

Reconfiguring Multicast Sessions in Elastic Optical Networks Adaptively with Graph-Aware Deep Reinforcement Learning

Xiaojian Tian, Baojia Li, Rentao Gu, and Zuqing Zhu, *Senior Member, IEEE*

Abstract—With the fast deployment of datacenters (DCs), bandwidth-intensive multicast services are becoming more and more popular in metro and wide-area networks, to support dynamic applications such as DC synchronization and backup. Hence, this work studies the problem of how to formulate and reconfigure multicast sessions in an elastic optical network (EON) dynamically. We proposed a deep reinforcement learning (DRL) model based on graph neural networks (GNNs) to solve the sub-problem of multicast session selection in a more universal and adaptive manner. The DRL model abstracts the topology information of the EON and the current provisioning scheme of a multicast session as graph-structured data, and analyzes the data to intelligently determine whether the session should be selected for reconfiguration. We evaluate our proposal with extensive simulations that consider different EON topologies, and the results confirm its effectiveness and universality. Specifically, the results show that it can balance the tradeoff between the number of reconfiguration operations and blocking performance much better than the existing algorithms, and the DRL model trained in one EON topology can easily adapt to solve the problem of dynamic multicast session reconfiguration in other topologies, without being redesigned or retrained.

Index Terms—Optical multicast, Elastic optical networks (EONs), Network reconfiguration, Deep reinforcement learning (DRL), Graph neural network (GNN).

I. INTRODUCTION

In recent years, the rising of cloud services and live video streaming has made multicast services more and more popular in the Internet [1]. This trend becomes even more remarkable since 2020, because of the surge in demands for video conferencing and online classroom services during the epidemic. Meanwhile, due to the fast deployment of datacenters (DCs) all over the world, the popularity of multicast services can also be seen in metro and wide-area networks [2], especially for bandwidth-intensive applications such as DC synchronization and backup, distributed scientific computing, *etc* [3]. This has put great pressure on DC interconnects (DCIs) and made multicast provisioning in DCIs an attractive research topic.

With the tremendous bandwidth in each optical fiber, optical networking plays an important role in DCIs, and a latest study [4] even suggested that an optical-circuit-switched architecture could be more scalable and cost-effective for regional DCIs

than a natural packet-switched network. More promisingly, the advances on the flexible-grid elastic optical networks (EONs) can further improve the performance of optical switching on spectrum-efficiency, adaptivity and application-awareness [5–7]. Note that, for bandwidth-intensive and long-lasting applications (*e.g.*, DC backup), realizing multicast directly in the optical domain has the benefits such as less bandwidth/protocol overheads and easier to obtain large throughputs [8]. The agility of EONs would further promote these benefits, which motivated people to study how to provision multicast services in EONs and proposed various algorithms [9–14].

Meanwhile, the semi-permanent optical layer in telecommunication networks might not adapt to the dynamic applications and traffic in DCIs [15]. Therefore, a dynamic optical layer with fast reconfiguration speed is desired. For instance, the standardization effort in [16] suggested that to properly support inter-DC communications, a dynamic optical network should be reconfigurable within a few milliseconds. Following this trend, researchers have considered different dynamic operation scenarios for EONs, *e.g.*, the reconfiguration to accommodate time-varying unicast traffic [17, 18], spectrum defragmentation [19], lightpath restoration [20], and spectrum retuning for bulk data transfers [21]. The dynamic nature of the multicast services in DCIs determines that each multicast session might also need to be updated consistently to maintain the optimality of its service provisioning scheme (*i.e.*, the one that consumes the least spectrum resources) [22]. For example, during a one-to-many DC backup, each destination DC joins the multicast session when the data of its interest starts to be transferred, and it will leave the session when its data transfer is done.

The problem of how to formulate and reconfigure multicast sessions in EONs dynamically was previously studied in [22]. Specifically, the authors divided the problem into two sub-problems, *i.e.*, session selection and session reconfiguration, and designed algorithms to solve them. The session selection algorithm finds the most “critical” multicast sessions whose provisioning schemes waste the most spectrum resources when being compared with the optimal ones (*i.e.*, off their optima the most), to reconfigure. After the sessions have been selected, they can be reconfigured with either full or partial rearrangements in the session reconfiguration, to free up unnecessary spectrum usages. Note that, the reconfiguration of multicast sessions should be evaluated from two perspectives, *i.e.*, the number of reconfiguration operations and overall blocking probability of multicast sessions. Specifically, by invoking more reconfiguration operations, we generally can

X. Tian, B. Li and Z. Zhu are with the School of Information Science and Technology, University of Science and Technology of China, Hefei, Anhui 230027, P. R. China (email: zqzhu@iee.org).

R. Gu is with the School of Information and Communication Engineering, Beijing University of Posts and Telecommunications, Beijing 100876, China.

Manuscript received on May 9, 2021.

readjust the provisioning schemes of multicast sessions better to save more spectrum resources, and thus a lower blocking probability will be get in the future. Hence, to maximize the efficiency of the reconfiguration, we should use the least reconfiguration operations to achieve the largest reduction on blocking probability. However, to the best of our knowledge, how to optimize this tradeoff has not been fully explored yet.

We can see that in the reconfiguration of multicast sessions, the sub-problem of session selection is more relevant to the aforementioned tradeoff. Nevertheless, the heuristic approaches developed in [22] (*i.e.*, the D-/Q-value based selection strategies) cannot universally adapt to dynamic EON environments, and the problem of how to select between them and determine their key parameters can only be tackled in an empirical manner. This motivates us to revisit the sub-problem in this work. Note that, deep reinforcement learning (DRL) can obtain statistically optimal solutions for complex and time-varying problems without explicit programming [23]. Hence, we try to replace the heuristic approaches for session selection with a DRL-based algorithm, and expect that it can balance the tradeoff between the number of reconfiguration operations and blocking probability better.

Note that, in order to select multicast sessions in an EON to reconfigure, we need to process data in graph structure, which can hardly be handled well by the neural networks (NNs) in linear structures. This is because certain important information buried in the graph-structured data can be lost, and the DRL models with NNs in linear structures need to be redesigned and retrained when the EON's topology changes. Fortunately, graph neural networks (GNNs) [24] can fulfill the requirements much better, as they can operate directly on graph-structured data to understand the complex relations in it for the applications related to networks [25].

In this work, we propose a DRL model based on GNNs to solve the sub-problem of multicast session selection in a more universal and adaptive way. The DRL model takes the topology information of the EON and the current provisioning scheme of a multicast session as the input, abstracts them as graph-structured data, and analyzes the data to intelligently determine whether the multicast session should be selected for reconfiguration. We evaluate the proposed graph-aware DRL model with extensive simulations that consider different EON topologies. The simulation results confirm the effectiveness and universality of our proposal, and show that it can balance the tradeoff between the number of reconfiguration operations and blocking probability much better than the existing heuristic approaches, without empirical parameter adjustments.

The rest of the paper is organized as follows. Section II briefly surveys the related work. We describe the network model and operation principle of the dynamic reconfiguration of multicast sessions in EONs in Section III. The graph-aware DRL model for session selection is designed in Section IV, and we discuss its performance evaluations in Section V. Finally, Section VI summarizes the paper.

II. RELATED WORK

Multicast in the optical domain has been studied since the inception of wavelength-division-multiplexing (WDM) net-

works, and Sahasrabudde *et al.* [8] first came up with the concept of light-tree for it. One can refer to the survey in [26] for a complete review of optical multicast in fixed-grid WDM networks. The proposals of flexible-grid EON [5–7] considered to leverage bandwidth-variable transponders (BVTs) and bandwidth-variable switches (BV-WSS) to manage the spectrum allocation in the optical layer with a fine granularity of 12.5 GHz or even less, and thus can make the optical layer more spectrum-efficient and adaptive. Meanwhile, the flexible spectrum management in EONs transforms the well-known routing and wavelength assignment (RWA) problem in WDM networks into a more complex one, *i.e.*, the routing and spectrum assignment (RSA) [27]. Hence, the provisioning of optical multicast should be revisited for EONs.

In [9], the authors proposed two multicast-capable RSA (MC-RSA) algorithms for EONs and analyzed their performance. Liu *et al.* [12] improved the performance of MC-RSA by leveraging layered auxiliary graphs. Nevertheless, these two studies did not consider the adaptive modulation-level selection in EONs. The multicast provisioning with impairment-aware routing, modulation and spectrum assignment (RMSA) was addressed in [11], where the authors designed two integer linear programming (ILP) models and a few heuristics. Then, the multicast-capable RMSA (MC-RMSA) algorithms to support distance-adaptive transmissions were developed in [13]. The authors of [14] introduced light-forest to further improve the performance of MC-RMSA and proposed a polynomial-time approximation algorithm. In addition to algorithmic contributions, people have also leveraged the idea of software-defined EON (SD-EON) to experimentally demonstrate the control plane operations for optical multicast in [28].

However, the aforementioned studies all assumed that the optical switches are multicast-capable (MC) (*i.e.*, supporting light-splitting). Note that, MC optical switches usually have complicated architectures and thus can be relatively expensive [26]. Therefore, it might not be cost-effective to build an EON with them, since the majority of the communications in the EON will still be for unicast services. This issue can be addressed by realizing multicast with multicast-incapable (MI) optical switches, *i.e.*, establishing a logic light-tree for each multicast session with multiple unicast lightpaths [10]. Specifically, the study in [10] proposed a spectrum-flexible member-only relay (OL-M-SFMOR) scheme for this purpose.

Another benefit of realizing multicast with MI optical switches is that the multicast sessions can be reconfigured in a local and easier manner. This is because the multicast with MC optical switches has the restriction that all the branches of a light-tree should have the same spectrum assignment, while this is not required by the OL-M-SFMOR scheme [10]. In [22], the authors studied how to formulate and reconfigure multicast sessions dynamically, assuming that OL-M-SFMOR is used in an EON built with MI optical switches. Nevertheless, as we have already explained, the algorithms proposed in [22] for multicast session selection still have a few drawbacks, which motivate us to revisit the sub-problem in this work and try to solve it better with a novel graph-aware DRL model.

Previously, Li *et al.* [29] designed a deep neural network (DNN) to predict the performance of multicast light-trees.

However, the DNN still uses a linear architecture, which is not good at processing graph-structured data, and the topic was not on multicast reconfiguration. Due to its promising performance on processing graph-structured data, GNN has attracted great attention nowadays [25], especially for the complex optimizations in networks [30, 31].

III. PROBLEM DESCRIPTION

In this section, we explain the network model and operation principle of the dynamic multicast reconfiguration in EONs.

A. Network Model

The topology of an EON for DCI is modeled as a directed graph $G(V, E)$, where V and E are the sets of DCs and fiber links, respectively. Here, similar to the case in [22], we assume that the EON is built with MI optical switches. On each link $e \in E$, there are F frequency slots (FS'), each of which has a bandwidth of 12.5 GHz. The BV-Ts that terminate each fiber link are assumed to be the sliceable ones [32], which means that as long as there are sufficient spectrum resources on a link, its BV-Ts can always be sliced to facilitate the requested lightpath transmissions.

We model each multicast session as $MR(s, D, b, t)$, where $s \in V$ denotes the source, D represents the set of destinations, b is the bandwidth demand in Gbps, and t stands for its life-time. In this work, we consider a dynamic EON environment that each multicast session $MR(s, D, b, t)$ can come and leave on-the-fly, and during its life-time t , the DCs in D can change over time too. Hence, when a new multicast session first comes in, we leverage the OL-M-SFMOR scheme in [10] to set up several lightpaths for establishing a logic light-tree, such that each destination in D can receive b Gbps from the source s through one or more lightpaths. Here, each lightpath for serving the multicast session can only start and end at its member nodes (*i.e.*, those in $s \cup D$) for saving BV-Ts, according to the principle of OL-M-SFMOR [10]. As the optical signal is only transmitted all-optically on each lightpath, the RSA schemes of different lightpaths in the logic light-tree are independent, *i.e.*, the spectrum assignments on different branches of the light-tree can be different.

After the initial provisioning of the multicast session, the DCs in D can change over time. Then, when a DC leaves the session or a new DC joins in, the lightpaths in the logic light-tree are updated, still with OL-M-SFMOR. Nevertheless, this might gradually degrade the optimality of the logic light-tree, the RSA schemes of certain lightpaths in it can be sub-optimal and waste spectrum resources. Therefore, we need to reconfigure the multicast session adaptively from time to time.

B. Dynamic Reconfiguration of Multicast Sessions

We use *Algorithm 1* to explain the operation principle of dynamic formulation and reconfiguration of multicast sessions [22]. *Lines 2-10* explains how to formulate multicast sessions dynamically. Then, the reconfiguration of multicast sessions is triggered periodically to maintain the optimality of the logic light-trees of in-service multicast sessions. Here, we need to

solve two sub-problems for the reconfiguration, *i.e.*, session selection (*Line 12*) and session reconfiguration (*Line 13*). The session selection needs to find the most ‘‘critical’’ multicast sessions whose logic light-trees are off their optima the most, to reconfigure. The session reconfiguration rearranges the logic light-trees of the selected sessions to save spectrum resources, which can be done with either full or partial rearrangements [22]. Specifically, the full rearrangement recalculates the logic light-tree of each selected session with OL-M-SFMOR, while the partial rearrangement only chooses certain lightpaths in the logic light-tree of each selected session to reconfigure, according to the average cost of the lightpaths in it¹.

Algorithm 1: Dynamic provisioning of multicast sessions

```

1 while the EON is operational do
2   for each newly-arrived session  $MR_i(s, D, b, t)$  do
3     try to set up a logic light-tree for it with
       OL-M-SFMOR;
4     if the light-tree cannot be established then
5       mark  $MR_i$  as blocked;
6     end
7   end
8   for each existing session  $MR_j(s, D, b, t)$  do
9     if  $t$  has been expired then
10      remove  $MR_j$  and free its resources;
11    end
12    if  $D$  has changed then
13      update its logic light-tree with
        OL-M-SFMOR;
14    end
15  end
16  if it is the time to reconfigure multicast sessions then
17    select existing multicast sessions to reconfigure;
18    reconfigure the selected multicast sessions;
19  end
20 end

```

Throughout the aforementioned process, we need to balance the tradeoff between the number of lightpath reroutings and overall blocking probability of multicast sessions. It can be seen that the sub-problem of session selection is more relevant to this tradeoff. Hence, in the following, we first review the D-/Q-value based selection strategies designed in [22], analyze their drawbacks, and then explain the principle of our graph-aware DRL based selection algorithm.

The D-value of a logic light-tree is actually the hop-count of its longest-destination branch [22]

$$\mathcal{D}(\mathcal{T}) = \max[\text{hops}(s \rightarrow d), \forall d \in D], \quad (1)$$

where \mathcal{T} is the light-tree for multicast session $MR(s, D, b, t)$ and $\text{hops}(\cdot)$ returns the hop-count of a routing path. With the definition in (1), the D-value based selection (DTS) strategy first calculates the average D-value of all the in-service multicast sessions, and then selects those whose D-values are larger

¹Here, the cost of a lightpath was defined in [22], which depends on the lightpath’s spectrum usage and the number of hops of its routing path.

than the average value to reconfigure. As DTS only considers the branch lengths of each logic light-tree but does not address the overall tree structure or the spectrum assignment on the links, it might not always select the most critical sessions to reconfigure.

The Q-value of a logic light-tree considers the overall tree structure and the spectrum assignment on its links [22]

$$Q(\mathcal{T}) = \frac{\text{hops}(\mathcal{T}^*) \cdot \text{hid}_x(\mathcal{T}^*)}{\text{hops}(\mathcal{T}) \cdot \text{hid}_x(\mathcal{T})}, \quad (2)$$

where \mathcal{T}^* is the logic light-tree that is calculated with OLM-SFMOR based on the current network status, and $\text{hid}_x(\cdot)$ returns the highest index of the used FS' on a light-tree. With (2), the Q-value based selection (QTS) strategy first chooses a threshold Q_{lb} , and then selects those whose Q-values are smaller than Q_{lb} to reconfigure. Although QTS considers more information of a logic light-tree than DTS, the information is still somehow limited, and the value of Q_{lb} can only be determined empirically, which is rather difficult in a dynamic EON or for EONs with various topologies.

To address the issues of DTS and QTS, this work proposes to select multicast sessions in a self-adaptive manner with graph-aware DRL. More specifically, the DRL model takes the topology information of the EON and the current provisioning scheme of a multicast session as the input, abstracts them as graph-structured data, and analyze the data with GNNs to intelligently determine whether the multicast session should be selected for reconfiguration. Meanwhile, after offline training, the DRL model is also trained in the online manner to make sure that it can optimize its decision-making automatically and adaptively according to the reward feedbacks from a dynamic EON environment, *i.e.*, its effectiveness and universality can be guaranteed without empirical parameter adjustments.

IV. GRAPH-AWARE DRL BASED APPROACH

This section elaborates on our graph-aware DRL model for multicast session selection. Note that, to determine whether a multicast session should be selected for reconfiguration or not, we need to process graph-structured data (*i.e.*, the tree topology and the spectrum usages on its links). This task is suitable for GNNs, because NNs in linear structures normally only deal with the data in Euclidean domains well [25].

A. System Architecture

We still assume that the EON for DCI is operated by leveraging software-defined networking (SDN), which means that the control plane consists of a centralized controller to handle the tasks for network control and management (NC&M). Our graph-aware DRL model obtains the information about the EON and the multicast sessions in it from the controller, and selects the most critical sessions to reconfigure.

Fig. 1 explains the operation principle of the graph-aware DRL model, and its work-flow is illustrated with step numbers. Dynamic requests regarding multicast sessions (*i.e.*, new multicast sessions and changes on in-service sessions) are first processed by the request handler, which dispatches them to both the traffic engineering database (TED) and the service

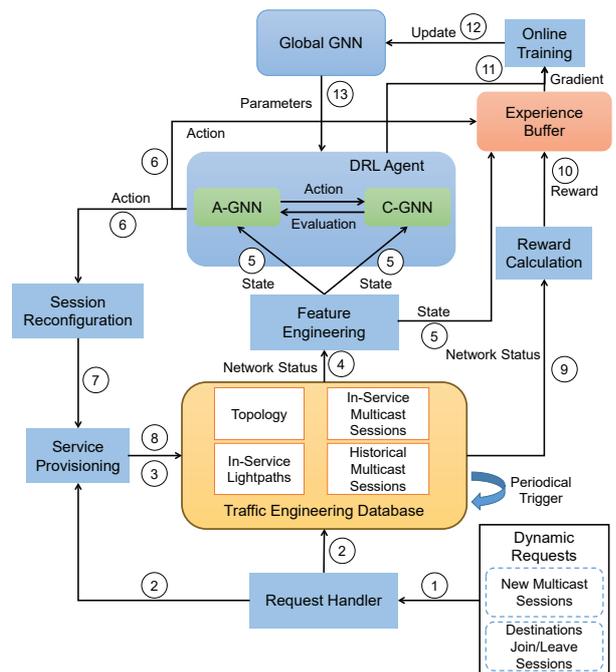


Fig. 1. Architecture and operation principle of our graph-aware DRL model.

provisioning module. As explained in *Algorithm 1*, the service provisioning module serves the dynamic requests and updates their provisioning results in the TED. Then, the reconfiguration of multicast sessions is triggered periodically, and it starts from TED sending the current network status to the feature engineering module, which abstracts the network status to a state that consists of graph-structured data. The DRL agent uses two GNNs, *i.e.*, the actor GNN (A-GNN) and critic GNN (C-GNN), to analyze the state of each in-service multicast sessions and select certain sessions to reconfigure.

Next, the action from the DRL agent (*i.e.*, the selected multicast sessions) is forwarded to the session reconfiguration module, which works with the service provisioning module to reconfigure the selected sessions. After this, the TED sends the new network status to the reward calculation module to obtain the reward of the last action conducted by the DRL agent. Then, we organize the state, action and reward as a training sample, and store it in the experience buffer. When enough entries of experience have been accumulated, the online training module invokes a training process to update the global GNN, which in turn updates the parameters of the A-GNN and C-GNN in the DRL agent accordingly.

B. Preprocessing of Data

To prepare the input to the graph-aware DRL model, we abstract the topology information of the EON and the current provisioning scheme of a multicast session as graph-structured data $\mathcal{G}(V, \tilde{V}, E, \tilde{E})$, where V and E still represent the sets of DC nodes and fiber links in the EON, respectively, while \tilde{V} and \tilde{E} denote the features of the nodes and links in V and E , respectively, regarding the current provisioning scheme a multicast session $MR(s, D, b, t)$. Specifically, according to the current logic light-tree \mathcal{T} of MR , we classify the nodes

in V into 5 categories, *i.e.*, the source s , destinations in D , intermediate node on \mathcal{T} that used to be a destination, normal intermediate node, and nodes that are not on \mathcal{T} . Then, the feature of a node $v \in V$ can be described with a corresponding vector $\tilde{v} \in \tilde{V}$, with one-hot coding. Here, we use 5 bits to represent the aforementioned node categories, respectively. For instance, if we have a node $v = s$, its feature vector \tilde{v} should be $[1, 0, 0, 0, 0]$, or if the node v is a destination, its feature vector \tilde{v} should be $[0, 1, 0, 0, 0]$. On the other hand, the feature of a link $e \in E$ is defined as $\tilde{e} = \frac{f}{F}$, where f is the number of unused FS' on link e and F is the total number of FS' there.

Here, for simplicity, we do not consider distance-adaptive modulation selection, and assume that all the lightpaths in each logic light-tree use the lowest modulation level (*i.e.*, BPSK). Note that, if we need to consider distance-adaptive modulation selection, the only difference is that we should let our DRL model learn the relation between the transmission distance of a lightpath and the number of FS' that it uses. Hence, when preprocessing the graph-structured data of $\mathcal{G}(V, \tilde{V}, E, \tilde{E})$, we need to include the length of each fiber link as an attribute, modify the feature of each link in E and \tilde{E} accordingly, and redesign the GNNs in the DRL model to accommodate the changes. This will be considered in our future work.

C. Structure of GNN

We design the GNNs used in our DRL model based on graph convolutional network (GCN) [33]. The GCN takes the graph-structured data $\mathcal{G}(V, \tilde{V}, E, \tilde{E})$ as the input, and performs two types of operations on the data, *i.e.*, the message transfer and information reduction. For the two types of operations, we define two functions as follows. The message function calculates the message to be sent from node v to node u

$$msg(v, u) = \tilde{v} \cdot \tilde{e}, \quad v, u \in V, e = (v, u) \in E, \quad (3)$$

where nodes v and u are connected with a link $e = (v, u)$ in $\mathcal{G}(V, \tilde{V}, E, \tilde{E})$. The reduction function reduces the messages that each node in V receives from its neighbors.

$$rdu(v) = \sum_{\{u:(u,v) \in E\}} msg(u, v), \quad v \in V. \quad (4)$$

Then, we send $\{rdu(v), \forall v \in V\}$ through a linear network in the GCN to obtain the new feature vector of v , and the transfer function from layer- l to layer- $(l+1)$ is defined as

$$\tilde{v}^{(l+1)} = \sigma(W \cdot rdu^{(l)}(v) + b), \quad (5)$$

where W and b denote the weight matrix and bias of the linear network, respectively, $\sigma(\cdot)$ is the nonlinear transfer function, and $rdu^{(l)}(v)$ represents the reduced information for node v obtained in layer- l of the linear network.

After several layers of GCNs, we introduce a pooling layer to aggregate the processed graph-structured data and get a vector for representing it. Specifically, we select the pooling layer that averages the feature vectors of each node as

$$\tilde{G} = \frac{1}{|V|} \sum_{v \in V} \tilde{v}^{(k)}, \quad (6)$$

where \tilde{G} is the obtained vector, $|V|$ is the number of nodes in V , and k is the number of GCN layers. Finally, we send \tilde{G} to go through several linear layers, for getting the final output.

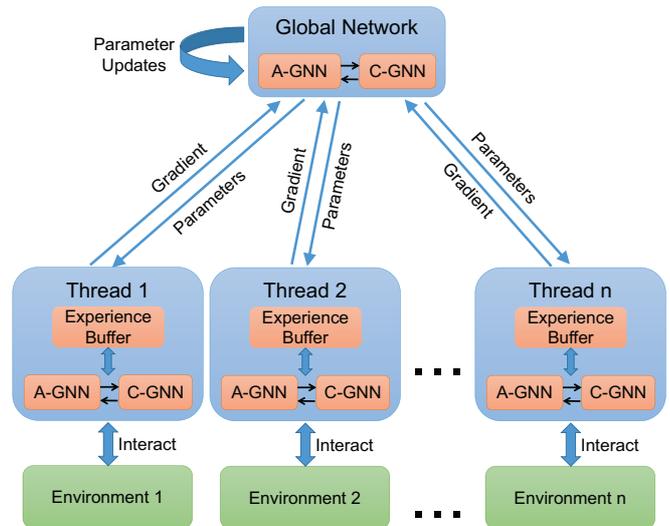


Fig. 2. Training of DRL model in A3C framework.

D. Design of Graph-aware DRL Model

We design the four basic elements of the DRL model as

- **Agent:** The DRL agent is based on the asynchronous advantage actor-critic (A3C) framework [34], which uses multiple pairs of A-GNN and C-GNN for parallel online training in several threads. For each pair of A-GNN and C-GNN, the A-GNN provides an action policy $\pi(S)$ based on the state S in graph structure, and chooses the appropriate action according to the policy $\pi(S)$. The C-GNN is responsible for learning the value of state S and evaluating the action from A-GNN based on it.
- **State:** The state S contains the topology information of the EON and the current provisioning scheme of a multicast session, and it is just the graph-structured data $\mathcal{G}(V, \tilde{V}, E, \tilde{E})$ obtained by the data preprocessing.
- **Action:** The action is modeled with a binary variable a , *i.e.*, if the multicast session should be selected for reconfiguration, we have $a = 1$, and $a = 0$, otherwise.
- **Reward:** We define the reward as follows

$$r = -k_1 \cdot N_{re} + k_2 \cdot [slots(\mathcal{T}) - slots(\mathcal{T}^*)] + k_3 \cdot [cuts(\mathcal{T}) - cuts(\mathcal{T}^*)], \quad (7)$$

where k_1 , k_2 and k_3 are the positive coefficients for normalization, N_{re} represents the number of lightpath reroutings to reconfigure the multicast session, \mathcal{T} and \mathcal{T}^* denote the logic light-trees for the multicast session before and after the reconfiguration, respectively, $slots(\cdot)$ returns the number of FS' used by a logic light-tree, and $cuts(\cdot)$ returns the number of spectrum cuts [19] caused by a logic light-tree. Hence, the reward in (7) decreases with the number of lightpath reroutings, and increases with the spectrum usage and spectrum cuts saved by the reconfiguration. In other words, by maximizing the reward, our graph-aware DRL model tries to invoke the smallest number of lightpath reroutings on a multicast session to achieve the largest savings on spectrum usage and spectrum cuts.

As shown in Fig. 2, we duplicate the A-GNN and C-GNN into several copies, use one copy as the global GNN, and put each of the others in a training thread to expedite the training process. Specifically, each training thread uses its A-GNN and C-GNN to interact with an EON environment independently to obtain training samples. In the iterative manner, the global GNN collects the gradients generated by the training threads, leverages them to update the parameters of its A-GNN and C-GNN, and synchronizes the updated parameters to the A-GNNs and C-GNNs in the training threads. As each thread is trained independently to obtain the gradients, the major benefit of this approach is that it effectively reduces the correlations among training samples. Meanwhile, the multi-thread training can make full use of available computing resources to accelerate the online training.

Algorithm 2 explains the training process in a thread in detail, where we use T to record the number of training iterations, and T_{max} is the upper-limit on training iterations. Lines 3-9 use the local A-GNN and C-GNN of the thread to interact with its own EON environment, for collecting training samples. Then, when enough training samples have been collected, Lines 11-17 perform one iteration of the training. Specifically, the gradients are first calculated locally with the obtained training sample (Lines 11-13), then they are forwarded to the global GNN (Line 14), and finally the thread updates the parameters of its A-GNN and C-GNN according to the feedback from the global GNN and prepares itself for the next iteration of training (Lines 15-17).

V. PERFORMANCE EVALUATION

In this section, we conduct extensive numerical simulations to evaluate our proposed approach based on graph-aware DRL.

A. Simulation Setup

The simulations use the four topologies in Fig. 3 for the EONs for DCIs, to confirm the universality of our proposal in terms of topologies. The capacity of each fiber link is assumed to be $F = 100$ FS', where each FS has a bandwidth of 12.5 GHz to deliver 12.5 Gbps throughput. For each multicast session $MR(s, D, b, t)$, s and D are randomly selected from the nodes in the EON, D contains $[2, 5]$ destinations initially, the bandwidth demand b is uniformly distributed within $[50, 200]$ Gbps, and the life-time t follows the exponential distribution with an average of 500 time-units. As the multicast sessions are dynamic, we generate new multicast sessions according to the Poisson traffic model, and for each in-service multicast, destinations can join or leave dynamically during its life-time. Specifically, the service time of each destination follows the exponential distribution, and it leaves its multicast session when the service time expires, while new destinations are generated with the Poisson distribution. In Section V-E, we will change the settings mentioned above and run more simulations to further verify the universality of our proposal.

The reconfiguration of multicast sessions is invoked every 100 time-units, and this interval is empirically set. The simulations compare our proposal based on graph-aware DRL with the heuristics for session selection in [22] (i.e., DTS and

Algorithm 2: Training process of a thread

```

1  $T = 0$ ;
2 while  $T < T_{max}$  do
3   if it is the time to reconfigure multicast sessions then
4     for each in-service multicast session  $MR_i$  do
5       get state  $S_i$  of  $MR_i$ ;
6       put  $S_i$  into the A-GNN to get an action  $a_i$ ;
7       apply  $a_i$  to the EON environment;
8       calculate reward  $r_i$ ;
9       push  $\{S_i, a_i, r_i\}$  to experience buffer;
10    end
11  end
12  if experience buffer is full then
13    reset the gradients as 0;
14    calculate the loss with A-GNN and C-GNN using
    the training samples in experience buffer;
15    get the gradients with the loss;
16    send the gradients to the global GNN;
17    update the parameters of A-GNN and C-GNN
    according to the feedback from the global GNN;
18    empty the experience buffer;
19     $T = T + 1$ ;
20  end
21 end

```

QTS), and consider both the partial and full rearrangements for session reconfiguration. To ensure sufficient statistical accuracy, we average the results from 5 independent runs to obtain each data point.

B. Training Performance

We first evaluate the training performance of our DRL model. Note that, the DRL model needs to first go through the offline training that optimizes its parameters initially, to make it suitable for being put into online operation/training. Hence, we study the performance of the offline training in this subsection, and will consider that of the online operation/training in subsequent ones. Figs. 4(a) and 4(b) show how the average number of lightpath reroutings per session and blocking probability change in the training process, respectively, for the case in the NSFNET topology with the traffic load at 25 Erlangs. For comparisons, we also plot the results from DTS-based and QTS-based algorithms in Fig. 4. Here, all the algorithms assume that full rearrangement is used to reconfigure the selected multicast sessions, and thus they are labeled with “-F”. In the following, the algorithms labeled with “-F” and “-P” mean that they accomplish multicast session reconfiguration with the full and partial rearrangements in [22], respectively. Meanwhile, for QTS-based algorithms, we can choose their thresholds on Q-value for session selection (i.e., Q_{lb}), and thus they are also labeled with their Q_{lb} values. For instance, the QTS-P-0.8 in Fig. 4 means that the multicast reconfiguration uses QTS to select multicast sessions with $Q_{lb} = 0.8$, and reconfigures them with partial rearrangement.

The results in Fig. 4(a) show that compared with QTS-P-0.8, DTS-P achieves a lower blocking probability by invoking

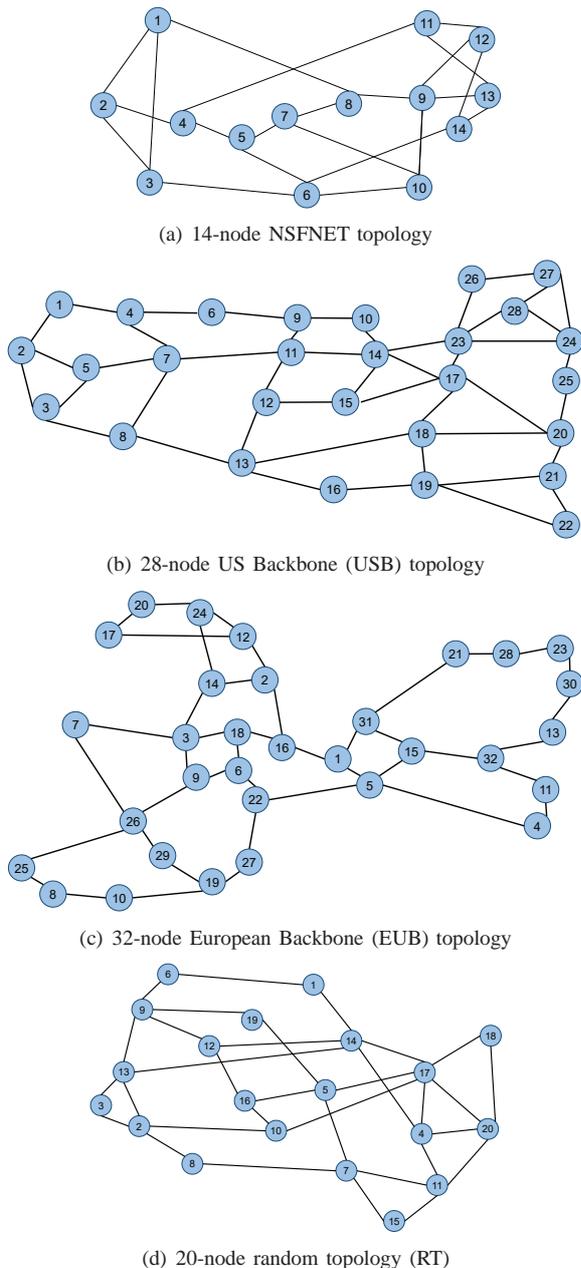


Fig. 3. EON topologies used in simulations.

much more lightpath reroutings per session. Meanwhile, after being trained with more than 5,000 episodes, our DRL-P can obtain a blocking probability that is as low as that of QTS-P-0.8, while its average lightpath reroutings per session is fewer than that of QTS-P-0.8 in Fig. 4(b). In other words, by utilizing its graph-aware intelligence, our DRL-P can balance the tradeoff between overall blocking probability and average lightpath reroutings per session much better than the two benchmarks that use deterministic strategies.

Moreover, to clearly see how the average value of DRL-P's reward correlates with the metrics in Figs. 4(a) and 4(b) in the training, we plot it in Fig. 4(c). Here, we empirically set the positive coefficients in (7) as $k_1 = 6.0$, $k_2 = 1.0$ and $k_3 = 2.0$. It can be seen that the average reward generally increases with the decreases of overall blocking probability

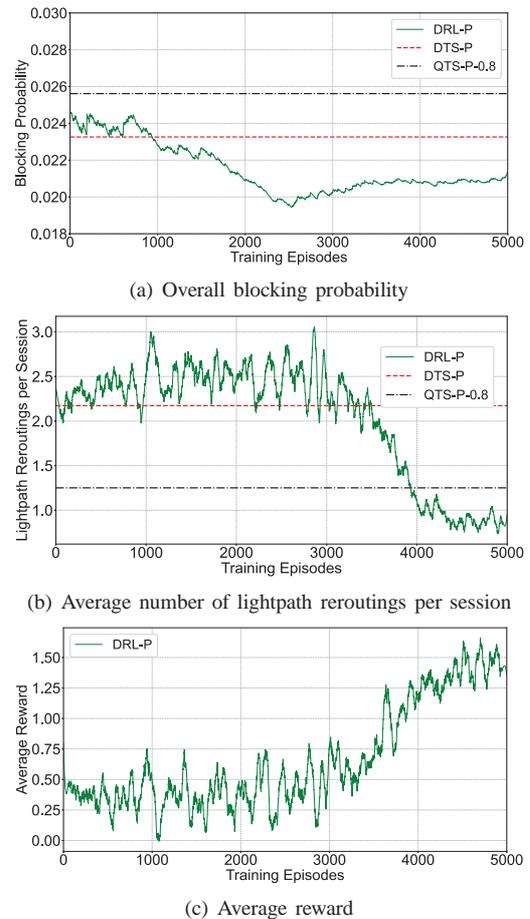


Fig. 4. Training performance (NSFNET, 25 Erlangs).

TABLE I
AVERAGE RUNNING TIME OF OFFLINE TRAINING (SECONDS)

Topology	NSFNET	USB	EUB	RT
DRL-P	25167.35	34698.01	34944.14	32342.45
DRL-F	31541.22	40423.58	45508.86	36843.80

and average number of lightpath reroutings per session in Figs. 4(a) and 4(b), respectively. Note that, we also check other traffic loads in NSFNET and the cases with full rearrangement, and confirm that our DRL-based approach can always achieve similar training performance as that in Fig. 4. Hence, the results are omitted due to the page limit.

Table I lists the running time of the offline training that makes our DRL model suitable for online operation/training. We observe that for the EONs with the NSFNET, US Backbone (USB), European Backbone (EUB), and random (RT) topologies, the running time actually increases with the size of the topology. Meanwhile, the DRL model that is for the full rearrangement scheme usually takes longer offline training time than that for the partial rearrangement scheme, regardless of the topology. These trends are expected, because when the topology of the EON becomes larger or the multicast reconfiguration changes from partial rearrangement to full rearrangement, the problem of multicast reconfiguration actually become more complex.

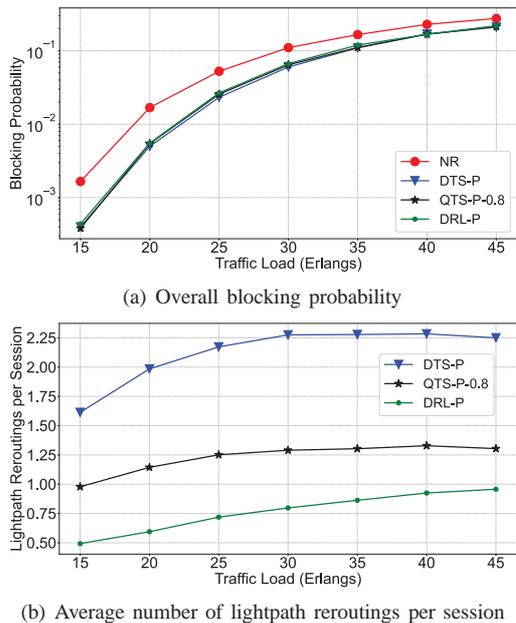


Fig. 5. Results of dynamic operations (NSFNET, partial rearrangement).

C. Performance in Dynamic Network Environments

Next, we evaluate the performance of our DRL-based approach by putting the graph-aware DRL model, which has passed the offline training in a dynamic network environment with the NSFNET topology, and compare its performance with DTS-based and QTS-based algorithms. Figs. 5 and 6 show the simulation results for the cases using partial and full rearrangements, respectively. Here, “NR” denotes the case without multicast session reconfiguration. Note that, in Figs. 5(a) and 6(a), when the traffic load is above 35 Erlangs, the blocking probabilities from the algorithms with multicast session reconfiguration can actually exceed the practical range of the blocking probability in a real-world EON. Although the traffic loads exceed what should be considered in a real-world EON, we still simulate them to get a complete picture about how the algorithms will perform at various traffic loads. The DTS-based algorithm still provides the lowest overall blocking probability with the largest number of lightpath reroutings per session. By combining the results in the figures, we can conclude that to keep the overall blocking probabilities comparable to those of DTS-based and QTS-based algorithms, our DRL model always requires the smallest number of lightpath reroutings per session effectively, for all the simulation scenarios considered in Figs. 5 and 6. Hence, our graph-aware DRL-based approach can effectively reduce the operational complexity of dynamic multicast session reconfiguration, without sacrificing much performance on request blocking.

Moreover, we notice that QTS-based algorithm can change the value of Q_{lb} to balance the tradeoff between blocking probability and average lightpath reroutings per session. Hence, we change Q_{lb} to obtain different sets of blocking probability and average lightpath reroutings per session, and plot the results in Fig. 7, when the traffic load is set as 40 Erlangs. Here, we take average lightpath reroutings per session and blocking probability as the X-axis and Y-axis, respectively, to illustrate

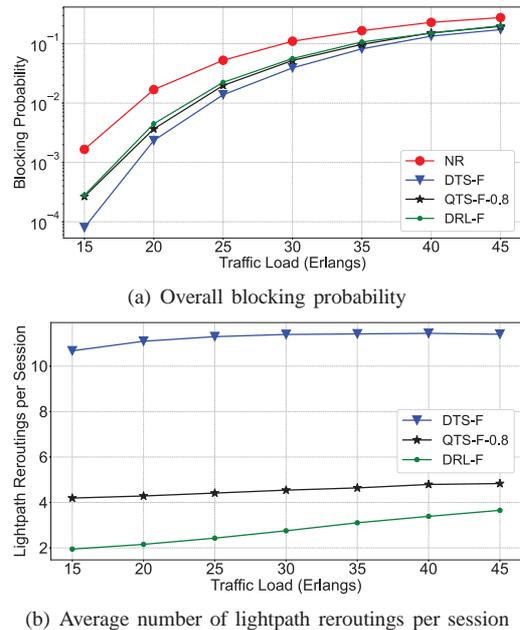


Fig. 6. Results of dynamic operations (NSFNET, full rearrangement).

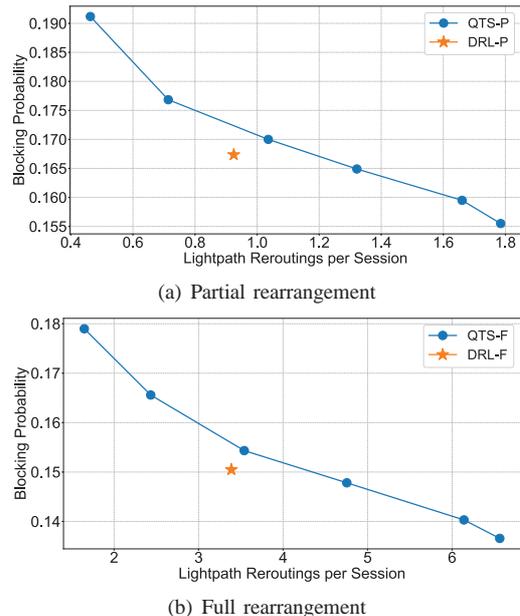


Fig. 7. Tradeoff between blocking probability and average lightpath reroutings per session (NSFNET, 40 Erlangs).

the tradeoff more clearly. It can be seen that no matter partial or full rearrangement is used, the data point for the results from the DRL model is always below the curve for the results from QTS-based algorithm. This verifies that the DRL model balances the tradeoff better than QTS, regardless of the choice of Q_{lb} . In addition to 40 Erlangs, the simulations also check other traffic loads, and similar trends can be obtained.

D. Universality across Different Topologies

We then evaluate the universality of our graph-aware DRL-based approach across different topologies. The operation principle of our graph-aware DRL model ensures that the

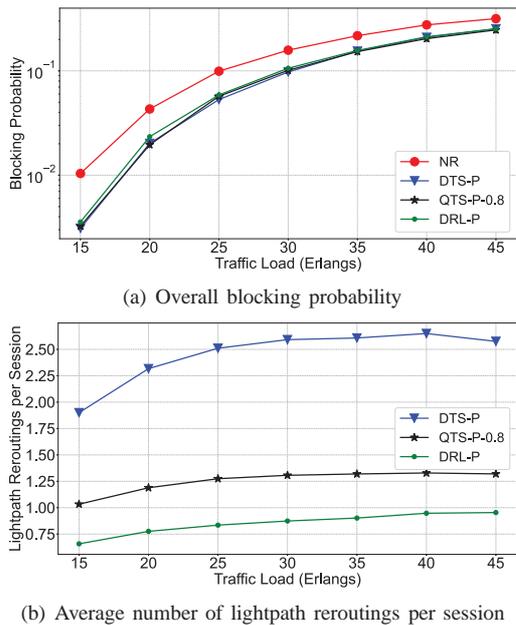


Fig. 8. Results of dynamic operations (USB, partial rearrangement).

DRL model trained in one EON topology can be directly applied to solve the problem of dynamic multicast session reconfiguration in others. Specifically, we only need to abstract the new topology information of the EON and the provisioning scheme of each multicast session as graph-structured data $\mathcal{G}(V, \tilde{V}, E, \tilde{E})$ and input the data to the trained DRL model, while the DRL model does not need to be redesigned or retrained. To verify this, the simulations apply the DRL model trained in NSFNET to solve the problem of dynamic multicast session reconfiguration in the other topologies in Fig. 3.

Fig. 8 shows the results for the dynamic operations in USB, when partial rearrangement is considered. We can see that the results follow the similar trends as those in Fig. 5. To further clarify the adaptability of our DRL model, we take the case of traffic load at 25 Erlangs in USB as an example, and plot how the performance metrics change over the simulation time in Fig. 9. As we directly apply the DRL model trained in NSFNET to the EON with the USB topology, a zero-shot transfer learning (*i.e.*, applying a trained DRL model to an unseen environment for the same task [35]) is actually considered. It can be seen that due to the superior adaptability of our DRL model, it achieves relatively good performance on the performance metrics at the beginning of the online operation/training, and both the overall blocking probability and average number of lightpath reroutings per session only changes slightly afterwards.

Figs. 10 and 11 illustrate the results obtained by directly applying the DRL model trained in NSFNET to the EONs with EUB and RT topologies, respectively. The results still follow the similar trends as those in Fig. 5. Note that, when the EON topology changes, we might need to change the value of Q_{lb} (*i.e.*, the threshold on Q-value for session selection) for QTS-based algorithms empirically. This is the reason why we simulate QTS-P-0.9 in EUB (as shown in Fig. 10). On the other hand, with its graph-aware intelligence, our DRL

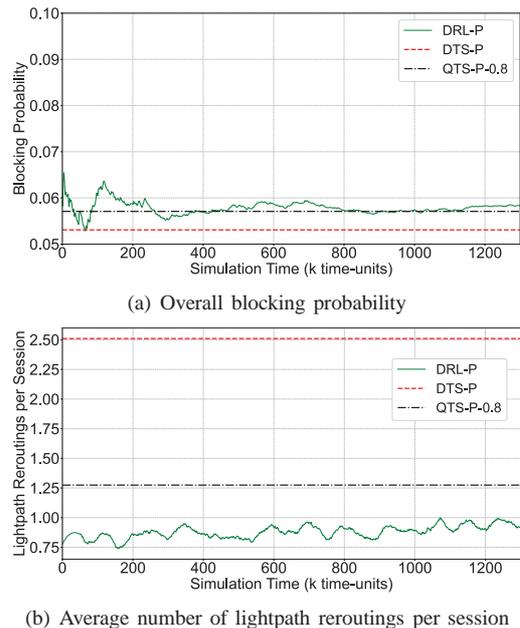


Fig. 9. Performance on zero-shot transfer learning (USB, 25 Erlangs).

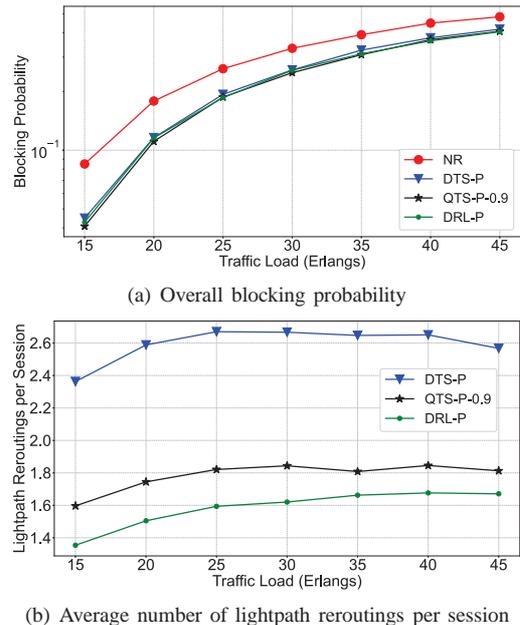


Fig. 10. Results of dynamic operations (EUB, partial rearrangement).

model can adapt to different topologies without such manual adjustments. Although the results in Figs. 8-11 are all about the cases that use partial rearrangement, we also check those with full rearrangement and confirm that our DRL-based approach achieves similar performance in them too. Therefore, we prove the universality of our DRL model across different topologies.

Table II lists the average running time per multicast session reconfiguration of the algorithms. Here, for our DRL model, the running time is only for its online operation/training, because the offline training should be finished before the DRL model can be put into operation and its running time has already been summarized in Table I. The results in Table II suggest that the running time of all the algorithms is

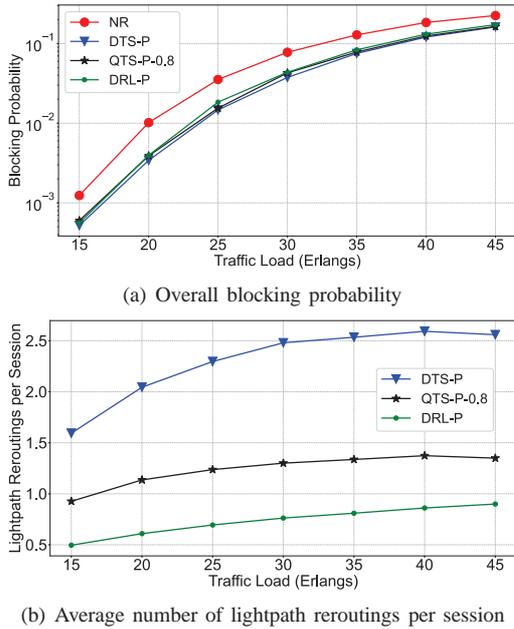


Fig. 11. Results of dynamic operations (RT, partial rearrangement).

TABLE II

AVERAGE RUNNING TIME PER MULTICAST RECONFIGURATION (SECONDS)

Topology	NSFNET	USB	EUB	RT
DRL-P	0.1943	0.2494	0.2383	0.2908
QTS-P	0.2964	0.3941	0.4040	0.3986
DTS-P	0.0693	0.0900	0.1016	0.0867
DRL-F	0.2104	0.2823	0.3022	0.3110
QTS-F	0.3620	0.4648	0.4927	0.4448
DTS-F	0.1995	0.2723	0.2882	0.2687

comparable and short enough to adapt to dynamic operations. The running time of our DRL model is less than that of the QTS-based algorithm in all the simulation scenarios, while as the DTS-based algorithm only makes decisions according to the depth of each logical light-tree, it runs the fastest. Meanwhile, the running time of each algorithm generally increases with the size of the topology, or from using partial rearrangement to using full rearrangement.

E. Generalization to Various EON Settings

Finally, we consider more simulation settings to verify that our proposed graph-aware DRL model can adapt to various EON settings. First of all, we notice that the assumption of Poisson traffic model might not hold in today's Internet. Hence, we design a new simulation scenario, in which the multicast sessions are generated dynamically in a bursty manner, *i.e.*, they come in according to the realistic ON/OFF pattern for bursty Internet traffic [36]. Note that, we still quantify the traffic load of the multicast sessions with Erlangs, *i.e.*, the production of the average number of new sessions per unit-time and the average lifetime of each session in time-units. The results of the simulations with NSFNET are shown in Fig. 12, and by comparing them with those in Fig. 5, we can see the similar trends. Meanwhile, as the bursty traffic model is more

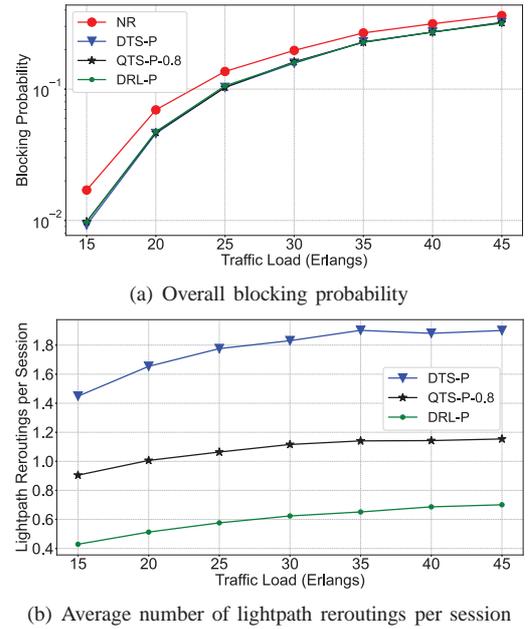
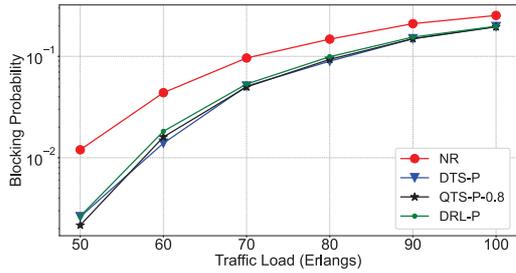


Fig. 12. Results of dynamic operations with bursty multicast sessions (NSFNET, partial rearrangement).

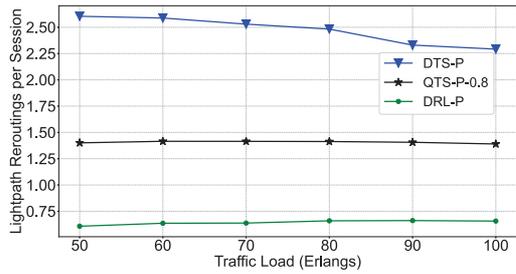
likely to cause session blockings, the blocking probability of each algorithm in Fig. 12 is higher. Nevertheless, our DRL model still retains its advantage of significantly reducing the number of reconfiguration operations without sacrificing the performance on blocking probability. With the bursty traffic model, we also simulate other EON topologies and test the algorithms with full rearrangement, while the results always follow the similar trends as those in Fig. 12.

Secondly, we increase the number of FS' on each fiber link to 200, for simulating the EONs with more spectrum resources. The results of the simulations with NSFNET are shown in Fig. 13, and by comparing them with those in Fig. 5, we still see the similar trends. Meanwhile, since there are more spectrum resources in the EON, we need to increase the traffic load to see the same blocking probability. Our DRL model still exhibits the advantages over the heuristics, which suggests that its performance is not affected by the change of spectrum resources in the EON. With the new setting of spectrum resources, we also simulate other EON topologies and test the algorithms with full rearrangement, and the results always follow the similar trends as those in Fig. 13.

Finally, considering the fact that in a real-world EON, there are unicast and anycast lightpaths coexisting with multicast sessions, we design a realistic simulation scenario that unicast and anycast lightpaths are used as the background traffic of multicast sessions. Specifically, to create a stressful scenario for our DRL model, we make the total bandwidth demands of unicast, anycast, and multicast account for 25%, 25% and 50% of the overall bandwidth usage in the EON, respectively. The results of the simulations with NSFNET are shown in Fig. 14, and by comparing them with those in Fig. 5, we can see that the blocking probability of multicast sessions becomes lower. This is because for the same traffic load, unicast and anycast lightpaths generally require less spectrum resources

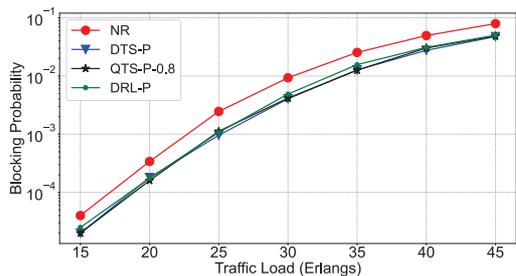


(a) Overall blocking probability

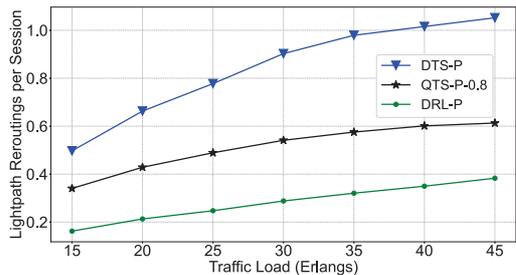


(b) Average number of lightpath reroutings per session

Fig. 13. Results of dynamic operations with 200 FS per fiber link (NSFNET, partial rearrangement).



(a) Overall blocking probability



(b) Average number of lightpath reroutings per session

Fig. 14. Results of dynamic operations with unicast/anycast background traffic (NSFNET, partial rearrangement).

than multicast sessions, and thus the total spectrum usage is actually smaller. Meanwhile, for the same reason, the gap on blocking probability between the multicast reconfiguration algorithms and the case without multicast reconfiguration (*i.e.*, NR) becomes smaller too. Note that, compared with the QTS-based and DTS-based algorithms, our DRL model still invokes a smaller number of lightpath reroutings per session to maintain almost the same blocking probability. This verifies the effectiveness of our DRL model in the EON environment that contains mixed types of traffic demands.

VI. CONCLUSION

In this work, we revisited the problem of how to formulate and reconfigure multicast sessions in an EON, and proposed a DRL model based on GNNs that can solve the sub-problem of multicast session selection in a more universal and adaptive way. Specifically, we abstracted the state information of each multicast session as graph-structured data, which can be directly analyzed by our graph-aware DRL model. Then, the graph-based reasoning capability of our proposal made sure that the state information of each multicast session can be analyzed in depth for dynamic reconfiguration, and facilitated the universality across different topologies. Hence, an important takeaway is that our graph-aware design of the DRL model made its architect and operation independent of the EON's topology, and thus avoided the hassle of redesigning its architecture to adapt to different EON topologies.

Simulation results verified that compared with the existing deterministic algorithms based on DTS and QTS, our graph-aware DRL based approach can significantly reduce the average lightpath reroutings per multicast session while maintaining the overall blocking probability approximately at the same level. This suggested that our proposal can balance the tradeoff between the number of reconfiguration operations and blocking performance much better than the existing algorithms. Moreover, our simulations also confirmed that the DRL model trained in one EON environment can easily adapt to solve the problem of dynamic multicast session reconfiguration in EONs with various settings (*e.g.*, different topologies, spectrum resources, traffic models and request types). Therefore, the universality of our proposal helped to effectively save the time and efforts that are needed to adjust the DRL model according to an EON's setting, and provided a more realistic solution for network automation.

ACKNOWLEDGMENTS

This work was supported in part by the National Key R&D Program of China (2020YFB1806400), NSFC project 61871357, SPR Program of CAS (XDC02070300), and Fundamental Funds for Central Universities (WK3500000006).

REFERENCES

- [1] Cisco Visual Networking Index, 2017-2022. [Online]. Available: <https://www.cisco.com/c/en/us/solutions/collateral/service-provider/visual-networking-index-vni/white-paper-c11-741490.html>.
- [2] P. Lu, L. Zhang, X. Liu, J. Yao, and Z. Zhu, "Highly efficient data migration and backup for Big Data applications in elastic optical inter-data-center networks," *IEEE Netw.*, vol. 29, pp. 36–42, Sept./Oct. 2015.
- [3] N. Laoutaris, M. Sirivianos, X. Yang, and P. Rodriguez, "Inter-datacenter bulk transfers with NetStitcher," in *Proc. of ACM SIGCOMM 2011*, pp. 74–85, Aug. 2012.
- [4] V. Dukic, C. Gkantsidis, T. Karagiannis, F. Parmigiani, A. Singla, M. Filer, J. Cox, A. Ptasznik, N. Harland, W. Saunders, and C. Belady, "Beyond the mega-data center: networking multi-data center regions," in *Proc. of ACM SIGCOMM 2020*, pp. 765–781, Aug. 2020.
- [5] O. Gerstel, M. Jinno, A. Lord, and B. Yoo, "Elastic optical networking: a new dawn for the optical layer?" *IEEE Commun. Mag.*, vol. 50, pp. s12–s20, Feb. 2012.
- [6] Z. Zhu, W. Lu, L. Zhang, and N. Ansari, "Dynamic service provisioning in elastic optical networks with hybrid single-/multi-path routing," *J. Lightw. Technol.*, vol. 31, pp. 15–22, Jan. 2013.

- [7] L. Gong and Z. Zhu, "Virtual optical network embedding (VONE) over elastic optical networks," *J. Lightw. Technol.*, vol. 32, pp. 450–460, Feb. 2014.
- [8] L. Sahasrabudde and B. Mukherjee, "Light trees: optical multicasting for improved performance in wavelength routed networks," *IEEE Commun. Mag.*, vol. 37, pp. 67–73, Feb. 1999.
- [9] Q. Wang and L. Chen, "Performance analysis of multicast traffic over spectrum elastic optical networks," in *Proc. of OFC 2012*, pp. 1–3, Mar. 2012.
- [10] X. Liu, L. Gong, and Z. Zhu, "On the spectrum-efficient overlay multicast in elastic optical networks built with multicast-incapable switches," *IEEE Commun. Lett.*, vol. 17, pp. 1860–1863, Sept. 2013.
- [11] L. Gong, X. Zhou, X. Liu, W. Zhao, W. Lu, and Z. Zhu, "Efficient resource allocation for all-optical multicasting over spectrum-sliced elastic optical networks," *J. Opt. Commun. Netw.*, vol. 5, pp. 836–847, Aug. 2013.
- [12] X. Liu, L. Gong, and Z. Zhu, "Design integrated RSA for multicast in elastic optical networks with a layered approach," in *Proc. of GLOBECOM 2013*, pp. 2346–2351, Dec. 2013.
- [13] K. Walkowiak, R. Goscien, M. Klinkowski, and M. Wozniak, "Optimization of multicast traffic in elastic optical networks with distance-adaptive transmission," *IEEE Commun. Lett.*, vol. 18, pp. 2117–2120, Dec. 2014.
- [14] Z. Zhu, X. Liu, Y. Wang, W. Lu, L. Gong, S. Yu, and N. Ansari, "Impairment- and splitting-aware cloud-ready multicast provisioning in elastic optical networks," *IEEE/ACM Trans. Netw.*, vol. 25, pp. 1220–1234, Apr. 2017.
- [15] A. Mahimkar, A. Chiu, R. Doverspike, M. Feuer, P. Magill, E. Mavrogiorgis, J. Pastor, S. Woodward, and J. Yates, "Bandwidth on demand for inter-data center communication," in *Proc. of HOTNETS 2011*, pp. 1–6, Nov. 2011.
- [16] A. Malis, B. Wilson, G. Clapp, and V. Shukla, "Requirements for very fast setup of GMPLS label switched paths (LSPs)," *RFC 7709*, Nov. 2015. [Online]. Available: <https://tools.ietf.org/html/rfc7709>.
- [17] A. Castro, L. Velasco, M. Ruiz, M. Klinkowski, J. Fernandez-Palacios, and D. Careglio, "Dynamic routing and spectrum (re) allocation in future flexgrid optical networks," *Comput. Netw.*, vol. 56, pp. 2869–2883, Aug. 2012.
- [18] M. Klinkowski, M. Ruiz, L. Velasco, D. Careglio, V. Lopez, and J. Comellas, "Elastic spectrum allocation for time-varying traffic in flexgrid optical networks," *IEEE J. Sel. Areas Commun.*, vol. 31, pp. 26–38, Dec. 2012.
- [19] Y. Yin, H. Zhang, M. Xia, Z. Zhu, S. Dahlfors, and S. Yoo, "Spectral and spatial 2D fragmentation-aware routing and spectrum assignment algorithms in elastic optical networks," *J. Opt. Commun. Netw.*, vol. 5, pp. A100–A106, Oct. 2013.
- [20] Y. Sone, A. Watanabe, W. Imajuku, Y. Tsukishima, B. Kozicki, H. Takara, and M. Jinno, "Bandwidth squeezed restoration in spectrum-sliced elastic optical path networks (SLICE)," *J. Opt. Commun. Netw.*, vol. 3, pp. 223–233, Mar. 2012.
- [21] W. Lu and Z. Zhu, "Malleable reservation based bulk-data transfer to recycle spectrum fragments in elastic optical networks," *J. Lightw. Technol.*, vol. 33, pp. 2078–2086, May 2015.
- [22] M. Zeng, Y. Li, W. Fang, W. Lu, X. Liu, H. Yu, and Z. Zhu, "Control plane innovations to realize dynamic formulation of multicast sessions in inter-DC software-defined elastic optical networks," *Opt. Switch. Netw.*, vol. 23, pp. 259–269, Jan. 2017.
- [23] R. Gu, Z. Yang, and Y. Ji, "Machine learning for intelligent optical networks: A comprehensive survey," *J. Netw. Comput. Appl.*, vol. 157, pp. 1–22, May 2020.
- [24] F. Scarselli, M. Gori, A. Tsoi, M. Hagenbuchner, and G. Monfardini, "The graph neural network model," *IEEE Trans. Neur. Net.*, vol. 20, pp. 61–80, Jan. 2009.
- [25] J. Zhou, G. Cui, S. Hu, Z. Zhang, C. Yang, Z. Liu, L. Wang, C. Li, and M. Sun, "Graph neural networks: A review of methods and applications," *AI Open*, vol. 1, pp. 57–81, Jan. 2021.
- [26] A. Ding and G. Poo, "A survey of optical multicast over WDM networks," *Comput. Commun.*, vol. 26, pp. 193–200, Feb. 2003.
- [27] K. Christodoulopoulos, I. Tomkos, and E. Varvarigos, "Elastic bandwidth allocation in flexible OFDM-based optical networks," *J. Lightw. Technol.*, vol. 29, pp. 1354–1366, May. 2011.
- [28] P. Zhu, J. Li, Y. Chen, X. Chen, Z. Wu, D. Ge, Z. Chen, and Y. He, "Experimental demonstration of EON node supporting reconfigurable optical superchannel multicasting," *Opt. Express*, vol. 23, pp. 20495–20504, Aug. 2015.
- [29] X. Li, L. Zhang, J. Wei, and S. Huang, "Deep neural network based OSNR and availability predictions for multicast light-trees in optical WDM networks," *Opt. Express*, vol. 27, pp. 10648–10669, 2020.
- [30] K. Rusek, J. Suarez-Varela, A. Mestres, P. Barlet-Ros, and A. Cabellos-Aparicio, "Unveiling the potential of graph neural networks for network modeling and optimization in SDN," in *Proc. of SOSR 2019*, pp. 140–151, Oct. 2019.
- [31] P. Sun, J. Lan, J. Li, Z. Guo, Y. Hu, and T. Hu, "Efficient flow migration for NFV with graph-aware deep reinforcement learning," *Comput. Netw.*, vol. 183, p. 107575, Sept. 2020.
- [32] N. Sambo, P. Castoldi, A. D'Errico, E. Riccardi, A. Pagano, M. Moreolo, J. Fabrega, D. Rafique, A. Napoli, S. Frigerio, E. Salas, G. Zervas, M. Nolle, J. Fischer, A. Lord, and J. Gimenez, "Next generation sliceable bandwidth variable transponders," *IEEE Commun. Mag.*, vol. 53, no. 2, pp. 163–171, 2015.
- [33] M. Defferrard, X. Bresson, and P. Vandergheynst, "Convolutional neural networks on graphs with fast localized spectral filtering," *arXiv preprint arXiv:1606.09375*, pp. 1–9, Feb. 2017.
- [34] V. Mnih, A. Badia, M. Mirza, A. Graves, T. Lillicrap, T. Harley, D. Silver, and K. Kavukcuoglu, "Asynchronous methods for deep reinforcement learning," in *Proc. of ICML 2016*, pp. 1928–1937, Jun. 2016.
- [35] Z. Zhang and V. Saligrama, "Zero-shot learning via semantic similarity embedding," in *Proc. of ICCV 2015*, pp. 4166–4174, Dec. 2015.
- [36] X. Yang, "Designing traffic profiles for bursty Internet traffic," in *Proc. of GLOBECOM 2002*, pp. 1–6, Dec. 2002.