Virtualized Controller Placement for Multi-domain Optical Transport Networks

Sabidur Rahman[1]($\boxtimes$), Tanjila Ahmed[1], Sifat Ferdousi[1], Partha Bhaumik[1], Pulak Chowdhury[1], Massimo Tornatore[1,2], Goutam Das[3], and Biswanath Mukherjee[1]

[1] University of California, Davis, USA
{krahman,tanahmed,sferdousi,pbhaumik,pchowdhury,bmukherjee}@ucdavis.edu
[2] Politecnico di Milano, Milan, Italy
massimo.tornatore@polimi.it
[3] Indian Institute of Technology, Kharagpur, India
gdas@gssst.iitkgp.ac.in

Abstract. Optical multi-domain transport networks are often controlled by a hierarchical distributed architecture of controllers. Optimal placement of these controllers is very important for their efficient management and control. Traditional SDN controller placement methods focus mostly on controller placement in datacenter networks. But the problem of virtualized controller placement for multi-domain transport networks needs to be solved in the context of geographically-distributed heterogeneous multi-domain networks. In this context, Edge Datacenters have enabled network operators to place virtualized controller instances closer to users, besides providing more candidate locations for controller placement. In this study, we propose a dynamic controller placement method for optical transport networks that considers the heterogeneity of optical controllers, resource limitations at edge hosting locations, latency requirements, and costs. Simulation studies considering practical scenarios show significant cost savings and delay reductions compared to standard placement approaches.

Keywords: SDN and optical controller · Optical transport network · Cost savings · Network virtualization · Edge computing

1 Introduction

Existing proposals for controller placement [1] have focused mostly on packet-switched Software-Defined Networks (SDNs) and they often ignore the complexity, heterogeneity, and vendor specificity of a transport-network control plane. Current technical solutions for transport-network control planes (e.g., Transport SDN, T-SDN) are designed for circuit-switched layer 0 (optical) and layer 1 (SONET/ SDH and OTN). T-SDN supports multi-layer, multi-vendor, circuit-oriented networks

that are different from packet-based SDN-controlled networks [2]. The control plane for optical transport networks employs a hierarchical distributed architecture [3] comprising heterogeneous (often vendor-specific) Optical Network (ON) Controllers (ONC) and SDN Controllers (SDNC).

Our study considers that T-SDN controllers can be deployed as virtualized controller instances (as in [5]). Virtualized controller placement has many benefits. First, manually deploying SDN and ON controllers in traditional 'hardware boxes' can take several days, compared to few minutes in case of virtualized instances (hosted on Virtual Machines (VMs), docker containers, etc.) in the cloud datecenter (DCs) or in computing nodes at edge datacenters (Edge-DCs) (such as Network Function Virtualization Infrastructure Points of Presence (NFVI-PoPs), metro datacenters (DCs), or Central Offices Re-architected as Datacenters (CORDs), etc.). Second, virtualized controllers can be easily recovered from failures or disasters using the backed-up/replicated virtual copy of the controllers. These instances can be easily moved from one location to another and can be redeployed [6] without significant down time. Third, operational cost savings for network operators and leasing cost savings for network leasers are other important motivations toward virtualization.

Prior studies exploring static [7,8], and dynamic [9] controller placement problems mostly focused on packet-based SDN controllers and DC networks [10,11]. But, as we discuss in Sect. 2, methods proposed in SDN and DC scenarios are often not applicable and not optimal for heterogeneous optical transport networks.

To the best of our knowledge, our study is the first to propose dynamic placement of controllers for heterogeneous, multi-domain transport networks comprising heterogeneous ON and SDN controllers, considering the complexity due to virtual instances hosted jointly on DCs and Edge (e.g. NFVi-PoPs), and inter-domain and intra-domain latency constraints. In addition, we observe that [2,4] both fixed-grid and flex-grid technologies might be required to co-exist (i.e., mixed-grid) with seamless interoperability (see Fig. 1). In this context, our proposed method can support the deployment of different controller types, by enforcing 'controller-type constraint' (e.g., flex-grid controllers for flex-grid domains, fixed-grid controllers for fixed-grid domains, etc.). In addition, we explore the technical details of the dynamic controller placement problem (e.g., latency requirements, resource limitations at Edge-DCs, controller capacity limitations, etc.), propose the Virtualized Controller Deployment Algorithm (VCDA), and report illustrative results comparing with prior studies.

This study is organized as follows. Section 2 reviews prior work on controller placement problems. Section 3 discusses the control-plane architecture. Section 4 provides a formal problem statement and describes the proposed solution method. Section 5 discussed numerical results on cost savings and delay minimization. Section 6 concludes the study.

2   Background and Related Work

In the context of SDN controllers, both static [7,8], and dynamic [9] placement problems have been explored. But most studies on SDN Controller Placement

Problems (CPPs), e.g., [7–9], consider placement of only controller 'middle-boxes', not virtualized instances. Ref. [15] considers recovery of SDN controllers in a disaster scenario. These early studies do not consider the additional complexities due to virtualization, delay constraints, hosting location constraints, etc.

Recently, control-plane architectures [12,16] in T-SDN paradigm are proposed to accommodate multiple heterogeneous network domains and associated domain-specific ON controllers. But, as recent studies [2,16] suggests, there is still no consensus on the design of the T-SDN control plane. Ref. [5] was among the first to propose a virtualized control plane architecture for transport networks, but not from placement perspective.

Recent studies [10,11] on Elastic Control Placement (ECP) for SDN controllers discuss threshold-based methods to dynamically resize the ECPs. Ref. [10] focuses on DC networks managed through homogeneous SDN controllers, minimizing control-plane resizing delay. Refs. [10,11] consider DC placement (of controller instances), which is practical for a DC network scenario. But, for a transport network with distributed heterogeneous domains, we also consider joint deployment in DCs and Edge-DCs. This introduces new constraints such as host location capacity, and inter-domain and intra-domain communication delays (which [10,11] do not consider).

3   Control Plane Architecture

Figure 1 shows an example of a hierarchical control plane for heterogeneous transport networks. Domain-specific controllers are connected with 'parent controller(s)', which are connected to the 'application plane' (e.g., Transport Network Orchestrator (TNO), Operations Support Systems (OSS), etc.).
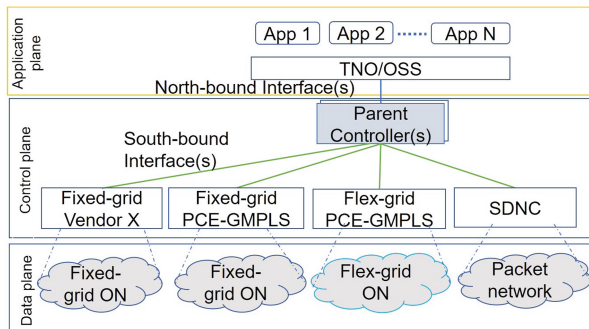
Fig. 1. Control-plane architecture for heterogeneous transport networks.

'Domain Controllers' are responsible for the communication between 'control plane' and 'data plane'. Different autonomous domains, depending on the underlying ONC/SDNC, use specific type of controllers and protocols to control the 'data plane' switches. Figure 1 shows three types of domains. This example architecture can be extended to support more variations of domain controllers and associated technologies.

4 Problem Statement and Solution Method

4.1 Problem Statement

The controller-placement problem is known to be NP-hard [1]. Traditional dynamic controller placement methods [9,10] focus on the 'switch-to-controller' mapping, ensuring that each switch (forwarding plane) is connected to at least one controller (control plane); incoming traffic flow requests originated inside a domain will be served by the same controller; and controller capacity limit is preserved. When a new traffic flow request arrives, the 'switch' depends on the controller for routing and path computation decisions.

The dynamic 'on-demand' controller deployment problem can be defined as follows: Given a topology, a set of controller hosting locations with limited capacity, arrival rate of traffic flows, a set of heterogeneous network domains, controller capacity, and constraints, deploy optimal number of controllers to satisfy all the domains, minimizing the leasing costs.

4.2 Input Parameters and Variables

- $G(V, E)$: Optical transport network topology where V is set of domains and E is set of links.
- M_v: set of controllers serving domain v.
- S_v: set of switches in domain v.
- T_v: domain-specific controller type.
- H_v: controller hosting location where $H_v \subseteq V$.
- χ_v^t: total compute capacity at v.
- χ_v^u: used compute capacity at v.
- ω_v^t: total memory capacity at v.
- ω_v^u: used compute capacity at v.
- T_v^μ: service capacity limit (i.e., maximum number of requests served per second) for controller type T_v.
- T_v^χ: compute resource requirement for controller-type T_v.
- T_v^ω: memory resource requirement for controller-type T_v.
- $\lambda(v)$: arrival rate of new traffic flows for a given domain (v), where r_v represents new arrival of flow routing request. S_r is the switch at which the request has arrived, and M_v gives the set of domain controllers the switch (S_r) is connected to.
- α: latency constraint (maximum allowed delay).
- C_T: variable containing total cost of running the controllers in all the domains (more details in Section IV.D).

4.3 Constraints

We consider the following constraints:

1. **Latency constraints**: Controllers must be placed within the allowed latency limit, i.e., switch-to-controller and controller-to-switch delay, including processing delay must not exceed the allowed delay limit:

$$D(s,h) + D_p + D(h,s) \leq \alpha; \forall s \epsilon S_v, \ \forall h \epsilon V \tag{1}$$

where function $D(x,y)$ represents transmission, propagation, and processing delay between points x and y, s is origin, h is controller hosting location, and D_p is processing delay at controller instances M_v.

2. **Controller type constraint**: To reflect heterogeneity, we enforce the following controller type constraints:

$$T_v == M_v^t; \ \forall v \epsilon V \tag{2}$$

where, the constraint enforces that all controller instances of v, (M_v^t) match the required controller type T_v.

3. **Controller capacity constraints**: Deployed controllers must have enough capacity to support domain switches:

$$\sum_{s \epsilon S_v} \lambda(s) \leq T_v^\mu * |M_v|; \ \forall v \epsilon V \tag{3}$$

4. **Controller host capacity limit**: Hosting location $(v.h)$ must have both compute and memory capacity:

$$\sum_{g \epsilon V} |M_g| * T_g^\chi \leq \chi_v^t; \ H_g == v; \forall v \epsilon V \tag{4}$$

$$\sum_{g \epsilon V} |M_g| * T_g^\omega \leq \omega_v^t; \ H_g == v; \forall v \epsilon V \tag{5}$$

4.4 Cost Models

Leasing Cost for Virtual Controller Instances. We consider that network operators lease capacity from DC operators. Virtual instance leasing cost C_C can be stated as:

$$C_C = \sum_{v \epsilon V} (|M_v| * \gamma * d) \tag{6}$$

where γ is per-unit compute per unit-time cost of leasing virtual instances and d is duration of operation.

Network Capacity Cost. Cost of network usage is often ignored in prior studies focusing on DC SDN scenarios. But the distributed nature of virtualized controller deployment for transport networks can add significant communication cost. First, switch-to-controller communication cost, C_N^S, is:

$$C_N^S = (\sum_{v \epsilon V} (B(S_v, M_v))) * \pi \tag{7}$$

where B(x,y) gives bandwidth consumption between x and y, and π gives the per-GBps per unit-time bandwidth price.

Similarly, controller-to-controller (same domain) communication cost, C_N^M, is:

$$C_N^M = \sum_{v \epsilon V} (B(M_v(i), M_v(j)); i \neq j)) * \pi \tag{8}$$

where $B(M_v(i), M_v(j)); i \neq j$ gives bandwidth consumption between controllers instances of the same domain.

Similarly, controller-to-controller (different domains) communication cost, C_N^P, is:

$$C_N^P = \sum_{p,q \epsilon V} (B(M_p, M_p; p \neq q)) * \pi \tag{9}$$

In addition, controllers may require to be migrated from one hosting location to other. This live VM migration process adds to the network cost (C_N^V) as follows:

$$C_N^V = \sum_{v \epsilon V} (B_{VM}(v)) * \pi \tag{10}$$

where $B_{VM}(.)$ is the bandwidth consumption due to migrations.

Hence, total network cost, C_N, is:

$$C_N = C_N^S + C_N^M + C_N^P + C_N^V \tag{11}$$

Delay Cost. We also consider the impact of controller delays on user experience (higher delay means unhappy user, leading to revenue penalty). This cost adds delay factor in decision making and encourages the algorithm to minimize delays. Similar to network cost, cost of switch-to-controller delay (C_U^S), controller-to-controller (same domain) communication delay (C_U^M), and controller-to-controller (different domains) delay (C_U^K) can be calculated by replacing B(x,y) with D(x,y) and π with σ in Eqs. (7), (8), and (9), respectively. Here, D(x,y) gives delay due to communication between sets x and y, and σ is cost (\$) associated to per unit-time delay. Also, our model considers cost due to processing delay at controllers, C_U^P.

Hence, total delay cost becomes:

$$C_U = C_U^S + C_U^M + C_U^K + C_U^P \tag{12}$$

Now, total cost is modeled as:

$$C_T = C_C + C_N + C_U \tag{13}$$

The objective of the proposed method is to minimize the cost of controller deployment:

$$Minimize(C_T) \tag{14}$$

4.5   Algorithm

We propose a polynomial-time heuristic, called Virtualized Controller Deployment Algorithm (VCDA), as a scalable solution for a heterogeneous optical transport network (see Algorithm 1). Since turning controllers on/off too often may make the network unstable, we introduce a decision epoch (e), a dynamic variable allowing network operators to tune the decision frequency. We also use two management entities: Network Management and Orchestration (NMO) and Distributed Cloud Management (DCM). NMO takes care of load balancing and assignment of switches and traffic with the controllers.

Algorithm 1. Virtualized Controller Deployment Algorithm (VCDA)

1: **Input:** $G(V,E), \lambda(v), \alpha, e$;
2: **for each** domain v in V **do**
3:                             ▷ Calculate required number of controllers
4:     $c \leftarrow (count(r_v))/T_v^\mu$;
5:     **if** $|M_v| \ == \ c$ **then**
6:         Consolidate and load balance traffic flows among the M_vs using Eqns. (1-5);
7:                            ▷ deploy more controllers
8:     **else if** $c \ > \ |M_v|$ **then**
9:         $h \leftarrow H_v$;
10:        $\delta \leftarrow c \ - \ M_v$;
11:                         ▷ enough resources at h
12:         **if** $(\chi_h^t - \chi_h^u) > \delta * T_v^\chi$ & $(\omega_h^t - \omega_h^u) > \delta * T_v^\omega$ **then**
13:            Turn on additional δ controllers (T_v type) at location h;
14:            Load balance and re-route among the M_vs using Eqns. (1-5);
15:                     ▷ not enough resources at h
16:         **else**
17:            $h^{'} \leftarrow$ find optimum location to host c using Eqns. (1-5) and (6-14);
18:            Allocate c controllers (T_v type) at $h^{'}$ via **DCM**;
19:            Migrate all M_v instances to $h^{'}$ via **DCM**;
20:            Turn on δ additional controller instances via **DCM**;
21:            Load balance and re-route flows among the controller instances;
22:            Turn off M_v controllers at h via **DCM**;
23:         **end if**
24:                         ▷ remove extra resources
25:     **else if** $c \ < \ |M_v|$ **then**
26:         $\delta \leftarrow c \ - \ M_v$;
27:         DCM finds optimum δ controllers to turn off;
28:         Reroute and load balance among the M_v using Eqns. (1-5);
29:         Turn off δ controllers;
30:     **end if**
31: **end for**
32: **if** e is expired **then**
33:     go to line 2;
34: **end if**

Our algorithm ensures that, for each domain, enough controllers are deployed to serve current load, observing the constraints. At a given load, if the controller capacity constraint holds, it means that we do not need additional controller instances (line 5). But, if the controller capacity constraint fails (line 9), the algorithm checks if the current hosting location (h) has enough compute and memory capacity to host the additional δ controllers (line 13). If yes, we turn on additional controllers and load balance the switches and traffic flows (line 14–17). If host location h does not have enough resources, the algorithm finds the next optimal location to host all the instances (line 20) following constraints as in Eqs. (1–5) and minimizing Eq. (14). In this step, we utilize the benefits of consolidation in computing. Placing controllers from the same domain closer to each other will reduce delay cost (Eq. (12)). We consider live VM migraion (line 24) to relocate the already-running controller instances to the new location with least interruption of services.

The algorithm turns off the extra controllers (line 32–37) to save operational cost. After each iteration, the algorithm waits for the epoch e to expire. The run-time complexity of VCDA depends on number of domains ($|V|$), maximum number of controllers ($max(|M_v|)$), maximum number of switches in a domain ($max(|S_v|)$), and number of host locations ($|H_v|$). The run-time complexity of VCDA can be expressed as $O(|V|*max(|M_v|)*max(|S_v|)+|V|*|H_v|*max(|M_v|)*max(|S_v|)$.

5   Illustrative Numerical Examples

We present illustrative results on a US-wide topology (see Fig. 2), with heterogeneous domains and Edge-DCs/DCs. Each network domain requires domain-specific controller(s), that are connected to other domains via backbone optical links. We consider two DC locations (2 and 13) and three Edge-DCs (6, 8, 10) to host controller instances. Capacities of DCs, racks, and servers vary significantly in practice. For Edge-DCs (6, 8, 10), we assume total compute capacity (χ_v^t) is 30 units and total memory capacity (ω_v^t) is 60 GB. For DCs, we consider 15000 compute and 30000 memory capacity (to represent virtually infinite capacity).

We observe that the size of the topology, number of domains, etc. have impact on the results from our study. Thus, we assume different domain scenarios using this topology: (a) for the results reported in Figs. 3 and 4, we assume that each of the nodes in topology is a domain. Hence, we have six individual fixed-grid domains (4, 5, 7, 9, 12), two packet-network domains (8 and 11), and six flex-grid domains (1, 2, 3, 10, 13, 14). and (b) Fig. 5 results consider larger six domains. Nodes 1, 2, and 3 create a flex-grid domain: 1-2-3. Similarly, the other domains are: 4-5-6-7 (fixed-grid), 8-11 (packet), 10-13-14 (flex-grid), 9 (fixed-grid), and 12 (fixed-grid).

Per census data, we have 47.6% of US population in Eastern time zone, 29.1% in Central, 6.7% in Mountain, and 16.6% in Pacific time zone. Our study uses this data to generate spatial variation of incoming load $\lambda(v)$. For example, full load ($\lambda = 1$) for the Eastern domains is 30,615 requests per seconds vs. 10,910
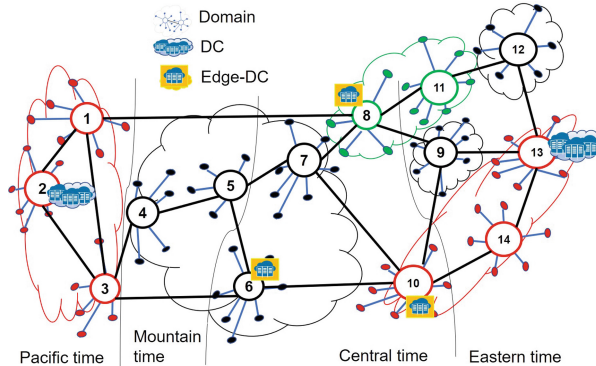
Fig. 2. Example optical network topology with controller host locations and heterogeneous domains (red denotes flex-grid, black denotes fixed-grid, and green denotes packet).

requests for Mountain domains. For illustrative examples, let $\alpha = 15$ ms, $T_v^\mu = 2500$ requests per second [9], per-controller instance compute requirement (T_v^χ) $= 2$ compute units [10], memory requirement (T_v^ω) $= 4$ GB, $\gamma = \$0.01$ per unit per hour, $\sigma = \$0.0001$ per minute, and $\pi = \$70$ per GBps per month [14].

Figure 3 compares the normalized cost (Eq. (14)) among three different methods: (1) DC-Only method mimics controller placement methods which focus on dynamic placement of controllers inside DCs only; (2) Edge-Greedy method considering both DCs and Edge-DCs, but instead of consolidated VM placement considering delays, this method places VMs in a greedy way ignoring network cost (Eq. 7–11) and delay cost (Eq. 12) (similar to [9], evolved to host controllers at Edge); and (3) our proposed VCDA method considering consolidated VM placement, delay and capacity constraints (Eqs. (1–5)), and cost minimization (Eqs. (6–14)).
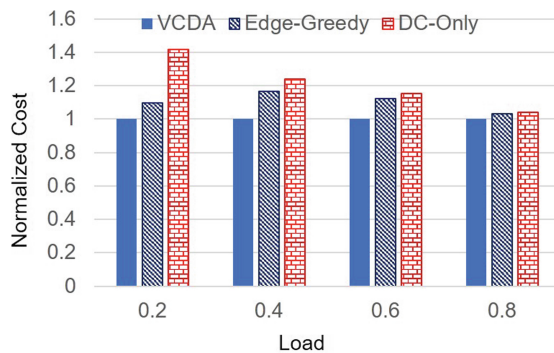
Fig. 3. Normalized cost versus load for topology scenario (a).

Figure 3 shows that our algorithm has lowest normalized cost among all three approaches. As expected, placing controllers only at DCs results in very high delays (see Fig. 4), causing higher delay cost. At lower load ($\lambda = 0.2$), Edge-Greedy placement saves cost compared to DC-Only (10% extra cost for Edge-Greedy vs. 41% extra cost for DC-Only), by placing the domain controllers closer to Edge. VCDA benefits more from Edge resources using consolidated placement and reducing inter-controller communication delays, resulting in cost saving. But, at higher load ($\lambda = 0.8$), VCDA saves less due to tighter capacity limits at Edge-DCs, as consolidated controllers need to move to DC locations, resulting in higher network and delay costs. Still, VCDA's cost savings is higher than the other two methods.

A major limitation of virtualized controller placement in transport networks is the additional delays. Edge-DCs helps to reduce those delays. But, as shown in Fig. 4, if the placement method is not aware of the delays (DC-Only) or takes a greedy placement approach (Edge-Greedy), the controllers will experience significant additional delays (resulting in higher delay cost). At lower load ($\lambda = 0.2$), DC-Only placement method places the controllers far from the domains, resulting in very high delays (60% extra delays). Even Edge-Greedy placement method (13.5% extra delays) reduces significant delays compared to DC-Only. But, as load increases ($\lambda = 0.6$), VCDA experiences higher delays as more controllers are now being placed at DCs (due to Edge-DC hosting capacity limit). Edge-Greedy also experience more delays at higher loads as, in addition to delays due to DC locations usage, more scattered controller instances lead to higher communication delays.
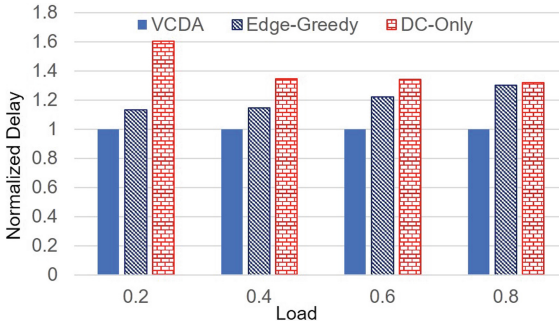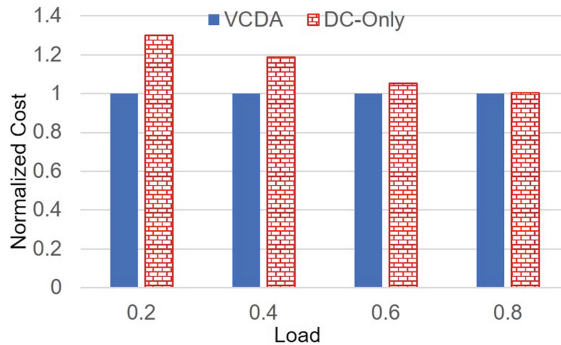

Fig. 4. Normalized delay versus load for topology scenario (a).

Figure 5 shows normalized cost for topology scenario (b) (only six domains, instead of 14). At lower load ($\lambda = 0.2$), cost savings is 30% compared to 'DC-Only'. VCDA looses some cost savings (compared to 41% for the same load in Fig. 3). This observation can be explained as follows. When node 10 was a separate domain, it's controllers were placed in the Edge-DC at 10. But, with the bigger domain 10-13-14, controllers for this domain are hosted in DC at

13, adding to the switch-to-controller network and delay cost. As load grows cost savings reduces and gets close to 'DC-Only' (0.358% cost savings at load $\lambda = 0.8$). This observation can be explained with the following example: for domain 4-5-6-7, initially at lower load, controllers were placed at Edge-DC at 6, but as load grew, requiring more resources for controllers, the controller instances were migrated to DC (reducing the cost gap).

Fig. 5. Normalized cost versus load for topology scenario (b).

More Edge-DC and DC locations can significantly change the cost and delay. More Edge-DC locations will allow to save more cost and reduce delay even in higher loads ($\lambda = 0.8$). In our study, we consider two DC locations, in two corners of the topology, from these intuitions: (1) DC operators in USA tend to place DCs in more populated east coast and west coast regions; and (2) Placing DCs apart allows us to demonstrate the impact of Edge-DCs. But if we add more DCs or change the DC locations, that will impact the cost and delays as well.

6   Conclusion

Virtualized controller placement in multi-domain heterogeneous optical transport networks introduces new challenges for network management. Our proposed method for controller placement considers transport-network-specific properties and constraints such as heterogeneous optical controller types, resource limitations at edge-hosting locations, cost from additional delays, etc. Illustrative examples show that our proposed method saves cost and reduces delays significantly, compared to prior studies. Future studies should explore variation of compute/memory requirements, variation of Edge-DC capacities, variation of Edge-DC and DC locations, temporal variation of load, and more detailed cost models.

References

1. Wang, G., et al.: The controller placement problem in software defined networking: a survey. IEEE Netw. **31**(5), 21–7 (2017)
2. Alvizu, R., et al.: Comprehensive survey on T-SDN: Software-defined networking for transport networks. IEEE Commun. Surv. Tutorials **19**(4), 2232–2283 (2017)
3. Lopez, V., et al.: Control plane architectures for elastic optical networks. J. Opt. Commun. Netw. **10**(2), 241–249 (2018)
4. Ahmed, T., et al.: Dynamic routing and spectrum assignment in co-existing fixed/flex grid optical networks. In: Proceedings of IEEE Advanced Networks and Telecom Systems, Indore, India (2018)
5. Munoz, R., et al.: Integrated SDN/NFV management and orchestration architecture for dynamic deployment of virtual SDN control instances for virtual tenant networks. J. Opt. Commun. Netw. **7**(11), 60–70 (2017)
6. Rahman, S., et al.: Dynamic workload migration over optical backbone network to minimize data center electricity cost. IEEE Trans. Green Commun. Netw. **2**(2), 570–579 (2017)
7. Heller, B., et al.: The controller placement problem. In: Proceedings of 1st Workshop on Hot Topics in Software Defined Networks, pp. 7–12 (2012)
8. Yao, G., et al.: On the capacitated controller placement problem in software defined networks. IEEE Commun. Lett. **18**(8), 1339–1342 (2014)
9. Sallahi, A., et al.: Optimal model for the controller placement problem in software defined networks. IEEE Comms. Lett. **19**(1), 30–33 (2015)
10. Kim, W., et al.: T-DCORAL: a threshold-based dynamic controller resource allocation for elastic control plane in software-defined data center networks. IEEE Commun. Lett. **23**(2), 198–201 (2018)
11. Potluri , A., et al.: An efficient DHT-based elastic SDN controller. In: Proceedings of 9th International Conference on Communication Systems and Networks (COMSNETS), pp. 267–273 (2017)
12. Aguado, A., et al.: ABNO: A feasible SDN approach for multivendor IP and optical networks. IEEE/OSA J. Opt. Commun. Netw. **7**(2), A356–A362 (2015)
13. Rahman, S., et al.: Dynamic controller deployment for mixed-grid optical networks. In: Proceedings of Communications and Photonics Conference (ACP) (2018)
14. Rahman, S., et al.: Auto-scaling VNFs using machine learning to improve QoS and reduce cost. In: Proceedings of IEEE International Conference on Communications (May 2018)
15. Savas, S.S., et al.: Disaster-resilient control plane design and mapping in software-defined networks. In: Proceedings of 16th IEEE International Conference on High Performance Switching and Routing, pp. 1–6 (July 2015)
16. Lourenco, R.B., et al.: Robust hierarchical control plane for transport software-defined networks. Opt. Switching Netw. **30**, 10–22 (2018)