Monte-Carlo sampling can be applied to approximate $J_L^{\pi}(\theta)$ by drawing B i.i.d samples $\mathcal{F}^1, \dots, \mathcal{F}^B \sim \mathcal{F}$, which can be written:

$$\nabla_{\theta} J_L^{\pi}(\theta) \approx \frac{1}{B} \sum_{i=1}^B (L(\mathcal{O}^i|\mathcal{F}^i) - b_{\theta_v}(\mathcal{F}^i)) \nabla_{\theta} \log \pi_{\theta}(\mathcal{O}^i|\mathcal{F}^i), \quad (28)$$

where $b_{\theta_v}(\mathcal{F}^i)$ is the baseline estimation given the state input of i -th batch, parameterized by a neural network structure with weights θ_v .

1) *Baseline Estimator*: The baseline choice can be from an exponential moving average of the reward over time that captures the improving policy in training. Although it succeeds in the Christofides algorithm, it does not perform well because it cannot differentiate between inputs [25]. To this end, we use a parametric baseline b_{θ_v} to estimate the expected total cost with penalization that typically improves the learning performance. We estimate the baseline through an auxiliary network built from an LSTM encoder connected to a multilayer perceptron output layer. The auxiliary network (parameterized by θ_v) that approximates the expected cost with penalization from input \mathcal{F} is trained with stochastic gradient descent. It employs a mean squared error (MSE) objective, calculated from the prediction of b_{θ_v} and the total cost with penalization $L(\mathcal{O}^i|\mathcal{F}^i)$, and sampled by the most recent policy (obtained from the environment). We formulate the auxiliary network goal is to minimize the below loss function:

$$\mathcal{L}(\theta_v) = \frac{1}{B} \sum_{i=1}^B \|b_{\theta_v}(\mathcal{F}^i) - L(\mathcal{O}^i|\mathcal{F}^i)\|_2^2. \quad (29)$$

Fig. 4 illustrates the architecture of the auxiliary network for estimating the baseline.

To sum up, our training procedures are summarized in Algorithm 1 and illustrated in Fig. 5, which run iteratively by K episodes (epochs) based on a single time-step Monte-Carlo policy gradient with a baseline estimator. The sequence of policy updates will converge to a locally optimal policy

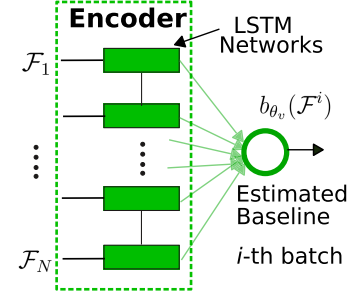


Fig. 4. **Baseline Estimator**. The self-competing baseline of CDRS is estimated using an auxiliary network constructed from an LSTM encoder connected to a multilayer perceptron output linear layer.

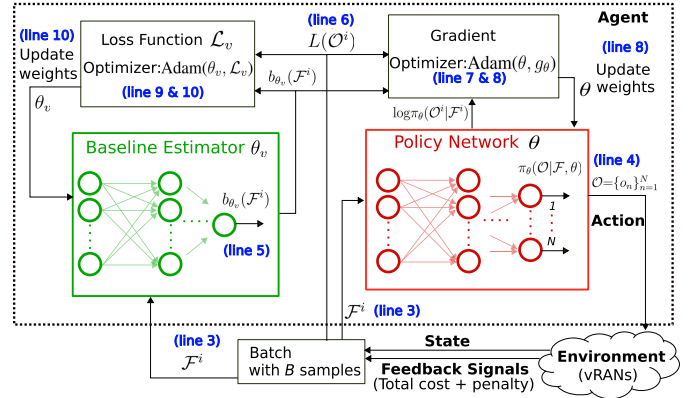


Fig. 5. **CDRS Diagram**. CDRS is trained using a single time step Monte-Carlo policy gradient algorithm, where at every epoch, the interactions with the environment are narrowed to a single time step. Our agent learns the policy iteratively over epochs.

Algorithm 1 CDRS Training

Input: K (Num of epoch), B (Batch size), \mathcal{F} (Learning set)

Initialize: assign agent and critic (baseline) networks with random weights θ and θ_v .

```

1 for  $k = 1, \dots, K$  do
2    $d\theta \leftarrow 0$  % Reset gradient
3    $\mathcal{F}^i \sim \text{SampleInput}(\mathcal{F})$  for  $i \in \{1, \dots, B\}$ .
4    $\mathcal{O}^i \sim \text{SampleSolution}(\pi_{\theta}(\cdot|\mathcal{F}))$  for  $i \in \{1, \dots, B\}$ .
5    $b^i \leftarrow b_{\theta_v}(\mathcal{F}^i)$  for  $i \in \{1, \dots, B\}$ .
6   Compute  $L(\mathcal{O}^i)$  for  $i \in \{1, \dots, B\}$ .
7    $g_{\theta} \leftarrow \frac{1}{B} \sum_{i=1}^B (L(\mathcal{O}^i) - b^i) \nabla_{\theta} \log \pi_{\theta}(\mathcal{O}^i|\mathcal{F}^i)$  from (28).
8    $\theta \leftarrow \text{Adam}(\theta, g_{\theta})$  %Run Adam algorithm
9    $\mathcal{L}_v \leftarrow \frac{1}{B} \sum_{i=1}^B \|b^i - L(\mathcal{O}^i)\|_2^2$  from (29).
10   $\theta_v \leftarrow \text{Adam}(\theta_v, \mathcal{L}_v)$  %Run Adam algorithm
11  Update  $\mu$  from (26) %CDRS-Ada
12  Set  $\mu = \max(0, \mu)$  %CDRS-Ada
13 end
14 return  $\theta, \theta_v, \mu$ 

```

and the penalty coefficient updates (e.g., CDRS-Ada) will converge to a fixed value when all constraints are satisfied; see also [30], [31].

⁹If Assumption 2 is satisfied, μ_{\max} can be set to ∞ [30].

C. Searching Strategy

At the test time, evaluating the total network cost is inexpensive as it only requires a forward pass from the policy network to decide the splits. Our agent can add a search procedure during the inference process by considering solution candidates from multiple pretraining models to select the splits. It can help to reduce the inferred policy suffering from a severe suboptimality. We employ two different search strategies by greedy decoding and temperature sampling [25].

1) *Greedy Decoding*: It is the simplest search strategy. The idea is to greedily select the splits with the highest probability for having the lowest cost from multiple pretraining models during the inference time. Then, we can extend CDRS to CDRS-Fixed-G, which uses a fixed penalty coefficient with greedy decoding and CDRS-Ada-G that uses an adaptive penalty coefficient with greedy decoding.

2) *Temperature Sampling*: This method samples through stochastic policy for each pretraining model to generate several candidate solutions then decides the splits with the lowest total cost among them [25], [29]. As opposed to the heuristic solvers, it does not sample the different split options. Instead, it samples through the stochastic policy and controls the sparsity of the output distribution by a temperature hyperparameter T . The softmax function in (19) is modified to

$$\mathbf{a}_{tk} = \frac{\exp(\text{score}(\mathbf{h}_t, \tilde{\mathbf{h}}_k)/T)}{\sum_{k'=1}^N \exp(\text{score}(\mathbf{h}_t, \tilde{\mathbf{h}}_{k'})/T)} \quad (\text{softmax temperature}).$$

In the training, the temperature hyperparameter T is set to 1. Meanwhile, we modify to $T > 1$ during the test, hence the output distribution becomes less step, which prevents the model from being overconfident. With this method, we can extend CDRS to CDRS-Fixed-T (fixed penalty coefficient, temperature sampling) and CDRS-Ada-T (adaptive penalty coefficient, temperature sampling). Note that this method requires additional time, which depends on the number of samples.

V. RESULTS AND DISCUSSION

In this section, we conduct several experiments to evaluate our approach using synthetic and real network datasets. We aim to examine our approach in regards to: (i) the behaviour during the training process, (ii) the accuracy and solution distributions to the optimality with different penalty coefficient and search strategy settings, (iii) the impact of routing costs and traffic loads on the optimality performance and total network cost, and (iv) the computational time.

A. Environment & Experiment Setup

We use synthetic (R1) and real (R2) network datasets to evaluate our approach. We generate R1 with stricter constraints and a larger scale environment than R2. R1 is generated using the Waxman algorithm [32] with parameters such as link probability (α) and edge length control (β). These respective parameters (α, β) are set to (0.5, 0.1). R1 has 1 CU and 99 DUs. In the case of R2, we utilize a real network dataset from [40], which has 1 CU and 63 DUs. We assume that the routers are co-located with the DUs. R1 and R2 differ in parameters, e.g., location, link capacity, weighted

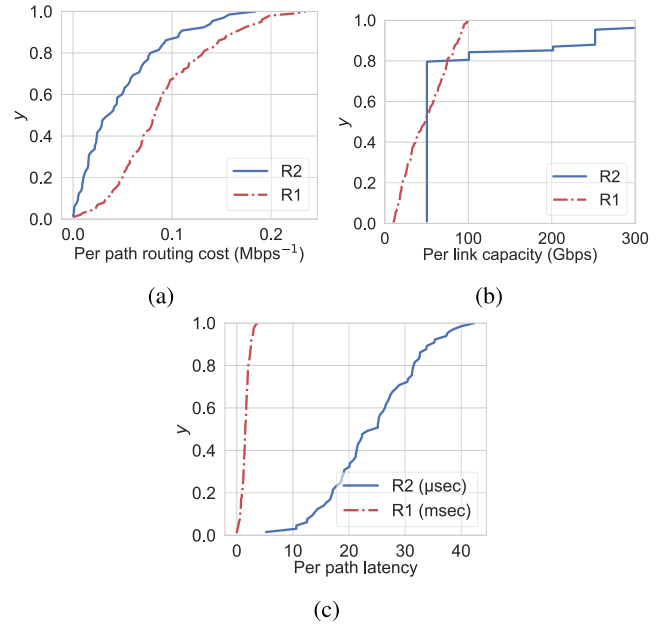


Fig. 6. **RANs dist.** eCDF of (a) per-path routing cost, (b) per-link capacity, (c) per-path latency for R1 and R2.

link, delay. We use a standard store-and-forward model to calculate the delay. It is from $12000/c_{ij}$, $4\mu\text{secs/Km}$ and $5\mu\text{secs}$ for transmission, propagation and processing delay, respectively; see [18]. The link capacity varies to 100 Gbps (R1) and 252 Gbps (R2). The path delay reaches to $3658\mu\text{s}$ (R1) and $42\mu\text{s}$ (R2). In R1, the routing cost per path is calculated from the total cost per link (randomly generated) which belongs to the selected path. A link with a routing cost of 1 monetary unit per Mbps means having the same cost as a DU computing cost. We consider the routing cost within a range of 0.001 – 0.01 times of DU computing cost (for the same network load) for each link in R1. In R2, we calculate the distance between nodes based on its geolocation dataset from [40] and charge the cost of 0.01 monetary units per Mbps/km. Fig. 6 depicts the parameter distributions of our RANs with eCDF.

In this experiment, all system parameters correspond to testbed measurements of previous studies [2], [3], [18], [41]. We assume a high load scenario $\lambda_n = 150$ Mbps for every DU. This setting is based on 1 user/TTI, 2×2 MIMO, 20 Mhz (100 PRB), 2 TBs of 75376 bits/subframe and IP MTU 1500B. We use an Intel Haswell i7-4770 3.40GHz CPU as the *reference core*, and set the maximum computing capacity to 75 RCs for CU and 7.5 RCs for each DU. Each split $o \in \{0, 1, 2, 3\}$ incurs computational load $\rho_o^d = \{0.05, 0.04, 0.00325, 0\}$ RCs per Mbps at each DU and $\rho_o^c = \{0, 0.001, 0.00175, 0.05\}$ RCs per Mbps at the CU. The VM instantiation cost at the CU is half of the DU ($\alpha_0 = \alpha_n/2$) and the processing cost is set to $\beta_0 = 0.017\beta_n$.

Our learning rate is initially set to $\eta_a = 0.0001$ (Agent) and $\eta_b = 0.005$ (Baseline) with the batch size: 128. Our neural network has the number of layers, hidden dimension and embedding size with 1, 32 and 32, respectively. The temperature hyperparameter is set to $T = 1$ by default,

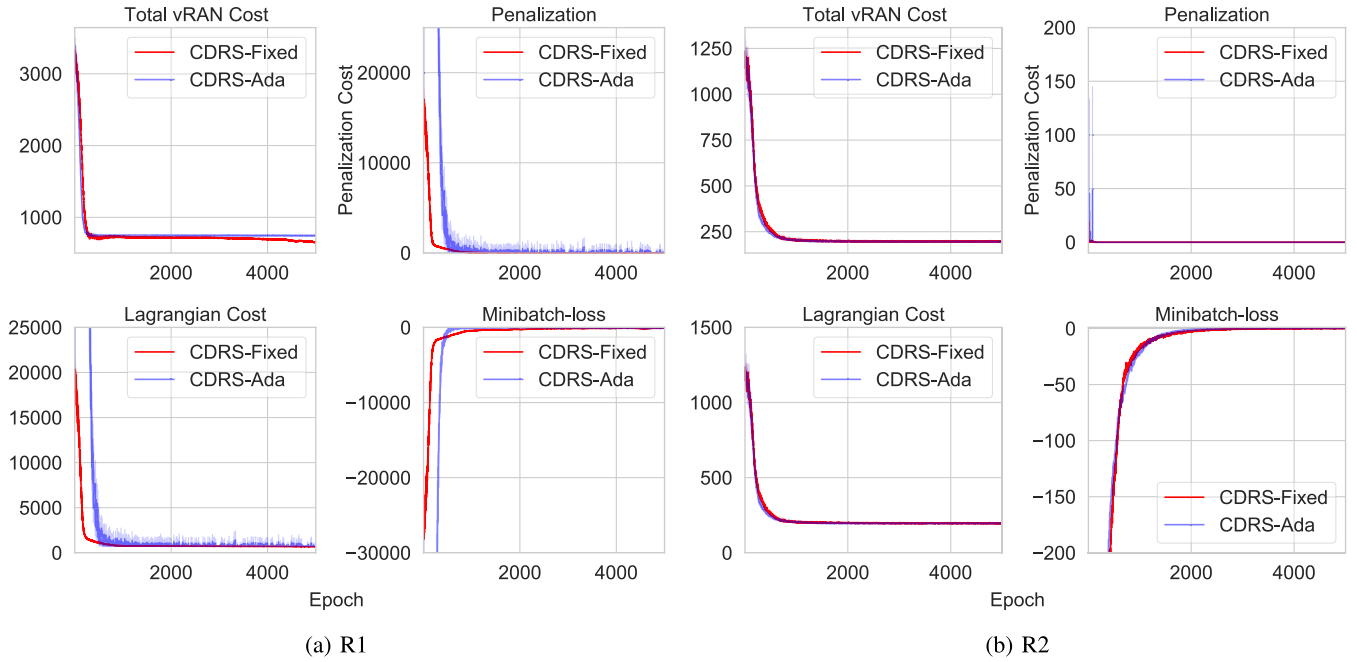


Fig. 7. **Training results of CDRS in (a) R1 and (b) R2.** CDRS-Fixed uses a fixed value of penalty coefficient (reward shaping) with $\mu_i = 1, \forall i$. CDRS-Ada utilizes an adaptive update of penalty coefficient.

so the model computes the softmax function directly. We scale all the original values of weighted paths and traffic loads randomly with uniform distribution $[0, 1]$ as in [25]. Then, we generate three models (RL-pretaining) as outputs of our training with 50000 (in R1) and 15000 (in R2) epochs each. CDRS-Fixed uses a fixed penalty coefficient with $\mu_i = 1, \forall i$ for all epochs while CDRS-Ada is set with initial penalty coefficient $\mu_i(0) = 1, \forall i$ and step-size $\eta_d = 0.001$. The training is performed with Tensorflow 1.15.3 and Python 3.7.4. In the test, the temperature sampling method uses 16 samples and $T = 15$ (softmax temperature).

B. Training Analysis

We aim to examine the behaviour of CDRS-Fixed and CDRS-Ada during the training process in R1 and R2. We focus on the mini-batch loss, reward (total network cost), Lagrangian cost and penalization.

Fig. 7 visualizes the training of CDRS-Fixed and CDRS-Ada in R1 and R2. We found additional costs because of penalization at the beginning of the training for both settings in R1 and R2. It occurs because CDRS-Fixed and CDRS-Ada try to find the solution, but violate the constraint sets (e.g., latency, bandwidth, computation). Fig. 7 also shows a significant difference in the cost of penalization in R1 compared to R2. The main reason is that R1 has stricter constraints, e.g., larger path delays, smaller link capacity than R2. We can also see that CDRS-Fixed and CDRS-Ada improve their policy by focusing on constraint satisfaction and then correcting the weights via stochastic gradient descent. It is proven from our agent's behaviour in R1 and R2, where each penalization cost keeps decreasing and turns to zero as soon as the training goes. CDRS-Ada sets the penalty coefficient

increasing in the ascent direction, causing a higher penalization value than CDRS-Fixed. However, it can help speed up the policy toward constraint satisfaction, i.e., CDRS-Ada penalization downs faster than CDRS-Fixed.

We also found that the policy of CDRS-Ada converges faster than CDRS-Fixed from the behaviour of mini-batch loss in R1. Despite the mini-batch loss decreases to near zero after several epochs, the mini-batch loss of CDRS-Ada diminishes faster than CDRS-Fixed. However, CDRS-Ada suffers from more severe sub-optimality. It is shown by the total vRAN cost of CDRS-Ada that converges to a fixed value but has a higher cost compared to CDRS-Fixed. Then, we have the Lagrangian cost from the sum of vRAN cost and penalization cost. It describes how our agent tries to minimize the primal problem \mathbb{P}_{IP} through the dual problem \mathbb{P}_{ID} . When our agent finally dismisses the penalization cost, it means that all constraints are satisfied. As a result, the Lagrangian cost becomes equal to the vRAN cost, and the penalty coefficient of CDRS-Ada converges to a fixed value. Although having different behaviours, CDRS-Ada and CDRS-Fixed can learn the solution and converge to the local minima or saddle point in R1 and R2.

1) *Findings:* 1) R1 has stricter constraint requirements than R2; hence, it produces a higher additional cost for penalization to CDRS-Fixed and CDRS-Ada. 2) CDRS-Fixed and CDRS-Ada improve the policy by focusing on the penalization; then, it adjusts the weights as the training goes. 3) CDRS-Ada receives higher penalization compared to CDRS-Fixed as a result of increasing the penalty coefficient in the ascent direction; however, it also helps speed up the policy to constraint satisfaction. 4) CDRS-Ada converges faster but has a higher cost than CDRS-Fixed in R1. 5) When all constraints are satisfied, the Lagrangian cost becomes equal

to the total vRAN cost, and the penalty coefficient of CDRS-Ada converges to a fixed value.

C. Accuracy of Solutions

In this part, we study the accuracy of CDRS over different penalty coefficient and search strategy settings: CDRS-Fixed-G, CDRS-Fixed-T, CDRS-Ada-G and CDRS-Ada-T. We conduct 128 tests with a distinct sequence order of the BSs in R1 and R2 to assess how accurate these four CDRS settings find the solution of the vRAN split problem. We utilize three pretraining models from our CDRS training.

Fig. 8 shows the distribution of the solutions from CDRS-Fixed-G, CDRS-Fixed-T, CDRS-Ada-G and CDRS-Ada-T in R1 and R2. Each bar counts the number of offered solutions resulting in some suboptimality, represented using the optimality gap (error). It shows that the distribution varies between four settings, especially in a stricter environment (R1). Still, all of these settings can guarantee less than 0.6% (R1) and 0.1% (R2) of the optimality gap. In R1, CDRS-Fixed-G and CDRS-Fixed-T perform better by offering lower solution errors ($\leq 0.05\%$ and $\leq 0.05\%$ of optimality gap) than CDRS-Ada-G and CDRS-Ada-T ($\leq 0.6\%$). It means that a fixed penalty coefficient setting can lead to a better optimality performance during the test than the adaptive one. However, CDRS-Fixed-G, CDRS-Ada-G and CDRS-Ada-T have a similar performance in R2. Regardless of R1 or R2, using a sampling method with a temperature hyperparameter can improve (or at least at same) the optimality performance than greedy decoding. It is shown from the higher total number of solutions (counts) for a sampling method that having a lower error. The combination of a fixed penalty in the training and temperature sampling method (CDRS-Fixed-T) can improve the solution performance significantly both in R1 and R2. It can achieve an optimal value (R2) and less than 0.05% of error for a more complex environment (R1). It is also shown that CDRS-Fixed-T is less affected to the stricter environment than any other settings where all of the distribution solutions are in less than 0.05%.

1) *Findings:* 1) CDRS-Fixed-G, CDRS-Fixed-T, CDRS-Ada-G and CDRS-Ada-T can guarantee the solution with very close to the optimal value offering less than 0.6% (R1) and 0.1% (R2) of the optimality gap over 128 tests. 2) CDRS-Fixed-T can significantly improve the optimality performance (offers $\leq 0.05\%$ of optimality gap) and outperforms the other settings.

D. Impact of Routing Cost

This part studies the impact of altering the routing cost to CDRS-Fixed-G, CDRS-Fixed-T, CDRS-Ada-G and CDRS-Ada-T. We aim to examine how the routing cost affects optimality performance and the total network cost. Hence, the default routing cost is changed within a range of $\gamma = 0.1$ to $\gamma = 1$. This change can arise due to increasing or decreasing the leasing agreement's price, maintenance, etc. The traffic load is fixed with $\lambda_n = 150$ Mbps. We utilize three pretraining

models, conduct 128 tests for each routing cost scale, and analyze the offered solutions' distribution. We also consider benchmarking with two extremes of RAN setups, fully D-RAN and C-RAN¹⁰ to assess how significant the routing cost affects the total network cost over various RAN setups.

Fig. 9 depicts how the routing cost affects the optimality performance of CDRS-Fixed-G, CDRS-Fixed-T, CDRS-Ada-G and CDRS-Ada-T. It shows that the overall optimality gap (error) diminishes as the routing cost increases; then, it converges to a specific value. In R1, we see a performance improvement as the errors decrease for CDRS-Fixed-G ($\approx 75\%$), CDRS-Fixed-T ($\approx 75\%$), CDRS-Ada-G ($\approx 78\%$) and CDRS-Ada-T ($\approx 75\%$) by the increase of routing cost. It also shows that CDRS-Ada-G gets the most impact while CDRS-Fixed-T is the least affected. In R2, all CDRS settings also have a similar trend in terms of error performance. Although we have changed the routing cost from the default parameter, we found that altering the routing cost gives relatively less effect to these settings where the errors are maintained under 1.8%. CDRS-Fixed-T even can guarantee the solution under 0.08% ($\gamma = 0.1$) of the optimality gap.

Fig. 10 shows the routing cost's effect on the total network cost of CDRS-Fixed-T and D-RAN, normalized to the C-RAN cost in R1 and R2. It shows that CDRS-Fixed-T can obtain a larger cost-saving than the D-RAN cost at a cheaper routing cost by up to 59.06% of cost-saving at $\gamma = 0.1$ while only 25.49% of cost-saving at $\gamma = 1$ in R1. Compared to C-RAN, CDRS-Fixed-T can save the cost by up to 92% at $\gamma = 1$ in R1. However, this gain diminishes as the routing cost decreases and eventually CDRS-Fixed-T will reach near the C-RAN cost if all constraint requirements are eligible. A similar trend also appears for R2. Moreover, CDRS-Fixed-T can offer the solution extremely close to the optimal solution by $\leq 0.09\%$ (R1) and $\leq 0.5\%$ (R2).

1) *Findings:* 1) The increase of routing cost reduces the optimality gap (error); then, it converges to a fixed value. 2) CDRS-Fixed-T is the least affected by the routing cost changes, while CDRS-Ada-G is the most affected. 3) Scaling the routing cost from $\gamma = 0.1$ to $\gamma = 1$ does not significantly degrade the optimality performance. 4) CDRS-Fixed-T has the lowest optimality gap than other CDRS settings, and becomes the most cost-effective setup in R1 and R2. 5) CDRS-Fixed-T can reach near the D-RAN cost at a high routing cost, while it can be near the C-RAN cost at a cheap routing cost if all constraint requirements are eligible.

E. Impact of Traffic Load

In this part, we assess how the traffic load affects the optimality performance and the total network cost. We change the traffic load from 10 Mbps to 150 Mbps. This evaluation is conducted using three pretraining models and examined over 128 tests.

¹⁰We practically cannot implement C-RAN because our RANs do not meet the constraint requirements of delay, bandwidth and CU capacity to deploy C-RAN. The presented C-RAN in this experiment is just for benchmarking; hence we also do not consider the penalization cost (constraints violation) for this case.

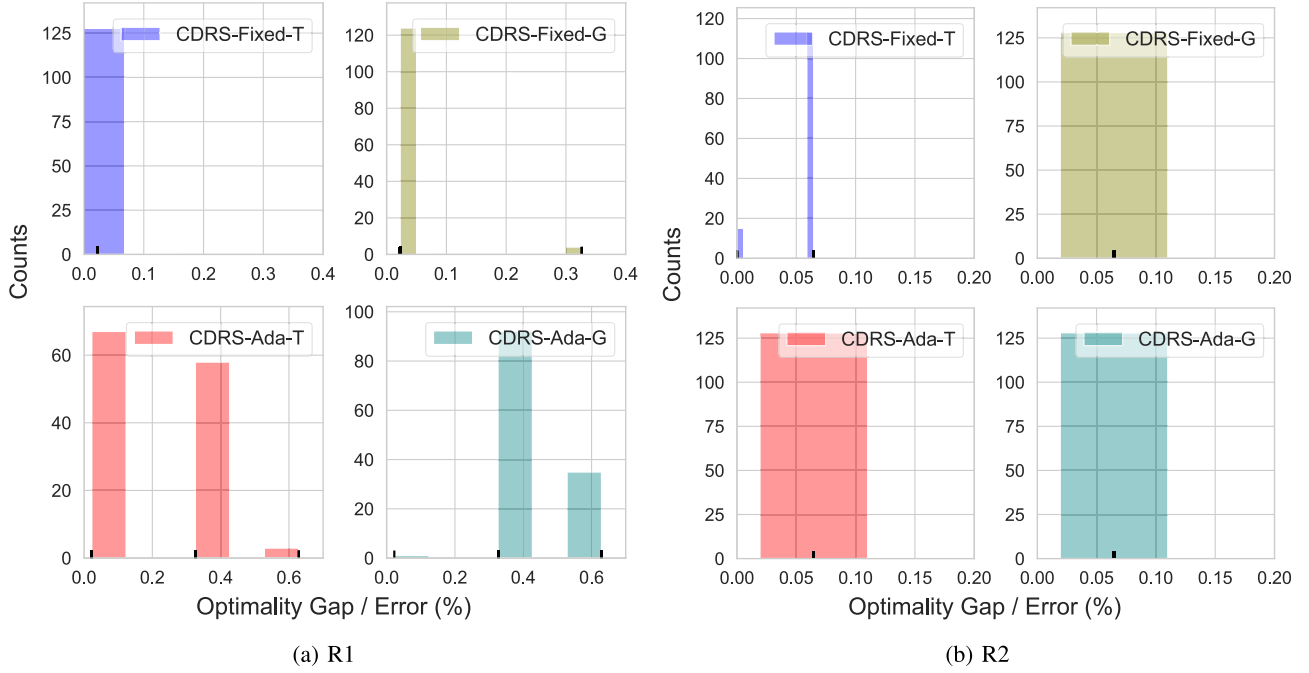


Fig. 8. **Histogram of CDRS accuracy in (a) R1 and (b) R2.** The accuracy is calculated over 128 tests. CDRS-Ada-T and CDRS-Fixed-T are set with $T = 15$ and 16 samples.

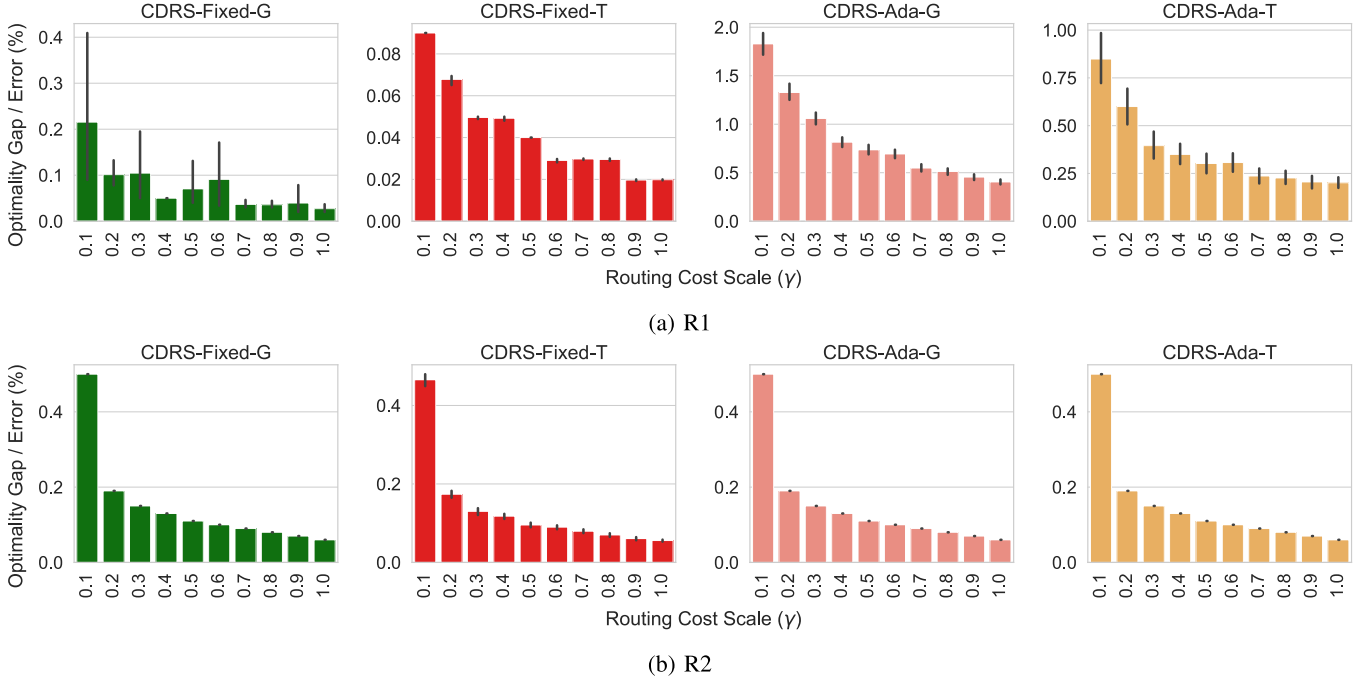


Fig. 9. **Impact of the routing cost to the accuracy in (a) R1 and (b) R2.** Study of altering the routing cost to the optimality performance with $\lambda_n = 150$ Mbps, $\forall n \in \mathcal{N}$. There are 128 tests for each routing scale $[0.1, 1]$.

Fig 11 shows the impact of altering the traffic load to the optimality performance of CDRS-Fixed-G, CDRS-Fixed-T, CDRS-Ada-G and CDRS-Ada-T. In R1, it shows that the increase of traffic load in line with the rise of the error to CDRS-Ada-G and CDRS-Ada-T, but it then diminishes to a fixed value, i.e., around 0.4% (CDRS-Ada-G) and 0.18% (CDRS-Ada-T). However, the traffic load does not

significantly affect CDRS-Fixed-G and CDRS-Fixed-T, where they stay at around 0.04% and 0.02% of errors, respectively, in R1. In R2, CDRS-Fixed-G, CDRS-Fixed-T, CDRS-Ada-G and CDRS-Ada-T have the same trend where the optimality gap increases with the traffic load; then, it diminishes at around 0.05%. We also found that CDRS-Fixed-T has a better optimality performance and a more stable solution.

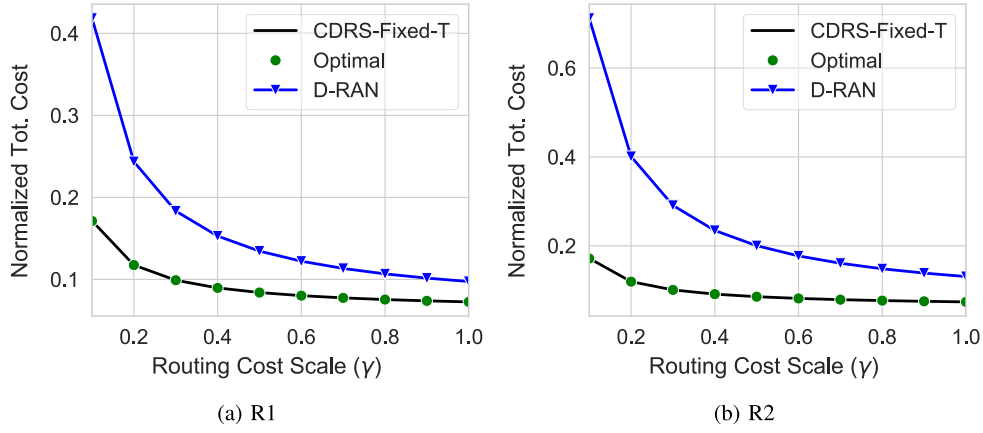


Fig. 10. **Impact of routing cost to the total cost in (a) R1 and (b) R2.** We also compare our approach (e.g., CDRS-Fixed-T) to two extreme cases: fully D-RAN and C-RAN, and the optimal value with the routing cost scaling from 0.1 to 1 of default R1 and R2. The presented cost above is normalized toward fully C-RAN cost.

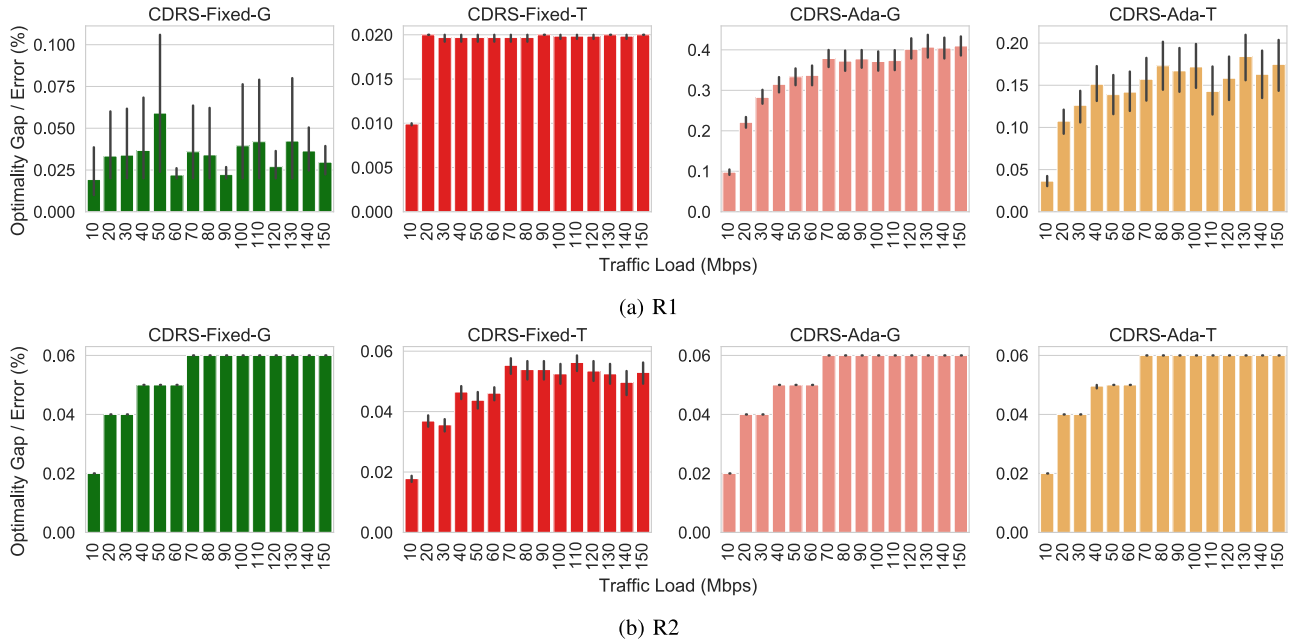


Fig. 11. **Impact of the traffic load to the accuracy in (a) R1 and (b) R2.** Study of traffic load to the optimality performance. There are 128 tests for each traffic load.

Fig 12 examines the impact of traffic load on CDRS-Fixed-T and D-RAN cost normalized to the C-RAN cost. Despite an increase in CDRS-Fixed-T cost as the traffic load rises, it shows that CDRS-Fixed-T is still the most cost-effective compared to D-RAN and C-RAN in R1 and R2. CDRS-Fixed-T almost has the same cost as D-RAN at the low traffic load with only 12.33% cost-saving. This cost-saving then increases for the higher traffic load settings by up to 25.5% at 150 Mbps in R1. This trend also happens in R2. Compared to C-RAN, CDRS-Fixed-T significantly outperforms at the low traffic load, but this gain then diminishes as the increase of the load. CDRS-Fixed-T can reach near the C-RAN cost when all constraint requirements are satisfied, and the traffic load is high, but the routing cost is significantly low.

1) *Findings:* 1) CDRS-Fixed-T can offer to better optimality performance and more stable solution than other CDRS settings. 2) In R2, all CDRS settings have similar trends where

the increase of traffic load can also increase the optimality gap, but it then diminishes and stays at around 0.05% for CDRS-Fixed-T and 0.06% for the others. 3) CDRS-Fixed-T is the most cost-efficient compared to C-RAN and D-RAN. 4) CDRS-Fixed-T can eventually almost have the same C-RAN cost when all constraint requirements are satisfied, and the traffic load is high, but the routing cost is significantly low.

F. Computational Time

Finally, we examine the computational time to solve a single instance of the vRAN split problem. We use a small laptop with an Intel Core i5-7300U CPU@2.60GHz and 8GB memory. The computational time for each CDRS setting is a result of averaging 128 executions with a trained model. We report this evaluation in Table II. Overall, our proposed CDRS settings: CDRS-Fixed-G, CDRS-Fixed-T,

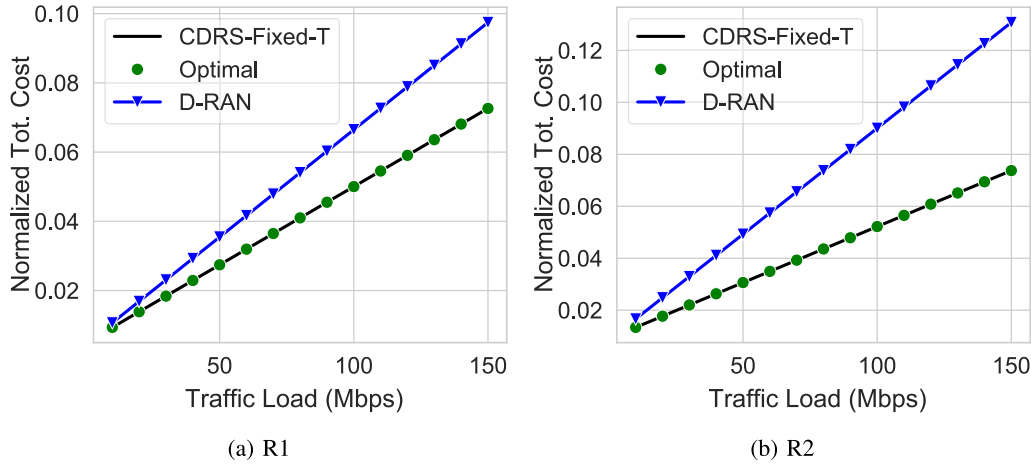


Fig. 12. **Impact of traffic load to total vRAN cost in (a) R1 and (b) R2.** On the comparison of our approach (e.g., CDRS-Fixed-T) to fully D-RAN. The presented cost above is normalized toward fully C-RAN cost.

TABLE II

COMPUTATIONAL TIME. STUDY OF COMPUTATIONAL TIME FOR SOLVING A SINGLE PROBLEM INSTANCE IN SECONDS. THE PRESENTED COMPUTATIONAL TIME IS A RESULT OF AVERAGING 128 EXECUTIONS

Topology	MIP solver	CDRS-Fixed-T	CDRS-Fixed-G	CDRS-Ada-T	CDRS-Ada-G
R1	0.2527	0.2026	0.0155	0.1985	0.0120
R2	0.1756	0.1240	0.0098	0.1207	0.0077

CDRS-Ada-G and CDRS-Ada-T, have a faster computational time than the MIP solver. CDRS-Ada-G is the fastest with 0.0120 secs and 0.0077 secs in R1 and R2 reaching to 22.82 times faster than the MIP solver. We also found that any CDRS settings with greedy decoding for the inference process, e.g., CDRS-Fixed-G, CDRS-Ada-G, is more time-efficient than a temperature sampling method with around 10-20 times faster. It is also shown that CDRS-Ada-G/T has a slightly faster computational time than CDRS-Fixed-G/T. Finally, we can sort from the fastest computational time as 1) CDRS-Ada-G, 2) CDRS-Fixed-G, 3) CDRS-Ada-T, 4) CDRS-Fixed-T, 5) the MIP solver.

1) *Findings:* 1) CDRS-Ada-G, CDRS-Fixed-G, CDRS-Ada-T, and CDRS-Fixed-T can reach up to 22.82, 17.99, 1.45 and 1.41 times faster than the MIP solver. 2) Greedy decoding is more time-efficient than a temperature sampling method for the inference process.

VI. CONCLUSION

In this paper, we have investigated the functional split optimization problem in which the BS functions can be deployed at the CU or DUs. We have formulated the problem mathematically and analyzed the complexity, which is shown to be combinatorial and NP-hard. Because finding the exact solution is computationally expensive and precise modelling the actual vRAN system is highly non-trivial, we have proposed CDRS as a solution framework to optimize the functional splits of the BSs amidst minimal assumptions about the underlying system. We have developed CDRS using a chain rule-based stochastic policy to handle the interdependence between split decisions and the large action space. We have applied LSTM networks-based sequence-to-sequence model to approximate

the policy. Since this policy is limited to an unconstrained problem, and vRAN's constraint requirements bound each function placement decision, we have leveraged a constrained policy gradient method to train the policy. We have also provided a search strategy by greedy decoding or temperature sampling to improve the optimality performance at the test time. The performance of CDRS has been extensively evaluated using synthetic and real network datasets. The results have shown that CDRS successfully learns the functional split decision with less than 0.05% optimality gap, attains considerable cost savings compared to C-RAN or D-RAN systems, and has a faster computational time than the optimal baseline.

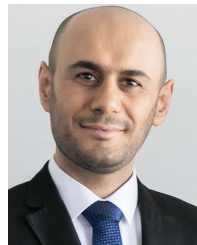
REFERENCES

- [1] F. W. Murti, S. Ali, and M. Latva-Aho, "Deep reinforcement based optimization of function splitting in virtualized radio access networks," in *Proc. IEEE Int. Conf. Commun. Workshops (ICC Workshops)*, Jun. 2021, pp. 1–6.
- [2] V. Suryaprakash, P. Rost, and G. Fettweis, "Are heterogeneous cloud-based radio access networks cost effective?" *IEEE J. Sel. Areas Commun.*, vol. 33, no. 10, pp. 2239–2251, Oct. 2015.
- [3] F. W. Murti, J. A. Ayala-Romero, A. Garcia-Saavedra, X. Costa-Perez, and G. Iosifidis, "An optimal deployment framework for multi-cloud virtualized radio access networks," *IEEE Trans. Wireless Commun.*, vol. 20, no. 4, pp. 2251–2265, Apr. 2021.
- [4] "Open and virtualized—The future radio access network," NEC, Tokyo, Japan, White Paper, 2020.
- [5] "5G immersive service opportunities with edge cloud and cloud RAN," Nokia, Espoo, Finland, White Paper, 2019.
- [6] L. M. P. Larsen, A. Checko, and H. L. Christiansen, "A survey of the functional splits proposed for 5G mobile crosshaul networks," *IEEE Commun. Surveys Tuts.*, vol. 21, no. 1, pp. 146–172, 1st Quart., 2019.
- [7] "Mobile anyhaul," Nokia, Espoo, Finland, White Paper, 2017.
- [8] S. Gonzalez-Diaz *et al.*, "Integrating fronthaul and backhaul networks: Transport challenges and feasibility results," *IEEE Trans. Mobile Comput.*, vol. 20, no. 2, pp. 533–549, Feb. 2021.

- [9] J. A. Ayala-Romero, A. Garcia-Saavedra, M. Gramaglia, X. Costa-Perez, A. Banchs, and J. J. Alcaraz, "vRAIn: Deep learning based orchestration for computing and radio resources in vRANs," *IEEE Trans. Mobile Comput.*, early access, Dec. 8, 2021, doi: [10.1109/TMC.2020.3043100](https://doi.org/10.1109/TMC.2020.3043100).
- [10] J. A. Ayala-Romero, A. Garcia-Saavedra, X. Costa-Perez, and G. Iosifidis, "Orchestrating energy-efficient vRANs: Bayesian learning and experimental results," *IEEE Trans. Mobile Comput.*, early access, Oct. 29, 2021, doi: [10.1109/TMC.2021.3123794](https://doi.org/10.1109/TMC.2021.3123794).
- [11] X. Foukas and B. Radunovic, "Concordia: Teaching the 5G vRAN to share compute," in *Proc. ACM SIGCOMM Conf.* New York, NY, USA: Association Computing Machinery, Aug. 2021, pp. 580–596.
- [12] *Study on New Radio Access Technology: Radio Access Architecture and Interfaces*, document TS 38.801, Version 14.0.0, 3GPP, Mar. 2017.
- [13] *NG-RAN: Architecture Description*, document TS 38.401, Version 16.8.0, 3GPP, Technical Specification Group Radio Access Network (NG-RAN), Mar. 2021.
- [14] *R6.0. Small Cell Virtualization Functional Splits and Use Cases*, document Rec. 159.07.02, Version 14.0.0, Release 7, Small Cell Forum, 2016.
- [15] H. Gupta, M. Sharma, A. A. Franklin, and B. R. Tamma, "Apt-RAN: A flexible split-based 5G RAN to minimize energy consumption and handovers," *IEEE Trans. Netw. Service Manage.*, vol. 17, no. 1, pp. 473–487, Mar. 2020.
- [16] A. Garcia-Saavedra, J. X. Salvat, X. Li, and X. Costa-Perez, "WizHaul: On the centralization degree of cloud RAN next generation fronthaul," *IEEE Trans. Mobile Comput.*, vol. 17, no. 10, pp. 2452–2466, Oct. 2018.
- [17] A. Garcia-Saavedra, X. Costa-Perez, D. J. Leith, and G. Iosifidis, "FluidRAN: Optimized vRAN/MEC orchestration," in *Proc. IEEE Conf. Comput. Commun.*, Apr. 2018, pp. 2366–2374.
- [18] A. Garcia-Saavedra, G. Iosifidis, X. Costa-Perez, and D. J. Leith, "Joint optimization of edge computing architectures and radio access networks," *IEEE J. Sel. Areas Commun.*, vol. 36, no. 11, pp. 2433–2443, Nov. 2018.
- [19] F. W. Murti, A. Garcia-Saavedra, X. Costa-Perez, and G. Iosifidis, "On the optimization of multi-cloud virtualized radio access networks," in *Proc. IEEE Int. Conf. Commun. (ICC)*, Jun. 2020, pp. 1–7.
- [20] F. Z. Morais *et al.*, "PlaceRAN: Optimal placement of virtualized network functions in beyond 5G radio access networks," *IEEE Trans. Mobile Comput.*, early access, Apr. 29, 2022, doi: [10.1109/TMC.2022.3171525](https://doi.org/10.1109/TMC.2022.3171525).
- [21] S. Ali *et al.*, "6G white paper on machine learning in wireless communication networks," 2020, *arXiv:2004.13875*.
- [22] S. Ali, W. Saad, and D. Steinbach, Eds., *White Paper on Machine Learning in Wireless Communication Networks* (6G Research Visions), no. 7. Oulu, Finland: Univ. Oulu, Jun. 2020.
- [23] S. Matoussi, I. Fajjari, N. Aitsaadi, and R. Langar, "Deep learning based user slice allocation in 5G radio access networks," in *Proc. IEEE 45th Conf. Local Comput. Netw. (LCN)*, Nov. 2020, pp. 286–296.
- [24] T. Pamuklu, M. Erol-Kantarci, and C. Ersoy, "Reinforcement learning based dynamic function splitting in disaggregated green open RANs," in *Proc. IEEE Int. Conf. Commun.*, Jun. 2021, pp. 1–6.
- [25] I. Bello, H. Pham, Q. V. Le, M. Norouzi, and S. Bengio, "Neural combinatorial optimization with reinforcement learning," in *Proc. Int. Conf. Learn. Represent. (ICLR)*, Nov. 2019, pp. 1–15.
- [26] I. Sutskever, O. Vinyals, and Q. V. Le, "Sequence to sequence learning with neural networks," in *Proc. 27th Int. Conf. Neural Inf. Process. Syst. (NIPS)*, vol. 2. Cambridge, MA, USA: MIT Press, 2014, pp. 3104–3112.
- [27] D. Bahdanau, K. Cho, and Y. Bengio, "Neural machine translation by jointly learning to align and translate," in *Proc. Int. Conf. Learn. Represent. (ICLR)*, 2015, pp. 1–15.
- [28] R. Solozabal, J. Ceberio, and M. Takáč, "Constrained combinatorial optimization with reinforcement learning," 2020, *arXiv:2006.11984*.
- [29] R. Solozabal, J. Ceberio, A. Sanchoyerto, L. Zabala, B. Blanco, and F. Liberal, "Virtual network function placement optimization with deep reinforcement learning," *IEEE J. Sel. Areas Commun.*, vol. 38, no. 2, pp. 292–303, Dec. 2020.
- [30] C. Tessler, D. J. Mankowitz, and S. Mannor, "Reward constrained policy optimization," in *Proc. Int. Conf. Learn. Represent. (ICLR)*, 2019, pp. 1–15.
- [31] Y. Chow, M. Ghavamzadeh, L. Janson, and M. Pavone, "Risk-constrained reinforcement learning with percentile risk criteria," *J. Mach. Learn. Res.*, vol. 18, no. 1, pp. 6070–6120, 2017.
- [32] B. M. Waxman, "Routing of multipoint connections," *IEEE J. Sel. Areas Commun.*, vol. SAC-6, no. 9, pp. 1617–1622, Dec. 1988.
- [33] S. Orlowski, R. Wessälly, M. Pióro, and A. Tomaszewski, "SNDlib 1.0 survivable network design library," *Special Issue, Netw. Optim.* vol. 55, no. 3, pp. 276–286, May 2010.
- [34] N. Nikaein, "Processing radio access network functions in the cloud: Critical issues and modeling," in *Proc. ACM Mobisys*, 2015, pp. 36–43.
- [35] A. M. Alba, J. H. G. Velásquez, and W. Kellerer, "An adaptive functional split in 5G networks," in *Proc. IEEE Conf. Comput. Commun. Workshops (INFOCOM WKSHPS)*, Apr. 2019, pp. 410–416.
- [36] M. M. Akbar, M. S. Rahman, M. Kaykobad, E. G. Manning, and G. C. Shoja, "Solving the multidimensional multiple-choice knapsack problem by constructing convex hulls," *Comput. Oper. Res.*, vol. 33, no. 5, pp. 1259–1273, May 2006.
- [37] A. Graves, A.-R. Mohamed, and G. Hinton, "Speech recognition with deep recurrent neural networks," in *Proc. IEEE Int. Conf. Acoust., Speech Signal Process.*, May 2013, pp. 6645–6649.
- [38] D. Bertsekas, *Nonlinear Programming*. Belmont, MA, USA: Athena Scientific, 1999.
- [39] R. Solozabal, J. Ceberio, and M. Takáč, "Constrained combinatorial optimization with reinforcement learning," 2020, *arXiv:2006.11984*.
- [40] S. Orlowski, R. Wessälly, M. Pióro, and A. Tomaszewski, "SNDlib 1.0—Survivable network design library," *Networks*, vol. 55, no. 3, pp. 276–286, 2010.
- [41] P. Rost, S. Talarico, and M. C. Valenti, "The complexity-rate tradeoff of centralized radio access networks," *IEEE Trans. Wireless Commun.*, vol. 14, no. 11, pp. 6164–6176, Jun. 2015.



learning and optimization techniques for intelligent wireless networks.



Fahri Wisnu Murti received the B.S. degree from Telkom University, Indonesia, and the master's degree from the WENS Laboratory, Department of IT Convergence Engineering, Kumoh National Institute of Technology, South Korea. He is currently pursuing the Ph.D. degree with the Centre for Wireless Communications (CWC), University of Oulu, Finland. Prior to his doctoral study, he has worked as a Research Assistant at the Department Computer Science, Trinity College Dublin, Ireland. His current research interests lie in the development of machine



Samad Ali received the B.S. degree in electrical engineering from the University of Tabriz, Tabriz, Iran, and the M.S. and Ph.D. degrees in wireless communications engineering from the University of Oulu, Oulu, Finland. He is currently a Post-Doctoral Researcher with the University of Oulu and a Senior Research Specialist at Nokia Standards. His research interests include machine learning in wireless communications, machine type communications, and RIS.

Matti Latva-aho (Senior Member, IEEE) received the M.Sc., Lic.Tech., and Dr.Tech. (Hons.) degrees in electrical engineering from the University of Oulu, Finland, in 1992, 1996, and 1998, respectively. From 1992 to 1993, he was a Research Engineer at Nokia Mobile Phones, Oulu, Finland, after which he joined the Centre for Wireless Communications (CWC), University of Oulu. He was the Director of CWC during the years 1998–2006 and the Head of the Department of Communication Engineering until August 2014. He currently serves as the Academy of Finland Professor for the period 2017–2022 and the Director for the National 6G Flagship Program for the period 2018–2026. He has published close to 500 conferences or journal articles in the field of wireless communications. His research interests are related to mobile broadband communication systems and his group focuses on beyond 5G systems research. He has received the Nokia Foundation Award in 2015 for his achievements in mobile communications research.