Understanding the Performance Guarantee of Physical Topology Design for Optical Circuit Switched Data Centers

SHIZHEN ZHAO, PEIRUI CAO, and XINBING WANG, Shanghai Jiao Tong University, China

As a first step of designing Optical-circuit-switched Data Centers (ODC), physical topology design is critical as it determines the scalability and the performance limit of the entire ODC. However, prior works on ODC have not yet paid much attention to physical topology design, and the adopted physical topologies either scale poorly, or lack performance guarantee.

We offer a mathematical foundation for the design and performance analysis of ODC physical topologies in this paper. We introduce a new performance metric $\beta(\mathcal{G})$ to evaluate the gap between a physical topology $\mathcal{G}$ and the ideal physical topology. We develop a coupling technique that bypasses a significant amount of computational complexity of calculating $\beta(\mathcal{G})$. Using $\beta(\mathcal{G})$ and the coupling technique, we study four physical topologies that are representative of those in literature, analyze their scalabilities and prove their performance guarantees. Our analysis may provide new guidance for network operators to design better physical topologies for their ODCs.

CCS Concepts: • **Networks** → **Network performance modeling**; **Network performance analysis**; **Data center networks**.

Additional Key Words and Phrases: Optical Circuit Switch, Data Center, Physical Topology, Logical Topology

ACM Reference Format:

Shizhen Zhao, Peirui Cao, and Xinbing Wang. 2021. Understanding the Performance Guarantee of Physical Topology Design for Optical Circuit Switched Data Centers. *Proc. ACM Meas. Anal. Comput. Syst.* 5, 3, Article 42 (December 2021), 24 pages. <https://doi.org/10.1145/3491054>

1 INTRODUCTION

As data center traffic doubles every year [24], building Clos topologies [1, 11, 24] for data centers using electrical switches is becoming more and more expensive and power prohibitive [2]. In order to meet the growing demand at reduced energy cost, building ODCs is becoming a promising alternative for future data centers. An optical circuit switch, e.g. Calient [3], could offer at least hundreds of times higher switching capacity than an electrical switch, while its energy cost is hundreds of times lower (less than 45 Watts for an optical circuit switch with 320 Tx/Rx pairs).

While the eventual goal of ODC design is the full-optical data center, due to technological immaturity, existing designs of ODC have mostly adopted a hybrid design that involves both electrical switching and optical switching. The design of an ODC typically covers the following four aspects:

- (1) Physical topology design: Determine how to interconnect hosts, electrical packet switches (EPS) and optical circuit switches (OCS).
- (2) OCS control: Determine the configurations of all the OCS nodes in an ODC.

Shizhen Zhao is the corresponding author.

Authors' address: Shizhen Zhao; Peirui Cao; Xinbing Wang, Shanghai Jiao Tong University, China.

This work is licensed under a Creative Commons Attribution International 4.0 License.

© 2022 Copyright held by the owner/author(s).

2476-1249/2021/12-ART42

<https://doi.org/10.1145/3491054>

- (3) EPS control: Set up the routing strategies for all the EPS nodes. Since OCS reconfiguration may affect the connected EPS nodes, the queuing and buffer management policies of the EPS nodes may also require careful redesign.
- (4) Host control: Modify the host protocol stacks at different layers in accordance with the OCS/EPS control strategies to attain better end-to-end performance.

Existing works on ODC, e.g., Helios [9], c-Through [25], Mordia [23], REACToR [18], OSA [6], Solstice [19], MegaSwitch [7], FireFly [12], ProjectToR [10], RotorNet [21], Opera [20], etc., have primarily focused on the OCS/EPS/Host control aspects, but their physical topology designs may not support large-scale data centers (see Section 2 for more details). Only until recently, three scalable physical topology designs are adopted by Flexfly [26], Sirius [2] and 3D-Hyper-FleX-LION [17]. However, the performance guarantee of these physical topologies remain unclear. To the best of our knowledge, there is a lack of rigorous performance metrics that evaluate the “goodness” of a physical topology design, and there is no systematic study on how to design a good physical topology for ODC.

In this paper, we focus on the physical topology design for ODC, and offer a rigorous performance evaluation for different physical topologies. One challenge of this work is to design a performance metric for physical topologies. While much of the literature on ODC has adopted network throughput, flow completion time, max link utilization, etc., to evaluate the end-to-end performance of an ODC, these metrics may not accurately reflect the “goodness” of a physical topology. The reason is that, both the physical topology design and the network control policies may affect these metrics, and thus it is hard to evaluate the exact contribution of a physical topology. To eliminate the impact of network control policies, we introduce $\mu(\mathcal{G}, f)$, which is defined as the optimal throughput across all possible network control policies for a given physical topology $\mathcal{G}$ and a given data center demand pattern f (see §3.2). Based on $\mu(\mathcal{G}, f)$, we then define “ β -optimality” for a given physical topology $\mathcal{G}$, i.e., $\mathcal{G}$ is said to be $\beta(\mathcal{G})$ -optimal as long as $\mu(\mathcal{G}, f) \geq \beta(\mathcal{G})\mu(\mathcal{G}', f)$ for any physical topology $\mathcal{G}'$ and any demand pattern f (see §3.3). A physical topology $\mathcal{G}$ is said optimal if $\beta(\mathcal{G}) = 1$.

We first study how to design the optimal physical topologies with $\beta(\mathcal{G}) = 1$. A naive approach requires calculating $\mu(\mathcal{G}, f)$ for different f 's and optimizing $\mu(\mathcal{G}, f)$ among different physical topologies. Unfortunately, this approach is computationally prohibitive. First, the formulation of $\mu(\mathcal{G}, f)$ is essentially an integer programming problem (see (5) in §3.2), which is NP-hard in general. Second, there are an uncountable number of different traffic patterns, and it is impossible to calculate $\mu(\mathcal{G}, f)$ one by one. To circumvent the above difficulty, we introduce a coupling technique (see Lemma 1 in §3.3) to compare two physical topologies. The intuition is that, if one physical topology can “imitate” all the control strategies of another physical topology, then the first physical topology will perform no worse than the second one. Using this coupling technique, we prove that the ideal physical topology and the uniform bipartite physical topology are both optimal. While these two physical topologies have been adopted by many existing ODC designs [4, 6, 9, 18–21, 23, 25] to build prototypes, they scale poorly to support large data centers.

We then study how to design scalable physical topologies for large ODCs. Along this direction, we first prove a negative result that no physical topology is optimal when the number of EPS nodes exceeds the number of ports of an OCS node. Hence, we have to seek for sub-optimal physical topologies, and calculate $\beta(\mathcal{G})$ as a performance guarantee of these physical topologies. We could still use the coupling technique to calculate $\beta(\mathcal{G})$ for these physical topologies. However, due to the lack of direct connectivity between certain node pairs, it may not always be possible for one physical topology to imitate the control decisions of another physical topology. To overcome this challenge, we enhance the above coupling technique using *overlay topologies*. If one physical topology allows constructing an overlay topology to “imitate” all the control strategies of another

physical topology, then the first physical topology will perform no worse than the second one (see Lemma 2 in §3.3). Using the enhanced coupling technique, we successfully calculated $\beta(\mathcal{G})$ for three representative physical topologies.

The contributions of this work are summarized below:

- (1) We introduce a new metric $\beta(\mathcal{G})$ to evaluate the goodness of a physical topology $\mathcal{G}$.
- (2) We develop a coupling technique to calculate $\beta(\mathcal{G})$. This technique circumvents the computational complexity of calculating $\mu(\mathcal{G}, f)$ for any traffic pattern f .
- (3) We design an optimal physical topology (with rigorous proof) for small-scale data centers.
- (4) We prove that no physical topology is optimal for large-scale data centers.
- (5) Motivated by Flexfly [26], 3D-Hyper-FleX-LION [17] and Sirius [2], we design three physical topologies that can scale to large ODCs, and prove their performance guarantees $\beta(\mathcal{G})$.

2 RELATED WORK

Existing works on optical circuit switched data centers have primarily focused on the overall network control. The adopted physical topology designs either scale poorly, or lack performance guarantee.

c-Through [25], Mordia [23], REACToR [18], OSA [6], Solstice [19] simply use one OCS to connect all the Top-of-Rack (ToR) switches. This physical topology design could only support tiny data centers with $\Theta(100)$ number of servers. To improve scalability, MegaSwitch [7] uses fiber rings for interconnection (a similar idea was also discussed in [23]). This physical topology architecture could support $\Theta(1000)$ number of servers, but still does not work for large-scale data centers with over 100k servers and thousands of ToRs.

Using free-space optics, FireFly [12] and ProjectToR [10] enable optical interconnection for thousands of ToRs. However, the free-space optical switching technology faces tremendous deployment complexity, as many environmental factors (e.g., vibration, dust, and humidity) may hinder the performance of the free-space optical links.

Helios [9], RotorNet [21], Opera [20], TROD [4] create a uniform bipartite graph between all the PoDs/ToRs and all the OCSs. This physical topology turns out to be optimal based on our analysis in §4. While this design scales well for the PoD-level interconnect, it scales poorly for the ToR-level interconnect. A large-scale data center may require hundreds of OCSs to connect over 100k servers. Unfortunately, a ToR switch only has tens of uplinks and thus it is impossible to establish a link between every pair of ToR and OCS. (In contrast, a PoD could have hundreds of uplinks.)

Recently, Flexfly [26], Sirius [2] and 3D-Hyper-FleX-LION [17] offered three physical topologies that can potentially scale to large-scale data centers. However, no performance guarantee is provided for any of the three physical topologies.

3 MATHEMATICAL MODEL

3.1 Basic Definitions

We study a network with N electrical-packet-switching (EPS) nodes $\mathcal{S} = \{S_1, S_2, \dots, S_N\}$ and K optical-circuit-switching (OCS) nodes $\mathcal{O} = \{O_1, O_2, \dots, O_K\}$. Each EPS node has L ports, each of which has a transceiver. Each OCS node has R ingress ports and R egress ports. Each EPS port can either connect to another EPS port, or connect to a pair of OCS ports. Note that when we connect an EPS port to OCS ports, we need to separate the transmitter and the receiver of the EPS port, have the transmitter connected to an OCS ingress port, and have the receiver connected to an OCS egress port. Since an OCS port capacity is typically much higher ($100\times$ and more) than an EPS port capacity, the capacity of each network link is dominated by the EPS port capacity B .

Remark on OCS-OCS connections: The optical transceivers commonly used in data centers, e.g., 100G QSFP28 CWDM4, typically have a link margin of 5dB and a maximum transmission distance of 2km. Note that signals traversing an OCS node could experience 3dB optical loss. To enable OCS-OCS connections, one may need to replace the 2km optical transceivers by 10km optical transceivers, e.g., QSFP28 LR4. However, the QSFP28 LR4 transceiver has a much higher consumption and laser cost than the QSFP28 CWDM4 transceiver. Due to the above cost reasons, we do not allow OCS-OCS connections in this paper. Nevertheless, the methodologies introduced in this paper can be easily generalized to incorporate OCS-OCS connections.

Traffic Matrix: Let f_{ij} be the relative traffic demand between the EPS node S_i and the EPS node S_j . Then, the network demand can be modeled using a traffic matrix $f = [f_{ij}, i = 1, 2, \dots, N, j = 1, 2, \dots, N]$.

Physical Topology: We use $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, where $\mathcal{V} = \mathcal{S} \cup \mathcal{O}$, to denote the physical topology among all the EPS nodes and all the OCS nodes. As shown in Fig. 1(a), the edges of $\mathcal{G}$ consist of two parts $\mathcal{E} = \mathcal{E}_1 \cup \mathcal{E}_2$, where $\mathcal{E}_1$ is the set of bidirectional links that interconnect two EPS nodes and $\mathcal{E}_2$ is the set of unidirectional links that interconnect one EPS node and one OCS node. The port count constraints of the EPS nodes and the OCS nodes are reflected by the degrees of the nodes in $\mathcal{G}$. Specifically, the in-degree and out-degree of each EPS node are no more than L and the in-degree and out-degree of each OCS node are no more than R . Note that there could be multiple physical links connecting two adjacent nodes.

OCS Configuration: Each OCS node has R ingress ports and R egress ports. As shown in Fig. 1(b), an OCS configuration is essentially a $1 - 1$ mapping from the R ingress ports to the R egress ports. In this paper, we denote by $\Pi_k, k = 1, 2, \dots, K$ as the OCS configuration of O_k , and let $\Pi = \{\Pi_1, \Pi_2, \dots, \Pi_K\}$. An OCS configuration helps create multiple *logical links* between EPS node pairs. For example, if an ingress port a is mapped to an egress port b in an OCS node, let S_i be the EPS node that connects to a and S_j be the EPS node that connects to b , then a unidirectional *logical link* is created between S_i and S_j . Since OCS nodes do not perform packet decoding, a packet traversing through a logical link will not be aware of the underlying OCS node.

Logical Topology: Given a physical topology $\mathcal{G}$ and an OCS configuration Π , a logical topology $\mathcal{G}_l = \mathcal{G}_l(\mathcal{G}, \Pi) = (\mathcal{V}_l, \mathcal{E}_l)$ is formed. In the logical topology $\mathcal{G}_l$ (also see Fig. 1(c)), the node set $\mathcal{V}_l = \mathcal{S}$ only contains the EPS nodes; the link set $\mathcal{E}_l = \mathcal{E}_1 \cup \mathcal{E}_3$, where $\mathcal{E}_1$ is the set of physical links among all the EPS nodes and $\mathcal{E}_3$ is the set of logical links formed by the OCS configuration Π . We use a non-negative integer $x_{ij}(\mathcal{G}_l)$ to denote the number of links from S_i to S_j in $\mathcal{G}_l$. Note that $x_{ij}(\mathcal{G}_l)$ could be larger than one.

3.2 Network Control for Optimal Throughput

Network performance are usually evaluated from different dimensions, including throughput, latency, flow completion time [8], deployment complexity [27], expansion complexity [27, 28], etc. We focus on throughput in this paper due to the following reasons. First, compared to latency and flow completion time, throughput can be mathematically formulated and thus is easy to evaluate even at the planning stage. Second, having a higher throughput could also help reduce latency, flow completion time, expansion complexity (the number of expansion stages can be reduced [27]), etc.

Given a traffic pattern f and a physical topology $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, network throughput can be affected by various control policies, including route selection, load balancing, congestion control, etc. For ODC, the OCS configurations can be also a determining factor for throughput. We consider an “oracle” that could choose the best control policies for every pair of $(\mathcal{G}, f)$. Specifically, this oracle performs topology engineering (ToE) and traffic engineering (TE) to attain the optimal network throughput, denoted by $\mu(\mathcal{G}, f)$. Admittedly, this oracle is hard to realize in practice. But on the other hand, this oracle allows us to eliminate the impact of different control policies on $\mu(\mathcal{G}, f)$,

Fig. 1. Physical topology and logical topology. When an OCS configuration acts on a physical topology, the logical topology will be changed. In (a) physical topology, red arrowlines represent the set of bidirectional links that interconnect two EPS nodes and green arrowlines represent the set of unidirectional links that interconnect one EPS node and one OCS node. In (c) logical topology, blue arrowlines represent the set of logical links formed by the OCS configuration. Note that each EPS port can either connect to one red bidirectional arrowline, or connect to two green unidirectional arrowlines with opposite directions.

and thus $\mu(\mathcal{G}, f)$ can be used to measure the goodness of the physical topology $\mathcal{G}$ under the traffic pattern f . Next, we will describe how this oracle calculates the optimal throughput $\mu(\mathcal{G}, f)$.

Topology Engineering (ToE): ToE controls the logical topologies used to serve different traffic patterns. For each traffic pattern f , ToE could either generate one logical topology, or generate multiple logical topologies and perform time sharing among these logical topologies. Let $M \geq 1$ be the number of logical topologies generated by ToE. Each logical topology, denoted by $\mathcal{G}_l^{(m)}$, $m = 1, 2, \dots, M$, corresponds to a set of K OCS configurations $\Pi^{(m)} = \{\Pi_1^{(m)}, \Pi_2^{(m)}, \dots, \Pi_K^{(m)}\}$, and lasts for $\Delta^{(m)}$ amount of time. Note that reconfiguring OCSs could incur a reconfiguration latency δ , which ranges from hundreds of nanoseconds to tens of milliseconds, depending on the optical switching technology being used. Hence, the average bandwidth allocated to (S_u, S_v) can be computed as

$$B_{uv} = \left(\sum_{m=1}^M x_{uv}(\mathcal{G}_l^{(m)}) \Delta^{(m)} B \right) / \left(M\delta + \sum_{m=1}^M \Delta^{(m)} \right). \quad (1)$$

In addition, traffic patterns may change in an ODC. To model this fact, we impose the following constraint on $\Delta^{(m)}$:

$$M\delta + \sum_{m=1}^M \Delta^{(m)} \leq \Delta, \quad (2)$$

where Δ is the duration of the traffic pattern f .

Traffic Engineering (TE): Given a ToE solution, network operators can then perform traffic engineering to maximize network throughput. Here we adopt the max-flow formulation. For every flow from S_i to S_j , let $f_{ij}(u, v)$ be the amount of traffic allocated to (S_u, S_v) . The traffic allocation $f_{ij}(u, v)$ must satisfy the following constraints:

- (1) Flow conservation constraints: The total ingress traffic equals to the total egress traffic at every EPS node, i.e.,

$$\begin{cases} \sum_{w_1 \neq v} f_{ij}(w_1, v) = \sum_{w_2 \neq v} f_{ij}(v, w_2), \forall v \neq i, j, \\ \mu f_{ij} + \sum_{w_1 \neq i} f_{ij}(w_1, i) = \sum_{w_2 \neq i} f_{ij}(i, w_2), \\ \sum_{w_1 \neq j} f_{ij}(u, j) = \mu f_{ij} + \sum_{w_2 \neq j} f_{ij}(j, w_2), \end{cases} \quad (3)$$

where μ is a scaling factor of the network traffic pattern f .

- (2) Network capacity constraints: For any EPS node pair (S_u, S_v) , the total amount of traffic allocated to (S_u, S_v) must be no larger than B_{uv} , i.e.,

$$\sum_{i=1}^N \sum_{j=1}^N f_{ij}(u, v) \leq B_{uv}. \quad (4)$$

Defining $\mu(\mathcal{G}, f)$: Note that we introduced a scaling factor μ in (3). We could optimize μ by solving the following optimization problem:

$$\begin{aligned} & \max \quad \mu \\ & \text{s.t. } \mathcal{G}_l^{(m)} = \mathcal{G}_l(\mathcal{G}, \Pi^{(m)}), \forall m = 1, 2, \dots, M, \\ & \quad \mu, B_{uv}, f_{ij}(u, v), \mathcal{G}_l^{(m)}, \Delta^{(m)} \text{ satisfy (1)(2)(3)(4)}. \end{aligned}$$

(5)

Then, $\mu(\mathcal{G}, f)$ is defined as the optimal objective value of the above optimization problem.

Remark on the complexity of calculating $\mu(\mathcal{G}, f)$: Even though we have defined $\mu(\mathcal{G}, f)$ rigorously, calculating $\mu(\mathcal{G}, f)$ turns out to be an NP-Complete problem (see Appendix A for the detailed proof). We could use Gurobi [22], the world's fastest linear programming solver, to solve (5). Unfortunately, (5) contains $\Theta(N^2KM)$ number of integer variables and $\Theta(N^4)$ fractional variables. As we will show in Section 7.2, Gurobi may fail to solve (5) even when there are only 9 EPS nodes.

Remark on TE: We have adopted an “edge formulation” for TE in (3) and (4). This edge formulation requires $\Theta(N^4)$ routing variables, one for every link & flow pair. Another widely used formulation for TE is the “path formulation” (see SWAN [13]). The path formulation pre-computes a set of paths for each flow, and then defines a traffic-spitting ratio variable for every flow and every path. This formulation generally requires fewer number of decision variables. However, its logical topology must be determined first in order to compute paths. In contrast, calculating $\mu(\mathcal{G}, f)$ requires a joint optimization over topology and routing, and thus cannot adopt the “path formulation”.

3.3 Objective of Physical Topology Design

DEFINITION 1. A physical topology $\mathcal{G}$ is said to be β -optimal if for every traffic pattern f ,

$$\mu(\mathcal{G}, f) \geq \beta \mu(\mathcal{G}', f)$$

holds for any possible physical topology $\mathcal{G}'$. β is called the **performance ratio** of $\mathcal{G}$.

Clearly, every physical topology $\mathcal{G}$ corresponds to a $\beta(\mathcal{G}) \in [0, 1]$. $\beta(\mathcal{G})$ can be viewed as a performance guarantee of the physical topology $\mathcal{G}$, and the larger the better. A physical topology is said to be optimal if its corresponding $\beta = 1$.

Unfortunately, directly calculating $\beta(\mathcal{G})$ based on (5) and Definition 1 is computationally intractable for large scale ODCs. First, calculating $\beta(\mathcal{G}, f)$ is NP-Complete. Second, there are uncountable number of traffic patterns f , and it is impossible to calculate $\mu(\mathcal{G}, f)$ one by one. Due to the above difficulties, we have to resort to a theoretical approach. The key to our theoretical approach is a coupling technique detailed in the following two lemmas. This coupling technique helps us circumvent the algorithmic complexity of calculating $\beta(\mathcal{G}, f)$ for all possible f 's.

LEMMA 1. *Given two physical topologies $\mathcal{G}$ and $\mathcal{G}'$. If any logical topology $\mathcal{G}_l$ formed on top of $\mathcal{G}$ can be also formed on top of $\mathcal{G}'$, then for any traffic pattern f , $\mu(\mathcal{G}', f) \geq \mu(\mathcal{G}, f)$.*

PROOF. Let $(\mathcal{G}_l^{*(m)}, f_{ij}^*(u, v))$ be the optimal topology+routing solution that attains the optimal throughput $\mu(\mathcal{G}, f)$. Since the logical topologies $\mathcal{G}_l^{*(m)}$ are also realizable on $\mathcal{G}'$, $(\mathcal{G}_l^{*(m)}, f_{ij}^*(u, v))$ must also satisfy the constraints in (5) when calculating $\mu(\mathcal{G}', f)$. Therefore, $\mu(\mathcal{G}', f) \geq \mu(\mathcal{G}, f)$. $\square$

Lemma 1 can be also generalized to the concept of overlay topology, which is defined below:

DEFINITION 2. *Given a logical topology $\mathcal{G}_l$. A topology $\mathcal{G}_{ol}$ can be formed as an overlay topology of $\mathcal{G}_l$ if we can reserve a number of paths for every source and destination pairs S_i, S_j , such that the following two conditions are met:*

- (1) *From the overlay topology $\mathcal{G}_{ol}$ aspect, for every pair of EPS nodes S_i, S_j , the total reserved capacity among all paths is equal to $x_{ij}(\mathcal{G}_{ol})B$.*
- (2) *From the underlay topology $\mathcal{G}_l$ aspect, for any EPS node pair (S_i, S_j) , the total capacity offered for reservation does not exceed $x_{ij}(\mathcal{G}_l)B$.*

LEMMA 2. *Given two physical topologies $\mathcal{G}$ and $\mathcal{G}'$. If for any logical topology $\mathcal{G}_l$ formed on top of $\mathcal{G}$, there is a logical topology $\mathcal{G}'_l$ formed on top of $\mathcal{G}'$, such that $\mathcal{G}_l$ can be formed as an overlay topology of $\mathcal{G}'_l$, then for any traffic pattern f , $\mu(\mathcal{G}', f) \geq \mu(\mathcal{G}, f)$.*

PROOF. See Appendix B. $\square$

3.3.1 An Ideal Physical Topology. Fix the number of links L and the link capacity B of an EPS node. Assume that each OCS node has the infinite number of ports. Then, we could connect all the EPS nodes to a single OCS node and obtain an ideal physical topology $\mathcal{G}_{L,B}^{\text{ideal}}$. It is easy to verify that any logical topology formed on top of any physical topology is realizable on top of $\mathcal{G}_{L,B}^{\text{ideal}}$. Then according to Lemma 1, $\mathcal{G}_{L,B}^{\text{ideal}}$ is optimal, as stated below:

THEOREM 3. *For any fixed number of links L and fixed link capacity B , $\mathcal{G}_{L,B}^{\text{ideal}}$ is optimal.*

Based on the ideal physical topology $\mathcal{G}_{L,B}^{\text{ideal}}$, we can then define the ideal throughput $\mu_{L,B}^{\text{ideal}}(f)$ for an arbitrary traffic pattern f as

$$\mu_{L,B}^{\text{ideal}}(f) = \mu(\mathcal{G}_{L,B}^{\text{ideal}}, f).$$

Clearly, $\mu_{L,B}^{\text{ideal}}(f)$ scales linearly with respect to B . Further, as the number of links L of an EPS node increases, the ideal throughput $\mu_{L,B}^{\text{ideal}}(f)$ would never decrease.

4 OPTIMAL PHYSICAL TOPOLOGY FOR SMALL DATA CENTERS WITH $N \leq R$

Building an ideal physical topology requires $R \geq NL$. This only applies to tiny data centers with at most tens of EPS nodes¹. In this section, we study a less restrictive case where $N \leq R$, which corresponds to hundreds of EPS nodes.

Before proposing our physical topology design, we first introduce the following lemma.

LEMMA 4. *Given an $N \times N$ non-negative integer matrix $\mathbf{x} = \{x_{i,j}\}$, for any integer $H \geq 1$ and mutually disjoint sets $\mathcal{A}_1, \mathcal{A}_2, \dots, \mathcal{A}_C$ satisfying $\bigcup_{c=1}^C \mathcal{A}_c = \{1, 2, \dots, N\}$ and $\mathcal{A}_{c_1} \cap \mathcal{A}_{c_2} = \emptyset, \forall c_1 \neq c_2$, there must exist H non-negative integer matrices $\mathbf{x}^{(1)}, \dots, \mathbf{x}^{(H)}$ satisfying*

- (1) $\mathbf{x} = \mathbf{x}^{(1)} + \dots + \mathbf{x}^{(H)}$;
- (2) $0 \leq \sum_{j=1}^N x_{i,j}^{(h)} \leq \left\lceil \frac{\sum_{j=1}^N x_{i,j}}{H} \right\rceil$, for any $i = 1, \dots, N$ and any $h = 1, 2, \dots, H$;
- (3) $0 \leq \sum_{i \in \mathcal{A}_c} \sum_{j=1}^N x_{i,j}^{(h)} \leq \left\lceil \frac{\sum_{i \in \mathcal{A}_c} \sum_{j=1}^N x_{i,j}}{H} \right\rceil$, for any $c = 1, \dots, C$ and any $h = 1, 2, \dots, H$;
- (4) $0 \leq \sum_{i=1}^N x_{i,j}^{(h)} \leq \left\lceil \frac{\sum_{i=1}^N x_{i,j}}{H} \right\rceil$, for any $j = 1, \dots, N$ and any $h = 1, 2, \dots, H$;
- (5) $0 \leq \sum_{i=1}^N \sum_{j \in \mathcal{A}_c} x_{i,j}^{(h)} \leq \left\lceil \frac{\sum_{i=1}^N \sum_{j \in \mathcal{A}_c} x_{i,j}}{H} \right\rceil$, for any $c = 1, \dots, C$ and any $h = 1, 2, \dots, H$.

Lemma 4 can be proved by transforming it into a sequence of max-flow problems. It is actually a direct consequence of Theorem 3 in [28] by setting $\mathcal{A} = \mathcal{B} = \{\mathcal{A}_1, \dots, \mathcal{A}_C, \{1\}, \dots, \{N\}\}$. See Appendix A in [28] for the detailed proof.

4.1 Physical Topology Design

Motivated by Lemma 4, we set the number of OCS nodes as $K = L$ and evenly distribute the L links of each EPS node across all the OCS nodes. In other words, the physical topology $\mathcal{G}$ becomes a uniform bipartite graph, with exactly one bidirectional physical link between every pair of EPS node and OCS node. (More specifically, the transmitter of the EPS port is connected to an OCS ingress port, and the receiver of the EPS port is connected to an OCS egress port.) Note that $N \leq R$ ensures that the number of ports of each OCS node is not exhausted. We denote such a physical topology by $\mathcal{G}_{\text{uniform}}$, as shown in Fig. 2. We can prove that $\mathcal{G}_{\text{uniform}}$ is optimal.

Fig. 2. $\mathcal{G}_{\text{uniform}}$ topology.

THEOREM 5. *If $N \leq R$, then $\mathcal{G}_{\text{uniform}}$ is optimal.*

PROOF. We only need to prove that any logical topology $\mathcal{G}_l$ is realizable on top of the physical topology $\mathcal{G}_{\text{uniform}}$.

¹A practical OCS node has hundreds of ports and a practical ToR switches has tens of ports.

Since each EPS node has L ports, we must have

$$\begin{cases} \sum_{i=1}^N x_{ij}(\mathcal{G}_l) \leq L, \forall j = 1, \dots, N, \\ \sum_{j=1}^N x_{ij}(\mathcal{G}_l) \leq L, \forall i = 1, \dots, N. \end{cases} \quad (6)$$

Then, according to Lemma 4, we can decompose $\{x_{ij}(\mathcal{G}_l)\}$ into $H = K = L$ integer matrices $\{x_{ij}^{(h)}(\mathcal{G}_l)\}$, $h = 1, 2, \dots, H$ satisfying the conditions in Lemma 4. Based on (6) and the conditions (2) and (4) of Lemma 4, it is easy to check that any $\{x_{ij}^{(h)}(\mathcal{G}_l)\}$, $h = 1, 2, \dots, H$, must satisfy

$$\begin{cases} \sum_{i=1}^N x_{ij}^{(h)}(\mathcal{G}_l) \leq 1, \forall j = 1, \dots, N, \\ \sum_{j=1}^N x_{ij}^{(h)}(\mathcal{G}_l) \leq 1, \forall i = 1, \dots, N. \end{cases}$$

Since $x_{ij}^{(h)}(\mathcal{G}_l)$'s are non-negative integers, $x_{ij}^{(h)}(\mathcal{G}_l)$ is actually a permutation matrix. Note that each EPS node has one bidirectional physical link connected to the h -th OCS node, it is easy to verify that the permutation matrix $x_{ij}^{(h)}(\mathcal{G}_l)$ is realizable on the h -th OCS node. $\square$

5 NEGATIVE RESULT ON PHYSICAL TOPOLOGY DESIGN

Having designed an optimal physical topology for the case $N \leq R$, we thus wonder if it is possible to design an optimal physical topology for large-scale data centers with $N > R$. Unfortunately, the following theorem demonstrates the impossibility of such a design.

THEOREM 6. *If $N > R$, then no physical topology is β -optimal for $\beta > \frac{N+R-1}{2N-2}$. In other words, for any physical topology $\mathcal{G}$, there exists a traffic pattern f and a physical topology $\mathcal{G}'$, such that*

$$\mu(\mathcal{G}, f) \leq \frac{N+R-1}{2N-2} \mu(\mathcal{G}', f).$$

PROOF. Let $\mathcal{Z} = (k_1, k_2, \dots, k_N)$ be a permutation of the N EPS nodes. In total, we obtain $N!$ different permutations. For each permutation $\mathcal{Z} = (k_1, k_2, \dots, k_N)$, we count the maximum number of unidirectional logical links that can be formed for $S_{k_1} \rightarrow S_{k_2}, S_{k_2} \rightarrow S_{k_3}, \dots, S_{k_N} \rightarrow S_{k_1}$ under a physical topology $\mathcal{G}$, and denote this number by $N(\mathcal{G}, \mathcal{Z})$.

Step 1: We would like to prove that for any physical topology $\mathcal{G}$, there must exist a permutation $\mathcal{Z}^*$, such that

$$N(\mathcal{G}, \mathcal{Z}^*) \leq \frac{NLR}{N-1}.$$

Let $N_{i \rightarrow j}$ be the maximum number of links that can be created from the EPS node S_i to the EPS node S_j . We compute

$$\sum_{\mathcal{Z}} N(\mathcal{G}, \mathcal{Z}) = \sum_{\mathcal{Z}} \sum_{i,j:i \neq j} 1_{\{j \text{ is after } i \text{ in } \mathcal{Z}\}} N_{i \rightarrow j} = \sum_{i,j:i \neq j} N_{i \rightarrow j} \sum_{\mathcal{Z}} 1_{\{j \text{ is after } i \text{ in } \mathcal{Z}\}}.$$

Here, " j is after i in $\mathcal{Z} = (k_1, k_2, \dots, k_N)$ " means that either there exists $r = 1, 2, \dots, N-1$, such that $k_r = i, k_{r+1} = j$, or $k_N = i, k_1 = j$. In order to compute the number of permutations in which j is after i , i.e., $\sum_{\mathcal{Z}} 1_{\{j \text{ is after } i \text{ in } \mathcal{Z}\}}$,

- (1) we first find a permutation with $N-2$ elements excluding i and j , and the total number of such permutations is $(N-2)!$;
- (2) we then insert i and j to the right places of the above permutation, and the number of different inserting positions is N . (For example, when $N = 4$, $ij **$, $*ij*$, $**ij$ and $j**i$ are the 4 generated permutations.)

Hence,

$$\sum_{\mathcal{Z}} 1_{\{j \text{ is after } i \text{ in } \mathcal{Z}\}} = N \times (N - 2)!$$

We then focus on $\sum_{i,j:i \neq j} N_{i \rightarrow j}$. We first fix i and compute $\sum_{j:j \neq i} N_{i \rightarrow j}$. Assume that the EPS node S_i has at least one transmitter connected to the OCS node O_k . Let R_k be the number of links between the egress ports of the OCS node O_k and the receivers of all the EPS nodes excluding S_i . Then it is easy to check that the contribution of the OCS node O_k to $\sum_{j:j \neq i} N_{i \rightarrow j}$ is at most $R_k \leq R$. Since the EPS node S_i can connect to at most L different OCS nodes, we must have

$$\sum_{j:j \neq i} N_{i \rightarrow j} \leq LR, \text{ for any fixed } i.$$

Therefore,

$$\sum_{i,j:i \neq j} N_{i \rightarrow j} = \sum_{i=1}^N \sum_{j:j \neq i} N_{i \rightarrow j} \leq NLR. \quad (7)$$

Note that there are $N!$ different permutations in total. According to the Pigeonhole Principle, there must exist a permutation $\mathcal{Z}^*$ such that

$$N(\mathcal{G}, \mathcal{Z}^*) \leq \frac{\sum_{\mathcal{Z}} N(\mathcal{G}, \mathcal{Z})}{N!} \leq \frac{N \times (N - 2)! \times NLR}{N!} \leq \frac{NLR}{N - 1}.$$

Step 2: We construct a permutation traffic pattern f^* based on $\mathcal{Z}^* = (k_1^*, k_2^*, \dots, k_N^*)$, i.e.,

$$f_{ij}^* = \begin{cases} 1, & \text{if } i = k_r^*, j = k_{r+1}^*, \text{ where } r = 1, 2, \dots, N - 1; \\ 1, & \text{if } i = k_N^*, j = k_1^*; \\ 0, & \text{otherwise.} \end{cases}$$

Clearly, if we create a physical topology $\mathcal{G}(\mathcal{Z}^*)$ such that the l -th ($l = 1, 2, \dots, L$) egress port of the EPS node $S_{k_i^*}$ (or $S_{k_N^*}$) and the l -th ingress port of the EPS node $S_{k_{i+1}^*}$ (or $S_{k_1^*}$) connect to the same OCS, then the resulting physical topology $\mathcal{G}(\mathcal{Z}^*)$ would be the optimal physical topology for f^* . Specifically, on top of the physical topology $\mathcal{G}(\mathcal{Z}^*)$, we can create L identical cycles $k_1^* \rightarrow k_2^* \rightarrow \dots \rightarrow k_N^* \rightarrow k_1^*$ as the logical topology. It is easy to check that this logical topology offers the highest throughput for the traffic matrix f^* over the physical topology $\mathcal{G}(\mathcal{Z}^*)$, and $\mu(\mathcal{G}(\mathcal{Z}^*), f^*) = BL$.

We then compute the throughput $\mu(\mathcal{G}, f^*)$ for f^* over $\mathcal{G}$. Consider the optimal ToE+TE strategy that attains the throughput value $\mu(\mathcal{G}, f^*)$. Then, within a time period of Δ ,

$$\mu(\mathcal{G}, f^*) \Delta \sum_{i=1}^N \sum_{j=1}^N f_{ij}^* = \mu(\mathcal{G}, f^*) \Delta N \quad (8)$$

amount of traffic is delivered from sources to destinations. The delivered traffic can be grouped into two classes:

- (1) Class-1: traffic that is delivered via only one hop;
- (2) Class-2: traffic that is delivered via at least two hops.

Class-1 traffic must occupy one of following types of logical links $S_{k_1^*} \rightarrow S_{k_2^*}, \dots, S_{k_{N-1}^*} \rightarrow S_{k_N^*}, S_{k_N^*} \rightarrow S_{k_1^*}$. Since the total number of such links is at most $N(\mathcal{G}, \mathcal{Z}^*)$, we must have

$$\text{Class-1 traffic} \leq N(\mathcal{G}, \mathcal{Z}^*) \Delta B. \quad (9)$$

Every byte of the Class-2 traffic occupies at least two bytes of the network capacity. Note that every logical topology has at most LN unidirectional links. Hence,

$$2 \times \text{Class-2 traffic} \leq (LN - N(\mathcal{G}, \mathcal{Z}^*)) \Delta B. \quad (10)$$

Based on (8)(9)(10), we then obtain

$$\begin{aligned}\mu(\mathcal{G}, f^*) &\leq \frac{1}{N} \left(BN(\mathcal{G}, \mathcal{Z}^*) + \frac{B}{2} (LN - N(\mathcal{G}, \mathcal{Z}^*)) \right) = \frac{BL}{2} + \frac{B}{2N} N(\mathcal{G}, \mathcal{Z}^*) \\ &\leq \frac{BL}{2} + \frac{BL}{2} \frac{R}{N-1} = \frac{N+R-1}{2N-2} \mu(\mathcal{G}(\mathcal{Z}^*), f^*).\end{aligned}$$

This completes the proof. $\square$

Discussion: Note that $\frac{N+R-1}{2N-2} < 1$ when $N > R + 1$. Then according to Theorem 6, no physical topology is optimal. Readers may wonder if it is possible to construct an optimal physical topology for the case where $N = R + 1$. The answer is no. When $N = R + 1$, $\frac{N+R-1}{2N-2} = 1$. This indicates that, if there exists an optimal physical topology, the “=” must hold for the inequality (7). This “=” requires that the $2R$ ports in each OCS (R ingress ports and R egress ports) must connect to $2R$ different EPS nodes. This is impossible because there are only $N = R + 1$ EPS nodes.

6 SCALABLE PHYSICAL TOPOLOGY DESIGNS WITH PERFORMANCE GUARANTEE

In this section, we propose three scalable physical topology designs and analyze their performance guarantees. The main results in this section are summarized in Table 1. Note that the maximum scale of different physical topologies depends on the OCS node port count R and the EPS node port count L . A commercially-available OCS node, e.g., the 3D-MEMS based OCS [3] and the Arrayed Wavelength Grating Routers [2], could have hundreds of ports. An off-the-shelf ToR EPS node, e.g., the CloudEngine 8800 [14], could have 32 ~ 64 ports. Hence, the physical topologies studied in this section could support large-scale ODCs with thousands or even tens of thousands of EPS nodes.

Table 1. The scalability and the performance ratios of different physical topologies.

Physical Topology	Maximum Number of EPS Nodes	Performance Ratio β
$\mathcal{G}_{\text{uniform}}$	R	1
$\mathcal{G}_{\text{eps-group}}$	$R\lceil L/3 \rceil$	$\min_f \left\{ \mu_{\lceil L/3 \rceil, B}^{\text{ideal}}(f) / \mu_{L, B}^{\text{ideal}}(f) \right\}$
$\mathcal{G}_{\text{eps-mesh}}$	R^2	$\min_f \left\{ \mu_{\lceil L/3 \rceil, B}^{\text{ideal}}(f) / \mu_{L, B}^{\text{ideal}}(f) \right\}$
$\mathcal{G}_{\text{ocs-mesh}}$	$\lfloor R/2 \rfloor \lfloor L/2 \rfloor$	$\min_f \left\{ \mu_{\lfloor L/2 \rfloor, B}^{\text{ideal}}(f) / \mu_{L, B}^{\text{ideal}}(f) \right\}$

6.1 EPS-Group based Physical Topology

Physical topology of $\mathcal{G}_{\text{eps-group}}$: Let $P = \lceil L/3 \rceil$. We create C EPS node groups, each of which contains P EPS nodes and is denoted by $\mathcal{A}_c$. Within each EPS node group, there are two links between every pair of EPS nodes. Each EPS node has $2(P-1)$ intra-group links, and the total number of links in each EPS node group is $P(P-1)$. Each EPS node also has P links connected to the OCS nodes to establish inter-group connections. It is easy to verify that the degree of each EPS node is $3P-2 \leq L$. We use $K = P^2$ OCS nodes and divide these OCS nodes into P equal-sized groups. We number the EPS nodes in a group from 1 to P . The p -th EPS node in a group has one link connected to every OCS node in the p -th OCS node group, where $p = 1, 2, \dots, P$. Since the number of ports of an OCS node is R , the number of EPS node groups of $\mathcal{G}_{\text{eps-group}}$ must satisfy $C \leq R$. Thus, the total number of EPS nodes in $\mathcal{G}_{\text{eps-group}}$ satisfies $N = CP \leq R\lceil L/3 \rceil$.

THEOREM 7. *For any traffic pattern f , the optimal throughput under $\mathcal{G}_{\text{eps-group}}$ satisfies*

$$\mu(\mathcal{G}_{\text{eps-group}}, f) \geq \mu_{\lceil L/3 \rceil, B}^{\text{ideal}}(f).$$

Fig. 3. Toy example of EPS-Group topology ($R = 2, N = 6$).

PROOF. Let $\mathcal{G}_{[L/3],B}^{\text{ideal}}$ be the ideal physical topology. According to Lemma 2, we need to show that any logical topology $\mathcal{G}_l$ formed over $\mathcal{G}_{[L/3],B}^{\text{ideal}}$ can be realized as an overlay topology over $\mathcal{G}_{\text{eps-group}}$.

Since each EPS node in $\mathcal{G}_{[L/3],B}^{\text{ideal}}$ has $P = \lceil L/3 \rceil$ number of uplinks, the following constraints must be satisfied

$$\begin{cases} \sum_{i=1}^N x_{ij}(\mathcal{G}_l) \leq \lceil L/3 \rceil = P, \forall j = 1, \dots, N, \\ \sum_{j=1}^N x_{ij}(\mathcal{G}_l) \leq \lceil L/3 \rceil = P, \forall i = 1, \dots, N. \end{cases} \quad (11)$$

Note that it may not always be possible to establish a logical link between two EPS nodes in $\mathcal{G}_{\text{eps-group}}$, especially when the two EPS nodes are in different EPS node groups and these two nodes have different intra-group indices. Hence, we are interested in constructing $\mathcal{G}_l$ as an overlay topology on top of $\mathcal{G}_{\text{eps-group}}$. For every pair of EPS nodes S_i and S_j in $\mathcal{G}_{\text{eps-group}}$, we would like to create $x_{ij}(\mathcal{G}_l)$ number of virtual logical links in between. Here, a **virtual logical link** is essentially a path between two EPS nodes, with each path segment occupying exactly one logical link.

For every EPS node pair (S_i, S_j) , we classify this pair's virtual logical links into $P = \lceil L/3 \rceil$ types:

- If S_i and S_j are in the same EPS node group, the type- p virtual logical links are two-hop paths with intermediate node being the p -th EPS node in this EPS node group. It is possible that the p -th EPS node is exactly S_i or S_j . In this case, this “two-hop path” degenerates to a single-hop path.
- If S_i and S_j belong to different EPS node groups, the type- p virtual logical links are three-hop paths with intermediate nodes being the p -th EPS node in the EPS node group of S_i and the p -th EPS node in the EPS node group of S_j . It is possible that the p -th EPS node in the EPS node group of S_i is exactly S_i , or the p -th EPS node in the EPS node group of S_j is exactly S_j . In this case, this “three-hop path” degenerates to a two/one-hop path.

We use $x_{ij}^p, p = 1, 2, \dots, P$ to denote the total number of type- p paths between S_i and S_j . According to Lemma 4, there exists an integer solution of $[x_{ij}^p]$ satisfying the following constraints:

- (1) $x_{ij} = \sum_{p=1}^P x_{ij}^p, \forall i, j = 1, \dots, N$;
- (2) $0 \leq \sum_{j=1}^N x_{ij}^p \leq 1, \forall i = 1, \dots, N, p = 1, \dots, P$;
- (3) $0 \leq \sum_{i \in \mathcal{A}_c} \sum_{j=1}^N x_{ij}^p \leq P, \forall c = 1, \dots, C, p = 1, \dots, P$;
- (4) $0 \leq \sum_{i=1}^N x_{ij}^p \leq 1, \forall j = 1, \dots, N, p = 1, \dots, P$;
- (5) $0 \leq \sum_{i=1}^N \sum_{j \in \mathcal{A}_c} x_{ij}^p \leq P, \forall c = 1, \dots, C, p = 1, \dots, P$.

Note that x_{ij}^p determines the underlying logical topology of the overlay topology $\mathcal{G}_l$. We need to show that this logical topology is compatible with the physical topology $\mathcal{G}_{\text{eps-group}}$.

Intra-group Logical Topology: If S_i and S_j are in the same EPS node group, then the number of links in between is $\sum_{u=1}^N x_{iu}^{(j \bmod P)} + \sum_{v=1}^N x_{vj}^{(i \bmod P)}$. Based on the constraints 2) and 4), it is easy to

check that the above number is no greater than 2, which is compatible with intra-group physical topology design of $\mathcal{G}_{\text{eps-group}}$.

Inter-group Logical Topology: Fix $p = 1, 2, \dots, P$. Let $S_{(c_1-1)P+p}$ and $S_{(c_2-1)P+p}$ be the p -th EPS node in the EPS node groups $\mathcal{A}_{c_1}$ and $\mathcal{A}_{c_2}$, respectively. Then the total number of links from $S_{(c_1-1)P+p}$ to $S_{(c_2-1)P+p}$ is $y_{c_1 c_2}^p = \sum_{i \in \mathcal{A}_{c_1}} \sum_{j \in \mathcal{A}_{c_2}} x_{ij}^p$. Based on the constraints 3) and 5), it is easy to verify that $\sum_{c_1=1}^C y_{c_1 c_2}^w \leq P$, $\sum_{c_2=1}^C y_{c_1 c_2}^w \leq P$. Then, using the same arguments in the proof of Theorem 5, we can prove that the logical topology $y_{c_1 c_2}^p$, $c_1, c_2 = 1, 2, \dots, C$ is realizable on top of $\mathcal{G}_{\text{eps-group}}$. This completes the proof. $\square$

Remark: Theorem 7 implies that $\beta(\mathcal{G}_{\text{eps-group}}) = \min_f \{\mu_{\lfloor L/3 \rfloor, B}^{\text{ideal}}(f) / \mu_{L, B}^{\text{ideal}}(f)\}$. $\mathcal{G}_{\text{eps-group}}$ is similar to the physical topology used by Flexfly [26]. In order to prove the performance guarantee in Theorem 7, we have restricted $\mathcal{G}_{\text{eps-group}}$ to use uniform mesh for its intra-group topology, which limits the scale of $\mathcal{G}_{\text{eps-group}}$. One could use flattened butterfly [16] for intra-group interconnect to increase the size of each EPS group of $\mathcal{G}_{\text{eps-group}}$, as suggested by Flexfly [26]. However, the performance guarantee would become much poorer.

6.2 EPS-Mesh based Physical Topology Design

Physical topology of $\mathcal{G}_{\text{eps-mesh}}$: Let $P = \lfloor L/3 \rfloor$. We arrange all the ESP nodes into a $C \times W$ mesh. The W EPS nodes in the c -th row forms an EPS node group, denoted by $\mathcal{A}_c$, $c = 1, 2, \dots, C$. We index all the EPS nodes row by row, from 1 to CW . There are $2P$ OCS nodes dedicated for each EPS node group $\mathcal{A}_c$. For every EPS node in $\mathcal{A}_c$, there is one bidirectional link between this EPS node and each of the $2P$ OCS nodes. There are also P OCS nodes dedicated for the EPS nodes in each column. For the w -th column, there is one bidirectional link between every EPS node in this column and every OCS node dedicated for this column. Since the number of ports of an OCS node is R , the number of EPS nodes in each row or each column must be no more than R . Thus, the total number of EPS nodes must satisfy $N = CW \leq R^2$.

Fig. 4. Toy example of EPS-Mesh topology ($R = 3, N = 9$).

THEOREM 8. For any traffic pattern f , the optimal throughput under $\mathcal{G}_{\text{eps-mesh}}$ satisfies

$$\mu(\mathcal{G}_{\text{eps-mesh}}, f) \geq \mu_{\lfloor L/3 \rfloor, B}^{\text{ideal}}(f).$$

PROOF. Let $\mathcal{G}_{\lfloor L/3 \rfloor, B}^{\text{ideal}}$ be the ideal physical topology. According to Lemma 2, we need to show that any logical topology $\mathcal{G}_l$ formed over $\mathcal{G}_{\lfloor L/3 \rfloor, B}^{\text{ideal}}$ can be realized as an overlay topology over $\mathcal{G}_{\text{eps-mesh}}$.

Since each EPS node in $\mathcal{G}_{[L/3],B}^{\text{ideal}}$ has $P = \lfloor L/3 \rfloor$ number of uplinks, the following constraints must be satisfied

$$\begin{cases} \sum_{i=1}^N x_{ij}(\mathcal{G}_l) \leq \lfloor L/3 \rfloor = P, \forall j = 1, \dots, N, \\ \sum_{j=1}^N x_{ij}(\mathcal{G}_l) \leq \lfloor L/3 \rfloor = P, \forall i = 1, \dots, N. \end{cases} \quad (12)$$

Similar to $\mathcal{G}_{\text{eps-group}}$, it may not always be possible to establish a logical link between two EPS nodes in $\mathcal{G}_{\text{eps-mesh}}$ as well, especially when the two EPS nodes are in different rows and different columns. Hence, we will construct $\mathcal{G}_l$ as an overlay topology on top of $\mathcal{G}_{\text{large}}$. For every pair of EPS nodes S_i and S_j in $\mathcal{G}_{\text{eps-mesh}}$, we would like to create $x_{ij}(\mathcal{G}_l)$ number of virtual logical links in between, and classify these virtual logical links into W types:

- If S_i and S_j are in the same row, the type- w virtual logical links are two-hop paths with intermediate node being the w -th EPS node in this row.
- If S_i and S_j belong to different rows, the type- w virtual logical links are three-hop paths with intermediate nodes being the w -th EPS node in the same row as S_i and the w -th EPS node in the same row as S_j .

We use x_{ij}^w , $w = 1, 2, \dots, W$ to denote the total number of type- w paths between S_i and S_j . According to Lemma 4, there exists an integer solution of x_{ij}^w , $w = 1, 2, \dots, W$ satisfying the following constraints:

- (1) $x_{ij}(\mathcal{G}_l) = \sum_{w=1}^W x_{ij}^w, \forall i, j = 1, \dots, N$;
- (2) $0 \leq \sum_{i \in \mathcal{A}_c} \sum_{j=1}^N x_{ij}^w \leq P, \forall c = 1, \dots, C, w = 1, \dots, W$;
- (3) $0 \leq \sum_{i=1}^N \sum_{j \in \mathcal{A}_c} x_{ij}^w \leq P, \forall c = 1, \dots, C, w = 1, \dots, W$.

Note that x_{ij}^w determines the underlying logical topology of the overlay topology $\mathcal{G}_l$. We need to show that this logical topology is compatible with the physical topology $\mathcal{G}_{\text{eps-group}}$.

Inter-group Logical Topology: Fix $w = 1, 2, \dots, W$. Let $S_{(c_1-1)W+w}$ and $S_{(c_2-1)W+w}$ be the w -th EPS node in the EPS node groups $\mathcal{A}_{c_1}$ and $\mathcal{A}_{c_2}$, respectively. Then the total number of links from $S_{(c_1-1)W+w}$ to $S_{(c_2-1)W+w}$ is $y_{c_1 c_2}^w = \sum_{i \in \mathcal{A}_{c_1}} \sum_{j \in \mathcal{A}_{c_2}} x_{ij}^w$. Based on the constraints 2) and 3), it is easy to verify that $\sum_{c_1=1}^C y_{c_1 c_2}^w \leq P, \sum_{c_2=1}^C y_{c_1 c_2}^w \leq P$. Then, using the same arguments in the proof of Theorem 5, we can prove that the logical topology $y_{c_1 c_2}^w, c_1, c_2 = 1, 2, \dots, C$ is realizable on the P OCS nodes dedicated for the w -th row.

Intra-group Logical Topology: Fix $c = 1, 2, \dots, C$. Let $S_{(c-1)W+w_1}$ and $S_{(c-1)W+w_2}$ be the w_1 -th and the w_2 -th EPS nodes in $\mathcal{A}_c$. Then the total number of links from $S_{(c-1)W+w_1}$ to $S_{(c-1)W+w_2}$ is $z_{w_1 w_2}^c = \sum_{u=1}^N x_{(c-1)W+w_1, u}^{w_2} + \sum_{v=1}^N x_{v, (c-1)W+w_2}^{w_1}$. According to Eqn. (13) and the constraints 2) and 3), we can verify that $\sum_{w_1=1}^W z_{w_1 w_2}^c \leq 2P, \sum_{w_2=1}^W z_{w_1 w_2}^c \leq 2P$. Take the first one for example:

$$\begin{aligned} \sum_{w_1=1}^W z_{w_1 w_2}^c &= \sum_{w_1=1}^W \sum_{u=1}^N x_{(c-1)W+w_1, u}^{w_2} + \sum_{w_1=1}^W \sum_{v=1}^N x_{v, (c-1)W+w_2}^{w_1} \\ &= \sum_{i \in \mathcal{A}_c} \sum_{u=1}^N x_{iu}^{w_2} + \sum_{v=1}^N x_{v, (c-1)W+w_2}(\mathcal{G}_l) \leq 2P. \end{aligned}$$

Again, using the same arguments in the proof of Theorem 5, we can prove that the logical topology $z_{w_1 w_2}^c, w_1, w_2 = 1, 2, \dots, W$ is realizable on the $2P$ OCS nodes dedicated for the EPS node group $\mathcal{A}_c$. This completes the proof. $\square$

Remark: Theorem 8 implies that $\beta(\mathcal{G}_{\text{eps-mesh}}) = \min_f \{\mu_{[L/3],B}^{\text{ideal}}(f) / \mu_{L,B}^{\text{ideal}}(f)\}$. The design of $\mathcal{G}_{\text{eps-mesh}}$ shares a similar idea as the physical topology used by 3D-Hyper-FleX-LION [17]. Although 3D-Hyper-FleX-LION uses a 3D-mesh instead, it is easy to generalize $\mathcal{G}_{\text{eps-mesh}}$ from 2D-mesh to 3D-mesh to further improve its scalability. Notably, one critical design of $\mathcal{G}_{\text{eps-mesh}}$ is that the

EPS ports used for intra-group interconnect is “two” times the EPS ports used for inter-group interconnect. This design allows us to attain the best performance guarantee.

6.3 OCS-Mesh based Physical Topology Design

Physical Topology of $\mathcal{G}_{\text{OCS-mesh}}$: We arrange C^2 OCS nodes into an $C \times C$ mesh, where $C = \lfloor L/2 \rfloor$. The EPS nodes are arranged into C groups, each of which contains W EPS nodes and is denoted by $\mathcal{A}_c, c = 1, 2, \dots, C$. For the EPS node group $\mathcal{A}_c$, we connect two directed links from the transmitters of each EPS node in $\mathcal{A}_c$ to the ingress ports of each OCS node in the c -th row; we also connect two directed links from the egress ports of each OCS node in the c -th column to the receivers of each EPS node in $\mathcal{A}_c$. Since each OCS node has R ingress/egress ports, the number of EPS nodes in each group must satisfy $W \leq \lfloor R/2 \rfloor$. Hence, the total number of EPS nodes must satisfy $N = CW \leq \lfloor L/2 \rfloor \lfloor R/2 \rfloor$.

Fig. 5. Toy example of OCS-Mesh topology ($R = 6, N = 9$). Some RX side links are omitted for clear visualization. Each EPS node in $\mathcal{A}_2$ is connected to two egress ports of OCS(1,2), OCS(2,2) and OCS(3,2). Each EPS node in $\mathcal{A}_3$ is connected to two egress ports of OCS(1,3), OCS(2,3) and OCS(3,3).

THEOREM 9. *For any traffic pattern f , the optimal throughput under $\mathcal{G}_{\text{OCS-mesh}}$ satisfies*

$$\mu(\mathcal{G}_{\text{OCS-mesh}}, f) \geq \mu_{\lfloor L/2 \rfloor, B}^{\text{ideal}}(f).$$

PROOF. Let $\mathcal{G}_{\lfloor L/2 \rfloor, B}^{\text{ideal}}$ be the ideal physical topology. According to Lemma 2, we need to show that any logical topology $\mathcal{G}_l$ formed over $\mathcal{G}_{\lfloor L/2 \rfloor, B}^{\text{ideal}}$ can be realized as an overlay topology over $\mathcal{G}_{\text{OCS-mesh}}$.

Since each EPS node in $\mathcal{G}_{\lfloor L/2 \rfloor, B}^{\text{ideal}}$ has $C = \lfloor L/2 \rfloor$ number of uplinks, the following constraints must be satisfied

$$\begin{cases} \sum_{i=1}^N x_{ij}(\mathcal{G}_l) \leq \lfloor L/2 \rfloor = C, \forall j = 1, \dots, N, \\ \sum_{j=1}^N x_{ij}(\mathcal{G}_l) \leq \lfloor L/2 \rfloor = C, \forall i = 1, \dots, N. \end{cases} \quad (13)$$

Again, we construct $\mathcal{G}_l$ as an overlay topology on top of $\mathcal{G}_{\text{OCS-mesh}}$. For every pair of EPS nodes S_i and S_j in $\mathcal{G}_{\text{OCS-mesh}}$, we would like to create $x_{ij}(\mathcal{G}_l)$ number of virtual logical links. Each virtual logical link is essentially a two-hop path $S_i \rightarrow S_k \rightarrow S_j, k = 1, 2, \dots, N$, and we use x_{ij}^k to denote the number of such virtual logical links.

Next, we would like to divide $x_{ij}(\mathcal{G}_l) = \sum_{k=1}^N x_{ij}^k$, such that the x_{ij}^k 's are compatible with the physical topology $\mathcal{G}_{\text{OCS-mesh}}$. Unlike $\mathcal{G}_{\text{eps-group}}$ and $\mathcal{G}_{\text{eps-mesh}}$, we need to apply Lemma 4 multiple times to obtain a valid decomposition of $x_{ij}(\mathcal{G}_l)$.

Step 1: Decompose $x_{ij}(\mathcal{G}_l) = \sum_{c=1}^C x_{ij}^{(c)}$ such that the following constraints are met:

- $0 \leq \sum_{j=1}^N x_{ij}^{(c)} \leq 1, \forall i = 1, \dots, N, c = 1, \dots, C;$
- $0 \leq \sum_{i=1}^N x_{ij}^{(c)} \leq 1, \forall j = 1, \dots, N, c = 1, \dots, C.$

Step 2: For every $c_0 = 1, 2, \dots, C$, decompose $x_{ij}^{(c_0)} = \sum_{w=1}^W x_{ij}^{(c_0-1)W+w}$ such that for any $w = 1, \dots, W$, the following constraints are met:

- $0 \leq \sum_{i \in \mathcal{A}_c} \sum_{j=1}^N x_{ij}^{(c_0-1)W+w} \leq 1, \forall c = 1, \dots, C;$
- $0 \leq \sum_{i=1}^N \sum_{j \in \mathcal{A}_c} x_{ij}^{(c_0-1)W+w} \leq 1, \forall c = 1, \dots, C.$

Here $\mathcal{A}_c = \{(c-1)W+1, (c-1)W+2, \dots, cW\}$.

Based on Step 1&2, we obtain an integer decomposition $x_{ij}(\mathcal{G}_l) = \sum_{c=1}^C x_{ij}^{(c)} = \sum_{k=1}^N x_{ij}^k$ satisfying

- (1) $0 \leq \sum_{j=1}^N \sum_{k \in \mathcal{A}_c} x_{ij}^k \leq 1, \forall i = 1, \dots, N, c = 1, \dots, C;$
- (2) $0 \leq \sum_{i=1}^N \sum_{k \in \mathcal{A}_c} x_{ij}^k \leq 1, \forall j = 1, \dots, N, c = 1, \dots, C;$
- (3) $0 \leq \sum_{i \in \mathcal{A}_c} \sum_{j=1}^N x_{ij}^k \leq 1, \forall c = 1, \dots, C, k = 1, \dots, N;$
- (4) $0 \leq \sum_{i=1}^N \sum_{j \in \mathcal{A}_c} x_{ij}^k \leq 1, \forall c = 1, \dots, C, k = 1, \dots, N.$

With x_{ij}^k , the total number of logical links needed from S_i to S_j can be computed as $y_{ij} = \sum_{k=1}^N x_{ij}^k + \sum_{k=1}^N x_{ik}^j$. Next, we will show that y_{ij} is compatible with $\mathcal{G}_{\text{ocs-mesh}}$.

Recall that the OCS nodes in $\mathcal{G}_{\text{ocs-mesh}}$ are arranged as a 2D mesh. Consider the OCS node in the c_1 -th row and c_2 -th column. This OCS node connects the transmitters of the EPS node in group $\mathcal{A}_{c_1}$ to the receivers of the EPS node in group $\mathcal{A}_{c_2}$. Then, if the following constraints are satisfied

$$\sum_{j \in \mathcal{A}_{c_2}} y_{ij} \leq 2, \sum_{i \in \mathcal{A}_{c_1}} y_{ij} \leq 2, \quad (14)$$

the logical topology $y_{ij}, i \in \mathcal{A}_{c_1}, j \in \mathcal{A}_{c_2}$ will be realizable on this OCS. (14) can be verified below:

$$\sum_{j \in \mathcal{A}_{c_2}} y_{ij} = \sum_{j \in \mathcal{A}_{c_2}} \sum_{k=1}^N x_{ij}^k + \sum_{j \in \mathcal{A}_{c_2}} \sum_{k=1}^N x_{ik}^j \leq 2; \sum_{i \in \mathcal{A}_{c_1}} y_{ij} = \sum_{i \in \mathcal{A}_{c_1}} \sum_{k=1}^N x_{ij}^k + \sum_{i \in \mathcal{A}_{c_1}} \sum_{k=1}^N x_{ik}^j \leq 2.$$

This completes the proof. $\square$

Remark: Theorem 9 implies that $\beta(\mathcal{G}_{\text{ocs-mesh}}) = \min_f \{\mu_{[L/2],B}^{\text{ideal}}(f) / \mu_{L,B}^{\text{ideal}}(f)\}$. The design of $\mathcal{G}_{\text{ocs-mesh}}$ shares a similar idea as the physical topology used by Sirius [2]. One critical difference between $\mathcal{G}_{\text{ocs-mesh}}$ and Sirius' physical topology is that we connect "two" links instead of one between each connected pair of EPS node and OCS node. This subtle design allows us to derive the performance guarantee in Theorem 9.

7 PERFORMANCE RATIO ANALYSIS

We have performed theoretical analysis for four physical topologies $\mathcal{G}_{\text{uniform}}$, $\mathcal{G}_{\text{eps-group}}$, $\mathcal{G}_{\text{eps-mesh}}$ and $\mathcal{G}_{\text{ocs-mesh}}$, and related their throughput metrics with that of the ideal physical topology. $\mathcal{G}_{\text{uniform}}$ is optimal but scales poorly. $\mathcal{G}_{\text{eps-group}}$, $\mathcal{G}_{\text{eps-mesh}}$ and $\mathcal{G}_{\text{ocs-mesh}}$ have better scalability but are sub-optimal. According to Theorem 7, Theorem 8 and Theorem 9, It is easy to obtain

$$\begin{cases} \beta(\mathcal{G}_{\text{eps-group}}) = \min_f \{\beta(\mathcal{G}_{\text{eps-group}}, f)\}, \text{ where } \beta(\mathcal{G}_{\text{eps-group}}, f) = \mu_{[L/3],B}^{\text{ideal}}(f) / \mu_{L,B}^{\text{ideal}}(f), \\ \beta(\mathcal{G}_{\text{eps-mesh}}) = \min_f \{\beta(\mathcal{G}_{\text{eps-mesh}}, f)\}, \text{ where } \beta(\mathcal{G}_{\text{eps-mesh}}, f) = \mu_{[L/3],B}^{\text{ideal}}(f) / \mu_{L,B}^{\text{ideal}}(f), \\ \beta(\mathcal{G}_{\text{ocs-mesh}}) = \min_f \{\beta(\mathcal{G}_{\text{ocs-mesh}}, f)\}, \text{ where } \beta(\mathcal{G}_{\text{ocs-mesh}}, f) = \mu_{[L/2],B}^{\text{ideal}}(f) / \mu_{L,B}^{\text{ideal}}(f). \end{cases} \quad (15)$$

In this section, we will calculate the approximated values of $\beta(\mathcal{G}_{\text{eps-group}})$, $\beta(\mathcal{G}_{\text{eps-mesh}})$ and $\beta(\mathcal{G}_{\text{ocs-mesh}})$. We focus on two cases below:

- (1) ToE can generate multiple logical topologies to serve each traffic matrix. This setting was adopted by [2, 19–21].
- (2) ToE generates only one logical topology (i.e., $M = 1$) for each traffic matrix. This setting was adopted by [4, 9, 12].

For the first case, we obtain some theoretical results in Section 7.1. The second case is hard to analyze, and we could only obtain some approximated numerical results in Section 7.2.

7.1 Calculating $\beta(\mathcal{G}_{\text{eps-group}})$, $\beta(\mathcal{G}_{\text{eps-mesh}})$ and $\beta(\mathcal{G}_{\text{ocs-mesh}})$ in Case 1

DEFINITION 3. A traffic pattern f is q -decomposable if there exists q permutation traffic matrices $Z_1, Z_2, \dots, Z_q$ and $\phi_1 + \phi_2 + \dots + \phi_q = 1, 0 \leq \phi_i \leq 1$, such that $f \leq \sum_{i=1}^q \phi_i Z_i$.

DEFINITION 4. A traffic pattern f is a normalized traffic pattern if the following condition is met

$$\max \left\{ \max_{i=1}^N \left\{ \sum_{j=1}^N f_{ij} \right\}, \max_{j=1}^N \left\{ \sum_{i=1}^N f_{ij} \right\} \right\} = 1.$$

LEMMA 10. For any normalized traffic pattern f , we must have $\mu_{L,B}^{\text{ideal}}(f) \leq BL$.

PROOF. Without loss of generality, we assume that the egress traffic from S_1 sums to 1, e.g., $\sum_{j=1}^N f_{1j} = 1$. Consider the network control policy that achieves the throughput value $\mu_{L,B}^{\text{ideal}}(f)$. Then, $\mu_{L,B}^{\text{ideal}}(f) \sum_{j=1}^N f_{1j} = \mu_{L,B}^{\text{ideal}}(f)$ amount of traffic can be delivered from the EPS node S_1 to other EPS nodes in a unit time period. On the other hand, the total egress bandwidth of S_1 is upper bounded by BL , because S_1 has L ports. Thus, $\mu_{L,B}^{\text{ideal}}(f) \leq BL$. $\square$

LEMMA 11. Given a q -decomposable traffic pattern f that lasts Δ amount of time, if the ToE policy allows generating M OCS configurations, then $\mu_{L,B}^{\text{ideal}}(f) \geq BL(1 - \frac{q}{ML})(1 - \frac{\delta M}{\Delta})$.

PROOF. We split Δ into M time slots, such that $\Delta^{(1)} = \Delta^{(2)} = \dots = \Delta^{(M)} = \Delta/M - \delta$. In each time slot, L permutations can be formed because each EPS node has L bidirectional links connected the OCS. In total, we can form ML permutations.

Let $f \leq \sum_{i=1}^q \phi_i Z_i$ be the q -decomposition of f . For every permutation $Z_i, i = 1, 2, \dots, q$, we configure OCSs such that Z_i is formed $\lceil (ML - q)\phi_i \rceil$ times. It is easy to verify that

$$\sum_{i=1}^q \lceil (ML - q)\phi_i \rceil \leq \sum_{i=1}^q ((ML - q)\phi_i + 1) = ML.$$

Hence, the above ToE strategy is feasible. Under this strategy, the bandwidth allocated to the ODC satisfies

$$\frac{1}{\Delta} \sum_{i=1}^q \lceil (ML - q)\phi_i \rceil B(\Delta/M - \delta) Z_i \geq \frac{(ML - q)B(\Delta/M - \delta)}{\Delta} \sum_{i=1}^q \phi_i Z_i \geq \frac{(ML - q)B(\Delta/M - \delta)}{\Delta} f.$$

Therefore, under the above ToE strategy and the one-hop routing, we can achieve a throughput value $\frac{(ML - q)B(\Delta/M - \delta)}{\Delta}$. Based on the definition of $\mu(\mathcal{G}, f)$, we must have

$$\mu_{L,B}^{\text{ideal}}(f) \geq \frac{(ML - q)B(\Delta/M - \delta)}{\Delta} = BL(1 - \frac{q}{ML})(1 - \frac{\delta M}{\Delta}).$$

$\square$

Note that Lemma 11 holds for any M . According to the definitions of $\beta(\mathcal{G}_{\text{eps-group}})$, $\beta(\mathcal{G}_{\text{eps-mesh}})$ and $\beta(\mathcal{G}_{\text{ocs-mesh}})$ in (15), it is sufficient to consider normalized traffic patterns, because $\mu_{L,B}^{\text{ideal}}(f)$ increases linearly if we scales f linearly. Further, based on the Birkhoff and von Neumann theorem [5], any normalized traffic pattern f is q -decomposable with $q \leq N^2 - 2N + 2$. Combining all the above analysis with Lemma 10, we then obtain

COROLLARY 11.1. Let $\gamma(M, L) = \left(1 - \frac{N^2 - 2N + 2}{ML}\right) \left(1 - \frac{\delta M}{\Delta}\right)$, we have

$$\frac{\lceil L/3 \rceil}{L} \max_M \{\gamma(M, \lceil L/3 \rceil)\} \leq \beta(\mathcal{G}_{\text{eps-group}}) \leq \frac{\lceil L/3 \rceil}{L} \left(\max_M \{\gamma(M, L)\} \right)^{-1},$$

$$\frac{\lfloor L/3 \rfloor}{L} \max_M \{\gamma(M, \lfloor L/3 \rfloor)\} \leq \beta(\mathcal{G}_{\text{eps-mesh}}) \leq \frac{\lfloor L/3 \rfloor}{L} \left(\max_M \{\gamma(M, L)\} \right)^{-1},$$

$$\frac{\lfloor L/2 \rfloor}{L} \max_M \{\gamma(M, \lfloor L/2 \rfloor)\} \leq \beta(\mathcal{G}_{\text{ocs-mesh}}) \leq \frac{\lfloor L/2 \rfloor}{L} \left(\max_M \{\gamma(M, L)\} \right)^{-1}.$$

When the OCS reconfiguration delay $\delta = 0$, $\gamma(M, l)$ can be arbitrarily close to 1 as we increase the number of OCS configurations M . In this case, we have

$$\text{COROLLARY 11.2. If } \delta = 0, \text{ then } \beta(\mathcal{G}_{\text{eps-group}}) = \frac{\lfloor L/3 \rfloor}{L}, \beta(\mathcal{G}_{\text{eps-mesh}}) = \frac{\lfloor L/3 \rfloor}{L}, \beta(\mathcal{G}_{\text{ocs-mesh}}) = \frac{\lfloor L/2 \rfloor}{L}.$$

7.2 Calculating $\beta(\mathcal{G}_{\text{eps-group}})$, $\beta(\mathcal{G}_{\text{eps-mesh}})$ and $\beta(\mathcal{G}_{\text{ocs-mesh}})$ in Case 2 Where $M = 1$

$M = 1$ means that only one logical topology is generated for each traffic pattern f (see the definition of ToE in Section 3.1). When $M = 1$, the lower bound of $\mu_{L,B}^{\text{ideal}}(f)$ given by Lemma 11 can be extremely loose, and thus may not be useful for calculating $\beta(\mathcal{G}_{\text{eps-group}})$, $\beta(\mathcal{G}_{\text{eps-mesh}})$ and $\beta(\mathcal{G}_{\text{ocs-mesh}})$. In this case, we perform numerical analysis instead. Note that exhaustively enumerating all the different traffic matrices is not feasible. We simply generate random traffic matrices to perform the calculation. We use two approaches to generate random traffic matrices:

Approach 1: Randomly generate Q permutation traffic matrices $f^{(1)}, f^{(2)}, \dots, f^{(Q)}$ and Q parameters $0 \leq \gamma_1, \gamma_2, \dots, \gamma_Q \leq 1$ satisfying $\sum_{q=1}^Q \gamma_q = 1$. Let $f = \sum_{q=1}^Q \gamma_q f^{(q)}$. It is easy to verify that the traffic matrix f generated with this approach is a normalized traffic pattern.

Approach 2: Generate a random non-negative number for each entry f_{ij} , $i \neq j$, and then normalize f such that $\max\{\max_{j=1}^N \{\sum_{i=1}^N f_{ij}\}, \max_{i=1}^N \{\sum_{j=1}^N f_{ij}\}\} = 1$.

For every N and every L , we generate a number of normalized random traffic patterns using the above two approaches, and calculate throughput values $\mu_{\lfloor L/3 \rfloor, B}^{\text{ideal}}(f)$, $\mu_{\lfloor L/3 \rfloor, B}^{\text{ideal}}(f)$, $\mu_{\lfloor L/2 \rfloor, B}^{\text{ideal}}(f)$, $\mu_{L, B}^{\text{ideal}}(f)$ for every traffic pattern f . We use three approaches to calculate the throughput values:

Joint Optimization: Directly solve the optimization problem (5) using Gurobi [22].

Decoupled Optimization: First, compute a logical topology $\bar{\mathcal{G}}_l^{(1)}$ that best approximates the traffic pattern f by solving the following optimization problem:

$$\begin{aligned} & \max \sum_{i=1}^N \sum_{j=1}^N (L f_{ij} - \lfloor L f_{ij} \rfloor) x_{ij}(\bar{\mathcal{G}}_l^{(1)}) \\ & \text{s.t. } \lfloor L f_{ij} \rfloor \leq x_{ij}(\bar{\mathcal{G}}_l^{(1)}) \leq \lceil L f_{ij} \rceil, \text{ for all } i, j = 1, 2, \dots, N, \\ & \sum_{k=1}^N x_{ik}(\bar{\mathcal{G}}_l^{(1)}) \leq L, \sum_{k=1}^N x_{kj}(\bar{\mathcal{G}}_l^{(1)}) \leq L, \text{ for all } i, j = 1, 2, \dots, N. \end{aligned} \quad (16)$$

Then, solve the optimization problem (5) with the logical topology $\mathcal{G}_l^{(1)}$ being fixed as $\bar{\mathcal{G}}_l^{(1)}$.

Fast Approximation: First, compute a logical topology $\bar{\mathcal{G}}_l^{(1)}$ by solving (16). Second, find all the shortest paths for every EPS node pair in $\bar{\mathcal{G}}_l^{(1)}$. Third, calculate the throughput value based on Equal-Cost Multi-Path (ECMP) routing.

The Joint Optimization approach solves topology and routing jointly, and gives the exact throughput values. However, this approach has the highest computational complexity. The Decoupled Optimization approach reduces algorithmic complexity by solving topology and routing separately. The Fast Approximation approach further reduces computational complexity by using a fixed ECMP routing. Notably, the second and the third approaches can only compute approximated throughput and performance ratio values.

7.2.1 Approximating $\beta(\mathcal{G}_{\text{eps-group}})$, $\beta(\mathcal{G}_{\text{eps-mesh}})$ and $\beta(\mathcal{G}_{\text{ocs-mesh}})$. We generate a list of (N, L) pairs: $(7, 10)$, $(8, 10)$, $(9, 10)$, $(15, 10)$, $(20, 10)$, $(30, 16)$, $(40, 16)$, $(128, 32)$, $(256, 32)$, $(512, 64)$ and $(1024, 64)$. For each (N, L) pair, we generate a set of 60 normalized traffic patterns, and denote this set by $\mathcal{F}_{N,L}$. For every traffic pattern $f \in \mathcal{F}_{N,L}$, we use the above three approaches to compute the per-traffic-pattern performance ratio $\beta(\mathcal{G}, f)$, $\mathcal{G} \in \{\mathcal{G}_{\text{eps-group}}, \mathcal{G}_{\text{eps-mesh}}, \mathcal{G}_{\text{ocs-mesh}}\}$ as defined in (15).

Fig. 6. Calculating $\beta(\mathcal{G}, f)$ using Joint Optimization, Decoupled Optimization and Fast Approximation.

Since the Joint Optimization approach can compute the exact per-traffic-pattern performance ratio values, we consider $\beta_{\text{joint}}(\mathcal{G}) = \min_{f \in \mathcal{F}_{N,L}} \beta_{\text{joint}}(\mathcal{G}, f)$ as the best approximation to $\beta(\mathcal{G})$ for every $\mathcal{G} \in \{\mathcal{G}_{\text{eps-group}}, \mathcal{G}_{\text{eps-mesh}}, \mathcal{G}_{\text{ocs-mesh}}\}$. However, calculating $\beta_{\text{joint}}(\mathcal{G}, f)$ is computationally expensive. The Joint Optimization approach can only compute $\beta_{\text{joint}}(\mathcal{G})$ for the (N, L) pairs $(7, 10)$, $(8, 10)$ and $(9, 10)$. When $N > 9$, the Joint Optimization approach may not be able to compute a solution even after running a few hours. Hence, in the following, we study how to approximate $\beta_{\text{joint}}(\mathcal{G})$ using the other two approaches.

We first use the Decoupled Optimization approach to approximate $\beta_{\text{joint}}(\mathcal{G})$. For every $f \in \mathcal{F}_{N,L}$, we compute $\beta_{\text{decoupled}}(\mathcal{G}, f)$ using the Decoupled Optimization approach. As shown in Fig. 6, the values of $\beta_{\text{decoupled}}(\mathcal{G}, f)$ exhibit much higher variance than that of $\beta_{\text{joint}}(\mathcal{G}, f)$. If we use $\min_{f \in \mathcal{F}_{N,L}} \beta_{\text{decoupled}}(\mathcal{G}, f)$ to approximate $\beta_{\text{decoupled}}(\mathcal{G})$, the gap between the estimated value and the true value would be too large. Hence, we adopt the following heuristic. Let $\alpha_{\text{decoupled}}(\mathcal{G}, p)$ be the p -th percentile value among all the values of $\beta_{\text{decoupled}}(\mathcal{G}, f)$. We aim to find p such that $\alpha_{\text{decoupled}}(\mathcal{G}, p) \approx \beta_{\text{joint}}(\mathcal{G})$. From Fig. 6, we can see that $p(\mathcal{G}_{\text{eps-group}}) \approx 25\text{th}$, $p(\mathcal{G}_{\text{eps-mesh}}) \approx 35\text{th}$ and $p(\mathcal{G}_{\text{ocs-mesh}}) \approx 25\text{th}$ when $(N, L) = (7, 10)$, $(8, 10)$ or $(9, 10)$. Then, we use $\alpha_{\text{decoupled}}(\mathcal{G}, p(\mathcal{G}))$ to approximate $\beta_{\text{joint}}(\mathcal{G})$ for larger data centers. However, even this Decoupled Optimization approach

Table 2. The values of $p'(\mathcal{G})$ for different (N, L) 's. When $N \leq 9$, we pick $p'(\mathcal{G})$ such that $\alpha_{\text{fast}}(\mathcal{G}, p'(\mathcal{G})) \approx \beta_{\text{joint}}(\mathcal{G})$. When $9 < N \leq 40$, we pick $p'(\mathcal{G})$ such that $\alpha_{\text{fast}}(\mathcal{G}, p'(\mathcal{G})) \approx \alpha_{\text{decoupled}}(\mathcal{G}, p(\mathcal{G}))$.

	(7,10)	(8,10)	(9,10)	(15,10)	(20,10)	(30,16)	(40,16)
$p'(\mathcal{G}_{\text{eps-group}})$	36th	26th	23th	35th	33th	58th	56th
$p'(\mathcal{G}_{\text{eps-mesh}})$	28th	31th	20th	36th	36th	56th	61th
$p'(\mathcal{G}_{\text{ocs-mesh}})$	48th	31th	26th	46th	45th	65th	63th

Table 3. Approximating $\beta(\mathcal{G}_{\text{eps-group}})$, $\beta(\mathcal{G}_{\text{eps-mesh}})$ and $\beta(\mathcal{G}_{\text{ocs-mesh}})$ using the Fast Approximation approach.

	(7,10)	(8,10)	(9,10)	(15,10)	(20,10)	(30,16)
$\beta(\mathcal{G}_{\text{eps-group}})$	0.34	0.34	0.33	[0.28,0.31]	[0.27,0.3]	[0.25,0.28]
$\beta(\mathcal{G}_{\text{eps-mesh}})$	0.24	0.23	0.22	[0.17,0.21]	[0.16,0.19]	[0.18,0.22]
$\beta(\mathcal{G}_{\text{ocs-mesh}})$	0.45	0.45	0.44	[0.39,0.43]	[0.36,0.42]	[0.36,0.42]
	(40,16)	(128,32)	(256,32)	(512,64)	(1024,64)	
$\beta(\mathcal{G}_{\text{eps-group}})$	[0.24,0.26]	[0.22,0.27]	[0.23,0.27]	[0.24,0.28]	[0.25,0.29]	
$\beta(\mathcal{G}_{\text{eps-mesh}})$	[0.18,0.2]	[0.2,0.25]	[0.19,0.24]	[0.22,0.26]	[0.21,0.28]	
$\beta(\mathcal{G}_{\text{ocs-mesh}})$	[0.35,0.4]	[0.37,0.45]	[0.39,0.44]	[0.38,0.43]	[0.37,0.45]	

cannot scale beyond 40 EPS nodes. When $N > 40$, the Decoupled Optimization approach may not produce any solution after running a few hours.

We then use the Fast Approximation approach to approximate $\beta_{\text{joint}}(\mathcal{G})$. As shown in Fig. 6, the values of $\beta_{\text{fast}}(\mathcal{G}, f)$ exhibit even higher variance, and the gap between $\min_{f \in \mathcal{F}_{N,L}} \beta_{\text{fast}}(\mathcal{G}, f)$ and $\beta_{\text{decoupled}}(\mathcal{G})$ would be even larger. To reduce the estimation gap, we let $\alpha_{\text{fast}}(\mathcal{G}, p')$ be the p' -th percentile value among all the values of $\beta_{\text{fast}}(\mathcal{G}, f)$. We aim to find p' such that $\alpha_{\text{fast}}(\mathcal{G}, p') \approx \beta_{\text{joint}}(\mathcal{G})$ when $N \leq 9$ and $\alpha_{\text{fast}}(\mathcal{G}, p') \approx \alpha_{\text{decoupled}}(\mathcal{G}, p(\mathcal{G}))$ when $9 < N \leq 40$. The values of $p'(\mathcal{G})$ vary for different physical topologies and different (N, L) pairs. The detailed values are summarized in Table 2. For every $\mathcal{G} \in \{\mathcal{G}_{\text{eps-group}}, \mathcal{G}_{\text{eps-mesh}}, \mathcal{G}_{\text{ocs-mesh}}\}$, we pick the minimum and the maximum values of $p'(\mathcal{G})$, and denote them by $p'_{\min}(\mathcal{G})$ and $p'_{\max}(\mathcal{G})$. Then, we can generate a range $[\alpha_{\text{fast}}(\mathcal{G}, p'_{\min}(\mathcal{G})), \alpha_{\text{fast}}(\mathcal{G}, p'_{\max}(\mathcal{G}))]$ to approximate $\beta_{\text{joint}}(\mathcal{G}) \approx \beta(\mathcal{G})$. We summarize the approximation results for all $N > 9$ in Table 3.

Note: The idea of using percentile values to estimate $\beta(\mathcal{G})$'s is just a heuristic. There is **no guarantee** that the true values of $\beta(\mathcal{G})$ are indeed within the ranges provided by Table 3. Nevertheless, we believe that our numerical estimation gives readers a sense on the magnitude of $\beta(\mathcal{G})$.

8 CONCLUSION

We study physical topology design for ODCs, and offer a novel methodology to analyze the performance guarantee for different physical topologies in this paper. Based on our methodology, we prove that the uniform bipartite graph based physical topology is optimal for small-scale data centers with $N \leq R$, while optimal physical topologies do not exist when $N > R$. We also design three physical topologies that support larger-scale ODCs, and prove their performance gaps with respect to the ideal physical topology. The methodology in this paper provides a theoretical foundation for ODC physical topology design, and may help network operators design physical topologies with a better guarantee.

There are a number of interesting directions for future work. First, we do not assume any prior knowledge on the traffic patterns while we define the performance metric $\beta(\mathcal{G})$. In practice, different ToR switches may belong to different production groups, and the intra-group traffic can be higher than the inter-group traffic. It would be interesting to study how such traffic knowledge affects physical topology design. Second, in addition to throughput, deployment/management/expansion complexities may also affect the physical topology design for ODC. It remains open to find a physical topology with the best tradeoff among different performance metrics for ODC.

9 ACKNOWLEDGEMENT

This work was supported by the NSF China under Grant 61902246. We also thank our shepherd Rachee Singh, and the SIGMETRICS reviewers.

REFERENCES

- [1] Mohammad Al-Fares, Alexander Loukissas, and Amin Vahdat. 2008. A scalable, commodity data center network architecture. In *SIGCOMM*.
- [2] Hitesh Ballani, Paolo Costa, Raphael Behrendt, Daniel Cletheroe, Istvan Haller, Krzysztof Jozwik, Fotini Karinou, Sophie Lange, Kai Shi, Benn Thomsen, and Hugh Williams. 2020. Sirius: A Flat Datacenter Network with Nanosecond Optical Switching. In *SIGCOMM*.
- [3] Inc. CALIENT Technologies. [n.d.]. <https://www.calient.net/>.
- [4] Peirui Cao, Shizhen Zhao, Min Yee Teh, Yunzhuo Liu, and Xinbing Wang. 2021. TROD: Evolving From Electrical Data Center to Optical Data Center. In *ICNP*.
- [5] Cheng-shang Chang, Wen-Jyh Chen, and Hsiang-Yi Huang. 2000. Birkhoff-von Neumann Input-Buffered Crossbar Switches. In *INFOCOM*.
- [6] Kai Chen, Ankit Singla, Atul Singh, Kishore Ramachandran, Lei Xu, Yueping Zhang, Xitao Wen, and Yan Chen. 2012. OSA: An optical switching architecture for data center networks with unprecedented flexibility. In *NSDI*.
- [7] Li Chen, Kai Chen, Zhonghua Zhu, Minlan Yu, George Porter, Chunming Qiao, and Shan Zhong. 2017. Enabling wide-spread communications on optical fabric with megaswitch. In *NSDI*.
- [8] Nandita Dukkkipati and Nick McKeown. 2006. Why Flow-Completion Time is the Right Metric for Congestion Control. In *SIGCOMM Review*.
- [9] Nathan Farrington, George Porter, Sivasankar Radhakrishnan, Hamid Hajabdolali Bazzaz, Vikram Subramanya, Yeshaiah Fainman, George Papen, and Amin Vahdat. 2010. Helios: a hybrid electrical/optical switch architecture for modular data centers. In *SIGCOMM*.
- [10] Monia Ghobadi, Ratul Mahajan, Amar Phanishayee, Nikhil Devanur, Janardhan Kulkarni, Gireeja Ranade, Pierre Alexandre Blanche, Houman Rastegarfar, Madeleine Glick, and Daniel Kilper. 2016. ProjecToR: Agile Reconfigurable Data Center Interconnect. In *SIGCOMM*.
- [11] Albert Greenberg, James R. Hamilton, Navendu Jain, Srikanth Kandula, and Sudipta Sengupta. 2009. VL2: A Scalable and Flexible Data Center Network. In *SIGCOMM*.
- [12] Navid Hamedazimi, Zafar Qazi, Himanshu Gupta, Vyas Sekar, Samir R. Das, Jon P. Longtin, Himanshu Shah, and Ashish Tanwer. 2014. FireFly: A Reconfigurable Wireless Data Center Fabric Using Free-Space Optics. In *SIGCOMM*.
- [13] Chi-Yao Hong, Srikanth Kandula, Ratul Mahajan, Ming Zhang, Vijay Gill, Mohan Nanduri, and Roger Wattenhofer. 2013. Achieving High Utilization with Software-Driven WAN. In *SIGCOMM*.
- [14] Co. Huawei Technologies. [n.d.]. <https://e.huawei.com/en/products/enterprise-networking/switches/data-center-switches/ce8800>.
- [15] Robert W. Irving and Mark R. Jerrum. 1992. Three-dimensional Statistical Data Security Problems. *SIAM J. Comput.* 23 (October 1992), 170–184.
- [16] John Kim, William J. Dally, and Dennis Abts. 2007. Flattened Butterfly : A Cost-Efficient Topology for High-Radix Networks. In *ISCA*.
- [17] Gengchen Liu, Roberto Proietti, Marjan Fariborz, Pouya Fotouhi, Xian Xiao, and S.J. Ben Yoo. 2020. Architecture and Performance Studies of 3D-Hyper-FleX-LION for Reconfigurable All-to-All HPC Networks. In *SC*.
- [18] He Liu, Feng Lu, Alex Forencich, Rishi Kapoor, Malveeka Tewari, Geoffrey M Voelker, George Papen, Alex C Snoeren, and George Porter. 2014. Circuit Switching Under the Radar with REACToR. In *NSDI*.
- [19] He Liu, Matthew K. Mukerjee, Conglong Li, Nicolas Feltman, George Papen, Stefan Savage, Srinivasan Seshan, Geoffrey M. Voelker, David G. Andersen, Michael Kaminsky, George Porter, and Alex C. Snoeren. 2015. Scheduling Techniques for Hybrid Circuit/Package Networks. In *CoNEXT*.

- [20] William M Mellette, Rajdeep Das, Yibo Guo, Rob McGuinness, Alex C Snoeren, and George Porter. 2020. Expanding across time to deliver bandwidth efficiency and low latency. In *NSDI*.
- [21] William M. Mellette, Rob McGuinness, Arjun Roy, Alex Forencich, and George Porter. 2017. RotorNet: A Scalable, Low-complexity, Optical Datacenter Network. In *SIGCOMM*.
- [22] Gurobi Optimization. [n.d.]. <https://www.gurobi.com/>.
- [23] George Porter, Richard Strong, Nathan Farrington, Alex Forencich, Pang Chen-Sun, Tajana Rosing, Yeshaiah Fainman, George Papen, and Amin Vahdat. 2013. Integrating microsecond circuit switching into the data center. In *SIGCOMM*.
- [24] Arjun Singh, Joon Ong, Amit Agarwal, Glen Anderson, Ashby Armistead, Roy Bannon, Seb Boving, Gaurav Desai, Bob Felderman, Paulie Germano, et al. 2015. Jupiter Rising: A Decade of Clos Topologies and Centralized Control in Google's Datacenter Network. In *SIGCOMM*.
- [25] Guohui Wang, David G Andersen, Michael Kaminsky, Konstantina Papagiannaki, TS Ng, Michael Kozuch, and Michael Ryan. 2010. c-Through: Part-time optics in data centers. In *SIGCOMM*.
- [26] Ke Wen, Payman Samadi, Sebastien Rumley, Christine P. Chen, Yiwen Shen, Meisam Bahadori, Keren Bergman, and Jeremiah Wilke. 2016. Flexfly: Enabling a Reconfigurable Dragonfly through Silicon Photonics. In *SC*.
- [27] Mingyang Zhang, Radhika Niranjana Mysore, Sucha Supittayapornpong, and Ramesh Govindan. 2019. Understanding lifecycle management complexity of datacenter topologies. In *NSDI*.
- [28] Shizhen Zhao, Rui Wang, Junlan Zhou, Joon Ong, Jeffery C. Mogul, and Amin Vahdat. 2019. Minimal Rewiring: Efficient Live Expansion for Clos Data Center Networks: Extended Version. In *NSDI* <https://ai.google/research/pubs/pub47492>.

A NP-COMPLETENESS OF CALCULATING $\mu(\mathcal{G}, f)$

We prove the following theorem in this section.

THEOREM 12. *The problem of calculating $\mu(\mathcal{G}, f)$ for a given physical topology $\mathcal{G}$ and a given traffic pattern f is NP-Complete.*

We prove Theorem 12 by reducing the following three-dimensional sum problem to the problem of calculating $\mu(\mathcal{G}, f)$. The three-dimensional sum problem was proven to be NP-Complete [15].

DEFINITION 5. (Three-dimensional Sum Problem) *Given a positive integer n and three non-negative integer matrices $D(i, k)$, $E(k, j)$ and $F(i, j)$, where $i, j, k = 1, 2, \dots, n$ and*

$$\sum_{i=1}^n \sum_{k=1}^n D(i, k) = \sum_{k=1}^n \sum_{j=1}^n E(k, j) = \sum_{i=1}^n \sum_{j=1}^n F(i, j),$$

does there exist a non-negative integer assignment of $A(i, j, k)$ such that

$$\begin{cases} \sum_{j=1}^n A(i, j, k) = D(i, k), \\ \sum_{i=1}^n A(i, j, k) = E(k, j), \\ \sum_{k=1}^n A(i, j, k) = F(i, j). \end{cases} \quad (17)$$

We are now ready to prove Theorem 12.

PROOF. Given an instance of the three-dimensional sum problem, we construct an ODC physical topology $\mathcal{G}$ and a traffic matrix f as follows. $\mathcal{G}$ has $2n$ EPS nodes and n OCS nodes. The $2n$ EPS nodes are separated into two groups, each with n EPS nodes. There are $D(i, k)$ number of bidirectional links between the i -th EPS node in Group 1 and the k -th OCS node; there are $E(k, j)$ number of bidirectional links between the k -th OCS node and the j -th EPS node in Group 2. Each link is assumed to have its link capacity $B = 1$. The traffic demand from the i -th EPS node in Group 1 to the j -th EPS node in Group 2 is $F(i, j)$. There is no traffic demand within each EPS group or from Group 2 to Group 1. In total, this network has n^2 flows, and the traffic matrix f is an $2n \times 2n$ matrix with at most n^2 non-zero entries (see Fig. 7). For topology engineering, we only allow generating one logical topology for each traffic matrix, i.e., $M = 1$. Next, we will prove that the three-dimensional sum problem is satisfiable if and only if $\mu(\mathcal{G}, f) = 1$.

Sufficiency ($\Rightarrow$): Let $A(i, j, k)$ be a solution of the three-dimensional sum problem. Then, we can set up the OCS configurations as follows: for the k -th OCS, establish $A(i, j, k)$ number of

Fig. 7. Network topology and traffic model.

bidirectional links between the i -th EPS node in Group 1 and the j -th EPS node in Group 2. Under this set of OCS configurations, we establish $\sum_{k=1}^n A(i, j, k) = F(i, j)$ number of logical links between the i -th EPS node in Group 1 and the j -th EPS node in Group 2. Then, if we route every flow along its shortest path, it is easy to check that the achievable throughput $\mu(\mathcal{G}, f) = 1$. Note that $\mu(\mathcal{G}, f) = 1$ is indeed the optimal throughput, because the bisection bandwidth from Group 1 to Group 2 $\sum_{i=1}^n \sum_{k=1}^n D(i, k) = \sum_{i=1}^n \sum_{j=1}^n \sum_{k=1}^n A(i, j, k)$ divided by the total traffic demand from Group 1 to Group 2 $\sum_{i=1}^n \sum_{j=1}^n F(i, j) = \sum_{i=1}^n \sum_{j=1}^n \sum_{k=1}^n A(i, j, k)$ is exactly 1.

Necessity ($\Leftarrow$): Let $\mathcal{G}_l$ be the logical topology such that the optimal throughput $\mu(\mathcal{G}_l, f) = 1$ is achieved, and let $\Pi = \{\Pi_1, \Pi_2, \dots, \Pi_n\}$ be the corresponding configurations of the n OCS nodes. Since the throughput is 1 under the logical topology $\mathcal{G}_l$, the bisection bandwidth from Group 1 to Group 2 in $\mathcal{G}_l$ must be at least $\sum_{i=1}^n \sum_{j=1}^n F(i, j)$. Note that the egress capacity of Group 1 and the ingress capacity of Group 2 are both at most

$$\sum_{i=1}^n \sum_{k=1}^n D(i, k) = \sum_{k=1}^n \sum_{j=1}^n E(k, j) = \sum_{i=1}^n \sum_{j=1}^n F(i, j).$$

Hence, in order to deliver $\sum_{i=1}^n \sum_{j=1}^n F(i, j)$ amount of traffic, the following statements must be true for $\mathcal{G}_l$:

- (1) Every egress EPS port of Group 1 must be connected with an ingress EPS port of Group 2 through an OCS node;
- (2) Every flow must be delivered directly from its source to its destination.

Now, let $A(i, j, k)$ be the number of links from the i -th EPS node in Group 1 to the j -th EPS node in Group 2 formed in the k -th OCS node. Note that there are $D(i, k)$ number of bidirectional links between the i -th EPS node in Group 1 and the k -th OCS node; there are $E(k, j)$ number of bidirectional links between the k -th OCS node and the j -th EPS node in Group 2. Combined with the first statement, we must have $\sum_{j=1}^n A(i, j, k) = D(i, k)$ and $\sum_{i=1}^n A(i, j, k) = E(k, j)$. According to the second statement and the fact that the bisection bandwidth from Group 1 to Group 2 is at most $\sum_{i=1}^n \sum_{j=1}^n F(i, j)$, the total number of links from the i -th EPS node in Group 1 to the j -th EPS node in Group 2 must be exactly $F(i, j)$. Hence, $\sum_{k=1}^n A(i, j, k) = F(i, j)$.

The above analysis indicates that all the three-dimensional sum problem can be reduced to a ODC throughput optimization problem. Since the three-dimensional sum problem is NP-Complete, the ODC throughput optimization problem must also be NP-Complete. $\square$

B PROOF OF LEMMA 2

PROOF. Let $(\mathcal{G}_l^{*(m)}, f_{ij}^*(u, v))$ be the optimal topology+routing solution that attains the optimal throughput $\mu(\mathcal{G}, f)$. Next we consider $\mu(\mathcal{G}', f)$.

Note that the logical topologies $\mathcal{G}_l^{*(m)}$ are realizable as overlay topologies of $\mathcal{G}'$. Let $\mathcal{G}_l'^{*(m)}$ be the corresponding underlay logical topology of $\mathcal{G}_l^{*(m)}$. Consider the following overlay-topology based routing allocation in $\mathcal{G}'$. For every flow from S_i to S_j , we allocate $f_{ij}^*(u, v)$ amount of traffic to the overlay paths from S_u to S_v . It is easy to verify that the resulting throughput is $\mu(\mathcal{G}, f)$. Note that restricting routing allocation to overlay paths reduces routing flexibility. In contrast, our TE formulation (3)(4) computes network throughput based on the underlay topologies $\mathcal{G}_l'^{*(m)}$ directly, and thus the resulting throughput must be no smaller than $\mu(\mathcal{G}, f)$. Combined with the definition of $\mu(\mathcal{G}', f)$, we must have $\mu(\mathcal{G}', f) \geq \mu(\mathcal{G}, f)$. $\square$

C NOTATIONS

The notations in this paper are summarized in Table 4.

Table 4. Notations used in this paper

$N, \mathcal{S}$	N electrical-packet-switching (EPS) nodes $\mathcal{S} = \{S_1, S_2, \dots, S_N\}$.
$K, \mathcal{O}$	K optical-circuit-switching (OCS) nodes $\mathcal{O} = \{O_1, O_2, \dots, O_K\}$.
$\mathcal{G}$	Physical topology $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, where $\mathcal{V} = \mathcal{S} \cup \mathcal{O}$. (see §3.1).
$\mathcal{G}_l$	Logical topology $\mathcal{G}_l = \mathcal{G}_l(\mathcal{G}, \Pi) = (\mathcal{V}_l, \mathcal{E}_l)$ (see §3.1).
x_{ij}	$x_{ij}(\mathcal{G}_l)$ denotes the number of links from S_i to S_j in $\mathcal{G}_l$.
$\beta(\mathcal{G})$	Performance ratio of the physical topology $\mathcal{G}$ (see §3.3).
f	Traffic matrix among the N EPS nodes (see §3.1).
$\mu(\mathcal{G}, f)$	Optimal throughput for the given physical topology $\mathcal{G}$ and traffic matrix f (see §3.2).
L	Each EPS node has L ports.
R	Each OCS node has R ingress ports and R egress ports.
B	EPS port capacity.
Π	OCS configuration $\Pi = \{\Pi_1, \Pi_2, \dots, \Pi_K\}$ (see §3.1).
M	The number of logical topologies generated by ToE for each traffic matrix (see §3.2).
$f_{ij}(u, v)$	Amount of traffic from S_i to S_j that is allocated to the link (S_u, S_v) (see §3.2).
$C, \mathcal{A}_c$	C mutually disjoint EPS node groups satisfying $\cup_{c=1}^C \mathcal{A}_c = \mathcal{S}$.

Received August 2021; revised October 2021; accepted November 2021