Towards Real-Time Routing Optimization with Deep Reinforcement Learning: Open Challenges

Paul Almasan*, José Suárez-Varela*, Bo Wu[†], Shihan Xiao[†], Pere Barlet-Ros*, Albert Cabellos-Aparicio*

*Barcelona Neural Networking Center, Universitat Politècnica de Catalunya, Spain

[†]Network Technology Lab., Huawei Technologies Co.,Ltd.

NOTE: This is a version of a paper accepted in IEEE HPSR SARNET 2021. Please use the following reference to cite the corresponding work: P. Almasan, J. Suárez-Varela, B. Wu, S. Xiao, P. Barlet-Ros and A. Cabellos-Aparicio, "Towards Real-Time Routing Optimization with Deep Reinforcement Learning: Open Challenges," 2021 IEEE 22nd International Conference on High Performance Switching and Routing (HPSR), 2021, pp. 1-6, doi: 10.1109/HPSR52026.2021.9481864.

Abstract—The digital transformation is pushing the existing network technologies towards new horizons, enabling new applications (e.g., vehicular networks). As a result, the networking community has seen a noticeable increase in the requirements of emerging network applications. One main open challenge is the need to accommodate control systems to highly dynamic network scenarios. Nowadays, existing network optimization technologies do not meet the needed requirements to effectively operate in real time. Some of them are based on hand-crafted heuristics with limited performance and adaptability, while some technologies use optimizers which are often too time-consuming. Recent advances in Deep Reinforcement Learning (DRL) have shown a dramatic improvement in decision-making and automated control problems. Consequently, DRL represents a promising technique to efficiently solve a variety of relevant network optimization problems, such as online routing. In this paper, we explore the use of state-of-the-art DRL technologies for real-time routing optimization and outline some relevant open challenges to achieve production-ready DRL-based solutions.

I. INTRODUCTION

The digital transformation is driving existing network technologies towards supporting emergent networking paradigms (e.g., vehicular networks [1], smart cities [2], satellite networks [3]). These modern applications impose new requirements for existing network optimization technologies (e.g., adaptability to dynamic scenarios, low latency). In this context, Machine Learning (ML) seems a good candidate to enable real-time operation and allow efficient optimization across dynamic systems [4]. In the present paper, we explore the application of modern ML-based solutions for routing optimization in dynamic network scenarios.

Traditional solutions for routing optimization can be mainly differentiated in two different categories: (i) Specific-purpose heuristics and (ii) Computationally-intensive mathematical solvers, which intrinsically are not adequate for efficient real-time optimization. Heuristic-based solutions (e.g., shortest path, load balancing) are often sufficiently fast, but they require expert knowledge to be adapted to specific optimization problems, and often do not adapt well to dynamic network scenarios. Solutions based on general-purpose mathematical optimizers (e.g., CPLEX [5]) are often too time-consuming to achieve real-time operation. This pushes the networking

community to explore new methods and technologies that can adapt well to the requirements of emerging network applications.

Recent advances in Deep Reinforcement Learning (DRL) showcased that this technology is capable to operate efficiently in complex optimization problems [4], [6]. However, state-of-the-art DRL-based solutions used in other fields are not directly applicable to modern networks, where it is needed to generalize across scenarios with dynamic changes (e.g., topology, traffic, link failures) [7]. One key technology that enables DRL to handle network changes is Graph Neural Network (GNN) [8]. GNN is a novel family of neural networks specifically intended to operate on graph-structured data. Particularly, these recent neural networks have a strong relational bias over graphs, which results in an unprecedented level of combinatorial generalization over graph-structured data [9]. As a result, GNNs achieve outstanding generalization capabilities over networks [7], [10] and, similarly to other deep learning methods, they achieve low execution times – often sub-second operation in commodity hardware.

In this paper, we explore the applicability of state-of-the-art solutions based on DRL and GNN to modern networking scenarios with highly dynamic topology changes. Specifically, we design a DRL+GNN architecture for online routing optimization, and discuss the open challenges of this technology. First, we train a GNN-based DRL agent in a single topology, and then we test if the agent is able to operate successfully on different topologies never seen before. This allows us to assess the capabilities of this solution to adapt to fundamental changes in the network topology. Lastly, we discuss the implications and challenges that represents the new paradigm of DRL+GNN.

II. RELATED WORK

Finding the optimal routing configuration given an estimated traffic matrix is a fundamental networking problem, which is known to be NP-hard [11], [12]. This problem has been largely studied in the past and we outline some of the most relevant works. In DEFO [12], the authors propose a solution that converts high-level optimization goals, indicated by the network operator, into specific routing configurations using

constraint programming. In [13], the authors propose to use local search, where they sacrifice space exploration to achieve lower execution times. A more recent work [14] leverages the column generation algorithm and dynamic programming to solve Traffic Engineering problems.

Recently, numerous DRL-based solutions have been proposed to solve network optimization problems. In [15], they propose a generic DRL framework for Traffic Engineering. Their solution is based on a DRL agent that defines the split ratios of a flow over a set of paths. In the field of optical networks, the work [16] proposes an elaborated representation of the network state to help a DRL agent learn to route traffic demands. A more recent work [7] proposes a different approach where the authors combine DRL and GNN to optimize the resource allocation in optical circuit-switched networks. The recent work from [17] proposes a scalable DRL-based solution where they use pinning control theory to select critical links in the network and optimize their weights assuming a weighted shortest path routing policy (e.g., OSPF and ECMP).

III. ROUTING OPTIMIZATION SCENARIO

In this paper we explore the challenges when using DRL-based solutions for routing optimization in dynamic network scenarios. Particularly, we consider a network optimization scenario based on Software-Defined Networking (SDN), where the DRL agent (located in the control plane) has a global view of the current network state, and has to optimize the routing configuration considering the current traffic matrix. This is a complex optimization problem that has been widely studied in the past [12], [14].

All the source-destination traffic demands are initially allocated using OSPF routing [18] and the DRL agent must find a new routing policy that minimizes the utilization of the most loaded link. This is a common scenario where overlay technologies (e.g., Segment Routing [19], LISP [20]) are used in the optimization layer [12]. Traffic demands are defined by the tuple $\{src, dst, bw\}$ where src and dst are the source and destination nodes respectively and bw represents the average traffic volume estimated for this demand. The DRL agent iterates over all demands and it has to decide the new routing path for each of them. Considering all the possible src - dst combinations in the network, there is a total of $N*(N-1)$ traffic demands, where N is the number of nodes (i.e., one demand per src - dst pair). A DRL episode ends when the agent has iterated over all traffic demands. Figure 1 shows a schematic representation of the optimization scenario where the DRL agent interacts with the environment. This figure also illustrates how the DRL agent changes the routing policy of a traffic demand.

Traffic demands do not expire during the episode. This implies a challenging task for the DRL agent, since it has not only to identify critical resources on networks (e.g., potential bottlenecks), but also to plan ahead and optimize considering future demands. We summarize below the main features of the SDN-based routing optimization scenario considered in this paper:

Fig. 1: Schematic representation of the DRL agent in the routing optimization scenario.

- The DRL agent makes sequential routing decisions for every traffic demand.
- It selects per-demand routing paths in order to minimize the most loaded link.
- Once a routing decision is made for a traffic demand, it cannot be re-routed, at the demand remains until the end of the episode.

IV. DRL+GNN ROUTING OPTIMIZATION SOLUTION

We implement a DRL-based solution for routing optimization that integrates a GNN. We adapt the DRL+GNN solution from [7] for solving the routing problem addressed in the present paper (Sec. III). While the solution from [7] is tailored for routing optimization in optical networks, in this work we design a DRL+GNN agent that is tailored for Traffic Engineering optimization in IP networks. Our agent implements the Proximal Policy Optimization (PPO) algorithm [21], which is a DRL on-policy algorithm based on actor-critic methods, in difference with the off-policy DQN algorithm used in the previous work. We also implement a simulation environment based on a fluid model, as the one used in [12].

The DRL agent's training process is based on a trial-and-error process. At each time step, the agent receives a network state observation from the simulation environment, including the current link's utilization. Then, the DRL agent uses a GNN to construct an internal graph representation with the topology links. Each link has associated a state represented by a fixed-size vector with some real values (Sec. IV-A). With this representation, a message passing algorithm runs between the graph elements (i.e., the links of the network topology) according to the topology structure, and updates the link states [22].

The outputs of this algorithm (i.e., the new links states) are aggregated into a global state, which encodes information about the whole network, and this global embedding is then processed by a fully-connected NN. Particularly, the GNN outputs a set of probabilities over the possible actions of the DRL agent. Finally, the output probability distribution is used by the DRL agent to sample the action to perform.

Fig. 2: Action representation in the link states. Nodes 1 and 5 are the source and destination nodes respectively, and node 3 is the midpoint selected.

A. Network Simulation Environment

The network state is defined by features on the links, which includes the link capacity and the current utilization. These features are stored in a fixed-size vector padded with zeros. At the beginning of a DRL episode, all demands are allocated according to the OSPF routing (i.e., shortest path policy). The links states change as the DRL re-allocates traffic demands to specific *src-dst* paths (i.e., sequences of links in the network). The final goal is to minimize the maximum link utilization at the end of the episode (i.e., when the DRL agent has iterated over all the traffic demands).

B. Action space

The number of possible routing combinations for each traffic demand (i.e., *src-dst* node pairs) results in a high dimensional state/action space, even in small networks [23]. This makes the routing problem complex for the DRL agent, since it should estimate for each action which is the one that will lead to lower maximum utilization in the long-term. In other words, which is the optimal routing configuration per traffic demand that leads to the minimum maximum link utilization. Moreover, to exploit the generalization capabilities of GNNs over graphs, we need to define the action in a way that is invariant to edge and node permutation, using only link-level features rather than specific identifiers or labels.

To overcome these problems, we use the midpoint routing (MR) model [12] to define the action space. This model abstracts the routing paths between nodes as middlepoints, where the paths are pre-configured. In our paper, the pre-configured paths between nodes and middlepoints are computed using the OSPF routing policy. Given a traffic demand with an associated *src-dst* node pair, we have as many different actions as number of middlepoints. Particularly, we limited the options to selecting only one midpoint. This means that to reach a destination node *dst*, the traffic demand can only cross at most one midpoint before reaching the destination node. In our DRL agent, all the nodes except the source node can be midpoint nodes, and *src-dst* traffic demands can be directed through specific middlepoints using Segment Routing, as in [12].

To represent the action, we add an additional feature in the link's state (see Figure 2). This feature includes the traffic

Fig. 3: Message passing architecture.

volume of the current traffic demand bw in case the link is within the selected path. Otherwise, the links that are not included in the path have this feature set to zero. Features x_1 and x_2 from Figure 2 are the link capacity and link utilization respectively. In this example, the traffic to allocate of the current demand is 7.

C. GNN architecture

The GNN we implemented is based on the Message Passing Neural Network model [22]. In our case, we consider the network links as the graph entities and perform an iterative message passing process between adjacent links in the topology. This process consists of combining the state of all links with those of their neighbours in the graph. Particularly, the states of connected links are processed by a fully-connected NN. Then, the results (called *messages*) are aggregated for each link using an element-wise summation. Finally, the links' hidden states are individually updated using a Recurrent Neural Network (RNN), whose input is the previous link state and the new aggregated message. This message passing process is performed T times and, at the end of this phase, the resulting link states are aggregated using an element-wise summation. The result is then a general graph embedding that is passed through a fully-connected NN, which models the Readout function of the GNN. The output of this latter function is the action probability. Thus, given a traffic demand, this process is executed for each action (or possible midpoint) and the final action of the DRL agent is sampled over the individual actions' probabilities. Figure 3 represents the internal architecture of the GNN, including the iterative message passing phase and the final readout.

D. DRL Agent operation

The DRL agent operates by interacting with the network simulation environment (Sec. IV-A), and the learning process is implemented using PPO [21]. First, we collect experiences from the DRL agent interaction with the environment. Then, using the collected experiences, we improve the current policy weights using mini-batch updates. Afterwards, the old experiences are removed and we collect new experiences using the new policy weights. This process is repeated until convergence or a pre-defined criteria is reached.

Fig. 4: DRL+GNN evaluation on GBN (a) and GEANT2 (b) using 50 TMs.

At the beginning of an episode, all traffic demands are allocated using OSPF routing. Then, the DRL agent iterates over all traffic demands and tries to change their current routing configuration. For each traffic demand, the DRL agent evaluates the action of re-allocating the demand on each possible midpoint using the GNN. Notice that each midpoint corresponds to an action and the traffic demand cannot be split between different middlepoints. Then, the GNN outputs a probability over each action, and the action performed by the agent is sampled from the resulting probability distribution of all actions. The chosen action is then applied to the network, leading to a new state (i.e., new links utilization), a reward, and a flag indicating if the agent finished iterating over all traffic demands. We set the reward to be the difference in the maximum link utilization between two steps.

V. EXPERIMENTAL RESULTS

A. Evaluation setup

We implemented the DRL+GNN solution of Section IV using Tensorflow [24] and evaluated it on a network simulator implemented using the OpenAI Gym framework [25]. All the experiments were executed on *off-the-shelf* hardware without any specific hardware accelerator (Ubuntu 20.04.1 LTS with processor AMD Ryzen 9 3950X 16-Core Processor). The simulation environment is configured by initializing the link features for each topology with the respective link capacities. All paths between nodes and middlepoints are pre-computed using OSPF, considering equal weights for all the links. The traffic demands are initially allocated following the pre-computed paths. Then, the environment sorts the demands in decreasing order of traffic volume, and the DRL agent sequentially re-routes them using MR.

Initial experiments were performed to choose an appropriate gradient-based optimization algorithm and hyperparameter

values for the DRL+GNN architecture. We defined the links' hidden states h_l as 20-element vectors (filled initially with the link features as shown in Figure 2). In every execution of the GNN, we run $T=5$ message passing steps, and use mini-batches of 50 samples. The optimizer used is the Adam Optimizer [26], with an initial learning rate of $2 \cdot 10^{-4}$ and following an exponential learning rate decay during training.

B. Methodology

In our experiments, we use the NSFNet [27], GBN [28] and GEANT2 [29] topologies. The traffic demands were generated using a realistic gravity model from [30]. We consider four possible different values for the link capacities for the NSFNet topology and four different ones for the GBN and GEANT2 topologies. Specifically, the NSFNet has 5K, 10K, 15K and 20K as possible link capacities, and GBN and GEANT2 have 25K, 50K, 75K and 100K possible link capacities. We set the capacity individually on each link to ensure that the maximum link utilization is between 0.95 and 1.3 when applying the OSPF routing policy.

We compare the DRL's agent performance against two baselines. The first one is the Shortest Available Path (SAP). This consists on starting with an empty network, and iterate over the demands to allocate them on the shortest path that has more available capacity [23]. The second baseline we implemented is based on the Simulated Annealing [31] algorithm. Starting from an initial routing configuration using OSPF, the algorithm performs an iterative process trying to decrease the initial state energy (i.e., minimize the maximum link utilization). In each step, the algorithm explores all possible next states and it decides whether to move to a new state or stay in the same using a temperature parameter. This parameter will guide the algorithm to move to states with lower maximum link utilization. As more iterations the algorithm performs, more chances it will have to approach to the global optimum. Therefore, we parameterized the algorithm to perform $4 \cdot 10^6$ steps. This makes Simulated Annealing not suitable for real-time operation since its cost (in time) is too high for finding good routing configurations. Specifically, Simulated Annealing spent around 30 minutes and 24 minutes for optimizing the GEANT2 and GBN topologies respectively.

We trained the DRL+GNN agent solely on the 14-node NSFNet topology [27]. During the training process, we used 100 different Traffic Matrices (TM) and we evaluated the agent performance on 25 TMs never seen during training.

C. Evaluation

In the evaluation experiments, we compare our DRL+GNN agent against the baseline solutions. Particularly, we pick the model with more performance during the training phase, and evaluate it in two real-world network topologies: GBN (17 nodes) and GEANT2 (24 nodes). We used 50 different TMs to evaluate the performance for each topology.

In Figures 4a and 4b we can observe the evaluation results on both networks. Particularly, these figures show the maximum link utilization for the different optimization strategies.

The OSPF label corresponds to the initial network state before optimization (i.e., shortest path routing). Likewise, in Figures 4c and 4d we plot the CDFs of the previous experiments, but in this case in relative values w.r.t. the Simulated Annealing baseline. Note that Simulated Annealing represents a near-optimal solution that is not suitable for real-time operation as it has a very high computational cost.

The experimental results indicate that our DRL+GNN agent outperforms the baseline heuristics (OSPF and SAP) on topologies not seen during the training phase (17 and 24 nodes), after being trained in a single 14-node topology. Specifically, the DRL+GNN agent reduces the maximum link utilization in average by $\approx 21.7\%$ and $\approx 14\%$ w.r.t. OSPF in the GBN and GEANT2 topologies respectively, while SAP only reduces it by $\approx 6.8\%$ and $\approx 10\%$. This reveals the ability of our DRL+GNN solution to adapt to different network scenarios – even to new network topologies –, which is an essential property to operate in highly dynamic networks (e.g., traffic changes, link failures, new nodes), such as vehicular networks [1]. In addition, the DRL+GNN agent was able to operate efficiently in scenarios with different link capacities than in the training scenario. This indicates that the DRL agent is able to adapt to different network topologies with different link features than those seen during training.

We also measured the average optimization time that each baseline spends on optimizing all the traffic demands. For the GBN topology we obtained an average time (in seconds) of ≈ 1476 s for the Simulated Annealing, ≈ 4 s for the DRL+GNN architecture and ≈ 0.05 s for the SAP heuristic. The optimization times for the GEANT2 topology are ≈ 1873 s for the Simulated Annealing, ≈ 12 s for the DRL+GNN and ≈ 0.15 s for the SAP. Notice that the larger is the topology (in number of nodes), the more traffic demands need to be optimized. This makes the problem complexity grow with the topology size. However, DRL is inherently a technology based on NNs whose operation process can be easily parallelized using commodity hardware accelerators (e.g., GPUs). In consequence, DRL is key on enabling real-time network optimization.

VI. DISCUSSION AND OPEN CHALLENGES

In this section we discuss which are the main challenges that need to be addressed to enable the operation of DRL on dynamic networks. We consider that the most relevant aspects to address are generalization, action space design, and training cost.

1) *Generalization*: We argue that *generalization* is an essential property for the successful adoption of DRL technologies in networks with dynamic topologies. In this context, generalization refers to the ability of the DRL agent to adapt to new network scenarios not seen during training (e.g., network topologies, traffic, configurations). With generalization, a DRL agent can be trained with a representative set of network topologies and configurations, and afterwards be applied to other scenarios never seen before.

In our work, we addressed generalization by designing a GNN-based DRL agent. In our experimental results, we

have seen that our DRL+GNN agent is able to operate successfully in network topologies never seen before. From a commercialization standpoint, such “universal” DRL agent can be trained in a laboratory and later on be incorporated in a product or a network device (e.g., router, load balancer). The resulting solution would be ready to be deployed in a production network, without requiring any further training or instrumentation of the network.

2) *Action space*: The definition of the DRL agent’s *action space* is fundamental to achieve a good optimization solution. An inadequate action space might explode as the topology size grows, having a direct impact in the complexity of the optimization problem to be addressed. Thus, it is important to find good network abstractions that can limit the dimension of the action space in order to facilitate the learning process to the DRL agent. In addition, the action space should offer enough flexibility and expressiveness to enable the finding of good routing configurations that lead to a high optimization performance.

In our work we leveraged MR [12] to limit the action space dimension (see Sec. IV-B). We consider that MR offers a good level of abstraction to represent routing, but other methods can be used to define DRL’s agent action space. For each traffic demand, we have as many actions as possible middlepoints. This means that the maximum number of actions are $N-1$, where N is the number of nodes (i.e., the action space grows linearly with the topology size). At the same time, this abstraction provides enough flexibility to achieve a near-optimal performance, as in [12].

3) *Training cost*: The *training cost* of the DRL agent must be taken into account, especially when it is trained in large topologies. The DRL training process is, by definition, extremely sequential. This significantly hinders the parallelization of the learning process. New methods should be explored to scale DRL solutions to larger optimization scenarios. For example, these methods could be to accelerate the training process [32], transfer learning [33], or problem reduction [17].

Besides the challenges mentioned previously, there are some open challenges that require further attention by the research community and here we mention some of them. One of the open challenges is that DRL doesn’t offer any kind of performance bounds. In other words, once a DRL agent is trained, there is no way to know exactly which are the upper or lower bounds in optimization performance. However, this problem is also present in many existing heuristic-based solutions.

VII. CONCLUSION

In this paper, we explored a DRL agent based on GNNs for solving a routing optimization problem. To do this, we trained a DRL agent in a single topology and we evaluated it on two topologies never seen before. The experimental results show that the DRL agent is able to operate unseen network configurations while still obtaining a high performance. These results indicate that the DRL agent is able to adapt to different network topologies, which is an essential property to operate

highly dynamic networks. In addition, the DRL agent has a low optimization cost, in the scale of seconds, and its performance is comparable to near-optimal iterative solutions. This makes DRL a key technology to enable modern network applications that require real-time efficient operation in highly dynamic network topologies.

ACKNOWLEDGMENT

This work has received funding from the European Union's Horizon 2020 research and innovation programme within the framework of the NGI-POINTER Project funded under grant agreement No 871528. This paper reflects only the author's view; the European Commission is not responsible for any use that may be made of the information it contains. This work was also supported by the Spanish MINECO under contract TEC2017-90034-C2-1-R (ALLIANCE), the Catalan Institution for Research and Advanced Studies (ICREA) and the Secretariat for Universities and Research of the Ministry of Business and Knowledge of the Government of Catalonia and the European Social Fund.

REFERENCES

- [1] S. Zeadally, R. Hunt, Y.-S. Chen, A. Irwin, and A. Hassan, "Vehicular ad hoc networks (vanets): status, results, and challenges," *Telecommunication Systems*, vol. 50, no. 4, pp. 217–241, 2012.
- [2] I. Yaqoob, I. A. T. Hashem, Y. Mehmood, A. Gani, S. Mokhtar, and S. Guizani, "Enabling communication technologies for smart cities," *IEEE Communications Magazine*, vol. 55, no. 1, pp. 112–120, 2017.
- [3] F. Alagoz, O. Korkak, and A. Jamalipour, "Exploring the routing strategies in next-generation satellite networks," *IEEE Wireless Communications*, vol. 14, no. 3, pp. 79–88, 2007.
- [4] V. Mnih, K. Kavukcuoglu, D. Silver, A. Graves, I. Antonoglou, D. Wierstra, and M. Riedmiller, "Playing atari with deep reinforcement learning," *arXiv preprint arXiv:1312.5602*, 2013.
- [5] "IBM CPLEX optimizer," <https://www.ibm.com/analytics/cplex-optimizer>.
- [6] J. Schrittwieser, I. Antonoglou, T. Hubert, K. Simonyan, L. Sifre, S. Schmitt, A. Guez, E. Lockhart, D. Hassabis, T. Graepel *et al.*, "Mastering atari, go, chess and shogi by planning with a learned model," *Nature*, vol. 588, no. 7839, pp. 604–609, 2020.
- [7] P. Almasan, J. Suárez-Varela, A. Badia-Sampera, K. Rusek, P. Barlet-Ros, and A. Cabellos-Aparicio, "Deep reinforcement learning meets graph neural networks: Exploring a routing optimization use case," *arXiv preprint arXiv:1910.07421*, 2019.
- [8] F. Scarselli, M. Gori, A. C. Tsoi, M. Hagenbuchner, and G. Monfardini, "The graph neural network model," *IEEE transactions on neural networks*, vol. 20, no. 1, pp. 61–80, 2008.
- [9] P. W. Battaglia, J. B. Hamrick, V. Bapst, A. Sanchez-Gonzalez, V. Zambaldi, M. Malinowski, A. Tacchetti, D. Raposo, A. Santoro, R. Faulkner *et al.*, "Relational inductive biases, deep learning, and graph networks," *arXiv preprint arXiv:1806.01261*, 2018.
- [10] J. Suárez-Varela, S. Carol-Bosch, K. Rusek, P. Almasan, M. Arias, P. Barlet-Ros, and A. Cabellos-Aparicio, "Challenging the generalization capabilities of graph neural networks for network modeling," in *Proceedings of ACM SIGCOMM Posters and Demos*, 2019, pp. 114–115.
- [11] D. Xu, M. Chiang, and J. Rexford, "Link-state routing with hop-by-hop forwarding can achieve optimal traffic engineering," *IEEE/ACM Transactions on networking*, vol. 19, no. 6, pp. 1717–1730, 2011.
- [12] R. Hartert, S. Vissicchio, P. Schaus, O. Bonaventure, C. Filsfils, T. Tekamp, and P. Francois, "A declarative and expressive approach to control forwarding paths in carrier-grade networks," *ACM SIGCOMM*, p. 15–28, 2015.
- [13] S. Gay, R. Hartert, and S. Vissicchio, "Expect the unexpected: Sub-second optimization for segment routing," in *IEEE Conference on Computer Communications (INFOCOM)*. IEEE, 2017, pp. 1–9.
- [14] M. Jadin, F. Aubry, P. Schaus, and O. Bonaventure, "Cg4sr: Near optimal traffic engineering for segment routing with column generation," in *IEEE Conference on Computer Communications (INFOCOM)*, 2019, pp. 1333–1341.
- [15] Z. Xu, J. Tang, J. Meng, W. Zhang, Y. Wang, C. H. Liu, and D. Yang, "Experience-driven networking: A deep reinforcement learning based approach," in *IEEE Conference on Computer Communications (INFOCOM)*, 2018, pp. 1871–1879.
- [16] J. Suárez-Varela, A. Mestres, J. Yu, L. Kuang, H. Feng, P. Barlet-Ros, and A. Cabellos-Aparicio, "Feature engineering for deep reinforcement learning based routing," in *ICC 2019-2019 IEEE International Conference on Communications (ICC)*. IEEE, 2019, pp. 1–6.
- [17] P. Sun, J. Lan, J. Li, J. Zhang, Y. Hu, and Z. Guo, "A scalable deep reinforcement learning approach for traffic engineering based on link control," *IEEE Communications Letters*, 2020.
- [18] J. Moy, "Rfc2328: Ospf version 2," 1998.
- [19] C. Filsfils, N. K. Nainar, C. Pignataro, J. C. Cardona, and P. Francois, "The segment routing architecture," in *2015 IEEE Global Communications Conference (GLOBECOM)*. IEEE, 2015, pp. 1–6.
- [20] A. Rodriguez-Natal, M. Portoles-Comeras, V. Ermagan, D. Lewis, D. Farinacci, F. Maino, and A. Cabellos-Aparicio, "Lisp: a southbound sdn protocol?" *IEEE Communications Magazine*, vol. 53, no. 7, pp. 201–207, 2015.
- [21] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov, "Proximal policy optimization algorithms," *arXiv preprint arXiv:1707.06347*, 2017.
- [22] J. Gilmer, S. S. Schoenholz, P. F. Riley, O. Vinyals, and G. E. Dahl, "Neural message passing for quantum chemistry," in *Proceedings of the International Conference on Machine Learning (ICML) - Volume 70*, 2017, pp. 1263–1272.
- [23] J. Suárez-Varela, A. Mestres, J. Yu, L. Kuang, H. Feng, A. Cabellos-Aparicio, and P. Barlet-Ros, "Routing in optical transport networks with deep reinforcement learning," *IEEE/OSA Journal of Optical Communications and Networking*, vol. 11, no. 11, pp. 547–558, 2019.
- [24] M. Abadi, P. Barham, J. Chen, Z. Chen, A. Davis, J. Dean, M. Devin, S. Ghemawat, G. Irving, M. Isard *et al.*, "Tensorflow: A system for large-scale machine learning," in *Proceedings of the 12th USENIX Symposium on Operating Systems Design and Implementation (OSDI)*, 2016, pp. 265–283.
- [25] G. Brockman, V. Cheung, L. Pettersson, J. Schneider, J. Schulman, J. Tang, and W. Zaremba, "Openai gym," *arXiv preprint arXiv:1606.01540*, 2016.
- [26] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," *arXiv preprint arXiv:1412.6980*, 2014.
- [27] X. Hei, J. Zhang, B. Bensaou, and C.-C. Cheung, "Wavelength converter placement in least-load-routing-based optical networks using genetic algorithms," *Journal of Optical Networking*, vol. 3, no. 5, pp. 363–378, 2004.
- [28] J. Pedro, J. Santos, and J. Pires, "Performance evaluation of integrated otn/dwdm networks with single-stage multiplexing of optical channel data units," in *2011 13th International Conference on Transparent Optical Networks*. IEEE, 2011, pp. 1–4.
- [29] F. Barreto, E. C. Wille, and L. Nacamura Jr, "Fast emergency paths schema to overcome transient link failures in ospf routing," *arXiv preprint arXiv:1204.2465*, 2012.
- [30] M. Roughan, "Simplifying the synthesis of internet traffic matrices," *ACM SIGCOMM Computer Communication Review*, vol. 35, no. 5, pp. 93–96, 2005.
- [31] S. Kirkpatrick, C. D. Gelatt, and M. P. Vecchi, "Optimization by simulated annealing," *science*, vol. 220, no. 4598, pp. 671–680, 1983.
- [32] A. Stooke and P. Abbeel, "Accelerated methods for deep reinforcement learning," *arXiv preprint arXiv:1803.02811*, 2018.
- [33] R. Glatt, F. L. Da Silva, and A. H. R. Costa, "Towards knowledge transfer in deep reinforcement learning," in *2016 5th Brazilian Conference on Intelligent Systems (BRACIS)*. IEEE, 2016, pp. 91–96.