

# FlagVNE: A Flexible and Generalizable Reinforcement Learning Framework for Network Resource Allocation

Tianfu Wang<sup>1</sup>, Qilin Fan<sup>2</sup>, Chao Wang<sup>3,1\*</sup>,

Long Yang<sup>4</sup>, Leilei Ding<sup>1</sup>, Nicholas Jing Yuan<sup>2</sup> and Hui Xiong<sup>6,7\*</sup>

<sup>1</sup>School of Computer Science and Technology, University of Science and Technology of China, China

<sup>2</sup>School of Big Data and Software Engineering, Chongqing University, China <sup>3</sup>Guangzhou HKUST Fok Ying Tung Research Institute, China <sup>4</sup>School of Artificial Intelligence, Peking University, China

<sup>5</sup>Microsoft <sup>6</sup>Thrust of Artificial Intelligence, The Hong Kong University of Science and Technology (Guangzhou), China <sup>7</sup>Department of Computer Science and Engineering, The Hong Kong University of Science and Technology, Hong Kong SAR, China

{tianfuwang,dingleilei}@mail.ustc.edu.cn, fanqilin@cqu.edu.cn, chadwang2012@gmail.com  
yanglong001@pku.edu.cn, nicholas.yuan@microsoft.com, xionghui@ust.hk

## Abstract

Virtual network embedding (VNE) is an essential resource allocation task in network virtualization, aiming to map virtual network requests (VNRs) onto physical infrastructure. Reinforcement learning (RL) has recently emerged as a promising solution to this problem. However, existing RL-based VNE methods are limited by the unidirectional action design and one-size-fits-all training strategy, resulting in restricted searchability and generalizability. In this paper, we propose a **FL**exible **AN**d **GE**neralizable RL framework for **VNE**, named **FlagVNE**. Specifically, we design a bidirectional action-based Markov decision process model that enables the joint selection of virtual and physical nodes, thus improving the exploration flexibility of solution space. To tackle the expansive and dynamic action space, we design a hierarchical decoder to generate adaptive action probability distributions and ensure high training efficiency. Furthermore, to overcome the generalization issue for varying VNR sizes, we propose a meta-RL-based training method with a curriculum scheduling strategy, facilitating specialized policy training for each VNR size. Finally, extensive experimental results show the effectiveness of FlagVNE across multiple key metrics. Our code is available at GitHub<sup>1</sup>.

## 1 Introduction

Network virtualization (NV) emerges as a pioneering technology that facilitates dynamic management of Internet architecture, which finds applications in 5G networks and cloud computing [Zhuang *et al.*, 2020]. Through network slicing and shared infrastructure, NV enables the deployment of multiple user-submitted virtual network requests (VNRs) within

the same physical network, thereby accommodating diverse network service requirements of users [Yang *et al.*, 2021; Chen *et al.*, 2022b]. The primary challenge in NV involves the embedding of VNRs within a physical network, known as virtual network embedding (VNE), an NP-hard combinatorial optimization problem [Rost and Schmid, 2020].

Effective resource allocation for VNRs is essential to improve the quality of service and the revenue of Internet service providers (ISPs) [Wang *et al.*, 2021b; Chen *et al.*, 2022a]. Regrettably, it is hard to address the VNE problem involving tackling combinatorial explosion and differentiated demands [Fischer *et al.*, 2013; Yang *et al.*, 2022b]. On the one hand, the solution space of VNE is extensive, encompassing vast permutations of VNRs within the underlying physical network. Consequently, a comprehensive exploration of this expansive solution space becomes imperative to determine superior solutions. On the other hand, due to specific requirements of user service, the integration of diverse VNR topologies and their associated resource demands is dynamic. VNRs of varying sizes manifest unique complexities, rendering a one-size-fits-all strategy inadequate to effectively manage the inherent variability in such circumstances.

Recently, reinforcement learning (RL) has shown promising potential for the VNE problem [Yan *et al.*, 2020; He *et al.*, 2023b; Zhang *et al.*, 2023]. RL approaches model the solution construction process of each VNR as Markov decision processes (MDPs), which can automatically build efficient solving policies. Unlike supervised learning relying on labeled data, RL facilitates the learning of effective heuristics through interactions with the environment. However, most existing RL-based VNE approaches are still plagued with some significant issues. Firstly, these approaches commonly adhere to a unidirectional action design within the MDP, i.e., presupposing a fixed decision sequence for virtual nodes, and subsequently designating a physical node to host each virtual node sequentially. Such unidirectional action schema significantly limits the available action space, consequently constraining the searchability of the agent and impeding the effi-

\*Chao Wang and Hui Xiong are both corresponding authors

<sup>1</sup><https://github.com/GeminiLight/flag-vne>

cacy of exploring solution space. Secondly, conventional RL-based methods usually just train a single general policy, disregarding the distinctive complexities of VNRs with varying sizes in practice. Treating variable-sized VNRs equally poses challenges in achieving balanced learning of cross-size strategic knowledge and hinders the ability to generalize across VNRs of differing sizes. Thirdly, the direct training of multiple policies tailored to different VNR sizes slowly adapts to unseen distributions. In particular, training specific policies for large-sized VNRs from scratch tends to be stuck in the local optimum, due to the high complexity and challenges in exploring feasible solutions. These difficulties inevitably exert negative impacts on overall system performance. We conduct a preliminary study to highlight our motivations and latent challenges, which is detailed in Appendix A.

In this paper, we propose a novel **FLexible And Generalizable RL** framework for the VNE problem, named **FlagVNE**. Our framework aims to enhance the searchability and generalizability of RL-based VNE methods while achieving rapid adaption to the unseen distribution of VNR sizes. Specifically, our contributions are summarized as follows. (1) We propose a bidirectional action-based MDP modeling approach to enable the joint selection of virtual and physical nodes, enhancing the flexibility of agent exploration and exploitation. This method offers superior searchability and is proven theoretically. To handle the resulting large and changeable action space, we abstract it as two dependent aspects and design a hierarchical decoder with a bilevel policy, ensuring adaptive action probability distribution generation and high training efficiency. (2) We propose a meta-RL-based training method to enable efficient acquisition of multiple size-specific policies and quick adaptation to new sizes. A meta-policy is trained to grasp cross-size knowledge for different VNR sizes and then fastly fine-tuned to develop size-specific policies for each VNR size, even unseen sizes. Specially, due to difficult exploration and prone to suboptimal convergence, using large-sized VNRs for initial meta-learning yields inferior knowledge, impairing the meta-policy and generalization. Thus, we develop a curriculum scheduling strategy that gradually incorporates larger VNRs, alleviating suboptimal convergence. (3) Finally, we conduct experiments on the simulation platform to mimic various network systems and extensive results demonstrate the superiority of FlagVNE in terms of multiple key indicators, compared to state-of-the-art (SOTA) heuristics and RL-based methods.

## 2 Related Work

In this section, we review the related work on VNE.

**Traditional Methods for VNE.** Initially, the VNE problem was tackled using exact methods such as integer linear programming [Shahriar *et al.*, 2018], which provides optimal solutions through exact solvers. However, these exact algorithms proved impractical for real-world scenarios due to their time-consuming nature. Thus, numerous heuristic algorithms have been proposed to find solutions in an acceptable time [Su *et al.*, 2014; Jin *et al.*, 2020; Fan *et al.*, 2023]. Among these approaches, node ranking is a prevalent strategy, which ranks virtual and physical nodes

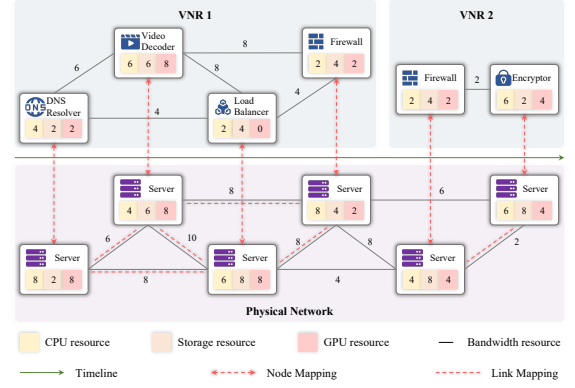


Figure 1: An example of the VNE problem with multidimensional resources. The numbers denote the unit counts of resources.

to determine the decision sequence and the matching priority, respectively. For example, [Zhang *et al.*, 2018] ranked nodes based on a node resource management (NRM) metric, and [Fan *et al.*, 2023] proposed a node essentiality assessment (NEA) metric considering topology connectivity. Additionally, [Dehury and Sahoo, 2019] designed VNE algorithms based on metaheuristics, such as particle swarm optimization (PSO). However, these algorithms heavily rely on manual designs and are usually tailored to specific scenarios, limiting their performance in general cases.

**Learning-based Methods for VNE.** Recently, machine learning techniques have been used to solve VNE, leading to faster and more efficient solutions [Blenk *et al.*, 2018; Geng *et al.*, 2023; He *et al.*, 2023b]. Particularly, RL has demonstrated significant potential as an intelligent decision-making framework [Liu *et al.*, 2023; Yang *et al.*, 2022a], which can effectively solve VNE with MDP modeling. In this paper, we unify most existing RL-based methods [Xiao *et al.*, 2019; Wang *et al.*, 2021c; Yan *et al.*, 2020; Yao *et al.*, 2020; Zhang *et al.*, 2022; Zhang *et al.*, 2023] into a general framework comprised of three key components: *MDP modeling*, *policy architecture*, and *training methods*. These methods model the process of VNE solution construction as unidirectional action-based MDPs, where a physical node is chosen to host a be-placing virtual node, and the decision sequence of virtual nodes is fixed. Then they build policy models with various neural networks and train a single general policy to deal with VNRs of varying sizes. For instance, [Xiao *et al.*, 2019] used multilayer perception (MLP) as a policy model and trained it with policy gradient (PG) algorithm, [Zhang *et al.*, 2023] designed a policy model with MLP and graph convolutional network (GCN) [Kipf and Welling, 2017] and trained it with asynchronous advantage actor-critic (A3C) [Mnih *et al.*, 2016]. However, existing RL-based VNE methods suffer from limited searchability and generalizability due to their unidirectional action design and one-size-fits-all training policy, ultimately affecting overall performance.

## 3 Preliminaries

### 3.1 Problem Definition

As shown in Fig. 1, in a practical network system, users' service requests arriving continuously are represented as VNRs. We collect all VNRs with a set  $\mathcal{V}$ . Mapping these VNRs onto

physical networks managed by ISPs is known as VNE, crucial in managing the quality of various network services [Chen *et al.*, 2020; Wang *et al.*, 2023a; Chen *et al.*, 2022c].

**System Modeling.** *Physical network* is formulated as a weighted undirected graph  $\mathcal{G}^p = (\mathcal{N}^p, \mathcal{L}^p)$ , where  $\mathcal{N}^p$  is the set of physical nodes,  $\mathcal{L}^p$  is the set of physical links. Each physical node  $n^p \in \mathcal{N}^p$  is equipped with multiple resource capacities  $\{C(n^p), \forall C \in \mathcal{C}\}$ , where  $\mathcal{C}$  is the set of node resource types, and each physical link  $l^p \in \mathcal{L}^p$  has bandwidth capacity  $B(l^p)$ . In this paper, we consider multidimensional node resources, including the central processing unit (CPU), storage resource, and graphics processing unit (GPU). Similarly, each VNR is modeled as a weighted undirected graph  $G^v = (\mathcal{N}^v, \mathcal{L}^v, d^v)$ , where  $\mathcal{N}^v$  is the set of virtual nodes and  $\mathcal{L}^v$  is the set of virtual links, and  $d^v$  denotes the lifetime of VNR. Once the VNR is accepted, it will be maintained for  $d^v$  time slots. Each virtual node  $n^v \in \mathcal{N}^v$  represents a virtual machine with resource demands  $\{C(n^v), \forall C \in \mathcal{C}\}$  and each virtual link  $l^v \in \mathcal{L}^v$  indicates the bandwidth demand  $B(l^v)$ .

**Objective.** Acknowledging the stochastic nature of online networking, most existing methods and this work aim to minimize the embedding cost of each VNR onto the physical network, which facilitates long-term performance. The quality of solutions is assessed using the revenue-to-cost ratio (R2C):

$$\text{R2C}(\mathcal{G}^v) = (\Psi \cdot \text{REV}(\mathcal{G}^v)) / \text{COST}(\mathcal{G}^v). \quad (1)$$

Here,  $\text{REV}(\mathcal{G}^v)$  denotes the revenue of the VNR  $G^v$  (i.e., the sum of VNR's resource requirements) and  $\text{COST}(\mathcal{G}^v)$  denotes the embedding cost resulting from the solution (i.e., the sum of ISP's resource consumption).  $\Psi$  is the binary variable that indicates the feasibility of a solution.

**Constraints.** The VNR embedding consists of two sub-processes. (1) *Node mapping* entails assigning each virtual node to a physical node with adequate resources, i.e.,  $C(n^p) \geq C(n^v), \forall C \in \mathcal{C}$ , while ensuring one-to-one placement and mutual exclusivity. (2) *Link mapping* involves finding a physical path for each virtual link, ensuring that the path connects the physical nodes hosting the virtual link endpoints and that each physical link  $l^p$  in the path has sufficient bandwidth, i.e.,  $B(l^p) \geq B(l^v)$ . A solution is deemed feasible ( $\Psi = 1$ ) only when all these constraints are satisfied.

Due to the space limit, we place detailed formulations of VNE's objective and constraints in Appendix B.

### 3.2 Motivations and Challenges

We conduct a preliminary study placed in Appendix A, and motivate our framework from the following two aspects.

**Flexibility of Action Space.** Most existing RL-based VNE approaches employ a unidirectional action design, assuming that the decision sequence of virtual nodes is predetermined. However, our analysis in Appendix A.1 reveals that varying the decision sequences of virtual nodes significantly impacts performance. This underscores the necessity of exploring different decision sequences for optimal solutions. Moreover, the fixed decision sequence of virtual nodes lacks the flexibility needed to adapt to the dynamic nature of exploration process. Thus, to enhance the flexibility of exploration and exploitation, we aim to achieve a joint selection of both physical and virtual nodes to eliminate the fixed decision sequence.

Nevertheless, it will pose some challenges, such as the difficulty of variable action probability distribution generation and the training efficiency issue caused by large action space.

**Generalization of Solving Policy.** VNRs of different sizes exhibit distinct complexities, necessitating varied solving strategies. Existing RL-based methods typically use a one-size-fits-all policy to tackle VNRs of varying sizes, leading to generalization issues. To address this, an intuitive approach might be to develop size-specific policies for different VNR sizes. Yet, as observed in Appendix A.2, specific policies for large-sized VNRs trained from scratch often get stuck in local optima due to their high complexity and the difficulty in exploring viable solutions. Their performance is even inferior to that of the general policy for all sizes. Furthermore, this strategy lacks the quick adaptability to handle previously unseen VNR sizes, since it requires extensive data demand.

## 4 FlagVNE Framework

In this section, we present the proposed RL-based framework for VNE, **FlagVNE**. As illustrated in Fig. 2, FlagVNE is designed to improve searchability and generalizability while achieving rapid adaptation to unseen distribution.

### 4.1 Bidirectional Action-based MDP

We formulate the solution construction process of each VNR as a bidirectional action-based MDP, allowing joint selection of virtual nodes and physical nodes. Specifically, at each decision timestep  $t$ , observing the state  $s_t$  of the environment, the agent takes an action  $a_t \sim \pi(\cdot | s_t)$  according to the policy  $\pi$ . Then, the environment will feedback a reward  $R(s_t, a_t)$  and transit to a new state  $s_{t+1} \sim P(s_t, a_t)$  following the transition probability function. During interactions, a trajectory memory  $\mathcal{D} = \{s_1, a_1, s_2, a_2, \dots\}$  collects state-action pairs. We present these notations in VNE as follows.

*State* represents the status of the network system at a specific decision timestep  $t$ , consisting of the current situation of VNR  $s_t^v$  and physical network  $s_t^p$ , i.e.,  $s_t = (s_t^v, s_t^p), s_t \in \mathcal{S}$ , where  $\mathcal{S}$  is the state space.

*Action* is defined as a pair of a virtual node to be placed and a physical node to host, denoted  $a_t = (n^v, n^p), a_t \in \mathcal{A}$ , where  $n^v \in \mathcal{N}^v, n^p \in \mathcal{N}^p$ , and  $\mathcal{A}$  is the action space.

*Transition*  $P(s_{t+1} | s_t, a_t)$  refers to the process of placing the virtual node  $n^v$  to the physical node  $n^p$  and routing the virtual links, resulting in the changes of state from  $s_t$  to  $s_{t+1}$ . Based on the selected bidirectional action, the environment attempts to place the virtual node  $n^v$  to the physical node  $n^p$ . If the node placement is successful, the link routing is executed based on the breadth-first search algorithm that finds the shortest physical paths meeting bandwidth demands from  $n^p$  to other physical nodes hosting the virtual node neighbors of  $n^v$ . If node placement and link routing are successful, the available resource of the physical network is updated with the VNR requirement. Otherwise, the current VNR is rejected.

*Reward*  $R$  measures the quality of agent's action at a given state. We define the reward function  $R$  as follows:

$$R(s_t, a_t) = \begin{cases} \text{R2C}(\mathcal{G}^v), & \text{if } \mathcal{G}^v \text{ is accepted at } t, \\ -1/|\mathcal{N}^v|, & \text{if } \mathcal{G}^v \text{ is rejected at } t, \\ 1/|\mathcal{N}^v|, & \text{otherwise.} \end{cases} \quad (2)$$

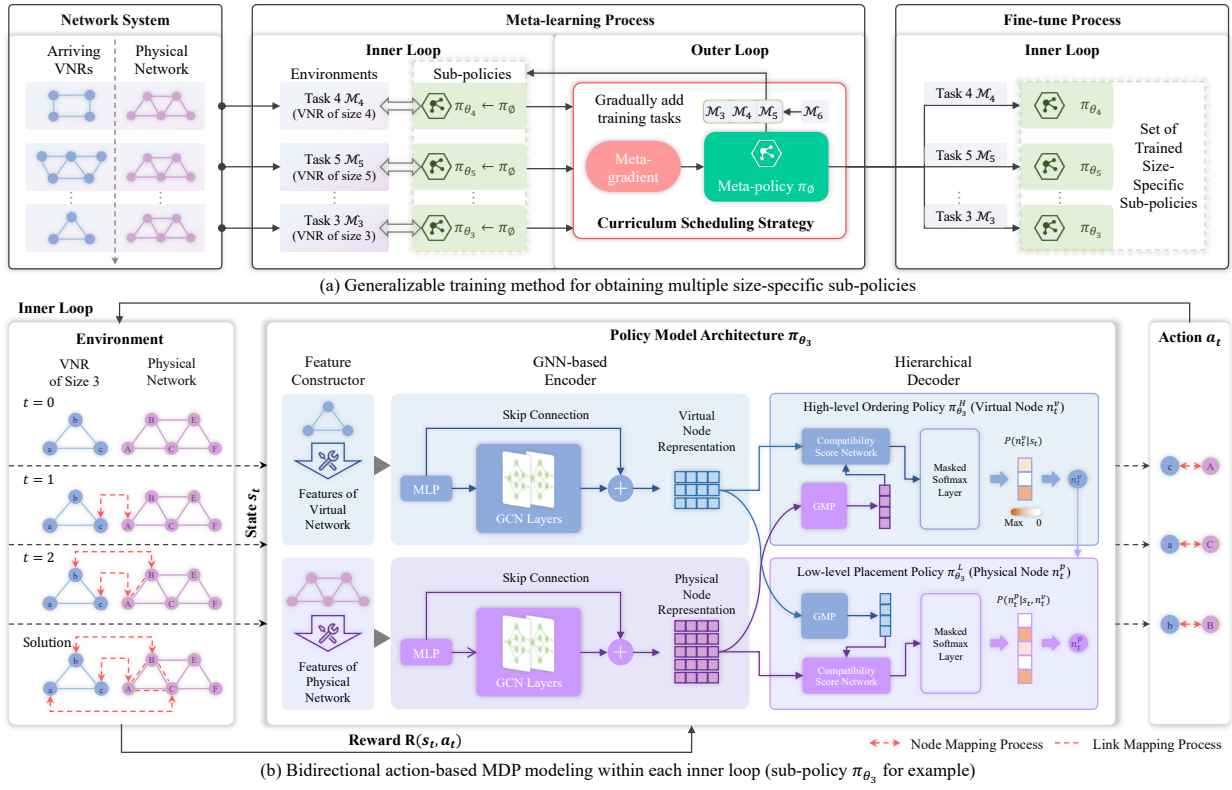


Figure 2: The overview of the FlagVNE framework. (a) For vary-sized VNRs that continuously arrive at the network system, we consider them as different tasks  $\mathcal{M}_i \sim p(\mathcal{M})$  based on their size. We first train a meta-policy  $\pi_\phi$  with cross-task knowledge in the meta-learning process, using a curriculum scheduling strategy. Then, we fine-tune it to obtain a set of size-specific sub-policies  $\pi_{\theta_i}$ . This generalizable training method effectively obtains refined solving policies for each VNR size. (b) Within each inner loop, we formulate the solution construction process of each VNR as a bidirectional action-based MDP, which enables the joint selection of virtual and physical nodes. We also design a hierarchical encoder with a bilevel policy to adaptively generate action probability distributions and ensure high training efficiency.

We design implicit rewards to encourage successful placement with  $1/|\mathcal{N}^v|$  and punish failure with  $-1/|\mathcal{N}^v|$ . Once the  $\mathcal{G}^v$  is completed embedding, we return R2C( $\mathcal{G}^v$ ).

*Policy* is parameterized by  $\theta$ , which denotes the distributions over the action space under a given state  $s_t$ :

$$\pi_\theta(a_t|s_t) = P(a_t|s_t). \quad (3)$$

*Discount factor*  $\lambda \in (0, 1)$  balances the importance of immediate rewards versus future rewards. Overall, the optimization objective of RL is to maximize the expected return, i.e., cumulative discounted rewards over timesteps  $T$ :

$$J_\pi = \mathbb{E}_{(s_t, a_t) \sim \mathcal{D}} \left[ \sum_{t=0}^T \lambda^t R(s_t, a_t) \right]. \quad (4)$$

If  $J_\pi \geq J_{\pi'}$ , then we denote it as  $\pi \succeq \pi'$ .

**Theorem 1.** Given two MDPs with bidirectional and unidirectional action,  $\mathcal{M}^b = \langle \mathcal{S}^b, \mathcal{A}^b, P^b, R, \lambda \rangle$  and  $\mathcal{M}^u = \langle \mathcal{S}^u, \mathcal{A}^u, P^u, R, \lambda \rangle$ , and their optimal policies denoted as  $\pi^{*,b}$  and  $\pi^{*,u}$ , respectively, we have  $\pi^{*,b} \succeq \pi^{*,u}$ .

See Appendix C for its proof. Our bidirectional action enhances flexibility and significantly expands the search space, allowing for a more comprehensive exploration of possible solutions, which offers superior MDP Optimality.

## 4.2 Hierarchical Policy Architecture

We construct raw features, encode them with a GCN-based encoder, and design a hierarchical decision module to ensure adaptive probability output and training efficiency.

**Feature Constructor.** We build the feature input for the subsequent encoder from the current state  $s_t$ , which includes the processing status of VNR  $s_t^v$  and the current physical network situation  $s_t^p$ . With comprehensive information on the current state, the agent gains deeper environmental insight, resulting in better decisions. For the VNR  $G_t^v$ , the feature constructor takes into account not only various node resource requirements denoted  $X_{t,N}^v$ , but also aggregates bandwidth resource requirements into the node features, represented as  $X_{t,L}^v$ . These features include essential bandwidth metrics, such as maximum, mean, and sum of bandwidth requirements of virtual links adjacent to one node. To indicate the embedding status, a placement flag  $X_{t,P}^v$  is designed for virtual nodes, with a value of 1 indicating that the virtual node has been placed and 0 otherwise. The VNR features  $X_t^v$  are organized as follows:  $X_t^v = (X_{t,N}^v, X_{t,L}^v, X_{t,P}^v) \in \mathbb{R}^{|\mathcal{N}^v| \times 7}$ .

Similarly, the physical network features  $X_t^p$  are constructed as follows:  $X_t^p = (X_{t,N}^p, X_{t,L}^p, X_{t,S}^p) \in \mathbb{R}^{|\mathcal{N}^p| \times 7}$ , where  $X_{t,N}^p$  denotes available resources of physical nodes,  $X_{t,L}^p$  similar to  $X_{t,L}^v$  denotes aggregated bandwidth availabil-

ity, and  $X_{t,S}^p$  is a selection flag indicating the status of physical nodes. A selection flag value of 1 indicates that a physical node has been selected to host a virtual node and 0 otherwise.

**GNN-based Encoder.** To encode the features of the virtual network  $X_t^v$  and the physical network  $X_t^p$ , into latent representations,  $Z_t^v$  and  $Z_t^p$ , respectively, we adopt a graph neural network (GNN) encoder. First, both  $X_t^v$  and  $X_t^p$  undergo the MLP to obtain the initial node representations, denoted  $I_t^v$  and  $I_t^p$ , respectively:  $I_t^v = \text{MLP}(X_t^v)$ ,  $I_t^p = \text{MLP}(X_t^p)$ .

Then, we consider multiple GCN [Kipf and Welling, 2017] layers as the GNN modules to obtain the latent representations of virtual nodes  $\tilde{Z}_t^v$  and physical nodes  $\tilde{Z}_t^p$ :  $\tilde{Z}_t^v = \text{GNN}(I_t^v, A^v)$ ,  $\tilde{Z}_t^p = \text{GNN}(I_t^p, A^p)$ , where  $A^v$  and  $A^p$  is adjacency matrixes of virtual and physical networks.

Furthermore, to enhance the feature representation ability, we also employ the residual connection method to combine the output of the GNN module with the initial representation:  $Z_t^v = \tilde{Z}_t^v + I_t^v$ ,  $Z_t^p = \tilde{Z}_t^p + I_t^p$ . Finally, we obtain the representation of each virtual and physical node.

**Hierarchical Decoder with Bilevel Policy.** In our bidirectional action-based MDP, the action space is represented by the matrix size  $|\mathcal{N}^v| \times |\mathcal{N}^p|$ , reflecting the number of virtual and physical nodes. The variable and often large size of VNRs contribute to the expansive and dynamic nature of the space. To effectively manage this, we develop a hierarchical decoder with a bilevel policy, ensuring high training efficiency and adaptive action probability generation. Specifically, we abstract this task into two dependent aspects: virtual node ordering and physical node placement. Our bilevel policy,  $\pi(a_t|s_t) = \pi^H(n_t^v|s_t) \cdot \pi^L(n_t^p|s_t, n_t^v)$ , consist of a high-level ordering policy  $\pi^H(n_t^v|s_t)$  and a low-level placement policy  $\pi^L(n_t^p|s_t, n_t^v)$ . This hierarchical approach reduces the size of policy distribution from  $|\mathcal{N}^v| \times |\mathcal{N}^p|$  to  $|\mathcal{N}^v| + |\mathcal{N}^p|$ , thus significantly enhancing training efficiency.

*High-level ordering policy* selects the appropriate virtual node  $n_t^v$  for placement. Concretely, we use an MLP-based compatibility scoring network to calculate the fitness between each virtual node representation and the graph-level representation of the physical network  $G_t^p = \text{GMP}(Z_t^p)$ . Here,  $\text{GMP}(Z) = \frac{1}{|Z|} \sum_{z \in Z} z$  denotes graph mean pooling (GMP), averaging all node representations. Then an MLP is applied to generate compatibility scores for each virtual node:

$$\tilde{Y}^H = \text{MLP}(Z_t^v + G_t^p) \in \mathbb{R}^{1 \times |\mathcal{N}^v|}. \quad (5)$$

Although the VNR's sizes are variable, this layer adaptively generates scores with the shape of  $(1, |\mathcal{N}^v|)$ . After masking virtual nodes already placed (i.e., setting their scores to  $-\infty$  on  $\tilde{Y}^H$ ), we apply a softmax function to the resultant score  $Y^H$  to produce the high-level action probability distribution.

$$\pi^H(n_t^v|s_t) = \text{softmax}(Y^H). \quad (6)$$

*Low-level placement policy* identifies a suitable physical node  $n_t^p$  for accommodating the to-be-placed virtual node  $n_t^v$ , which is selected by  $\pi^H$ . Similarly, we adopt an MLP-based compatibility scoring network to calculate the fitness between the representation of each physical node and the current context representation of virtual network, including the graph-level representation of virtual network  $G_t^v = \text{GMP}(Z_t^v)$  and

to-be-placed virtual node's representation  $z_{n_t^v}$ :

$$\tilde{Y}^L = \text{MLP}(Z_t^p + G_t^v + z_{n_t^v}) \in \mathbb{R}^{1 \times |\mathcal{N}^p|}. \quad (7)$$

To avoid unnecessary exploration, we mask the physical nodes that do not have enough resources or have been selected to obtain the final scores  $Y^L$ . Then, the low-level action probability distribution is generated:

$$\pi^L(n_t^p|s_t, n_t^v) = \text{softmax}(Y^L). \quad (8)$$

For both two-level probability distributions, we employ the sampling and greedy strategy to select actions during the training and inference phases, respectively.

### 4.3 Generalizable Training Method

Training a general policy for VNRs of varying sizes leads to imbalanced learning of cross-size strategy and generalization issues. Conversely, individualized training of multiple policies for each size is slow to adapt to new sizes, in which policies for large-sized VNRs are prone to suboptimal. To address this, we develop a meta-RL-based training method with a curriculum scheduling strategy. As illustrated in Algorithm 1 (see Appendix D), our method enables efficient training of multiple size-specific policies and quick adaptation to new sizes, while balancing the learning process across tasks of varying difficulty and avoiding suboptimal convergence.

**Meta-RL for VNE.** We treat VNRs of different sizes as distinct tasks and formulate them as multiple MDPs following a distribution  $\mathcal{M}_i \sim p(\mathcal{M})$ . Note that this distribution of VNR size is bounded and always obviously smaller than the number of physical nodes, following the network service orchestration standards [Zhuang *et al.*, 2020]. We adopt model-agnostic meta-learning (MAML) as the basic training method [Finn *et al.*, 2017]. MAML facilitates the learning of a meta-policy that can be swiftly fine-tuned on new tasks with only a few training samples, which improves generalizability and adaptability. This training process comprises two stages as follows. Firstly, during the meta-learning process, we iteratively execute the inner loops and outer loops to derive a well-trained meta-policy  $\pi_\phi$  with cross-task knowledge. Secondly, in the fine-tuning process, we leverage task-specific experiences to fine-tune the meta-policy to a set of size-specific policies  $\pi_{\theta_i}$  solely through inner loops.

Concretely, in the inner loop, the meta-policy  $\pi_\phi$  is updated to accommodate a specific task  $\mathcal{M}_i$  by performing gradient descents with the learning rate  $\alpha$  and task-specific data  $\mathcal{D}_i$ :

$$\theta_i = f(\phi, \mathcal{D}_i) = \phi - \alpha \nabla_\phi \mathcal{L}_{\mathcal{D}_i}(\phi). \quad (9)$$

Here,  $\mathcal{L}(\cdot)$  follows the objective of proximal policy optimization (PPO) algorithm [Schulman *et al.*, 2017]:

$$\mathcal{L}_{\mathcal{D}_i}(\phi) = \mathbb{E}_{(s_t, a_t) \sim \mathcal{D}_i} \left[ \min \left( r_\phi \hat{A}, \text{clip}(r_\phi, \epsilon) \hat{A} \right) \right], \quad (10)$$

where  $\hat{A}$  denotes the estimated advantage of taking an action.  $r_\phi = \frac{\pi_\phi(a_t|s_t)}{\pi_{\phi_{\text{old}}}(a_t|s_t)}$  denotes the ratio between the current policy  $\pi_\phi$  and the last updated policy  $\pi_{\phi_{\text{old}}}$ . The clip function with a hyperparameter  $\epsilon$  is used to limit  $r_\phi$  within the range of  $[1-\epsilon, 1+\epsilon]$ , improving the stability of policy updates. In PPO,

the critic uses a GNN-based encoder and GMPs, then inputs concatenated virtual and physical graph representations into an MLP-based decoder to estimate value.

In the outer loop, our objective is to find a meta-policy  $\pi_\phi$  that learns balanced strategy knowledge required by VNRs of different sizes and exhibits superior generalizability, enabling it to quickly learn optimal task-specific policies:

$$J_\phi = \mathbb{E}_{\mathcal{M}_i \sim P(\mathcal{M})} \left[ \mathbb{E} \left[ \sum_{t=0}^T \lambda^t R(s_t, a_t) | \theta_i, \mathcal{D}_i \right] \right]. \quad (11)$$

We update  $\phi$  with a meta-learning rate  $\beta$  according to average second-order meta-gradient over task-specific policies:

$$\phi \leftarrow \phi - \beta \nabla_\phi \left( \frac{1}{|\mathcal{M}|} \sum_{i=1}^{|\mathcal{M}|} \mathcal{L}(\theta_i) \right). \quad (12)$$

**Curriculum Scheduling Strategy.** In our preliminary study (see Appendix A.2), we observed that training specific policies for large VNRs often leads to suboptimal convergence. This issue stems from the complexity of large-sized VNRs and the challenges of exploring the solution space to find feasible solutions. This tendency also towards local optima adversely impacts the meta-learning process. Specifically, using large-sized VNRs in the initial stages of meta-learning results in low-quality gradients, which negatively affects the convergence and generalizability of the meta-policy.

To address this challenge, we draw inspiration from curriculum learning [Wang *et al.*, 2021d] and propose a curriculum scheduling strategy to gradually integrate larger VNRs into the meta-learning process. This strategy enables high-quality initializations for sub-policies of large-sized VNRs, alleviating the problems of suboptimal convergence and compromising meta-policy. We implement this by maintaining a training task list  $\mathcal{I}$ , initially containing the smallest VNR size. The meta-learning process begins by focusing on tasks with smaller VNR sizes, which are inherently easier and provide beneficial foundational knowledge for tackling more complex tasks. Policies adeptly trained on these smaller tasks serve as effective initializations for larger VNR tasks, facilitating to mitigating local optima issues.

To achieve a gradual increase in task complexity, we use the entropy metric  $H(\pi)$  to evaluate the stability of policy. For our bilevel policy, we approximate it with  $H(\pi) = H(\pi^H) + H(\pi^L)$ . A lower entropy suggests that the policy is making more confident decisions. When the policy entropy  $H(\pi_{\theta_k})$  for the largest size  $k = \max(\mathcal{I})$  currently on the training task list falls below a specified threshold  $\delta$ , we consider the policy ready to handle more complex tasks. At this point, we introduce the next larger VNR size to the training task list  $\mathcal{I}$ . This progressive approach allows the meta-policy to adapt and generalize effectively to larger VNRs.

## 5 Performance Evaluation

In this section, we evaluate the effectiveness of FlagVNE.

### 5.1 Experiment Setup

**Simulations.** Following the latest works [He *et al.*, 2023b; Wang *et al.*, 2023b], we conduct experiments on the simulation platform to mimic various realistic network systems.

We adopt two topologies, GEANT (40 nodes and 61 links) and WX100 (100 nodes and 500 links) [Waxman, 1988], as physical networks. See Appendix E.1 for these topologies' descriptions. The multiple-type resources (i.e., CPU, storage, GPU) of physical nodes and bandwidth resources of physical links are uniformly generated within the range of [50, 100] units. In each simulation run, we randomly generate 1000 VNRs with varying sizes ranging from 2 to 10. The virtual nodes within each VNR are randomly interconnected with a probability of 50%. Additionally, resource demands of each VNR's node and link requirements are uniformly generated within the range of [0, 20] and [0, 50] units, respectively. The lifetime of each VNR is exponentially distributed with an average of 500 time units. The arrival of these VNRs follows a Poisson process with an average rate  $\eta$ , wherein  $\eta$  VNRs are received per unit of time. In subsequent experiments, we first train models with  $\eta = 0.001$  on GEANT and  $\eta = 0.08$  on WX100, due to their different capacities of physical resources. Then we manipulate the value of  $\eta$  to emulate network systems with different traffic throughputs and infer with trained models to study the sensitivity of algorithms.

**Implementations.** During training, we first conduct meta-learning in the initial 20 simulations and then focus on fine-tuning in the subsequent 10 simulations. We set the policy entropy threshold  $\delta$  to 2. We implement neural network models with PyTorch and decide reasonable values for hyperparameters following the guide of related studies [Huang *et al.*, 2022; Zhou *et al.*, 2023; Wang *et al.*, 2021a; He *et al.*, 2023a; Kingma and Ba, 2014; Joshi *et al.*, 2022]. See Appendix E.2 for hyperparameter settings on neural networks and meta-RL.

**Baselines.** To validate the effectiveness of FlagVNE, we compare it with the following SOTA heuristics (NRM-VNE [Zhang *et al.*, 2018]; NEA-VNE [Fan *et al.*, 2023]; PSO-VNE [Jiang and Zhang, 2021]) and RL-based baselines (MCTS-VNE [Haeri and Trajković, 2017]; PG-CNN [Zhang *et al.*, 2022]; A3C-GCN [Zhang *et al.*, 2023]; DDPG-Attention [He *et al.*, 2023b]). See Appendix E.3 for their descriptions.

**Metrics.** The following metrics are widely used to evaluate the long-term operational status of network systems over a period  $\mathcal{T}$  [Fischer *et al.*, 2013]: *request acceptance rate* (RAC), *long-term average revenue* (LAR) and *long-term revenue-to-cost* (LT-R2C). See Appendix E.3 for their definitions.

### 5.2 Results and Analysis

**Overall Performance.** To simulate diverse and complex scenarios with varying traffic throughputs, we manipulate the arrival rate of VNRs in two settings due to the difference in physical resource capacity: in GEANT, we explore a range of [0.001, 0.006] with a step of 0.001, and in WX100, we investigate a range of [0.08, 0.18] stepped by 0.02.

Fig. 3(a)(b)(c) and (d)(e)(f) illustrate the performance of all algorithms in GEANT and WX100, respectively. As the arrival rate  $\eta$  increases, all algorithms experience a decline in RAC on both topologies, attributed to heightened competition for limited physical resources among VNRs. Despite the variability in algorithm performance across different network topologies, influenced by the varying abundance of physical bandwidth resources, FlagVNE consistently achieves the best performance in all scenarios. We observe that the improve-

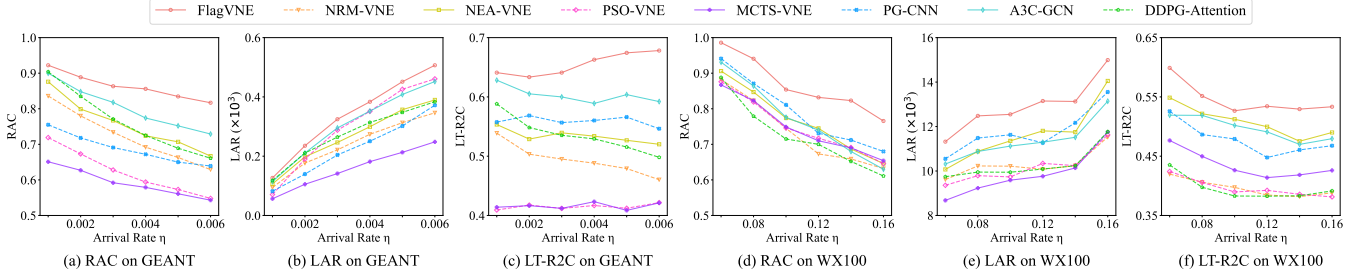


Figure 3: Experimental results in traffic throughput test.

	GEANT			WX100		
	RAC $\uparrow$	LAR $\uparrow$	LT-R2C $\uparrow$	RAC $\uparrow$	LAR $\uparrow$	LT-R2C $\uparrow$
FlagVNE-UniActionNEA	0.781	475.335	0.637	0.724	14334.671	0.493
FlagVNE-MetaFree-SinglePolicy	0.758	472.455	0.614	0.712	14170.514	0.501
FlagVNE-MetaFree-MultiPolicy	0.746	435.502	0.593	0.685	14069.938	0.472
FlagVNE-MetaPolicy	0.773	478.646	0.634	0.717	14292.962	0.485
FlagVNE-NoCurriculum	0.787	485.267	0.643	0.708	14144.234	0.509
<b>FlagVNE</b>	<b>0.804</b>	<b>499.303</b>	<b>0.668</b>	<b>0.754</b>	<b>14769.080</b>	<b>0.526</b>

Table 1: Results on ablation study. ( $\eta = 0.006$  on GEANT and  $\eta = 0.18$  on WX100).

ments of FlagVNE are more pronounced at higher values of  $\eta$ , corresponding to heightened resource competition. This underscores the importance of searchability and generalizability in network environments with limited resources. Specifically, at  $\eta = 0.006$  on GEANT, FlagVNE surpasses A3C-GCN, NEA-VNE and NRM-VNE by margins of 10.4%, 20.7% and 27.9% on RAC, 10.5%, 28.1%, and 44.2% on LAR, and 12.8%, 28.4%, and 45.1% on LT-R2C. On WX100, compared to A3C-GCN, NEA-VNE and NRM-VNE, FlagVNE shows average improvements over different  $\eta$  of 12.4%, 12.5% and 17.4% in RAC, 12.8%, 10.4% and 24.3% on LAR, and 9.1%, 6.7% and 36.7% on LT-R2C, respectively. Overall, FlagVNE demonstrates exceptional performance across various network system conditions.

**Ablation Study.** To verify the effectiveness of each proposed component, we build several variations of FlagVNE: (1) *FlagVNE-UniActionNEA* replaces the bidirectional action with the unidirectional one and sorts the decision sequence of virtual nodes with NEA [Fan *et al.*, 2023]. (2) *FlagVNE-MetaFree-SinglePolicy* trains a single general policy with valina PPO, without the help of Meta-RL. (3) *FlagVNE-MetaFree-MultiPolicy* directly trains a set of sub-policies from scratch, without using Meta-RL. (4) *FlagVNE-MetaPolicy* only uses the meta-policy to handle variable-sized VNRs. (5) *FlagVNE-NoCurriculum* discards the curriculum scheduling strategy during the meta-learning process.

We examine their performance under arrival rate settings of  $\eta = 0.006$  on GEANT and  $\eta = 0.18$  on WX100. These cases exhibit more intense competition for resources, accentuating the performance differentials stemming from the algorithms’ searchability and generalizability. As shown in Table 1, FlagVNE outperforms all variations on three metrics, demonstrating that each component of FlagVNE contributes to the improvement in the final performance. Notably, we observe significant performance declines in FlagVNE-

MetaFree-MultiPolicy and FlagVNE-MetaFree-SinglePolicy compared to FlagVNE, which shows the effectiveness of our meta-RL training method with a curriculum scheduling strategy in achieving generalization.

**Additional Evaluation.** Due to the space limit, we place more experiments and analyses in Appendix F. Concretely, in Appendix F.1, we conduct the running time test to verify the efficiency of FlagVNE, and the results show that FlagVNE strikes a better balance between performance and running time. We also provide the adaptation and convergence analysis in Appendix F.2. Results show that FlagVNE can efficiently learn a meta-policy with cross-size knowledge in known distributions and quickly adapt to unseen sizes through fine-tuning. Besides, in Appendix F.3, we evaluate FlagVNE’s scalability on large-scale network systems, and results demonstrate that FlagVNE consistently outperforms all baseline models, even in this large-scale network system scenario. Finally, in Appendix F.4, we explore the impact of the key hyperparameter  $\epsilon$  on the performance of FlagVNE.

## 6 Conclusion

In this paper, we proposed FlagVNE, a flexible and generalizable RL framework for VNE. Specifically, we developed a bidirectional action MDP modeling approach to enable the joint selection of virtual nodes and physical nodes, which expands the agent’s search space. Additionally, we designed a hierarchical recorder with a bilevel policy to ensure adaptive output and high training efficiency. Furthermore, we presented a generalizable training method based on meta-RL that efficiently trains a set of size-specific policies to tackle VNRs of varying scales. We also developed a curriculum scheduling strategy that gradually incorporates larger VNRs, thus alleviating suboptimal convergence. Finally, we conducted extensive experiments to verify the effectiveness of FlagVNE.

## Acknowledgements

This work was partially supported by National Natural Science Foundation of China (Grant No.92370204), Guangzhou-HKUST(GZ) Joint Funding Program (Grant No.2023A03J0008), Education Bureau of Guangzhou Municipality, Guangdong Science and Technology Department Project funded by China Postdoctoral Science Foundation (Grant No.2023M730785), National Natural Science Foundation of China (Grant No. 62102053).

## References

- [Blenk *et al.*, 2018] Andreas Blenk, Patrick Kalmbach, Johannes Zerwas, Michael Jarschel, Stefan Schmid, and Wolfgang Kellerer. NeuroViNE: A neural preprocessor for your virtual network embedding algorithm. In *IEEE International Conference on Computer Communications*, pages 405–413, 2018.
- [Chen *et al.*, 2020] Liyi Chen, Zhi Li, Yijun Wang, Tong Xu, Zhefeng Wang, and Enhong Chen. Mmea: Entity alignment for multi-modal knowledge graph. In *International Conference on Knowledge Science, Engineering and Management*, pages 134–147, 2020.
- [Chen *et al.*, 2022a] Jin Chen, Guanyu Ye, Yan Zhao, Shuncheng Liu, Liwei Deng, Xu Chen, Rui Zhou, and Kai Zheng. Efficient join order selection learning with graph-based representation. In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, page 97–107, 2022.
- [Chen *et al.*, 2022b] Liyi Chen, Zhi Li, Weidong He, Gong Cheng, Tong Xu, Nicholas Jing Yuan, and Enhong Chen. Entity summarization via exploiting description complementarity and salience. *IEEE Transactions on Neural Networks and Learning Systems*, 2022.
- [Chen *et al.*, 2022c] Liyi Chen, Zhi Li, Tong Xu, Han Wu, Zhefeng Wang, Nicholas Jing Yuan, and Enhong Chen. Multi-modal siamese network for entity alignment. In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, 2022.
- [Dehury and Sahoo, 2019] Chinmaya Kumar Dehury and Prasan Kumar Sahoo. DYVINE: Fitness-based dynamic virtual network embedding in cloud computing. *IEEE Journal on Selected Areas in Communications*, 37(5):1029–1045, 2019.
- [Fan *et al.*, 2023] Weibei Fan, Fu Xiao, Mengjie Lv, Lei Han, Junchang Wang, and Xin He. Node essentiality assessment and distributed collaborative virtual network embedding in datacenters. *IEEE Transactions on Parallel and Distributed Systems*, 34(4):1265–1280, 2023.
- [Finn *et al.*, 2017] Chelsea Finn, Pieter Abbeel, and Sergey Levine. Model-agnostic meta-learning for fast adaptation of deep networks. In *International Conference on Machine Learning*, pages 1126–1135, 2017.
- [Fischer *et al.*, 2013] Andreas Fischer, Juan Felipe Botero, Michael Till Beck, Hermann de Meer, and Xavier Hesselbach. Virtual network embedding: A survey. *IEEE Communications Surveys Tutorials*, 15(4):1888–1906, 2013.
- [Geng *et al.*, 2023] Haoyu Geng, Runzhong Wang, Fei Wu, and Junchi Yan. Gal-vne: Solving the vne problem with global reinforcement learning and local one-shot neural prediction. In *Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, KDD ’23, page 531–543, 2023.
- [Haeri and Trajković, 2017] Soroush Haeri and Ljiljana Trajković. Virtual network embedding via Monte Carlo tree search. *IEEE Transactions on Cybernetics*, 48(2):510–521, 2017.
- [He *et al.*, 2023a] Huiguo He, Tianfu Wang, Huan Yang, Jianlong Fu, Nicholas Jing Yuan, Jian Yin, Hongyang Chao, and Qi Zhang. Learning profitable nft image diffusions via multiple visual-policy guided reinforcement learning. In *Proceedings of the 31st ACM International Conference on Multimedia*, page 6831–6840, 2023.
- [He *et al.*, 2023b] Nan He, Song Yang, Fan Li, Stojan Trajanovski, Liehuang Zhu, Yu Wang, and Xiaoming Fu. Leveraging deep reinforcement learning with attention mechanism for virtual network function placement and routing. *IEEE Transactions on Parallel and Distributed Systems*, 34(4):1186–1201, 2023.
- [Huang *et al.*, 2022] Shengyi Huang, Rousslan Fernand Julien Dossa, Antonin Raffin, Anssi Kanervisto, and Weixun Wang. The 37 implementation details of proximal policy optimization. In *ICLR Blog Track*, 2022.
- [Jiang and Zhang, 2021] Chunxiao Jiang and Peiying Zhang. VNE-HPSO Virtual Network Embedding Algorithm Based on Hybrid Particle Swarm Optimization, pages 129–152. Springer Singapore, Singapore, 2021.
- [Jin *et al.*, 2020] Panpan Jin, Xincan Fei, Qixia Zhang, Fangming Liu, and Bo Li. Latency-aware VNF chain deployment with efficient resource reuse at network edge. In *IEEE International Conference on Computer Communications*, pages 267–276, 2020.
- [Joshi *et al.*, 2022] Chaitanya K Joshi, Quentin Cappart, Louis-Martin Rousseau, and Thomas Laurent. Learning the travelling salesperson problem requires rethinking generalization. *Constraints*, 27(1-2):70–98, 2022.
- [Kingma and Ba, 2014] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [Kipf and Welling, 2017] Thomas N. Kipf and Max Welling. Semi-supervised classification with graph convolutional networks. In *International Conference on Learning Representations*, 2017.
- [Liu *et al.*, 2023] Shuncheng Liu, Yuyang Xia, Xu Chen, Jiandong Xie, Han Su, and Kai Zheng. Impact-aware maneuver decision with enhanced perception for autonomous vehicle. In *2023 IEEE 39th International Conference on Data Engineering (ICDE)*, pages 3255–3268, 2023.
- [Mnih *et al.*, 2016] Volodymyr Mnih, Adria Puigdomenech Badia, Mehdi Mirza, Alex Graves, Timothy Lillicrap, Tim Harley, David Silver, and Koray Kavukcuoglu. Asynchronous methods for deep reinforcement learning. In

- International Conference on Machine Learning*, pages 1928–1937, 2016.
- [Rost and Schmid, 2020] Matthias Rost and Stefan Schmid. On the hardness and inapproximability of virtual network embeddings. *IEEE/ACM Transactions on Networking*, 28(2):791–803, 2020.
- [Schulman *et al.*, 2017] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.
- [Shahriar *et al.*, 2018] Nashid Shahriar, Shihabur Rahman Chowdhury, Reaz Ahmed, Aimal Khan, Siavash Fathi, Raouf Boutaba, Jeebak Mitra, and Liu Liu. Virtual network survivability through joint spare capacity allocation and embedding. *IEEE Journal on Selected Areas in Communications*, 36(3):502–518, 2018.
- [Su *et al.*, 2014] Sen Su, Zhongbao Zhang, Alex X. Liu, Xi-ang Cheng, Yiwen Wang, and Xinchao Zhao. Energy-aware virtual network embedding. *IEEE/ACM Transactions on Networking*, 22(5):1607–1620, 2014.
- [Sutton and Barto, 2018] Richard S. Sutton and Andrew G. Barto. *Reinforcement Learning: An Introduction*. A Bradford Book, Cambridge, MA, USA, 2018.
- [Wang *et al.*, 2021a] Chao Wang, Hengshu Zhu, Qiming Hao, Keli Xiao, and Hui Xiong. Variable interval time sequence modeling for career trajectory prediction: Deep collaborative perspective. In *Proceedings of the Web Conference 2021*, pages 612–623, 2021.
- [Wang *et al.*, 2021b] Chao Wang, Hengshu Zhu, Peng Wang, Chen Zhu, Xi Zhang, Enhong Chen, and Hui Xiong. Personalized and explainable employee training course recommendations: A bayesian variational approach. *ACM Transactions on Information Systems*, 40(4):1–32, 2021.
- [Wang *et al.*, 2021c] Tianfu Wang, Qilin Fan, Xiuhua Li, Xu Zhang, Qingyu Xiong, Shu Fu, and Min Gao. DRL-SFCP: Adaptive service function chains placement with deep reinforcement learning. In *IEEE International Conference on Communications*, pages 1–6, 2021.
- [Wang *et al.*, 2021d] Xin Wang, Yudong Chen, and Wenwu Zhu. A survey on curriculum learning. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(9):4555–4576, 2021.
- [Wang *et al.*, 2023a] Chao Wang, Hengshu Zhu, Chen Zhu, Chuan Qin, Enhong Chen, and Hui Xiong. Setrank: A setwise bayesian approach for collaborative ranking in recommender system. *ACM Transactions on Information Systems*, 42(2):1–32, 2023.
- [Wang *et al.*, 2023b] Tianfu Wang, Li Shen, Qilin Fan, Tong Xu, Tongliang Liu, and Hui Xiong. Joint admission control and resource allocation of virtual network embedding via hierarchical deep reinforcement learning. *IEEE Transactions on Services Computing*, 2023.
- [Waxman, 1988] B.M. Waxman. Routing of multipoint connections. *IEEE Journal on Selected Areas in Communications*, 6(9):1617–1622, 1988.
- [Xiao *et al.*, 2019] Yikai Xiao, Qixia Zhang, Fangming Liu, Jia Wang, Miao Zhao, Zhongxing Zhang, and Jiaxing Zhang. Nfvdeep: Adaptive online service function chain deployment with deep reinforcement learning. In *IEEE/ACM International Symposium on Quality of Service*, 2019.
- [Yan *et al.*, 2020] Zhongxia Yan, Jingguo Ge, Yulei Wu, Liangxiong Li, and Tong Li. Automatic virtual network embedding: A deep reinforcement learning approach with graph convolutional networks. *IEEE Journal on Selected Areas in Communications*, 38(6):1040–1057, 2020.
- [Yang *et al.*, 2021] Song Yang, Fan Li, Stojan Trajanovski, Ramin Yahyapour, and Xiaoming Fu. Recent advances of resource allocation in network function virtualization. *IEEE Transactions on Parallel and Distributed Systems*, 32(2):295–314, 2021.
- [Yang *et al.*, 2022a] Long Yang, Jiaming Ji, Juntao Dai, Linrui Zhang, Binbin Zhou, Pengfei Li, Yaodong Yang, and Gang Pan. Constrained update projection approach to safe policy optimization. *Advances in Neural Information Processing Systems*, 35:9111–9124, 2022.
- [Yang *et al.*, 2022b] Long Yang, Yu Zhang, Gang Zheng, Qian Zheng, Pengfei Li, Jianhang Huang, and Gang Pan. Policy optimization with stochastic mirror descent. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, pages 8823–8831, 2022.
- [Yao *et al.*, 2020] Haipeng Yao, Sihan Ma, Jingjing Wang, Peiying Zhang, Chunxiao Jiang, and Song Guo. A continuous-decision virtual network embedding scheme relying on reinforcement learning. *IEEE Transactions on Network and Service Management*, 17(2):864–875, 2020.
- [Zhang *et al.*, 2018] Peiying Zhang, Haipeng Yao, and Yunjie Liu. Virtual network embedding based on computing, network, and storage resource constraints. *IEEE Internet of Things Journal*, 5(5):3298–3304, 2018.
- [Zhang *et al.*, 2022] Peiying Zhang, Yu Su, Jingjing Wang, Chunxiao Jiang, Ching-Hsien Hsu, and Shigen Shen. Reinforcement learning assisted bandwidth aware virtual network resource allocation. *IEEE Transactions on Network and Service Management*, 19(4):4111–4123, 2022.
- [Zhang *et al.*, 2023] Xiaoquan Zhang, Lin Cui, Fung Po Tso, Zhetao Li, and Weijia Jia. Dapper: Deploying service function chains in the programmable data plane via deep reinforcement learning. *IEEE Transactions on Services Computing*, pages 1–14, 2023.
- [Zhou *et al.*, 2023] Jianan Zhou, Yaoxin Wu, Wen Song, Zhiguang Cao, and Jie Zhang. Towards omnigeneralizable neural methods for vehicle routing problems. In *the 40th International Conference on Machine Learning*, 2023.
- [Zhuang *et al.*, 2020] Weihua Zhuang, Qiang Ye, Feng Lyu, Nan Cheng, and Ju Ren. SDN/NFV-empowered future IoV with enhanced communication, computing, and caching. *Proceedings of the IEEE*, 108(2):274–291, 2020.

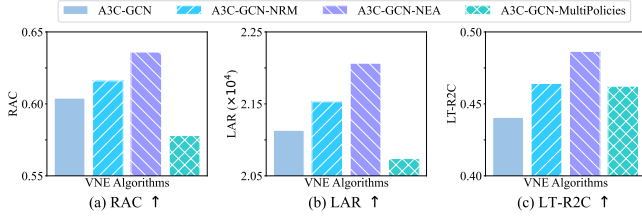


Figure 4: Comparative performance of A3C-GCN variants on three metrics: Impact of decision sequence and size-specific policies on VNE. (We conduct experiments using WX100 as the physical network, with a VNR arrival rate of 0.18. All other settings remained consistent with those described in Section 5.)

## A Preliminary Study

We conduct a preliminary analysis to reveal the limitations of existing RL-based works and highlight our motivations and latent challenges. Concretely, we investigate a widely adopted RL-based VNE algorithm, A3C-GCN, used in [Yan *et al.*, 2020; Zhang *et al.*, 2023]. We extend its capabilities by incorporating several methods. Instead of using the default decision sequence based on virtual node ID numbers, we introduce two variants: A3C-GCN-NRM and A3C-GCN-NEA. These variants utilize NRM [Zhang *et al.*, 2018] and NEA [Fan *et al.*, 2023] metrics, respectively, to rank virtual nodes and rearrange the decision sequence. Furthermore, we develop A3C-GCN-MultiPolicy that fine-tune the pretrained A3C-GCN model to obtain multiple policies tailored to different VNR sizes, rather than use a single one-size-fits-all to accommodate all sizes. Through the analysis of experimental results, our motivations are attributed to the following two key aspects.

### A.1 The Flexibility of Action Space

Most existing approaches based on RL suffer from a unidirectional action design, assuming that the decision sequence of virtual nodes is fixed. They severely limit the action search space from  $|\mathcal{N}^p| \times |\mathcal{N}^v|$  to  $|\mathcal{N}^p| \times 1$ , where  $|\mathcal{N}^p|$  and  $|\mathcal{N}^v|$  denote the number of physical nodes and virtual nodes, respectively. However, as shown in Fig. 4, A3C-GCN, A3C-GCN-NRM and A3C-GCN-NEA exhibit different performance due to their distinct decision sequence of virtual nodes, which highlights the importance of sufficiently exploring different decision sequences to achieve better solutions. Moreover, the fixed decision sequence of virtual nodes fails to perceive the dynamic nature of the environment and the potential interdependencies between decisions made at distinct time steps. Therefore, due to this action design, the exploration and exploitation of the solution space by the RL agent are inflexible, which restricts the searchability of the algorithm and reduces the probability of discovering high-quality solutions.

To enhance exploration and exploitation and consider the dynamic nature of solving process, an accessible way is to achieve a joint selection of physical and virtual nodes to eliminate the fixed decision sequence. But it will pose some challenges, such as the difficulty of variable action probability distribution generation and the training efficiency issue caused by large action space.

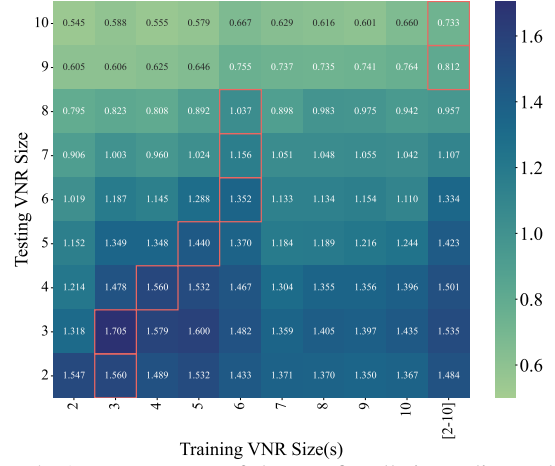


Figure 5: Average returns of the one-fits-all-size policy and each size-specific policy on all testing VNR sizes. The red boxes indicate the best performance results for test sizes. In the horizontal axis, [2-10] indicates a well-trained A3C-GCN policy while a single number represents a size-specific policy derived from well-trained A3C-GCN-MultiPolicy. (We use WX100 as the physical network and all training settings are the same as those mentioned in Section 5. For testing data of each VNR size, to exclude network system dynamics for a fairer comparison, we randomly generated 1000 static instances, including VNR and physical networks, as the benchmark. The performance metric is defined as the average episode return over 1000 instances.)

### A.2 The Generalization of Solving Policy

As revealed in studies [Joshi *et al.*, 2022; Zhou *et al.*, 2023], combinatorial optimization problems on graphs demonstrate variable complexity as graph sizes change. A one-size-fits-all strategy for handling these instances across different sizes often leads to generalization issues. VNE is no exception. Existing methods ignore this diversity and train a universal solving policy with vanilla RL approaches, which fails to effectively balance the strategic knowledge required for varying sizes of VNRs. Additionally, training within a limited range of VNR sizes also incurs a deficiency in generalization and adaptability for previously unseen VNR sizes. Furthermore, training a universal solving policy may perform reasonably well on average but not be optimal for specific VNR sizes. This lack of specialization for different sizes results in sub-optimal solutions when dealing with VNRs of varying sizes. Therefore, addressing this generalization problem is crucial to achieving high-quality solutions.

One intuitive solution is to train a set of sub-policies directly to handle VNRs of different sizes. However, as shown in Fig. 4, A3C-GCN-MultiPolicy that have multiple size-specific policies are even inferior to A3C-GCN on some performance metrics. To further explore this phenomenon, we test the performance of each size-specific strategy on different sizes and provide the results in Fig. 5. We can observe that specific policies trained for VNRs of small sizes often achieve the best performance in their corresponding testing dataset of the same size. Meanwhile, the specific policies of large-sized VNRs exhibit unsatisfactory performance on nearly all testing sizes. Even on the same test size as the training size, their performance is worse than the general policy. This indicates

they are stuck in local optima.

This issue arises due to the more complex solution spaces and stringent constraints associated with larger VNRs, which significantly complicate the exploration of high-reward states and the discovery of feasible solutions. Consequently, when these sub-policies for large-sized VNRs are trained from scratch, they tend to settle into suboptimal solutions and lead to diminished performance. Furthermore, this method based on direct training from scratch lacks quick adaptability to handle situations where unseen VNR sizes occur. When faced with new sizes, the system needs to collect ample data to train policies for the new sizes from scratch, which significantly increases the demand for data.

## B Problem Formulation

### B.1 Optimization Objectives

Acknowledging the stochastic nature of online networking, most existing methods and this work aim to minimize the embedding cost of each arriving VNR onto the physical network. This facilitates the improvement of resource utilization and VNR acceptance rate. To evaluate the quality of solutions, the revenue-to-cost ratio (R2C) serves as a crucial indicator, defined as follows:

$$\text{R2C}(\mathcal{G}^v) = (\Psi \cdot \text{REV}(\mathcal{G}^v)) / \text{COST}(\mathcal{G}^v). \quad (13)$$

Here,  $\Psi$  is the binary variable that indicates the feasibility of a solution.  $\Psi = 1$  if the solution of  $\mathcal{G}^v$  is accepted, and  $\Psi = 0$  if it is not.  $\text{REV}(\mathcal{G}^v)$  denotes the revenue of the VNR  $\mathcal{G}^v$  and  $\text{COST}(\mathcal{G}^v)$  denotes the embedding cost resulting from the solution, which are calculated as follows:

$$\text{REV}(\mathcal{G}^v) = \frac{1}{|\mathcal{C}|} \sum_{n^v \in \mathcal{N}^v} \sum_{C \in \mathcal{C}} C(n^v) + \sum_{l^v \in \mathcal{L}^v} B(l^v), \quad (14)$$

$$\text{COST}(\mathcal{G}^v) = \frac{1}{|\mathcal{C}|} \sum_{n^v \in \mathcal{N}^v} \sum_{C \in \mathcal{C}} C(n^v) + \sum_{l^v \in \mathcal{L}^v} H_{l^v} B(l^v), \quad (15)$$

where  $H_{l^v}$  denotes the hop count of the physical path that routes the virtual link  $l^v$ .

### B.2 Constraint Conditions

The process of embedding one VNR  $\mathcal{G}^v \in \mathcal{V}$  onto the physical network is formulated as a mapping function  $f: \mathcal{G}^v \rightarrow \mathcal{G}^p$ . This process utilizes two types of boolean variables to make decisions: (1)  $x_i^m = 1$  indicating that virtual node  $n_m^v$  is placed in physical node  $n_i^p$ , and 0 otherwise; (2)  $y_{i,j}^{m,w} = 1$  indicating that virtual link  $l_{m,w}^v = (n_m^v, n_w^v)$  traverses through physical link  $l_{i,j}^p = (n_i^p, n_j^p)$ , and 0 otherwise. Here, we use the  $m$  and  $w$  as identifiers of physical nodes and  $i$  and  $j$  as identifiers of virtual nodes. A VNR is successfully embedded when a feasible mapping solution is found, i.e., the following constraints are satisfied:

$$\sum_{n_i^p \in \mathcal{N}^p} x_i^m = 1, \forall n_m^v \in \mathcal{N}^v, \quad (16)$$

$$\sum_{n_i^p \in \mathcal{N}^p} x_i^m \leq 1, \forall n_i^p \in \mathcal{N}^p, \quad (17)$$

$$x_i^m C(n_m^v) \leq C(n_i^p), \forall n_m^v \in \mathcal{N}^v, n_i^p \in \mathcal{N}^p, C \in \mathcal{C}, \quad (18)$$

$$\sum_{n_i^p \in \Omega(n_k^p)} y_{i,k}^{m,w} - \sum_{n_j^p \in \Omega(n_k^p)} y_{k,j}^{m,w} = x_k^m - x_k^w, \forall l_{m,w}^v \in \mathcal{L}^v, n_k^v \in \mathcal{N}^p, \quad (19)$$

$$y_{i,j}^{m,w} + y_{j,w}^{m,w} \leq 1, \forall l_{m,w}^v \in \mathcal{L}^v, l_{i,j}^p \in \mathcal{L}^p, \quad (20)$$

$$\sum_{l_{m,w}^v \in \mathcal{L}^v} (y_{i,j}^{m,w} + y_{j,i}^{m,w}) B(l_{m,w}^v) \leq B(l_{i,j}^p), \forall (l_{i,j}^p) \in \mathcal{L}^p. \quad (21)$$

Here,  $\Omega(n_k^p)$  denotes the neighbors of  $n_k^p$ . Constraint (16) ensures that each virtual node is assigned to exactly one physical node, while the constraint (17) ensures that each physical node accommodates at most one virtual node, thus enforcing a one-to-one mapping between them. Constraint (18) ensures that each virtual node is assigned to a physical node with sufficient resources where available resources of all types exceed the demands. Constraint (19) ensures that each virtual link  $(n_m^v, n_w^v)$  is routed through a connective physical path from the physical node  $n_i^p$  (where virtual node  $n_m^v$  is mapped, i.e.,  $x_i^m = 1$ ) to the physical node  $n_j^p$  (where virtual node  $n_w^v$  is mapped, i.e.,  $x_j^w = 1$ ). This constraint follows the flow conservation law. Constraint (20) prevents routing loops and guarantees acyclic routing of physical paths for virtual links. Constraint (21) ensures that the bandwidth consumption of each physical link does not exceed its available capacity.

## C Proof of the Superiority of MDP Optimality

**Theorem.** Given two MDPs with the bidirectional and unidirectional action,  $\mathcal{M}^b = \langle \mathcal{S}^b, \mathcal{A}^b, P, R, \lambda \rangle$  and  $\mathcal{M}^u = \langle \mathcal{S}^u, \mathcal{A}^u, P, R, \lambda \rangle$ , and their optimal policy denote as  $\pi^{*,b}$  and  $\pi^{*,u}$ , respectively, we have  $\pi^{*,b} \succeq \pi^{*,u}$ .

*Proof.* In  $\mathcal{M}^b$ , the bidirectional action at the timestep  $t$ ,  $a_t = (n_t^p, n_t^v)$ , is the joint selection of virtual and physical nodes. The optimal policy for  $\mathcal{M}^b$  can be defined as:

$$\pi^{*,b}(a_t | s_t) = \pi^{H,*}(n_t^v | s_t) \cdot \pi^{L,*}(n_t^p | s_t, n_t^v), \quad (22)$$

where  $\pi^{H,*}(n_t^v | s_t)$  optimally orders virtual nodes and  $\pi^{L,*}(n_t^p | s_t, n_t^v)$  optimally selects physical nodes given the virtual node selection.

In  $\mathcal{M}^u$ , the predefined policy  $\pi^{H,u}(n_t^v | s_t)$  governs the decision sequence of virtual nodes. Similarly, The optimal policy for  $\mathcal{M}^u$  can be similarly defined as

$$\pi^{*,u}(a_t | s_t) = \pi^{H,u}(n_t^v | s_t) \cdot \pi^{L,*}(n_t^p | s_t, n_t^v). \quad (23)$$

Notably,  $\pi^{H,u}$  may not be optimal.

Following the standard reinforcement learning framework [Sutton and Barto, 2018], the value of a state  $s$  under a policy  $\pi$ , denoted  $V_\pi(s)$  is the expected return when starting in  $s$  and following  $\pi$  thereafter. For our bilevel policy,  $V_\pi(s)$  can be formulated as:

$$V_\pi(s) = \mathbb{E}_{n^v \sim \pi^H, n^p \sim \pi^L} \left[ \sum_{t=0}^T \lambda^t R(s_t, a_t) | s_0 = s \right]. \quad (24)$$

If and only if  $V_\pi(s) \geq V_{\pi'}(s)$  for all state,  $\pi \succeq \pi'$ .

---

**Algorithm 1:** Training Process of FlagVNE

---

**Input** : Initial meta-policy  $\phi$ ; Policy set  $\Theta = \{\phi\}$ ;  
Meta learning rate  $\beta$ ; Task learning rate  $\alpha$ ;  
Policy entropy threshold  $\delta$

**Output** : Trained policies set  $\Theta$ ;

```

1 // Meta-learning Process
2 Initialize the training task ID list  $\mathcal{I} = \{1\}$ ;
3 while not done do
4   Collect the trajectory memory  $\mathcal{D}$  by interactions;
5   Split  $\mathcal{D}$  into  $\{\mathcal{D}_1, \dots, \mathcal{D}_{|\mathcal{M}|}\}$  based on VNR' size;
6   Analyze the task distribution  $\mathcal{M}_i \sim p(\mathcal{M})$ ;
7   for  $i = 1, 2, \dots, |\mathcal{M}|$  do
8     if  $i \notin \mathcal{I}$  then
9       continue
10    end
11     $\theta_i \leftarrow \text{DeepCopy}(\phi)$ ;
12    Update  $\theta_i$  with Eq. (9) and (10); // Inner loop
13  end
14  Update  $\phi$  with Eq. (12); // Outer loop
15  // Curriculum Scheduling Strategy
16  Get the current most complex task ID  $k = \max(\mathcal{I})$ ;
17  if  $H(\pi_{\theta_k}) < \delta$  and  $k < |\mathcal{M}|$  then
18     $\mathcal{I} \leftarrow \mathcal{I} \cup \{k + 1\}$ 
19  end
20 end
21 // Fine-tuning Process
22 for  $i = 1, 2, \dots, |\mathcal{M}|$  do
23    $\theta_i \leftarrow \text{DeepCopy}(\phi)$ ;
24 end
25 while not done do
26   Collect task trajectory memories  $\{\mathcal{D}_1, \dots, \mathcal{D}_{|\mathcal{M}|}\}$ ;
27   for  $i = 1, 2, \dots, |\mathcal{M}|$  do
28     Update  $\theta_i$  with Eq. (9) and (10); // Inner loop
29   end
30 end
31 for  $i = 1, 2, \dots, |\mathcal{M}|$  do
32    $\Theta \leftarrow \Theta \cup \{\theta_i\}$ ;
33 end

```

---

Since  $\pi^{H,u}$  in  $\pi^{*,u}$  is not necessarily optimal, for some states, it might not maximize the expected return, i.e.,  $\pi^{H,u}(n^v|s_t) \neq \pi^{H,*}(n^v|s_t)$ . In contrast,  $\pi^{H,*}$  in  $\pi^{*,b}$  is optimal by definition. Thus, for some states, we have

$$V_{\pi^{*,b}}(s) > V_{\pi^{*,u}}(s). \quad (25)$$

This inequality demonstrates the potential sub-optimality of  $\pi^{*,u}$  compared to  $\pi^{*,b}$ .

Given the above, we conclude that  $\pi^{*,b} \succeq \pi^{*,u}$ .  $\square$

## D Pseudocode of FlagVNE Training

We describe FlagVNE's training process in Algorithm 1.

## E Experimental Details

### E.1 Topology Descriptions

The descriptions of used realistic topologies are as follows:

- *GEANT* is a network infrastructure linking Europe's national research and education networks. It acts as a pan-European backbone, connecting European researchers,

academics, and students, and extends globally to over half the world's countries. This network topology has 40 nodes and 64 lines, with a density of 0.0821.

- *WX100* is generated following the Waxman graph method [Waxman, 1988]. The Waxman graph is a type of random graph model, widely used in the field of communication and network simulation. WX100 has 100 nodes and 500 links, with a density of 0.1010.

### E.2 Implementation Details

Detailed hyperparameter settings are as follows: each neural network has a hidden dimension  $H$  of 128 and the GNN module consists of  $K = 3$  GCN layers. We train this model using Adam optimizer [Kingma and Ba, 2014], with learning rates  $\alpha$  and  $\beta$  of 0.001 and a batch size of 128. In PPO, the RL discounted factor  $\lambda$  is set to 0.99, the coefficient of critic loss is set to 0.5, and the number of repeat times is set to 10. All experiments are carried out on a server equipped with 24 Intel Xeon E5-2650 v4 @ 2.20GHz CPUs and 4 V100 GPUs.

### E.3 Baseline Descriptions

The descriptions of compared baselines are as follows:

- NRM-VNE [Zhang *et al.*, 2018] sorts virtual and physical nodes based on a node resource management metric. Greedy matching and breadth-first search are then used for node mapping and link mapping, respectively.
- NEA-VNE [Fan *et al.*, 2023] ranks virtual and physical nodes with an essentiality assessment metric and then performs node mapping and link mapping similar to [Zhang *et al.*, 2018].
- PSO-VNE [Jiang and Zhang, 2021] uses particle swarm optimization (PSO) to explore the VNE's solution space.
- MCTS-VNE [Haeri and Trajković, 2017] uses the Monte Carlo tree search algorithm to search solutions.
- PG-CNN [Zhang *et al.*, 2022] builds a policy network with convolutional neural network (CNN) and trains it with PG algorithm.
- A3C-GCN [Zhang *et al.*, 2023] constructs a policy network with GCN and MLP and uses A3C algorithm as the training method. It extends the method proposed in [Yan *et al.*, 2020] with an action mask mechanism.
- DDPG-Attention [He *et al.*, 2023b] builds an attention-based policy and trains it with DDPG algorithm.

### E.4 Performance Metrics

The following metrics are widely used to evaluate the long-term operational status of the network system over a period  $\mathcal{T}$  [Fischer *et al.*, 2013; Yan *et al.*, 2020]:

*Request Acceptance Rate* (RAC) quantifies the ratio of accepted VNRs to the total number of VNRs arrived, indicating the system's ability to meet user service demands, defined as:

$$\text{RAC} = \frac{\sum_{t=0}^{\mathcal{T}} |\mathcal{V}_s(t)|}{\sum_{t=0}^{\mathcal{T}} |\mathcal{V}(t)|}, \quad (26)$$

where  $\mathcal{V}(t)$  and  $\mathcal{V}_s(t)$  denote the set of totally arrived and accepted VNRs at the unit of time slot  $t$ , respectively.

	Average Running Time (s) ↓	
	GEANT	WX100
NRM-VNE	10.079	28.079
NEA-VNE	31.011	238.403
PSO-VNE	1330.706	1516.340
MCTS-VNE	240.195	679.007
PG-CNN	75.259	203.965
A3C-GCN	47.079	204.073
DDPG-Attention	81.713	164.355
FlagVNE	84.987	239.251

\* The average simulation time (seconds) over various  $\eta$

Table 2: Average running time in traffic throughput test.

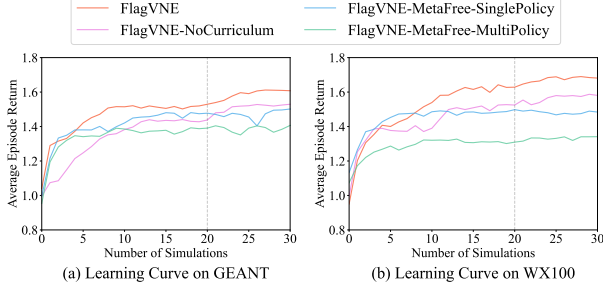


Figure 6: Learning curves over simulations. Within each simulation, there are 1000 VNRs, i.e., 1000 episodes and we average their returns. ( $\eta = 0.001$  on GEANT and  $\eta = 0.06$  on WX100.)

*Long-term Average Revenue (LAR)* quantifies the total gained revenue, directly assessing the financial performance of the ISP, defined as:

$$\text{LAR} = \left( \sum_{t=0}^{\mathcal{T}} \sum_{\mathcal{G}^v \in \mathcal{V}_s(t)} \text{REV}(\mathcal{G}^v) \times d^v \right) / \mathcal{T}. \quad (27)$$

where  $d^v$  denotes the lifetime of VNR  $\mathcal{G}^v$ . *Long-term R2C (LT-R2C)* quantifies the overall revenue-to-cost ratio, evaluating the solution quality of all accepted VNRs and resource utilization, defined as:

$$\text{LT-R2C} = \frac{\sum_{t=0}^{\mathcal{T}} \sum_{\mathcal{G}^v \in \mathcal{V}_s(t)} \text{REV}(\mathcal{G}^v) \times d^v}{\sum_{t=0}^{\mathcal{T}} \sum_{\mathcal{G}^v \in \mathcal{V}_s(t)} \text{COST}(\mathcal{G}^v) \times d^v}. \quad (28)$$

## F Additional Validation

### F.1 Running Time Test

Due to the low-latency requirements of network systems, VNE algorithms should provide solutions within an acceptable time. The average running time over different traffic throughputs same to the above settings of each algorithm is illustrated in Table 2. Compared to PSO-VNE and MCTS-VNE which are quite time-consuming, FlagVNE and other algorithms can offer solutions more efficiently. While NRM-VNE achieves the fastest solving speeds, it is noteworthy that FlagVNE markedly outperforms NRM-VNE on evaluation metrics. Overall, FlagVNE strikes a better balance between performance and running time.

### F.2 Adaptation and Convergence Analysis

We provide the analysis of learning curves on both known distribution and unseen size.

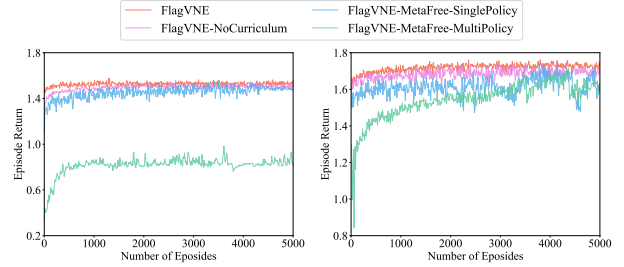


Figure 7: Learning curves for the unseen size. (There are 100 VNRs in one simulation and each VNR’s size is set to 12).

**Training on Known Distribution.** Fig. 6 presents the learning curves of FlagVNE alongside three other variations, trained on a known distribution of VNR’s size [2-10]. We observe that FlagVNE achieves higher average returns at convergence compared to both FlagVNE-MetaFree-SinglePolicy and FlagVNE-MetaFree-MultiPolicy. This outcome highlights the significance of generalization in training methods and validates the efficacy of our generalizable training approach. Additionally, FlagVNE’s superior performance over FlagVNE-NoCurriculum underscores the effectiveness of our curriculum scheduling strategy in mitigating suboptimal convergence and enhancing the overall learning process.

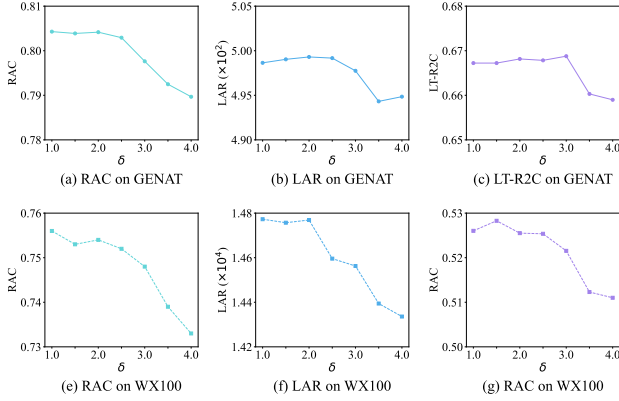
**Adaptation to Unseen Size.** To assess FlagVNE’s adaptability to previously unseen VNR sizes, we conduct experiments with all incoming VNRs in the network system set to size 12, a size not included in earlier training phases. Given the heightened resource demands of these larger-sized VNRs, we limited the number of simultaneous VNRs to 1 per simulation, maintaining all other simulation parameters as described in Section 5.1. During this phase, we fine-tuned the meta-policies of both FlagVNE and FlagVNE-NoCurriculum, as well as FlagVNE-MetaFree-SinglePolicy. Meanwhile, FlagVNE-MetaFree-MultiPolicy underwent training to develop a new size-specific policy from scratch. As shown in Fig. 7, FlagVNE achieves rapid convergence within a small amount of training data, significantly reducing the data requirements. This highlights the benefits of utilizing the initial meta-policy that acquires cross-task knowledge, which enables FlagVNE to rapidly adapt to unseen-size-specific sub-policies. Notably, in the case of GEANT, FlagVNE-MetaFree-MultiPolicy fails to converge within 5000 epochs. This outcome is mainly due to GEANT’s smaller topology and constrained resources, which limit the exploration of feasible solutions. Such conditions pose significant challenges for a policy that lacks prior experience or foundational knowledge, further emphasizing the effectiveness of our meta-learning approach in environments with limited resource scenarios.

### F.3 Scalability Validation

To further validate the effectiveness of FlagVNE in large-scale network systems, we study the performance of all algorithms in large-scale networks. Following previous works [Yan *et al.*, 2020; Geng *et al.*, 2023], we generate a random Waxman topology with 500 nodes and nearly 13000 links, named WX500 [Waxman, 1988]. We also increase the VNE size distribution to a uniform distribution from 2 to 20,

Table 3: Results on Scalability Validation.

Algorithm	RAC $\uparrow$	LAR ( $\times 10^6$ ) $\uparrow$	LT-R2C $\uparrow$
NRM-VNE	0.631	0.710772	0.507
NEA-VNE	0.857	1.186615	0.690
PSO-VNE	0.805	1.042604	0.537
MCTS-VNE	0.782	0.968175	0.563
PG-CNN	0.851	1.046523	0.548
A3C-GCN	0.869	1.147116	0.715
DDPG-Attention	0.796	1.013670	0.617
<b>FlagVNE</b>	<b>0.932</b>	<b>1.347162</b>	<b>0.744</b>

Figure 8: The impact of  $\delta$  on FlagVNE’s performance. ( $\eta = 0.001$  on GEANT and  $\eta = 0.06$  on WX100.)

and the arrival rate of VNR  $\eta$  is set to 3. For the training of FlagVNE, we execute the meta-learning in the initial 40 simulations and then conduct the fine-tuning with 10 simulations. All other simulation and hyperparameter settings remain consistent with those outlined in Section 5.1. The results of all algorithms are shown in Table 3. We observe that FlagVNE and NRM-VNE achieve the best and worst performance, respectively. Compared to other baselines, FlagVNE presents clear performance advantages, which demonstrates its effectiveness in large-scale network systems.

#### F.4 Hyperparameter Sensitivity

We investigate the impact of the policy stability threshold  $\delta$  on FlagVNE’s performance. We keep all other simulation and training parameters consistent with those detailed in Section 5.1. The experimental results are shown in Fig. 8, where (a)(b)(c) and (e)(f)(g) show the testing results on GEANT at  $\eta = 0.001$  and WX100 at  $\eta = 0.18$ , respectively. We observe that FlagVNE exhibit relative stability within a  $\delta$  range of [1.0, 2.5]. This stability is attributed to the higher  $\delta$  value signals more consistent and confident decision-making by the current meta-policy for the currently largest size VNRs. Such policies can serve as a better initialization of specific policies tailored to larger VNRs. To trade off performance and efficiency, we set this parameter to 2. Notably, FlagVNE’s performance on WX100 demonstrates a higher sensitivity to  $\delta$  compared to GEANT. This increased sensitivity can be ascribed to the more challenging exploration environment in WX100, where the policy is more prone to suboptimal con-

vergence. This observation underscores the value of our curriculum scheduling strategy, which aids in mitigating the issues associated with such exploration challenges by gradually increasing task complexity.