

Managing Recurrent Virtual Network Updates in Multi-Tenant Datacenters: A System Perspective

Zhuotao Liu¹, Yuan Cao¹, *Member, IEEE*, Xuwu Zhang, Changping Zhu, and Fan Zhang, *Member, IEEE*

Abstract—With the advent of software-defined networking, network configuration through programmable interfaces becomes practical, leading to various on-demand opportunities for network routing update in multi-tenant datacenters, where tenants have diverse requirements on network routings such as short latency, low path inflation, large bandwidth, high reliability, etc. Conventional solutions that rely on topology search coupled with an objective function to find desired routings have at least two shortcomings: (i) they run into scalability issues when handling consistent and frequent routing updates and (ii) they restrict the flexibility and capability to satisfy various routing requirements. To address these issues, this paper proposes a novel search and optimization decoupled design, which not only saves considerable topology search costs via search result reuse, but also avoids possible sub-optimality in greedy routing search algorithms by making decisions based on the global view of all possible routings. We implement a prototype of our proposed system, OpReduce, and perform extensive evaluations to validate its design goals.

Index Terms—Recurrent virtual network, multi-tenant datacenters, routing management

1 INTRODUCTION

ALONG with the rise of cloud computing, multi-tenant datacenters grow into the scale of tens thousands of servers hosting millions of tenant VMs [1]. Managing multi-tenancy at such a scale to ensure efficiency, scalability and agility is a challenging problem drawing considerable research and engineering attention. Prior proposals [2], [3] place major management effort at hypervisors for the sake of easy configuration and implementation, whereas little effort has been made to manage *in-network* routings, particularly at the granularity of tenants. As a result, although the VM locations of each tenant are decided, datacenter network operators lose the visibility of actual in-network traffic forwarding for each tenant.

The lack of tenant-level traffic accountability and routing control could lead to various limitations. From the

perspective of business, network operators cannot customize tenant routings in accord with tenants' service level agreements (SLA) (e.g., latency, bandwidth, reliability or security requirements), which may close the door for such a business model in the virtual private cloud (VPC) market. From the perspective of management, in case of resolving hot spots in datacenters, for instance, it is difficult to determine the affected tenants and effectively re-route their traffic around congestion in time.

Thus, it is desirable to have explicit tenant routing control. The traditional way to achieve tenant routing management relies on topology search coupled with an objective function to greedily find the desired overlay network for each tenant. Although instant network configuration is technically enabled by SDN [4], the conventional approach still has at least two shortcomings. First, tenant routing updates in our production datacenters are recurrent, which can be triggered by various reasons including network load dynamics, tenant arrivals/departures, hot spots, link failures and so on. Repeatedly performing topology search for each routing update will impose significant search cost. Second, a tenant routing is updated typically to fulfill certain goals. However, it is uncertain that topology search coupled with an objective function is sufficient to achieve any goal. For instance, optimizing a performance metric depending on virtual links (i.e., VM pair communications) is subject to sub-optimality since the mapping between physical links and virtual links is unknown during topology search (Section 2.3).

To address these issues, we propose OpReduce, a novel search and optimization *decoupled* design for routing management. For each tenant-routing update, OpReduce first comprehensively searches all desired overlay candidates considering only the tenant's VM locations. Then it applies a global objective function over these candidates to finalize

- Z. Liu is with the College of Internet of Things Engineering, Hohai University, Changzhou 213022, China, and also with Electrical and Computer Engineering Department, University of Illinois at Urbana-Champaign, Urbana, IL 61801, and also with Google Inc., Mountain View, CA 94043. E-mail: zliu48@illinois.edu.
- Y. Cao and X. Zhang are with the College of Internet of Things Engineering, Hohai University, Changzhou 213022, China. E-mail: {20161965, zhangxw}@hhu.edu.cn.
- C. Zhu is with the Key Laboratory of Unmanned Aerial Vehicle Development and Data Application of Anhui Higher Education Institutes, Wanjiang University of Technology, Maanshan 243031, China. E-mail: cpzhu5126081@163.com.
- F. Zhang is with College of Information Science and Electronic Engineering, Zhejiang University, Alibaba-Zhejiang University Joint Institute of Frontier Technologies, Institute of Cyberspace Research, and Zhejiang Lab, Hangzhou 310027, China. E-mail: Fanzhang@zju.edu.cn.

Manuscript received 28 Mar. 2018; revised 1 Dec. 2018; accepted 4 Jan. 2019.
Date of publication 16 Jan. 2019; date of current version 8 July 2019.
(Corresponding author: Yuan Cao.)

Recommended for acceptance by Thu D. Nguyen.

For information on obtaining reprints of this article, please send e-mail to: reprints@ieee.org, and reference the Digital Object Identifier below.

Digital Object Identifier no. 10.1109/TPDS.2019.2893239

the most desired one. OpReduce’s decoupled model offers at least two advantages. First, since topology search is not directed by any objective function, search results are general and can be reused across routing updates that share the same VM locations: i.e., regardless of their goals, topology search for future routing updates is saved as long as their VM locations have been explored. Second, with the global view of all routing candidates, objective functions are not limited to be greedy and local. This eliminates the possibility of sub-optimality even when optimizing complex performance metrics, for instance, involving virtual links.

One concern of the decoupled design may be that comprehensively finding all desired overlay candidates for a tenant can be expensive. However, since the VMs of one tenant are typically spanning across a few racks, we can reduce the search space to a subgraph of the entire topology. Further, most data center topologies (e.g., VL2 [5], fattree [6]) are hierarchically organized into several layers and network traffic is never forwarded back and forth between different layers. We can adopt these properties to further refine the search space. Our proposed algorithm can perform comprehensive topology search in polynomial time (Section 3.1).

We implement a prototype of OpReduce and perform extensive evaluations to validate its design goals. On the one hand, we show that OpReduce greatly reduces the search cost for tenant-routing updates, meanwhile imposing small system overhead for managing large scale datacenters (e.g., small routing cache size and agile network configuration). On the other hand, we demonstrate that OpReduce is able to achieve complex optimization goals for routing updates, yielding significant networking performance gain over common practice.

2 BACKGROUND AND MOTIVATION

2.1 Managing Per-Tenant Routing

For configuration simplicity, encapsulations protocols such as VXLAN [7] and NVGRE [8] are widely adopted in cloud [3]. However, by simply tunneling tenant traffic, network operators lose the traffic visibility and accountability, i.e., they are not able to identify the origin of network traffic since packets are sent on behalf of hypervisors and one hypervisor may host VMs for multiple tenants. Such invisibility becomes even worse when datacenters run link aggregation and/or perform load balancing (i.e., ECMP, Hedera [9], or CONGA [10]) to spread traffic across redundant physical links. As a result, even if one tenant occupies only a small fraction of computation resources in the datacenter, its traffic could appear on many physical links, making network operators unaware of actual in-network traffic forwarding for the tenant.

The lack of in-network tenant-routing management causes various operation restrictions. For instance, it is difficult for network operators to perform monitoring, measurement and trouble-shooting for a tenant as its traffic may spread across many network links. Further, as virtual private cloud gains popularity, more tenants have incentives to customize their private cloud based on their own needs, e.g., one tenant may want to promote latency performance of web servers, whereas someone else may want to optimize bandwidth for MapReduce tasks. Without explicit tenant-level routing control, it is difficult to customize routings for individual tenants.

Thus, it is desirable for network operators to have explicit control over tenant routings. In particular, for each tenant, network operators explicitly configure a Virtual Tenant Network (VTN), which is an overlay network connecting the tenant’s VMs. The tenant’s traffic is confined within its VTN and no other physical links besides the ones in its VTN can carry the tenant’s traffic. As tenants can be identified by their VTNs, traffic is accountable, which eliminates these management limitations. Further, network operators can customize a tenant’s routing based on its requirement by embedding its VTN into an overlay that can best satisfy the requirement.

The traditional way of embedding VTN for a tenant relies on topology search coupled with an objective function to find the desired overlay. Such an approach, however, has at least two shortcomings: (i) it runs into scalability issues when handling frequent VTN embedding requests and (ii) it may be insufficient to find the optimal overlay for complex embedding goals. Next, we elaborate on the two issues.

2.2 Frequent VTN Updating Requests

In production multi-tenant datacenters, managing VTN embeddings/updates are all-time tasks on the basis of individual tenants, which are up to tens of thousands [1]. This is because that in practice, VTN embedding requests are triggered by many reasons, including network load dynamics, link congestion and failures, hot spots, tenant demand change, network capacity augment and so on. Although efficient graph search algorithms, such as Prim’s and Kruskal’s, have been proposed for decades, consistently and frequently performing overlay search in large topology for a large number of VTN updates still imposes significant cost. In fact, failing to accurately manage recurrent VTN updates is one of the primary reason why the today’s cloud offers limited SLAs for tenant networks. Therefore, resolving the scalability problem is the very first step towards efficient tenant-routing management.

2.3 Fulfilling VTN Embedding Goals

A VTN embedding/update is as simple as finding a random overlay network for the tenant. Instead, the VTN may need to satisfy various requirements such as guaranteed bandwidth [11], [12], [13], bounded latency [14], [15], required security appliances [16], [17], and other performance metrics depending on the tenant’s SLA. As these embedding goals become more complex, it is uncertain that the conventional search & optimization coupled solution can find the real optimal routing.

We use an illustrative example in Fig. 1 to demonstrate such sub-optimality. Network operators perform VTN update for a tenant whose VMs are hosted by the set of hypervisors $[S_1, S_2, S_3]$. The embedding goal is to minimize the communication cost of all virtual links. For simplicity, we define the communication cost of a virtual link as the number of physical hops on the path connecting the two VMs of the virtual link. Thus, if the virtual link between VM1 and VM2 takes path $S_1-A-C-S_2$, the communication cost is 3.

This embedding is a minimum spanning tree problem. A classic way of solving the problem is using a greedy topology search algorithm coupled with an objective function, which proceeds as follows. At certain intermediate state of

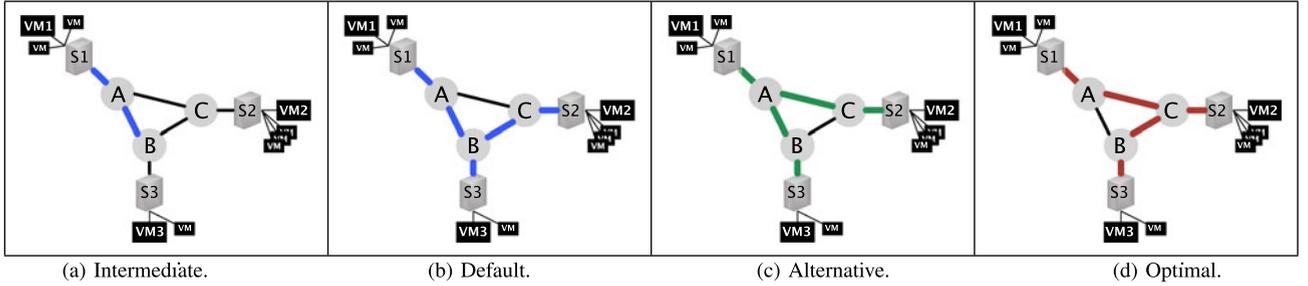


Fig. 1. Due to the lack of mapping between virtual links and physical links, conventional search & optimization coupled solution is subject to sub-optimality when optimizing performance metrics depending on virtual links.

the overlay search (Fig. 1a), AB has been appended to the overlay. Since BC and AC are equally good to be selected, based on the tie breaker in the search algorithm, BC may be appended to the ongoing overlay, and the final overlay returned by the algorithm is S_1-A , S_2-C , S_3-B , AB , BC (Fig. 1b). However, given another greedy search algorithm with a different tie breaker, AC could be the link to append after state (a) rather than BC , and the final overlay returned will be the one shown in Fig. 1c. However, no matter how the search algorithm is designed, the real optimality of an overlay is decided by the distribution of VMs. Although VM locations are known a priori, in an intermediate state, the search algorithm cannot determine exactly which set of virtual links a physical link will carry traffic for. Thus, no greedy search algorithm can guarantee to capture the real optimal overlay (Fig. 1d in this example). In general, due to the lack of mapping between virtual links and physical links, the greedy search & optimization coupled solution is always subject to sub-optimality when optimizing any performance metric involving virtual links.

We note that providing rigid theoretical proof to show the sub-optimality of the conventional search & optimization approach is beyond the scope of this paper. Instead, this paper focuses on offering a system perspective of how tenant routings are managed based on our experience with enterprise datacenters.

2.4 Decoupling Search and Optimization

To address the above challenges, OpReduce proposes a search and optimization decoupled design. For a routing update request, OpReduce first comprehensively finds all desired overlay candidates considering only VM locations. Then it applies an objective function, designed with the global view on all routing candidates, to finalize the optimal one. In the above example, OpReduce will first obtain all 3 overlay candidates (figures (b), (c) and (d) in Fig. 1) and then evaluate them to find the optimal one (Fig. 1d). With the global view on all candidates, OpReduce knows the exact mapping between virtual links and physical links. Thus, it can design a global objective function in evaluation to guarantee optimality. Further, as topology search is not directed by any objective function, search results are general so that they can be reused for future VTN updates that have the same VM locations. In addition, when handling VTN updates for existing tenants triggered by either link utilization change or demand change (two major reasons triggering VTN updates), no additional search effort is required. As a result, OpReduce saves considerable cost for performing frequent VTN updates in large scale datacenters.

3 SYSTEM DESIGN

In this section, we elaborate on the design of OpReduce. Fig. 2 illustrates the architecture of OpReduce. OpReduce is built upon a network information database and a controller. The network information database allows OpReduce to retrieve network related information, such as network topology and link utilization. The controller is used to manage both computation and networking resources, such as assigning VMs to tenants and configuring VTNs for tenants. OpReduce's decoupled design is achieved by its VTN embedding module which includes routing search engine, routing cache, objective functions and network action container.

We now describe the workflow of performing routing update using OpReduce. Upon receiving a request, OpReduce first obtains a list of desired routing candidates based on the VM placement. These routing candidates may come from either the routing search engine or the routing cache if the VM placement has been explored before. Then OpReduce evaluates each candidate to determine the most desired one for achieving a specific goal. To assist evaluation, network operators propose evaluation methods, defined as objective functions, to score each routing candidate. The design of objective functions is facilitated by the global view of all possible routing candidates. Further, useful network information (i.e., link loads) is retrieved from the network information database. Finally, after determining the most desired routing,

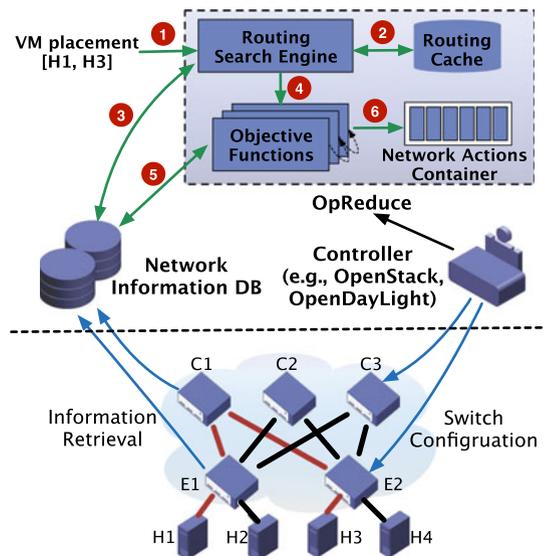


Fig. 2. The architecture of OpReduce.

OpReduce enforces the VTN embedding inside the network by executing network tasks (i.e., switch configuration) via the network action container.

Although setting up both the network information database and the network/computation controller takes considerable implementation efforts, we omit these details to focus on the research part of OpReduce in this paper. Next, we describe each individual component of OpReduce's VTN embedding module. Hereafter, we use *routing* and *VTN* interchangeably to indicate the overlay network carrying a tenant's traffic.

3.1 Routing Search Engine

Given a certain VM placement, the routing search engine is used to produce a list of *desired* routings (or routing candidates). In the typical layered datacenter network layout (e.g., fattree [6], Clos [5], [18]), a routing is desired if traffic between two VMs does not bounce back and forth more than once between two different layers. For random topology (e.g., [19], [20]), we can bound the maximum number of hops on VM-pair paths to exclude undesired routings. As datacenter fabric is often built with a high level of path redundancy, exploring all desired routings will provide sufficient candidates to achieve the embedding goal of the routing update request.

Next, we propose an algorithm that can produce a comprehensive list of all desired routings in layered data center topology. In the network, we associate each node with a *height*. All core switches are assigned height 1 and the initial height for other nodes is infinity. Then we trace down from core switches towards other nodes to assign them a height. The height of a node is the minimum height among all its neighbours plus one. All nodes with the same height compose a layer whose layer number is the node height. We define a path as *straight* if all nodes on the path have different heights. Straight path is either upward (from a source at a lower layer to a destination at a higher layer) or downward.

In a desired routing, to connect two hypervisors h_1 and h_3 , h_1 needs to have a straight path which shares the same endpoint with one of h_3 's straight paths. The shared endpoint is defined as a *common node* for the hypervisor set $[h_1, h_3]$. Formally, a node is defined as the common node for a set of hypervisors if it has at least one straight path to reach each hypervisor in the set. Then traffic from h_1 can first take one straight path to reach the common node and then bounces back to reach h_3 via another straight path. Thus, a routing connecting $[h_1, h_3]$ is composed of the above two straight paths. Similarly, a routing connecting all VMs of a tenant is composed of several straight paths, among which each straight path is originated from one hypervisor and all these paths are ended at one same common node.

We formulate our search procedure in Algorithm 1. At the very high level, the algorithm works as follows: (i) find all common nodes for the hypervisor set (line 5); (ii) for each common node, find one downward straight path from the common node to each hypervisor (line 8) and (iii) combine these paths to produce one routing candidate (line 9). Since a routing candidate can only contain one common node (having more will produce loops), Algorithm 1 provably finds all desired routings in the network. Even though the algorithm comprehensively finds all desired routings, the algorithm complexity is polynomial (not exponential). We defer detailed complexity analysis in Section 8.

Algorithm 1. Routing Search Algorithm

```

1: Input: VM placement  $H$ .
2: Output: All desired routing candidates.
3: for  $h_i \in$  hypervisor set  $H$  do
4:    $T_i \leftarrow$  GetUpwardGraph( $h_i$ );  $T \leftarrow \bigcup_i T_i$ ;
5: CommonNodes  $\leftarrow \bigcap_{T_i \in T} T_i.nodes()$ ;
6: foreach  $c_j \in$  CommonNodes do
7:   for  $h_i \in$  hypervisor set  $H$  do
8:      $N_j \leftarrow \bigcup_i$  GetDownwardPath( $c_j, h_i, T_j$ );
9: return  $N \leftarrow \bigcup_i N_i$ ;
10: Function: GetUpwardGraph( $h_i$ )
11:   return The graph containing all upward straight paths
        starting from  $h_i$  and ending at core switches.
12: Function: GetDownwardPath( $c_j, h_i, T_j$ )
13:   return The straight path from  $c_j$  to  $h_i$  in graph  $T_j$ .

```

For a random data center topology [19], [20], OpReduce can first adopt the k -shortest path algorithm [21] to obtain a set of paths between each hypervisor pair and then combine these paths to produce routing candidates. k can be parameterized to exclude undesired routings.

3.2 Tenant-Routing Cache

The search results for routing candidates are cached using a dictionary (hash table) data structure, where the key is VM placement and the value is a list of all desired routings for the VM placement. Each routing is stored as a list of physical link IDs, which can be used to retrieve link related information (e.g., utilization, status) from the network information database. Generally, entries in the dictionary are valid as long as the network topology remains the same. In OpReduce's prototype, we associate each entry in the dictionary with a relatively long validation period (e.g., few weeks), and re-perform routing search after an entry is expired.

In OpReduce's implementation, we allocate 32 bits for both the hash key and physical link IDs. In Section 5.2.2, we show that even in large scale $k = 32$ fat-tree datacenter with over 10 thousands servers (millions of VMs), the cache size for managing 10 thousand tenants is about 100 MB, which can be easily managed by commodity servers. Intuitively, the cache size is small because the topology size of a single VTN is small. Therefore, caching the routing search results would be not a bottleneck in practice.

3.3 Objective Functions

The objective functions are used to evaluate routing candidates so as to determine the most desired one for fulfilling embedding goals. For instance, a valid objective function can be as sophisticated as a combination function balancing latency, bandwidth, the number of hops and so on. Or it can be as detailed as optimizing specific virtual links for latency and other virtual links for bandwidth. In general, OpReduce is open to accept any objective function. But OpReduce offers global views on all possible routing candidates so that objective functions are not limited to be greedy and local. Relevant network information, available in OpReduce's network information database, can be applied to evaluate these routing candidates. By the time of this writing, an instance of OpReduce deployed in our production datacenter has pre-installed tens of abstract objective functions that can be

parameterized, allowing network operators to achieve a wide variety of VTN embedding goals.

3.4 Network Action Container

After determining the most desired routing, the final step is enforcing the routing inside network. The enforcement process is essentially configuring switches on the routing (e.g., configuring VLAN tags [22] or adding OpenFlow rules [23]) to guide the switches to perform desired traffic forwarding. To facilitate routing management in large scale datacenters, we develop a network action container to perform network configurations in a batch. For instance, configurations on different ports of one switch are aggregated in one thread and configurations on different switches are paralleled via multi-threading. As shown in Section 5.2.3, the network action container significantly reduces the overall configuration latency.

4 IMPLEMENTATION

We have a full implementation of OpReduce. We use OpenStack to create VMs on hypervisors and assign VMs to tenants. On our testbed dedicated for experiments, the OpenStack environment has 4 compute nodes as hypervisors that can host about one hundred VMs, one network node, and one controller node. Because OpenStack is limited to computation virtualization, we unify it with our physical network via an OpenDayLight [24] controller that uses OpenFlow and *netconf* [25] protocols to manage OpenFlow and legacy switches, respectively.

On our testbed, we allocate a dedicate VLAN tag for each VTN. Then embedding a VTN is about configuring relevant switch interfaces with corresponding VLAN tags to achieve reachability. Certainly there could be other network virtualization solutions, e.g., slices in sliceable switch [26], VLAN tag stacking [27]. OpReduce's network controller is extendible to support these virtualization solutions. We are aware that VLAN tags are limited to 4096 values which may be insufficient in large scale datacenters. To resolve this scalability issue, OpReduce implements a helper component, based on Panopticon [28], to achieve VLAN tag reuse in hybrid datacenters with both legacy and OpenFlow switches.

5 EVALUATION

Our evaluation centers around the followings. (i) In Section 5.1, we show OpReduce is guaranteed to find the least congested routings for VTN embeddings under various settings, which yields significant networking performance improvement over common practice. (ii) In Section 5.2, we show that OpReduce greatly reduces the search cost for managing numerous routing updates and imposes small system overhead while managing routing updates in large scale datacenters.

5.1 Congestion-Aware Routing Updates

One representative goal for routing updates in multi-tenant datacenters is to find the least congested routing to re-accommodate a tenant. With a suite of experiments (on testbeds and simulations) in this section, we aim to demonstrate that OpReduce is able to find more desired routings to embed

VTNs, compared with a couple of common practice. As a result, OpReduce yields improved networking performance.

To find a desired routing, the traditional search & optimization coupled solution uses Prim's minimal spanning tree algorithm with physical link utilization as the edge weight. However, since network traffic is generated by VM pairs, minimizing the congestion experienced by virtual links is the more direct and therefore more accurate goal for finding the least congested routing. Thus, OpReduce uses the following algorithm to determine the most desired routing. Assume OpReduce produces m routing candidates $\{\mathcal{T}_1, \dots, \mathcal{T}_m\}$ for a routing update request that has n virtual links $\{l_1, \dots, l_n\}$. Then the most desired routing is

$$\mathcal{T}^* = \min_{\mathcal{T}_k} \frac{\sum_{i=1}^n \lambda_i \cdot \mathcal{F}_k(l_i)}{n}, \quad (1)$$

where $k \in [1, m]$, $\mathcal{F}_k(l_i)$ is the congestion level experienced by virtual link l_i in routing candidate \mathcal{T}_k and λ_i is the weight of l_i . In OpReduce's instance deployed in our production datacenters, both \mathcal{F} and λ_i are configurable to enable highly flexible routing customization. For evaluation purpose in this paper, $\mathcal{F}_k(l_i)$ is defined as the highest congestion level of all physical links in \mathcal{T}_k that carry l_i , where the congestion level of a physical link is estimated as its average link utilization. Further, all virtual links are equally weighted.

Besides the conventional search & optimization coupled solution (referred to as the *local* solution), we also compare OpReduce with flow-level ECMP, the common practice for load balancing and congestion reduction in datacenters, and the *bottomline* solution which embeds VTN to a randomly selected routing candidate.

For each embedding request, we compare the routing determined by OpReduce and the ones selected by the other three solutions. We use average flow completion time (FCT) as the metric to quantify the congestion experienced by virtual links (similar to [10], [29]), expecting that a better routing will have shorter average FCT. *Note that we do not claim that FCT is either the only or the optimal metric. We use FCT for the sake of quantifiable results.*

5.1.1 Testbed Experiments

We start the evaluation on our physical testbed. We build a network topology illustrated in Fig. 4a. The datacenter has pre-embedded tenants to generate traffic using our client/server programs. The clients initiate long-lived TCP connections to randomly selected servers to request flow transfers. All VMs run both the client and server programs. Only intra-tenant communication is allowed.

In this experiment, we perform routing update for a tenant who has 10 VMs that are randomly distributed in all four hypervisors on our testbed. In total, we perform 40 sets of experiments under different network utilization. Fig. 4a shows the snapshot of network utilization in one experiment set. In each experiment set, we perform 3 individual routing updates using OpReduce, ECMP and the bottomline solution. After the tenant VTN is embedded, its virtual links (VM pairs) start to generate traffic using our client/server programs. Flow sizes are randomly sampled from the empirical datacenter workload [10]. To ensure fair comparison, the set of flow sizes used for the 3 experiments in

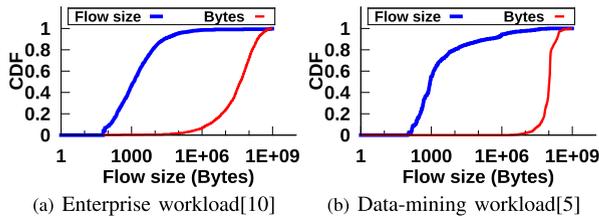


Fig. 3. Empirical workloads used in evaluation.

each experiment set is identical. We compute the average FCT among all flow transfers as our performance metric.

Fig. 4b illustrates the desired routing produced by each solution for the network utilization shown in Fig. 4a. Since OpReduce first outputs all routing candidates (i.e., $E_1-C_1-E_2$, $E_1-C_2-E_2$, $E_1-C_3-E_2$) and then applies the objective function (Equation (1)) to determine the most desired one, it can always find the least congested routing ($E_1-C_1-E_2$ in this case). However, both ECMP and bottomline are unaware of the link utilization. Consequently, their static hashing could result in overwhelming these more congested links.

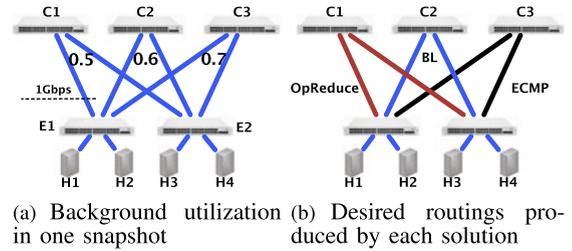
Fig. 4c plots the average FCT in each experiment set when using enterprise workload [10] (results for using data-mining workload [5] are similar, and we omit them for brevity). Although the bottomline solution and ECMP may produce different FCTs in one set, their average FCTs over all sets are close. Comparing with these two solutions, OpReduce can consistently find the optimal routing, yielding ~ 30 percent less FCTs on average and up to $\sim 5x$ reduced FCTs in some sets (outside the plot scope of Fig. 4c).

Note we do not compare OpReduce with the local solution in our testbed experiments. This is because the local solution and OpReduce will produce the same routing in such a small topology. Next we show that given large network topologies, it is very likely for the local solution to return sub-optimal overlays, which results in significant performance degradation.

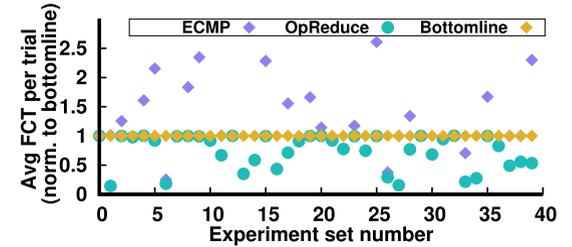
5.1.2 Large Scale Simulations

In order to investigate how OpReduce’s performance is affected by various factors that are not covered by our small-scale testbed experiments, we perform detailed simulations using large scale datacenter topologies and empirical workloads obtained from production datacenters (Fig. 3). As summarized in Table 1, we thoroughly evaluate the performance of OpReduce under the impact of five factors: the topology, the network/fabric load, the traffic workload, the average number of VMs occupied by a tenant, and the average VTN scale.

Various Fabric Loads (Exp NO. 1). In multi-tenant datacenters, fabric load can be quantified as the VM over-subscription ratio, which is defined as the ratio of worst-case achievable aggregate bandwidth among VM pairs to the total capacity of the topology. A fat-tree fabric has an hypervisor over-subscription of 1:1 because all hypervisors may potentially send at the full bandwidth of their network interfaces. Thus, if the average number of VMs hosted by a hypervisor is N , the VM over-subscription ratio is $N:1$. Hereafter, we use the VM over-subscription ratio as *load degree*.



(a) Background utilization (b) Desired routings produced by each solution in one snapshot



(c) Average FCT in each experiment set

Fig. 4. Testbed experiments for improving FCTs.

We tune the datacenter load degree in the range of 1:1-7:1 by varying the number of embedded tenants. For each load, we create 5 different snapshots to avoid bias. For each snapshot, we randomly pick a tenant to perform routing update using all four solutions: OpReduce, the local solution, ECMP and the bottomline solution. Thus, we perform 4 independent experiments for the tenant. In each solution, after the desired routing is finalized, we assign each virtual link a randomly sampled flow size to generate traffic and compute the average FCT among all flow transfers. Among all 4 experiments for the same tenant, we use the same set of flow sizes. After finishing the current tenant, we recover the network snapshot and re-sample another tenant to continue evaluation. In total, 100 routing updates are performed for each snapshot.

Fig. 5 plots the min-median-max distribution of the averaged FCT across all 5 snapshots for each load degree. As illustrated in the Fig. 5a, OpReduce significantly outperforms ECMP and bottomline: up to 80 percent FCT reduction for small fabric loads and at least 40 percent FCT reduction for all loads. Again, this is because ECMP and bottomline are static solutions that are unaware of the network utilization. Further, as load degree increases, OpReduce offers less FCT reductions since there is less routing optimization space in heavily utilized network. Since production datacenters are typically over-provisioned, we can expect large performance benefits offered by OpReduce in production datacenters.

Due to the lack of exact mapping between virtual links and physical links, the local solution often produces sub-optimal

TABLE 1
Experimental Settings

Exp. NO.	Topology	Fabric load	Workload	# VMs per tenant	Average VTN scale
1	$k=8$ FT	*	enterprise	20	4
2	$k=8$ FT	3:1	*	20	4
3	*	3:1	enterprise	20	4
4	$k=8$ FT	3:1	enterprise	20	*
5	$k=8$ FT	3:1	enterprise	*	4

“*” refers to that we vary the factor to isolate its impact in the experiment.

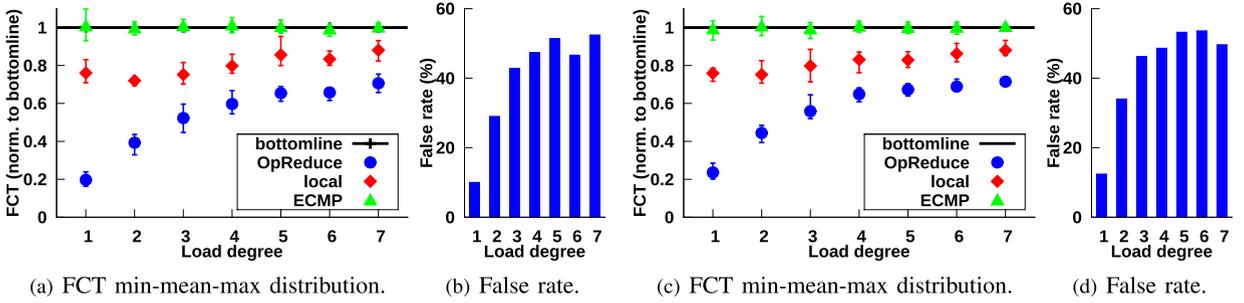


Fig. 5. Impact of the fabric load on OpReduce's performance. The normalized FCT is used as the performance metric. Part (a) and (b) use the enterprise workload [10]. Part (c) and (d) use the data-mining workload [5].

routings. As shown in Fig. 5b, with high probabilities (up to ~ 60 percent), the local solution returns an overlay different from the one produced by OpReduce. Consequently, these sub-optimal routings result in 1.2–4 \times FCT inflation, compared with the routings produced by OpReduce (Fig. 5a).

Datacenter Traffic Workload (Exp NO. 2). In Figs. 5a and 5b, we use the enterprise datacenter workload. We repeat the same experiment using the data-mining workload to learn the impact of traffic workload on OpReduce's performance. The results, plotted in Figs. 5c and 5d, show that OpReduce provide similar benefits for data-mining workload.

Fabric Topologies (Exp NO. 3). We investigate four datacenter topologies: three organized topologies (Clos, $k = 8$ and $k = 16$ fat-tree) and another organized topology added with random short-cuts. The Clos topology has the same number of hypervisors as the $k = 8$ fat-tree topology except that the over-subscription ratio is 2:1. All three organized topologies have different routing redundancy. In particular, they have 8, 16 and 64 shortest paths between two randomly chosen hypervisors in different pods, respectively. The short-cut topology is built by adding random links into the $k = 8$ fat-tree topology. The links are randomly added between ToR switches to inter-connect different pods so that besides traversing through the core switches, the inter-pod communication can alternatively use the short-cut bridges as well. We set the load degree as 3:1 in all these topologies.

For each topology, we perform the similar experiments as in evaluating the impact of fabric loads. Fig. 6 plots the CDF of all obtained FCT reductions (compared with the baseline solution). By comparing the results of all 3 fat-tree topologies, we conclude that OpReduce offers more benefits for topologies built with more redundant links. This is because OpReduce has larger optimization space in more redundant topologies.

We further investigate how topology short-cuts may affect the performance of OpReduce. In organized datacenter

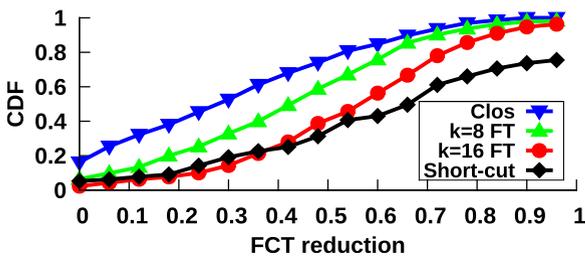


Fig. 6. Impact of datacenter topology on OpReduce's performance of FCT reduction.

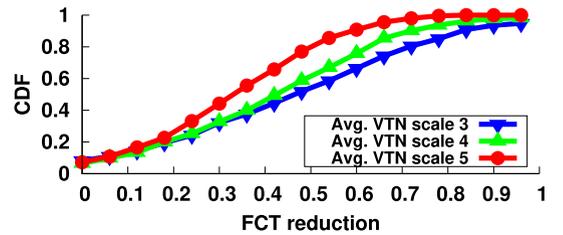
topologies, alternative routing options typically have less diversity in the sense that they all have the same number of hops, although they are varying in terms of utilization. In contrast, short-cut topology can create more diverse routings with different numbers of hops as well as different utilization. Thus, OpReduce has even larger optimization space in the short-cut topology so as to produce more performance benefits.

We also investigate the average VTN scale \mathcal{N}_h (the average number of hypervisors used by a tenant's VMs) and the average per-tenant VM count \mathcal{N}_v . We consider $\mathcal{N}_h = [3, 4, 5]$ and $\mathcal{N}_v = [10, 20, 30]$ since they are the common cases in our production datacenters. We find that OpReduce consistently provides performance benefits in these settings (Fig. 7).

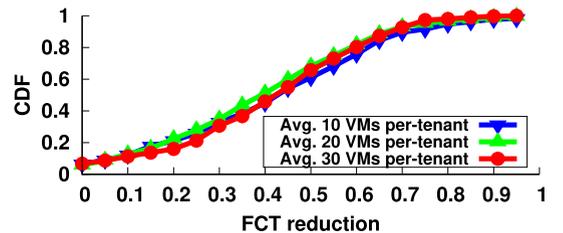
5.2 System Properties

5.2.1 Search Cost Reduction

The number of routing update requests is affected by the amount of embedded tenants, i.e., the load degree. Fig. 8a plots search cost reduction under various loads for $k = 16$ fat-tree topology. Although the absolute number of routing updates increase as load degree increases, the normalized search cost actually reduces. This is because OpReduce's knowledge base about routing candidates for various VM locations also expands as the load degree increases. As a result, the cache hit ratio dramatically increases as well



(a) Average VTN scale (Exp NO. 4)



(b) Average per-tenant VM count (Exp NO. 5)

Fig. 7. OpReduce offers consistent benefits for different VTN scales and different per-tenant VM counts.

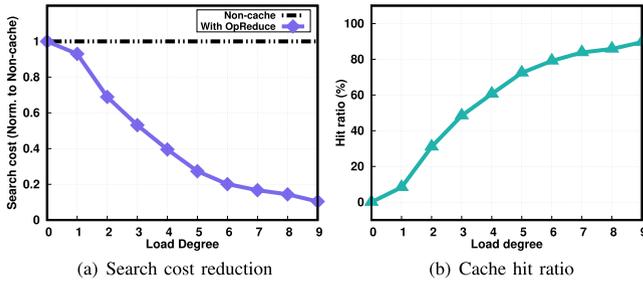


Fig. 8. In spite of the numerous routing updates as load degree increases, OpReduce introduces small search cost since the cache hit ratio increases dramatically.

(Fig. 8b). Thus, in spite of the numerous routing updates, OpReduce introduces small topology search cost.

5.2.2 Routing Cache Size

The routing cache size is affected by two factors: the topology and the average VTN scale. The topology affects the cache size because it affects the number of routing candidates for one VM placement. The VTN scale affects the cache size since it determines the number of physical links in one routing candidate.

Fig. 9a plots cache sizes with respect to the number of tenants and the average VTN scales in $k = 16$ fat-tree topology. The cache size increases with the average VTN scale and linearly grows with the number of tenants. However, even with very sporadically distributed VTNs, i.e., large average VTN scales, caching routing candidates for 10K tenants consumes no more than 40 MB memory. Even if we consider a much larger $k = 32$ fat-tree datacenter with over 8K servers, the cache size for 10K tenants is about hundreds of Megabytes (Fig. 9b), which can be easily managed by commodity servers with gigabytes of memory.

5.2.3 Switch Configuration Latency

Both SDN switches and legacy switches may co-exist in today's datacenters [30]. Configurations on OpenFlow Switches can be finished almost in real time via SDN controllers such as OpenDayLight [24], but it takes non-trivial time to configure legacy switches. Thus, to ensure that routing enforcement does not become the bottleneck for routing update, OpReduce designs a network action container to properly aggregate configuration tasks so as to reduce the overall configuration delay.

In OpReduce's prototype, one switch configuration task is about associating one VLAN tag on a certain port of a switch. We notice that configuring a single port on a legacy switch

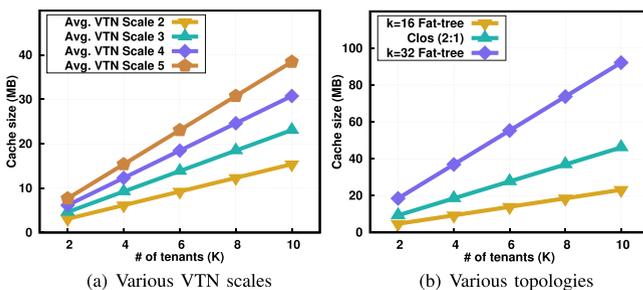


Fig. 9. The routing cache size under various settings.

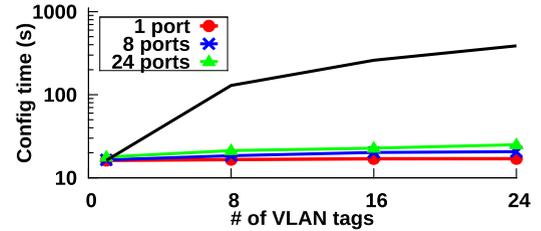


Fig. 10. Delay measurements for configuring legacy switches. The dash line indicates the configuration time without using the network action container.

takes almost the same amount of time as configuring multiple ports on the switch. Thus, our network action container aggregates all configuration tasks on the same switch together to perform batch configuration so as to reduce the overall configuration delay. Meanwhile, batch configurations for different switches are executed simultaneously via multithreading. Fig. 10 plots the measured configuration delay on our testbed. The results show that even if we simultaneously configure 24 ports on one switch and on each port we configure 24 VLAN tags (576 single configuration tasks), the overall configuration time is less than 1.5 times the delay for configuring just one VLAN tag on a single port. Thus, with the network action container, routing enforcement can be finished timely so that it will not be a bottleneck in practice.

6 RELATED WORK AND DISCUSSION

Multi-Tenancy Management. Prior designs for multi-tenant datacenters, such as NVP [3] and Netlord [2], focus on multi-tenancy management at hypervisors. For instance, NVP maintains virtual switches on each hypervisor and leverages a set of tunnels between each pair of hypervisors to deliver traffic for tenant VMs. The actual tenant traffic forwarding in the physical network is not managed. Several prior works have considered to perform one-time in-network routing management to achieve various goals, such as guaranteed bandwidth [11], [12], [13], [31], low or bounded latency [14], [15], [32] and user/service isolation [33]. OpReduce, instead, focuses on managing *recurrent* routing updates in multi-tenant datacenters.

A Wide Variety of Performance Enhancement. Many approaches have been proposed to improve datacenter networking performance. For instance, load balancing approaches [9], [10], [34], [35], priority queuing approaches [36], [37], [38], [39], deadline-aware approaches [40] and DCTCP [41] are proposed to improve latency performance. Portland [42] and fat-tree [6] propose scalable datacenter architectures to support high bandwidth between servers whereas VL2 [5] virtualizes datacenters into server pools to allow applications to obtain high throughput. Although OpReduce is not proposed to explicitly improve certain performance, its efficient tenant routing management and decoupled design allow network operators to enhance a wide variety of *customer-interested* performance metrics, and some of these metrics cannot be optimized using prior approaches.

Achieving Agile Routing Updates. To be readily deployable, OpReduce is augmented by SDN to perform agile in-network routing updates. B4 [43] and SWAN [44] also adopt SDN to perform traffic engineering in wide area networks to achieve high inter-datacenter throughput.

TABLE 2
One-time Comprehensive Search Cost for Fattree Topologies

Topo.	Intra-Pod	Inter-Pod	(nodes,edges)
$k=8$	$5\mathcal{N}_h$	$21\mathcal{N}_h$	(200,384)
$k=16$	$9\mathcal{N}_h$	$71\mathcal{N}_h$	(1296,3072)
$k=32$	$17\mathcal{N}_h$	$273\mathcal{N}_h$	(9248,24567)
k	$(1 + \frac{k}{2}) \cdot \mathcal{N}_h$	$(1 + \frac{k}{2} + (\frac{k}{2})^2) \cdot \mathcal{N}_h$	$(\sim \frac{k^3}{4}, \frac{3k^3}{4})$

\mathcal{N}_h is the average VTN scale.

7 CONCLUSION

In this paper, we present OpReduce, a system for managing virtual tenant network update in multi-tenant datacenters. Conventional solutions that rely on topology search coupled with an objective function to find desired routings have at least two shortcomings: scalability issue for handling recurrent routing updates and the inefficiency for satisfying various routing requirements. To address these issues, OpReduce proposes a novel search and optimization decoupled design, which enables routing search result reuse and guaranteed routing optimality. We implement a prototype of OpReduce and perform extensive evaluations to valid OpReduce's design goals. Evaluation results show that (i) Even for complex VTN embedding goals, OpReduce ensures routing optimality which yields significant networking performance improvement over conventional approaches; (ii) OpReduce greatly reduces search cost for managing numerous routing updates and imposes small system overhead.

8 APPENDIX

In this section, we analyze the complexity of Algorithm 1 in Section 3.1. The first step of the algorithm is to find the upward graph for each hypervisor in H . Essentially, it is a breadth-first search process. Thus the search complexity is $O(|E|)$ where $|E|$ is the number of edges in the upward graph. Note that the upward graph is much smaller than the entire network. Specifically, $|E| = \sum_{k=2}^L R_k \cdot N_k$, where R_k denotes the link redundancy ratio at layer k , N_k is the number of nodes at layer k in the upward graph and L is the number of layers in the upward graph. The link redundancy ratio at layer k is defined as the number of links that one node in layer k can has to reach nodes in layer $k-1$. For instance, in $k=8$ fat-tree topology with 4 layers, $R_4 = 1, R_3 = R_2 = 4$. There is no R_1 since layer 0 does not exist. Thus $|E| = 21$, which is very small compared with entire network size ($|E| = 384$). For simplicity of presentation, we assume all nodes in layer k have the same R_k . However, both our search algorithm and complexity analysis are not restricted by this assumption.

The second search step is that in upward graph \mathcal{T}_j , finding the downward straight paths from the common node to hypervisor h_i . It is depth-first search, which has $O(|E|)$ worst-case complexity. However, in topologies like fat-tree and Clos [18], \mathcal{T}_j turns out to be a tree rooted at h_i with all the common nodes as its leaves. Thus only $L-1$ edges need to be visited to find a downward straight path from the common node (leaf) to h_i (root), which introduces constant (negligible) overhead.

Table 2 summarizes the search cost (the number of visited edges) for finding all desired routing candidates in different

fat-tree topologies. We list the number of nodes and edges in the network to show that the search space is much smaller. \mathcal{N}_h is the average number of hypervisors used by one tenant. We consider the complexity for both intra-pod (all the hypervisors are located within one pod) and inter-pod (the tenant spreads across at least two pods). For the special case that the all hypervisors in H are sharing the same ToR switches, there will only one routing connecting them. It is clear that even in large scale datacenters with thousands of servers, the one-time comprehensive search cost is small and acceptable. Mathematically, we derive that algorithm complexity is $O(V^{\frac{3}{3}})$, where V is the number of nodes in the fattree topology. Thus, although our algorithm performs comprehensive topology search, the complexity is still polynomial.

For other topologies with larger over-subscription ratio (e.g., 2:1 Clos topology) than the fat-tree topology (1:1), search cost will further be reduced due to the smaller number of redundant paths in these topologies.

ACKNOWLEDGMENTS

We thank the authors listed in [45] who made substantial contributions to the work initially. This work was supported by the National Natural Science Foundation of China (61601168, 61671202), National Key R&D Program of China (2016YFC0401606), the Fundamental Research Funds for the Central Universities (2019B22714), Ministry of Education of the Peoples Republic of China, the Fundamental Research Foundation of Shenzhen (JCYJ20170302151209762), Major Scientific Research Project of Zhejiang Lab (Grant No. 2018FD0ZX01) and Alibaba-Zhejiang University Joint Institute of Frontier Technologies.

REFERENCES

- [1] Amazon Data Center Size. (2012). [Online]. Available: <http://huanliu.wordpress.com/2012/03/13/amazon-data-center-size/>
- [2] J. Mudigonda, P. Yalagandula, J. Mogul, B. Stiekes, and Y. Pouffary, "NetLord: A scalable multi-tenant network architecture for virtualized datacenters," *ACM SIGCOMM Comput. Commun. Rev.*, vol. 41, no. 4, pp. 62–73, 2011.
- [3] T. Koponen, K. Amidon, P. Baland, M. Casado, A. Chanda, B. Fulton, I. Ganichev, J. Gross, N. Gude, P. Ingram, et al., "Network virtualization in multi-tenant datacenters," in *Proc. 11th USENIX Conf. Networked Syst. Des. Implementation*, 2014, pp. 203–216.
- [4] N. McKeown, T. Anderson, H. Balakrishnan, G. Parulkar, L. Peterson, J. Rexford, S. Shenker, and J. Turner, "OpenFlow: Enabling innovation in campus networks," *ACM SIGCOMM Comput. Commun. Rev.*, vol. 38, pp. 69–74, 2008.
- [5] A. Greenberg, J. R. Hamilton, N. Jain, S. Kandula, C. Kim, P. Lahiri, D. A. Maltz, P. Patel, and S. Sengupta, "VL2: A scalable and flexible data center network," in *Proc. ACM SIGCOMM Conf. Data Commun.*, 2009, pp. 51–62.
- [6] M. Al-Fares, A. Loukissas, and A. Vahdat, "A scalable, commodity data center network architecture," *ACM SIGCOMM Comput. Commun. Rev.*, vol. 38, pp. 63–74, 2008.
- [7] T. Sridhar, L. Kreeger, D. Dutt, C. Wright, M. Bursell, M. Mahalingam, P. Agarwal, and K. Duda, "Virtual eXtensible local area network (VXLAN): A framework for overlaying virtualized layer 2 networks over layer 3 networks," *IETF*, 2014.
- [8] Y.-S. Wang and P. Garg, "NVGRE: Network virtualization using generic routing encapsulation," RFC 7637, Oct. 2015.
- [9] M. Al-Fares, S. Radhakrishnan, B. Raghavan, N. Huang, and A. Vahdat, "Hedera: Dynamic flow scheduling for data center networks," in *Proc. 7th USENIX Conf. Networked Syst. Des. Implementation*, 2010, p. 19.
- [10] M. Alizadeh, et al., "CONGA: Distributed congestion-aware load balancing for datacenters," in *Proc. ACM Conf. SIGCOMM*, 2014, pp. 503–514.

- [11] H. Ballani, P. Costa, T. Karagiannis, and A. Rowstron, "Towards predictable datacenter networks," in *Proc. ACM SIGCOMM Conf.*, 2011, pp. 242–253.
- [12] C. Guo, G. Lu, H. J. Wang, S. Yang, C. Kong, P. Sun, W. Wu, and Y. Zhang, "Secondnet: A data center network virtualization architecture with bandwidth guarantees," in *Proc. Proc. 6th Int. Conf.*, 2010, Art. no. 15.
- [13] J. Lee, Y. Turner, M. Lee, L. Popa, S. Banerjee, J.-M. Kang, and P. Sharma, "Application-driven bandwidth guarantees in datacenters," *ACM SIGCOMM Comput. Commun. Rev.*, vol. 44, pp. 467–478, 2014.
- [14] K. Jang, J. Sherry, H. Ballani, and T. Moncaster, "Silos: Predictable message latency in the cloud," in *Proc. ACM Conf. Special Interest Group Data Commun.*, 2015, pp. 435–448.
- [15] M. P. Grosvenor, M. Schwarzkopf, I. Gog, R. N. Watson, A. W. Moore, S. Hand, and J. Crowcroft, "Queues dont matter when you can JUMP them!" in *Proc. USENIX NSDI*, 2015, pp. 1–14.
- [16] S. Shin and G. Gu, "CloudWatcher: Network security monitoring using OpenFlow in dynamic cloud network," in *Proc. 20th IEEE Int. Conf. Network Protocols*, 2012, pp. 1–6.
- [17] S. Shin, H. Wang, and G. Gu, "A first step toward network security virtualization: From concept to prototype," *IEEE Trans. Inf. Forensics Secur.*, vol. 10, no. 10, pp. 2236–2249, Oct. 2015.
- [18] W. J. Dally and B. P. Towles, *Principles and Practices of Interconnection Networks*. Amsterdam, The Netherlands: Elsevier, 2004.
- [19] A. Singla, C.-Y. Hong, L. Popa, and P. B. Godfrey, "Jellyfish: Networking data centers randomly," in *Proc. USENIX NSDI*, 2012, pp. 1–6.
- [20] J.-Y. Shin, B. Wong, and E. G. Sirer, "Small-world datacenters," in *Proc. 2nd ACM Symp. Cloud*, 2011, Art. no. 2.
- [21] D. Eppstein, "Finding the k shortest paths," in *Proc. IEEE Annu. Symp. Found. Comput. Sci.*, 1994, pp. 652–673.
- [22] P. J. Frantz and G. O. Thompson, "VLAN frame format," U.S. Patent 5,959,990, Sep. 1999.
- [23] OpenFlow specification. [Online]. Available: <https://www.opennetworking.org/sdn-resources/onf-specifications/openflow>, Accessed on: 2016.
- [24] OpenDaylight. [Online]. Available: <http://www.opendaylight.org/>, Accessed on: 2016.
- [25] R. Enns, M. Bjorklund, and J. Schoenwaelder, "NETCONF configuration protocol," RFC 6241, 2011.
- [26] Trema: Full-stack OpenFlow framework in ruby and C. [Online]. Available: <http://trema.github.io/trema/>, Accessed on: 2015.
- [27] Cisco stacked VLAN processing. [Online]. Available: https://www.cisco.com/c/en/us/td/docs/ios/12_0s/feature/guide/qinq.html, Accessed on: 2016.
- [28] D. Levin, M. Canini, S. Schmid, F. Schaffert, A. Feldmann, et al., "Panopticon: Reaping the benefits of incremental SDN deployment in enterprise networks," in *Proc. USENIX Annu. Tech. Conf.*, 2014, pp. 333–345.
- [29] N. Dukkupati and N. McKeown, "Why flow-completion time is the right metric for congestion control," *ACM SIGCOMM Comput. Commun. Rev.*, vol. 36, no. 1, pp. 59–62, 2006.
- [30] H. Lu, N. Arora, H. Zhang, C. Lumezanu, J. Rhee, and G. Jiang, "HybNET: Network manager for a hybrid network infrastructure," in *Proc. Ind. Track 13th ACM/IFIP/USENIX Int. Middleware Conf.*, 2013, Art. no. 6.
- [31] Z. Liu, K. Chen, H. Wu, S. Hu, Y.-C. Hut, Y. Wang, and G. Zhang, "Enabling work-conserving bandwidth guarantees for multi-tenant datacenters via dynamic tenant-queue binding," in *Proc. IEEE Conf. Comput. Commun.*, 2018, pp. 1–9.
- [32] S. Hu, W. Bai, K. Chen, C. Tian, Y. Zhang, and H. Wu, "Providing bandwidth guarantees, work conservation and low latency simultaneously in the cloud," in *Proc. 35th Annu. IEEE Int. Conf. Comput. Commun.*, 2016, pp. 1–9.
- [33] R. Sherwood, G. Gibb, K.-K. Yap, G. Appenzeller, M. Casado, N. McKeown, and G. M. Parulkar, "Can the production network be the testbed?" in *Proc. 9th USENIX Conf. Operating Syst. Des. Implementation*, 2010, pp. 365–378.
- [34] T. Benson, A. Anand, A. Akella, and M. Zhang, "MicroTE: Fine grained traffic engineering for data centers," in *Proc. 7th Conf. Emerging Netw. Exp. Technol.*, 2011, Art. no. 8.
- [35] S. Kandula, D. Katabi, S. Sinha, and A. Berger, "Dynamic load balancing without packet reordering," *ACM SIGCOMM Comput. Commun. Rev.*, vol. 37, pp. 51–62, 2007.
- [36] C.-Y. Hong, M. Caesar, and P. Godfrey, "Finishing flows quickly with preemptive scheduling," in *Proc. ACM SIGCOMM Conf. Appl. Technol. Archit. Protocols Comput. Commun.*, 2012, pp. 127–138.
- [37] M. Alizadeh, S. Yang, M. Sharif, S. Katti, N. McKeown, B. Prabhakar, and S. Shenker, "pFabric: Minimal near-optimal datacenter transport," in *Proc. ACM SIGCOMM Conf. SIGCOMM*, 2013, pp. 435–446.
- [38] D. Zats, T. Das, P. Mohan, D. Borthakur, and R. Katz, "DeTail: Reducing the flow completion time tail in datacenter networks," in *Proc. ACM SIGCOMM Conf. Appl. Technol. Archit. Protocols Comput. Commun.*, 2012, pp. 139–150.
- [39] W. Bai, L. Chen, K. Chen, D. Han, C. Tian, and H. Wang, "Pias: Practical information-agnostic flow scheduling for commodity data centers," *IEEE/ACM Trans. Netw.*, vol. 25, no. 4, pp. 1954–1967, Aug. 2017.
- [40] B. Vamanan, J. Hasan, and T. Vijaykumar, "Deadline-aware datacenter TCP (D2TCP)," *Proc. ACM SIGCOMM Conf. Appl. Technol. Archit. Protocols Comput. Commun.*, 2012, pp. 115–126.
- [41] M. Alizadeh, et al., "Data center TCP (DCTCP)," *ACM SIGCOMM Comput. Commun. Rev.*, vol. 40, pp. 63–74, 2011.
- [42] N. Mysore, et al., "Portland: A scalable fault-tolerant layer 2 data center network fabric," in *Proc. ACM SIGCOMM Conf. Data Commun.*, 2009, pp. 39–50.
- [43] C.-Y. Hong, et al., "B4 and after: Managing hierarchy, partitioning, and asymmetry for availability and scale in Google's software-defined WAN," in *Proc. Conf. ACM Special Interest Group Data Commun.*, 2018, pp. 74–87.
- [44] C.-Y. Hong, S. Kandula, R. Mahajan, M. Zhang, V. Gill, M. Nanduri, and R. Wattenhofer, "Achieving high utilization with software-driven WAN," in *Proc. ACM SIGCOMM Conf. SIGCOMM*, 2013, pp. 15–26.
- [45] Q. Xu, C. Lumezanu, Z. Liu, N. Arora, A. Sharma, H. Zhang, and G. Jiang, "Optimization framework for multi-tenant data centers," U.S. Patent 9,813,301, Nov. 2017.

Zhuotao Liu received the BS degree from Shanghai Jiaotong University, in 2012 and the PhD degree from the University of Illinois at Urbana-Champaign, in 2017. Currently, he works in Network Infrastructure Team at Google, maintaining Google's global-scale private WAN. His research interests include Internet security & privacy, data center networking and blockchain infrastructure.

Yuan Cao (S'09-M'2014) received the BS degree from Nanjing University, the ME degree from Hong Kong University of Science and Technology, and the PhD degree from Nanyang Technological University in 2008, 2010 and 2015, respectively. Currently he works as an assistant professor with the College of Internet of Things Engineering, Hohai University. His research interests include hardware security, silicon physical unclonable function, and analog/mixed-signal VLSI circuits and systems. He is a member of the IEEE.

Xuewu Zhang received the ME and PhD degrees from Hohai University, in 2006 and 2011, respectively. Currently he works as a full professor with the College of Internet of Things Engineering, Hohai University. His research interests include information security and cryptographic engineering.

Changqing Zhu received the MS degree from Huazhong University of Science and Technology, in 2002 and PhD degree in communication and information system from China University of Mining and Technology, in 2010, respectively. He is a professor with Hohai University and Wanjiang Institute of Technology. His research interests include communication circuit and power ultrasonic technology. His research interests include underwater security.

Fan Zhang received the PhD degree from the Department of Computer Science and Engineering, University of Connecticut, in 2012. He is an associate professor with the College of Information Science and Electronic Engineering, Zhejiang University, China. He is also jointly appointed as an associate professor with the Institute of Cyberspace Research, Zhejiang University, an associate professor with Zhejiang Lab, and a visiting professor in Pillar of ISTD of SUTD Singapore. He also works for State Key Laboratory of Cryptology and Alibaba-Zhejiang University Joint Institute of Frontier Technologies. His main interests include general information security, cryptography, computer architecture, and wireless network. He is a member of the IEEE.

► For more information on this or any other computing topic, please visit our Digital Library at www.computer.org/publications/dlib.