DRL-based Network Orchestration to Realize Cooperative, Distributed and Tenant-driven Virtual Network Slicing

(*Invited Paper*)

Xu Zhang, Wei Lu, Baojia Li, Zuqing Zhu

School of Information Science and Technology, University of Science and Technology of China, Hefei, China, Email: zqzhu@ieee.org

Abstract: We propose a cooperative, distributed and tenant-driven framework for realizing virtual network embedding (VNE) in optical datacenter interconnections (O-DCIs), by leveraging a deep reinforcement learning (DRL) module.

OCIS codes: (060.1155) Software-defined optical networks; (060.4251) Networks, assignment and routing algorithms.

1. Introduction

Nowadays, the rapid development of cloud computing has made both the architecture of optical datacenter interconnections (O-DCIs) [1] and the network virtualization in them [2, 3] attractive research topics. This is because network virtualization is a must-have enabling technology for cloud computing [4, 5], to divide the role of conventional Internet service provider (ISP) into infrastructure provider (InP) and service provider (SP). Specifically, an InP manages a substrate network (SNT) and accepts virtual network (VNT) requests from multiple SPs (*i.e.*, tenants), while each tenant can submit requests to lease substrate resources from the InP and customize its own VNT with them in the pay-as-you-go manner. Recent advances on flexible-grid elastic optical networks (EONs) [6–8] have made the optical layer of O-DCIs more adaptive for network virtualization. However, how to provision the VNT requests from tenants time-efficiently and cost-effectively is still challenging. This is because the problem of virtual network embedding (VNE) is known to be NP-hard [9], and to the best of our knowledge, most of the existing approaches to solve it rely solely on the computing power of the InP. Moreover, the existing approaches obtain VNE schemes based on the current network status without the intelligence of looking ahead. These dilemmas inspire us to propose a cooperative, distributed and tenant-driven VNE framework, which not only relieves the computing pressure on an InP but also motivates the tenants to request for substrate resources in the load-balanced manner.

In this paper, we describe our VNE framework that was initially proposed in [10], and elaborate on our new efforts to enhance its performance. Simulation results show that by leveraging deep reinforcement learning (DRL), our VNE framework can utilize substrate resources wisely to provision VNT requests time-efficiently and cost-effectively.

2. Cooperative, Distributed and Tenant-driven VNE Framework

Fig. 1(a) illustrates our proposed VNE framework, whose basic idea is to involve tenants in the VNE calculation. Specifically, in each service cycle, the framework operates in four steps. Firstly, the InP collects the current network status regarding the in-service VNTs and the resource usages on substrate fiber links (SLs) and substrate DC nodes (SNs), and utilizes a DRL module to price the substrate resources on SLs (*i.e.*, optical spectra) and SNs (*i.e.*, IT resources on DCs) based on the network status. Here, the resource pricing is introduced to motivate the tenants to

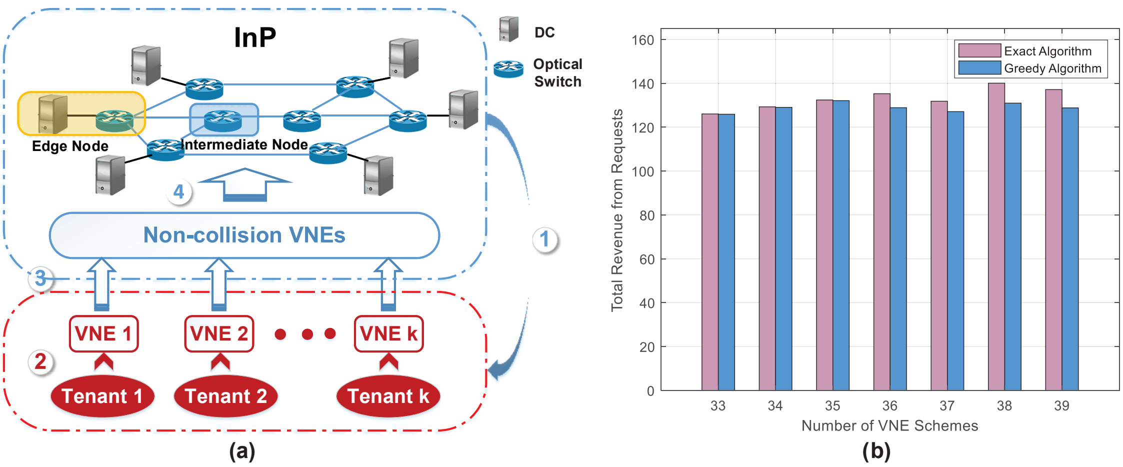

Fig. 1. (a) Architecture and operation principle of our VNE framework, and (b) Performance comparison of greedy algorithm and exact algorithm for selecting VNE schemes to provision.

demand for substrate resources in the load-balanced way. Note that, both the algorithmic [2–5, 9, 11] and experimental [12–14] studies on the VNE problem have concluded that balancing the resource utilization in an SNT can effective reduce the blocking probability of VNT requests. The information about the available substrate resources and their prices is then broadcasted to all the tenants. Secondly, based on the information provided by the InP, the tenants distributedly calculate the VNE schemes of their VNTs with the objective of minimizing VNE cost, and then submit their results to the InP. Thirdly, after collecting all the VNE schemes, the InP uses an algorithm to accept some or all of them under the resource capacity constraint. Finally, before moving to the next service cycle, the InP provisions all the collision-free VNE schemes obtained in the previous step, and updates the network status of the SNT.

3. Algorithm Design

We model the topology of an O-IDC as $G(V,E)$ (*i.e.*, the SNT), where V and E represent the sets of SNs and SLs, respectively. As shown in Fig. 1(a), there are two kinds of SNs in the O-IDC, which are the edge and intermediate nodes, respectively. Here, an edge node includes a DC and an optical switch, while an intermediate node only contains an optical switch. Then, the VNE calculation of a VNT needs to solve two subproblems: 1) selecting a group of edge nodes in the SNT to embed the virtual nodes (VNs) such that all the IT resource requirements of the VNT can be satisfied, and 2) mapping the virtual links (VLs) to substrate paths and satisfying their bandwidth requirements with proper spectrum allocations on the related SLs. Our framework lets each tenant to calculate its own VNE scheme distributedly, and this effectively restricts the size of each VNE problem. Therefore, we can leverage an existing integer linear programming (ILP) model [2] to obtain the exact solution for each tenant.

For the third step in Fig. 1(a), the InP needs to select the most profitable VNE schemes to provision, if not all the VNE schemes submitted by the tenants can be accommodated in the SNT due to resource capacity constraint. To solve this problem, we first construct an auxiliary graph (AG), in which each node corresponds to a VNE scheme, two nodes are connected if the corresponding VNE schemes cannot be simultaneously served because of the resource capacity constraint, and the weight of each node is just the cost of its VNE scheme. Therefore, the original problem gets transformed to finding the maximum weighted independent set in the AG, which can solved by utilizing a time-efficient greedy algorithm [15]. Fig. 1(b) and Table 1 compare the performance of the greedy algorithm and an exact algorithm based on exhaustive searching, in terms of total revenue from requests and running time. We can see that the greedy algorithm provides similar results on the total revenue with much shorter running time.

# of VNE schemes	33	34	35	36	37	38	39
Exact algorithm (sec)	0.00103	0.00442	0.02089	0.56306	19.567	144.19	1031.1
Greedy algorithm (sec)	0.00005	0.00005	0.00006	0.00032	0.00109	0.00056	0.0016

Table 1. Running time of algorithms to select VNE schemes for provisioning.

After finalizing and provisioning the collision-free VNE schemes, the InP utilizes a DRL module to adjust the prices of substrate resources based on the network status before the next service cycle starts. Here, we design the DRL module based on deep deterministic policy gradient algorithm (DDPG) [16], which is known to be powerful on optimizing actions in states with high-dimensional and continuous space. More specifically, the DRL is designed as follows.

- **State**: the resource usage on each substrate element in the SNT (*i.e.*, either the spectrum usage on an SL or the IT resource usage on a DC), and the abstracted information of pending VNT requests in the next service cycle.
- **Action**: increasing or decreasing the unit price of the resource on each substrate element by a margin which values from 0 to 1 , *i.e.*, the action-space is continuous within $[-1,1]$ for each substrate element.
- **Reward**: the acceptance ratio of VNT requests, which is the ratio of the total resources of the accepted requests to the total resources of all the requests coming in the previous service cycle.

4. Performance Evaluation

We run simulations with the 8-node O-IDC topology in Fig. 1(a) to evaluate the performance of our proposed framework. Here, we assume that at beginning of each service cycle, the number of tenants to submit VNT requests is uniformly distributed within $[20,25]$, the number of VNs in each VNT request is randomly selected within $[2,4]$, each VN pair is connected with a VL with a probability of 0.6, and the hold-on time of each VNT request is uniformly distributed within $[1,3]$ service cycles. As the benchmark, we also consider an adaptive price adjusting scheme, which sets the resource price of a substrate element according to $p_n = \frac{p_{n-1}}{1.5-u_{n-1}}$, where p_n and p_{n-1} are the resource prices in this and previous service cycle, respectively, and u_{n-1} is the resource utilization on the substrate element at the beginning of this service cycle. Each simulation runs for 100 service cycles, and we consider two scenarios, *i.e.*, 1) O-IDC with sufficient substrate resources, and 2) reducing substrate resources of the O-IDC in 1) by half.

Fig. 2 shows the algorithms' performance on cumulative blocking probability over time. It can be seen that the DRL-based pricing scheme achieves much better performance on blocking probability in both scenarios, and its performance is more stable too, *i.e.*, its cumulative blocking probability does not exhibit dramatic variations as that from the adaptive pricing scheme. This is because our DLR module possesses more intelligence to adjust resource prices such that the tenants are motivated to request for substrate resources in a more load-balanced manner.

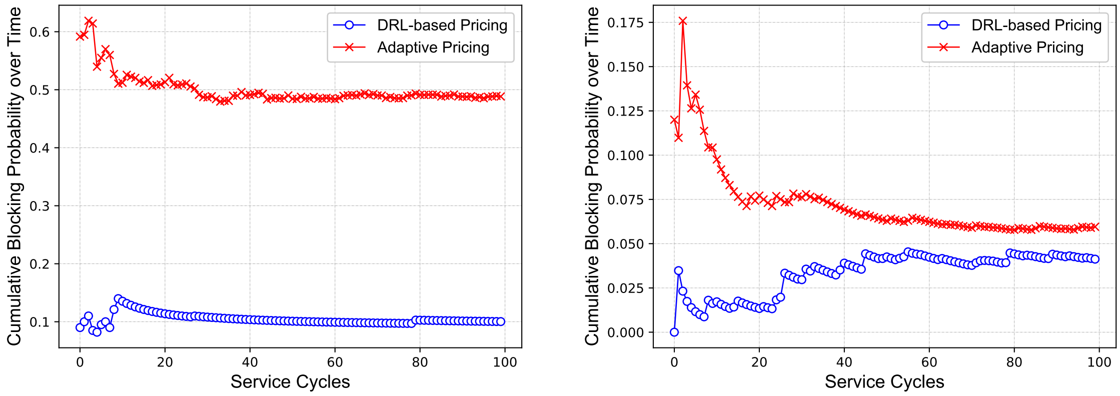

(a) O-IDC with sufficient substrate resources

(b) Reducing substrate resources of the O-IDC in (a) by half

Fig. 2. Cumulative blocking probability *versus* service cycles.

5. Conclusion

In this work, we proposed a cooperative, distributed and tenant-driven VNE framework, which not only relieves the computing pressure on an InP but also motivates the tenants to request for substrate resources in the load-balanced manner. Both the system architecture and the algorithm design of the framework were discussed to explain our efforts on ensuring high service provisioning performance. Simulation results indicated that by leveraging DRL, our proposal can utilize substrate resources wisely to provision VNT requests time-efficiently and cost-effectively.

Meanwhile, we hope to point out that even though our proposal introduces distributed calculation for VNE schemes, the InP still acts as a centralized arbiter to manage the whole O-IDC. Therefore, our proposed framework can easily fit into the appealing software-defined networking (SDN) architecture for O-IDCs [17, 18].

References

[1] P. Lu *et al.*, "Highly-efficient data migration and backup for Big Data applications in elastic optical inter-data-center networks," *IEEE Netw.*, vol. 29, no. 5, pp. 36-42, Sept./Oct. 2015.

[2] L. Gong and Z. Zhu, "Virtual optical network embedding (VONE) over elastic optical networks," *J. Lightw. Technol.*, vol. 32, no. 3, pp. 450-460, Feb. 2014.

[3] L. Gong *et al.*, "Toward profit-seeking virtual network embedding algorithm via global resource capacity," *in Proc. of INFOCOM 2014*, pp. 1-9, Apr. 2014.

[4] H. Jiang *et al.*, "Availability-aware survivable virtual network embedding (A-SVNE) in optical datacenter networks," *J. Opt. Commun. Netw.*, vol. 7, no. 12, pp. 1160-1171, Dec. 2015.

[5] L. Gong *et al.*, "Novel location-constrained virtual network embedding (LC-VNE) algorithms towards integrated node and link mapping," *IEEE/ACM Trans. Netw.*, vol. 24, no. 6, pp. 3648-3661, Dec. 2016.

[6] Z. Zhu *et al.*, "Dynamic service provisioning in elastic optical networks with hybrid single-/multi-path routing," *J. Lightw. Technol.*, vol. 31, no. 1, pp. 15-22, Jan. 2013.

[7] L. Gong *et al.*, "Efficient resource allocation for all-optical multicasting over spectrum-sliced elastic optical networks," *J. Opt. Commun. Netw.*,, vol. 5, no. 8, pp. 836-847, Aug. 2013.

[8] Y. Yin *et al.*, "Spectral and spatial 2D fragmentation-aware routing and spectrum assignment algorithms in elastic optical networks," *J. Opt. Commun. Netw.*, vol. 5, no. 10, pp. A100-A106, Oct. 2013.

[9] M. Chowdhury *et al.*, "Vineyard: Virtual network embedding algorithms with coordinated node and link mapping," *IEEE/ACM Trans. Netw.*, vol. 20, no. 1, pp. 206-219, Feb. 2012.

[10] W. Lu *et al.*, "AI-assisted resource advertising and pricing to realize distributed tenant-driven virtual network slicing in inter-DC optical networks," *in Proc. of ONDM 2018*, pp. 1-6, May 2018.

[11] W. Fang *et al.*, "Joint defragmentation of optical spectrum and IT resources in elastic optical datacenter interconnections," *J. Opt. Commun. Netw.*, vol. 7, no. 3, pp. 314-324, Mar. 2015.

[12] J. Yin *et al.*, "Experimental demonstration of building and operating QoS-aware survivable vSD-EONs with transparent resiliency," *Opt. Express*, vol. 25, no. 13, pp. 15468-15480, 2017.

[13] Z. Zhu *et al.*, "Build to tenants' requirements: On-demand applicationdriven vSD-EON slicing," *J. Opt. Commun. Netw.*, vol. 10, no. 2, pp. A206-A215, Feb. 2018.

[14] B. Kong *et al.*, "Demonstration of application-driven network slicing and orchestration in optical/packet domains: On-demand vDC expansion for Hadoop MapReduce optimization," *Opt. Express*, vol. 26, no. 11, pp. 14066-14085, 2018.

[15] S. Sakai *et al.*, "A note on greedy algorithms for the maximum weighted independent set problem," *Discrete Appl. Math.*, vol. 126, no. 2-3, pp. 313-322, Mar. 2003.

[16] T. Lillicrap *et al.*, "Continuous control with deep reinforcement learning," *arXiv:1509.02971*, Feb. 2016. [Online]. Available: https://arxiv.org/abs/1509.02971

[17] Z. Zhu *et al.*, "Demonstration of cooperative resource allocation in an OpenFlow-controlled multidomain and multinational SD-EON testbed," *J. Lightw. Technol.*, vol. 33, no. 8, pp. 1508-1514, Apr. 2015.

[18] X. Chen *et al.*, "Flexible availability-aware differentiated protection in software-defined elastic optical networks," *J. Lightw. Technol.*, vol. 33, no. 18, pp. 3872-3882, Sept. 2015.