LightFDG: An Integrated Approach to Flow Detection and Grooming in Optical Wireless DCNs

Amer AlGhadhban, *Member, IEEE*, Abdulkadir Celik, *Senior Member, IEEE*,
Basem Shihada, *Senior Member, IEEE*, and Mohamed-Slim Alouini, *Fellow, IEEE*.

Abstract—LightFDG is an integrated approach to flow detection (FD) and flow grooming (FG) in optical wireless data center networks (DCNs), which is interconnected via wavelength division multiplexing (WDM) based free-space optical (FSO) links. Since forwarding bandwidth-hungry elephant flows (EFs) and delay-sensitive mice flows (MFs) on the same path can cause severe performance degradation, the LightFDG optically grooms flows of each class into rack-to-rack (R2R) flows. Then, R2R-MF and R2R-EF flows are separately forwarded over lightpaths of separate MF and EF virtual topologies, respectively. Lightpaths are provisioned by jointly determining the capacity and route based on flows' arrival rate, size, and completion time request. To prevent EFs from congesting the MF lightpaths, high speed and accurate flow-detection mechanisms are also necessary for classifying EFs as soon as possible. Therefore, a fast-lightweight-and-accurate flow detection framework is developed by leveraging the transmission control protocol (TCP) behaviors. The proposed FD scheme has the flexibility of being implemented as in-network or centralized to classify flows of modifiable and unmodifiable hosts, respectively. Since the centralized scheme incurs considerable overhead, the processing and communication overhead is also mitigated by proposed techniques. Numerical results show that LightFDG outperforms traditional load balancers by about $3\times$ for EFs and $10\times$ for MFs. Along with the developed overhead mitigation methods, the centralized scheme is shown to provide up to $62\times$ lower overhead with 100% accuracy and with about $224\times$ higher detection speeds than the existing centralized solutions.

Index Terms—Flow classification; FSO; wavelength division multiplexing; lightpath provisioning; flow grooming.

I. INTRODUCTION

A. Preliminaries

The cloud service providers, mobile operators, electronic commerce companies, and content delivery services strive to enhance the performance of their data center networks (DCNs) for keeping their services up with the recent technological trends [1]. These include popular bandwidth-hungry services such as the fifth generation (5G) networks and beyond, the Internet of things, smart cities, big data, artificial intelligence, etc., that continuously exchange huge volume of diverse and complex data. To store and process such a deluge of data, DCNs are required to interconnect a plethora of servers

with high-speed low-latency links, which eventually causes a significant cost and complexity related to wire deployment, maintenance, and modifications [2].

B. Problem Statement and Motivation

Traditional wired DCN topologies with fixed and uniform link capacities have been shown to be deficient in supporting the optimal capacity allocation mechanisms for aforementioned applications [3]–[6]. This is mainly because of the fact that the majority of inter-rack traffic is exchanged between few racks whereas remaining racks exchange less traffic or no traffic at all. For instance, Hadoop servers of Facebook DCNs [5] communicate to only 1.5% of the other inter-rack servers, which comprises only 17% of the racks while approximately contributing to the 85% of entire DCN traffic. This pattern is aligned with previous findings in [4] where inter-rack traffic is carried out by up to 10% of the servers. Thereby, the DCN topologies with uniform link capacities are either under-utilized compared to the rack pairs that exchange less traffic or over-utilized in comparison with the rack pairs that generate larger traffic sets [7], [8], which is also known as *over-subscription*.

As a remedy wireless DCNs (WDCNs) can provide necessary reconfigurability to alleviate over-subscription, deliver a higher throughput thanks to efficient bandwidth utilization, eliminate the need for cost and power inefficient switches, and reduce the cabling cost & complexity [9]. Required high-speed wireless links can be obtained by state-of-the-art technologies such as millimeter wave [10] and free-space optical (FSO) communications [6], [11]. In particular, extremely high link capacities of FSO links can be further improved by wavelength division multiplexing (WDM) methods which yield a high fan-out (i.e., large number of channels per FSO link) as well as elastic topologies [12]. Recently, Ciaramella et al. [13] achieved a total of 1.28 Tbps speed in an outdoor experiment over 212 meters distance by using WDM-FSO link of 32 wavelengths (32×40 Gbps). Thanks to the acclimatized environment of DCNs, indoor FSO links can attain even better performances at longer distances since they are less vulnerable to the optical channel impairments such as scintillation, alignment errors, atmospheric turbulence, etc.

In this paper, a spine-leaf WDCN architecture is considered which is inter-connected via WDM-FSO links. Since bandwidth demands of flows can be much lower than the available capacity on WDM channels, aggregating (i.e., grooming) sub-wavelength flows onto high-speed lightpaths can offer desirable link utilizations [14], [15]. Considering the enormous

The authors are with Computer, Electrical, and Mathematical Sciences and Engineering (CEMSE) Division at King Abdullah University of Science and Technology (KAUST), Thuwal, 23955-6900, KSA. Amer AlGhadhban is also with Electrical Engineering Department at University of Hail, Hail, 55474, KSA. (Corresponding author: Abdulkadir Celik, email: abdulka-dir.celik@kaust.edu.sa) This article was presented in part at the IEEE GLOBECOM 2018, Abu Dhabi, UAE [1].

amount of flows generated across the DCNs, flow grooming (FG) is also useful for avoiding unnecessary topology reconfiguration delays caused by the computational and control overhead of handling flows individually. Therefore, merging FG with the flexibility of high fan-out WDM-FSO links may provide necessary performance enhancement required by DCNs.

QoS provisioning is another crucial task to meet the diversified demands of different flow classes. It has been shown that the majority of bytes are carried by bandwidth-hungry elephant flows (EFs), while the majority of flows are delay-sensitive mice flows (MFs) [3]–[6]. To reap the real benefits of reconfigurability, the FG policy should also adapt the network topology to provision the diverging QoS requirements of varying flow classes. From a load balancing point of view, MFs experience a severe delay if they are routed along the same path of EFs [16]. Therefore, grooming and forwarding the same class of flows altogether is a promising solution to handle heavy load and highly dynamic traffic characteristics of DCNs [17], [18]. In this regard, fast and precise flow classification schemes are necessary to groom the same class of flows and then steer them to their dedicated virtual-topology. Unfortunately, existing flow classification mechanisms have major drawbacks in terms of detection speed, controlling overhead, and accuracy.

Moreover, a DCN is a heterogeneous environment where the hosts can be modifiable or unmodifiable as per the DC administrator's right and capability to make changes [19]. In the former, system administrator can make modifications at user or kernel levels, which is subject to an agreement between the service provider and the subscriber. In the latter, the system modifications may not be possible for two reasons: 1) Changes are prohibited if a server and/or virtual machines (VMs) on the server belong to an identity other than the system administrator, and 2) Changes are not possible if a VM runs on a closed-source operating system (OS). Since the system administrator cannot modify all hosts, both centralized and in-network flow detection mechanisms might be needed.

Accordingly, this paper proposes LightFDG; an integrated flow detection, grooming, and forwarding approach. The LightFDG first determines the flow classes in a fast and accurate fashion, then grooms the same flow classes to obtain R2R flows, and finally forwards these aggregated flows over dedicated virtual topologies.

C. Paper Objectives and Achievements

Main contributions of this work can be summarized as follows:

- In order to avoid congestion and performance degradation, the proposed flow grooming and forwarding (FGF) policy is designed to handle the two main flow classes (i.e., MFs and EFs) by separately routing them over disjoint virtual topologies. To overcome the complexity of handling each flow individually, a three-step flow grooming approach is adopted to obtain rack-to-rack (R2R) MFs and EFs. By jointly determining the route & capacity of R2R MFs/EFs, the proposed FGF policy

innately balances the traffic load across the WDM-FSO links and efficiently utilizes the available bandwidth while provisioning the QoS demands of each flow class. The emulation results show that the proposed FGF policy outperforms traditional load balancing schemes by about $3\times$ for EFs and $10\times$ for MFs.

- Instead of using traditional sampling methods, LightFDG employs a fast, lightweight, and accurate flow detection (FD) framework that collects packets from every flow at a pace proportional to the flow rate. To do so, the proposed FD scheme leverages low-frequency (i.e., three-way handshaking packets) and high-frequency (i.e., acknowledgment packets) phases of the transmission control protocol (TCP). This categorization enables speeding up the classification process and minimizes the associated delays. Furthermore, it has flexibility to be implemented either in a distributed (in-network) or centralized manner to classify flows of modifiable and unmodifiable hosts, respectively.
- The in-network scheme is a module installed in virtual-switches/hypervisors, where the network is relaxed by utilizing location privileges of virtual-switches/hypervisors. On the other hand, the centralized scheme is implemented in the central unit (CU) by preconfiguring edge switches (ESs) to capture indicative packets (e.g., ACK, SYN, FIN, RST) which contain general information about a flow such as its starting time, source, and destination. The evaluation results show that the proposed centralized FD mechanism is about $224\times$ faster than the existing centralized solutions with a 100% accuracy with very low computational and communication overhead. Nonetheless, capturing and forwarding all the packets naturally introduce extensive processing and communication overhead. As a remedy, the ESs are configured to capture only indicative TCP packets. Two techniques are developed to further reduce the overhead: *Stop Useless Notifications* and *Sample Only ACK Packets*, where we tackle the challenge of differentiating between ACK and non-ACK packets. Noting that overhead of the centralized scheme without proposed techniques are close to that of the well-known sampled flow (sFlow) standard [20], developed methods provide up to $62\times$ lower overhead.

D. Paper Organization

The rest of the paper is organized as follows: The background and related work is presented in Section II. The node architecture, network topology, and an overview of LightFDG is introduced in Section III. Flow grooming and forwarding policy is explained in Section IV. Then, centralized and in-network FD schemes are provided in Section V. The Section VI presents the overhead minimization methods for the centralized scheme. Thereafter, the performance evaluations are presented in Section VII. Finally, conclusions are drawn in Section VIII.

II. BACKGROUND AND RELATED WORK

A. Literature Survey

Even though traffic grooming has been extensively studied for passive optical networks [21], it was first introduced for wired DCNs in [22], [23] where flows are simply groomed into three classes of wavelengths. The groomed wavelengths are broadcasted within racks and higher-layer switches. However, these works only consider fixed wavelength capacity and do not deal with the flow characteristics of real-life DCNs. In [17], [18], the authors conceptualized FG for MFs in WDM-FSO based wireless DCNs consisting of optoelectronic switches. After formulating the optimal FG problem, a sub-optimal FG policy is designed for MFs, whereas EFs are carried out separately via server-to-server express lightpaths without going through any grooming operation. The results showed that the proposed FG method provides superior performance in both throughput and flow completion times.

Existing FD schemes can be categorized into in-host [24], in-network [25]–[27] and centralized [28]–[31]. The in-host scheme is implemented in the host kernel where packets are counted before being injected into the network. When the counter exceeds a pre-defined threshold value, the flow is considered as an EF. However, in-host solutions need to modify the kernel network stack of the DCN hosts and hence require higher privileges that cannot always be granted.

Unlike the in-host approaches, in-network schemes detect EFs at network switches. The flow information along with their counters are maintained at the switch memory whereas the comparison task is performed internally at the switch hardware (e.g., HashPipe [25] and OpenSketch [27]) or software (e.g., Software-Defined Counter [26]), or otherwise is exposed to a centralized entity (e.g., FlowRadar [32]). One of the main challenges of the in-network mechanism is the limited resources of network switches (e.g., memory, processing delay, and power consumption). For instance, the time budget of DCN switches is too small (e.g., 25 ns) for a 128 bytes packet in a 40Gbps switch. Moreover, flow arrival-rates are in the order of milliseconds [3] and the number of flow-entries in switching/routing tables is in the order of a hundred thousand. As a remedy, the researchers introduced hybrid solutions where some of the classification functions are migrated into cheaper hardware [27] or into a centralized entity [32].

The centralized schemes employ the available network features (e.g., packet sampling, port mirroring, or statistical polling) to collect network statistics and forward them to a centralized collector, which is usually programmed to perform the flow classification functions. However, the existing centralized schemes have limitations such as high monitoring overhead, the necessity for hardware modifications, high latency, and low accuracy [20], [29]–[31]. For example, OpenSample measures the link utilization from the sequence numbers in the TCP header and uses sFlow to collect the TCP packets [29]. Planck [30] utilizes the port-mirroring, which causes extra processing overhead and cabling from ESs to the centralized collector. Moreover, it is possible to miss some EFs due to the limited buffer size. Similar to sFlow,

Planck can also attempt to detect already classified flows. To overcome these limitations, a hierarchical statistics pulling (HSP) mechanism is developed by utilizing a combination of aggregate and individual statistical messages in the OpenFlow protocol [33]. In FLight [34], centralized classification idea is introduced with preliminary results extracted from small-scale experiments. Unfortunately, FLight fell short of convincing the idea can overcome the practical limitations of Planck.

B. Distinctive Features of LightFDG

LightFDG can be distinguished from above works in the following aspects:

- Instead of assuming flow classes are known a priori, LightFDG provides a holistic approach by integrating FGF with effective FD mechanisms. Numerical results show that classification accuracy and speed has a significant impact on the FGF performance.
- Unlike the above centralized and in-network FD schemes, LightFDG also provides a generic flow classification mechanism that is capable of operating either in centralized or in-network fashion. In addition to being fast-accurate-and-lightweight, the proposed detection scheme is favorable for two reasons: 1) To strike a balance between communication and computation overhead, and 2) To be able to compatible with both modifiable and unmodifiable hosts.

III. AN OVERVIEW OF THE LIGHTFDG

This section first explains the node architecture and DCN topology, then outlines the proposed LightFDG approach, and finally briefly introduce its algorithmic implementation.

A. Node Architecture and Network Topology

A two-tier DCN architecture has been considered where each of N leaf layer ESs is connected with ηN core switches (CSs) in the spine layer, where η is the spine/leaf ratio. That is, leaf and spine layers are interconnected in a full-mesh manner as shown in Fig. 1. LoS connectivity between these layers can be realized by using the physical topology in Fig. 1 where optical transceivers of ESs are connected to CS transceivers located at the top. WDM-FSO links can be implemented by photonic integrated circuits (PIC) [35] or photo-diode array [36] based transmitters to steer the lightbeam to the photo-diode receivers. Since transceivers are not required to change their direction after the setup, an alternative approach is setting gimbals atop the racks and directing transmitter and receiver towards each other¹.

Each ES is considered as an optical cross-connect (OXC) with N I/O ports such that received optical beam at each input port is first demultiplexed into W wavelengths then fed into the connection matrix which determines the connectivity between OXC's I/O ports. After that output lines of the connection matrix are groomed via a multiplexer at each

¹We should note that our contributions do not involve in hardware topology of WDCNs. We are rather interested in virtual topology design to investigate the flexibility and reconfigurability advantages of wireless DCNs.

Algorithm 1 - LightFDG

Input: $\mathcal{F}(\mathcal{V}, \mathcal{W}, \mathcal{E})$, th

```

1:  $\mathcal{F}(\mathcal{V}, \mathcal{W}, \mathcal{E}) \leftarrow$  Initialize the physical topology
2:  $\mathcal{G}_m(\mathcal{V}_m, \mathcal{W}_m, \mathcal{E}_m) \leftarrow$  Initialize the MF virtual topology
3:  $\mathcal{G}_e(\mathcal{V}_e, \mathcal{W}_e, \mathcal{E}_e) \leftarrow$  Initialize the EF virtual topology
4:  $(\mathcal{L}_i^j, \mathcal{G}_m(\mathcal{V}_m, \mathcal{W}_m, \mathcal{E}_m)) \leftarrow$  LIGHTPATH PROVISIONING ( $m$ )
5:  $(\tilde{\mathcal{L}}_i^j, \mathcal{G}_e(\mathcal{V}_e, \mathcal{W}_e, \mathcal{E}_e)) \leftarrow$  LIGHTPATH PROVISIONING ( $e$ )
6: for each flow  $f$  between racks  $i$  and  $j$ ,  $\forall i, j$ , do
7:    $t_f \leftarrow$  Record the initial sequence number of flow  $f$ .
8:   while  $f$  is not completed do
9:      $t_c \leftarrow$  Update the sequence number
10:    if  $t_c - t_f < th$  then (Flow Detection)
11:      Flow Grooming: Perform three-step MF grooming
12:      Flow Forwarding: Forward the R2R-MF over  $\mathcal{L}_i^j$ 
13:    else
14:      Flow Grooming: Perform three-step EF grooming
15:      Flow Forwarding: Forward the R2R-EF over  $\tilde{\mathcal{L}}_i^j$ 
16:    end if
17:  end while
18: end for
19: procedure LIGHTPATH PROVISIONING ( $x$ )
20:   for  $i=1:N$  do
21:     for  $j=1:N$  do
22:        $E \leftarrow$  Determine the required intensity as per (2)
23:        $\mathcal{P}, \mathcal{C} \leftarrow$  Obtain K shortest paths and costs using K-SPR
24:        $\mathcal{P}_{ij}^x \leftarrow$  Decide on the path  $\text{Rack}_i \rightarrow \text{Rack}_j$ .
25:        $\mathcal{W}_{ij}^x \leftarrow$  Assign the wavelength to  $\mathcal{P}_{ij}^x$ 
26:        $\mathcal{L}_{ij}^x \leftarrow (E_{ij}^x, \mathcal{P}_{ij}^x, \mathcal{W}_{ij}^x)$  Record the lightpath
          $\text{Rack}_i \rightarrow \text{Rack}_j$ .
27:        $\mathcal{G}_x(\mathcal{V}_x, \mathcal{W}_x, \mathcal{E}_x) \leftarrow$  Update the virtual topology.
28:        $\mathcal{F}(\mathcal{V}, \mathcal{W}, \mathcal{E}) \leftarrow$  Update the physical topology.
29:     end for
30:   end for
31:   return  $\mathcal{G}(\mathcal{V}, \mathcal{W}, \mathcal{E})$ ,  $\mathcal{L}_{ij}^x, \forall i, j$ .
32: end procedure

```

link cannot be shared by different lightpaths and 2) *Wave-length continuity constraint* requires lightpaths to operate on the same wavelength along the routing path since OXCs are assumed to be not capable of wavelength conversion. Based on the lightpath provisioning procedure of Algorithm 1, the CU manipulates virtual topologies to provision R2R lightpaths, which consists of three tasks: light intensity (i.e., capacity) allocation, wavelength assignment, and routing.

C. Algorithmic Implementation of LightFDG

The pseudo-code of the LightFDG is given in Algorithm 1 where the first line initializes the virtual topology which is always kept updated by the CU over a wavelength dedicated for broadcasting the control signals. The MF (EF) virtual topology is characterized by a weighted graph $\mathcal{G}_x(\mathcal{V}_x, \mathcal{W}_x, \mathcal{E}_x)$ where $x = m$ for MFs, $x = e$ for EFs, $\mathcal{V}_x$ is the set of nodes, $\mathcal{E}_x$ presents available light intensity, and $\mathcal{W}_x$ symbolizes wavelength availability of the FSO links. Similarly, physical topology is denoted by $\mathcal{F}(\mathcal{V}, \mathcal{W}, \mathcal{E}) \supseteq \bigcup_x \mathcal{G}_x(\mathcal{V}_x, \mathcal{W}_x, \mathcal{E}_x)$. After that, the algorithm first provisions the R2R-MF lightpaths between the rack pairs (line 2). The residual light intensity and wavelengths are then exploited to provision the R2R-EF lightpaths between the racks. Since LightFDG guarantees the QoS demands of delay-sensitive MFs first, they are prioritized over EFs for protection against performance degradation in heavy traffic conditions. The details of lightpath provisioning are provided in the next sections.

For each flow arrival, an initial ACK sequence number, t_f , is recorded along with the current sequence number, t_c (lines 6-8). As long as the packet size is below a certain threshold, $t_c - t_f < th$, LightFDG regards flow f as an MF (line 10), grooms it with other MFs destined to the same rack (line 11), and forwarded over predetermined R2R-MF lightpaths (line 12). Otherwise, f is regarded as an EF, groomed with other EFs destined to the same rack (line 14), and forwarded over predetermined R2R-EF lightpaths (line 15). In the remainder, we present more technical details regarding each step of Algorithm 1.

IV. FLOW GROOMING AND FORWARDING POLICY

FGF policy comprises of three joint subproblems: 1) Designing grooming strategy, 2) Virtual topology design, i.e., provisioning lightpaths' capacity and routing the aggregated traffic over the physical topology, and 3) Assignment of wavelengths to the lightpaths. Since each of these subproblems are NP-hard [39], FGF is also an NP-hard problem which belongs to the class of mixed-integer linear programming (MILP) problems. Due to dynamic and heavy traffic characteristics of DCNs, obtaining an optimal FGF policy requires impractical time complexity even for small-scale DCNs. Hence, fast yet high-performance suboptimal solutions are necessary to implement FGF for WDCNs.

A. Flow Grooming: A Three-Step Strategy

The three-step optical grooming policy aggregates each class of flows in the optical domain as follows [17], [18] :

- 1) *Server-to-server (S2S) Grooming*: Each server combines all flow arrivals destined to the same server into a single flow. Denoting the set of hosts/servers within rack i as $\mathcal{R}_i$, arrival rates of S2S-MFs and S2S-EFs from $s \in \mathcal{R}_i$ to $t \in \mathcal{R}_j$ are assumed to follow Poisson distribution with arrival rates λ_{ij}^{st} and $\tilde{\lambda}_{ij}^{st}$, respectively.
- 2) *Server-to-Rack (S2R) Grooming*: Servers further groom S2S flows according to common destination rack such that all flows destined to the same rack is groomed into a single S2R flow. Assuming that flow arrivals are independent from each other, S2R-MF and SR-EF arrival rates also follow a Poisson distribution with the composite rates of $\lambda_{ij}^s = \sum_{t \in \mathcal{R}_j} \lambda_{ij}^{st}$ and $\tilde{\lambda}_{ij}^s = \sum_{t \in \mathcal{R}_j} \tilde{\lambda}_{ij}^{st}$, respectively.
- 3) *Rack-to-Rack (R2R) Grooming*: The ES finally grooms all S2R flows received from different servers according to their destination racks to obtain R2R flows. Similarly, overall arrival rate for the R2R-MF and R2R-EF flows have the composite rates of $\Lambda_{ij}^m = \sum_{s \in \mathcal{R}_i} \lambda_{ij}^s$ and $\Lambda_{ij}^e = \sum_{s \in \mathcal{R}_i} \tilde{\lambda}_{ij}^s$, respectively.

Following the three-step grooming, R2R-MF and R2R-EF flows are then directed to the relevant optical transmitter based on the predetermined routing paths. Therefore, CSs are not required to implement any optical grooming since they merely forward the lightpaths to the ES transceivers of the destination racks.

B. Flow Forwarding: Lightpath Provisioning Procedure

We denote the R2R lightpaths between racks i and j by $\mathcal{L}_{ij}^x = \{E_{ij}^x, \mathcal{P}_{ij}^x, \mathcal{W}_{ij}^x\}$ where $x = m$ ($x = e$) represents MFs (EFs), E_{ij}^x is the allocated light intensity, $\mathcal{P}_{ij}^x$ is the routing path, and $\mathcal{W}_{ij}^x$ is the assigned wavelength. As shown between lines 19 and 28, we iteratively provision lightpaths for the rack pairs as follows.

The first step of the lightpath provisioning (line 22) is determining the required light intensity of the wavelengths along the R2R paths, which is explained as follows: Based on the flow completion time τ_m and flow size F_m of MFs, required capacity for R2R-MF between Rack $_i$ and Rack $_j$ is given by, $C_{ij}^m = \Lambda_{ij}^m F_m / \tau_m$, which must be guaranteed at each hop along the routing path. By substituting C_{ij}^j into (1), required light intensity for wavelength and link pairs of $\mathcal{L}_{ij}^m$ is given by

$$E_{k,l}^{\omega} = \left(\frac{\left(2^{\frac{2C_{ij}^m}{B}} - 1 \right) 2\pi}{e(h_k^l)^2} \right)^{1/2}, \forall k, l \in \mathcal{P}_{ij}^m, \omega \in \mathcal{W}_{ij}^m. \quad (2)$$

Likewise, required light intensity for wavelength and link pairs of $\mathcal{L}_{ij}^e$ can be obtained by substituting $C_{ij}^e = \Lambda_{ij}^e F_e / \tau_e$ into (1).

In Line 23, we eliminate edges with insufficient light intensity (i.e., availability is less than the required amount in (2)) and run a modified version of the loop-less k-shortest path (KSP) algorithm [40] which not only finds the path with the maximum available light intensity (i.e., capacity), but also $k - 1$ other paths in non-increasing order of the available intensity. Thereafter, the available light intensity of these k lightpaths are weighted with the available number of wavelengths on the route. Finally, line 24 selects the lightpath with the maximum intensity and number of wavelength product as the dedicated R2R path for the racks i and j . In this way, we ensure at each iteration that the number of available wavelengths and light intensity on FSO links are evenly distributed across the entire topology.

In order to ensure that there exists no less than two R2R paths between each rack pairs (i.e., one for MFs and other for EFs), the required least number of wavelengths is given by

$$W \geq \left\lceil \frac{2(N-1)}{\eta N} \right\rceil. \quad (3)$$

Assuming (3) holds, line 25 randomly assigns wavelengths to lightpaths since they are identical once the lightpath capacity and route is determined. Nonetheless, we must note that wavelength continuity and collision constraints may cause connectivity issues between the rack pairs if the number of wavelengths is limited (i.e., $W < \lceil 2(N-1)/(\eta N) \rceil$) [41]. In this case, lightpath provisioning procedure must allow multiple lightpaths to share the same wavelength based on a certain multiple access scheme, e.g., time-division multiple access [42]. After that, line 26 records the lightpath and line 27 (28) updates the virtual (physical) topology for the next iteration. Following the MF lightpath provisioning in line 3, line 4 also employs the above LIGHTPATH PROVISIONING

procedure for R2R-EF lightpaths by exploiting the residual light intensity and available wavelengths on FSO links.

V. FLOW DETECTION: A FAST, ACCURATE, AND LIGHTWEIGHT FRAMEWORK

As the TCP carries out almost 99% of the DCN traffic [3], its attributes are leveraged by many existing flow classification methods to build accurate and fast flow detection mechanisms [24], [29]. A single packet of every TCP flow offers partial information about the flow size which is useless if it is individually considered. When two packets are captured from the same TCP flow, one can discover the number of bytes that have been transmitted between the two captured packets. This technique is fundamental idea of the existing flow detection solutions (e.g., OpenSample-TCP and Planck) where the difference between the TCP sequence numbers of captured packets and time of capture is used to measure the link utilization.

A. Operational Principles of the Proposed FD Schemes

LightFDG divides the TCP communication into low-frequency and high-frequency phases. The former consists of three-way handshaking packets (i.e., SYN, SYN+ACK, and ACK) that appears once to set the initial sequence number and prepare the connection parameters at the beginning of each flow. On the other hand, the latter comprises of ACK packets sent by the receiver to acknowledge the transmitter about the successful reception of the flow packets. LightFDG also benefits from the privileges of its position, i.e., the CU in the centralized scheme or the virtual-switch/hypervisor in the in-network scheme.

To this end, LightFDG employs two main components: collector and classifier. The collector reads the initial sequence number (ISN) of every TCP flow, which is the first half of the required information to classify a flow, and stores it in the flow-information table. Likewise, the classifier needs to compare the ISN with every captured packet from the same flow until a threshold value is reached. However, instead of reading/capturing³ the data packets themselves, LightFDG is programmed to read/capture the headers of ACK packets. For instance, if the classifier reads ACK $_1$ and then ACK $_t$ after a while, the number of transmitted bytes between these two ACKs is calculated as $t - 1$. The ACK sequence number only presents the bytes that have been successfully received, that is, the lost or out-of-order packets are not counted. Therefore, the classifier needs only to compare the number of bytes with a predefined threshold value th to classify a flow as an EF. Once $t - 1 > th$ is satisfied, the flow is considered as an EF and the subsequent packets are marked to be forwarded via the virtual topology of EFs and the CU is reported about this detection.

Even though we consider one-way flows throughout the paper, the LightFDG is also capable of handling two-way flows where the receiver temporarily delays sending ACKs

³We use the term *reading* for in-network scheme and *capturing* for the centralized scheme.

Fig. 3: The in-network FD scheme installed in the hypervisor.

until it has data to be sent towards the transmitter. This is also known as piggybacking and favored by its efficient bandwidth utilization [43]. In principle, one-way LightFDG can still operate under piggybacking communication by individually handling flows at both directions. That is, the CU updates two records in the flow information table, i.e., one for each direction. Whenever a direction exceeds the predefined threshold value, its packets are marked to be forwarded through the EF virtual topology. However, delaying ACKs may cause the receiver jamming the service if it has no data to send back. This delay will naturally reduce the flow detection speed of LightFDG. Fortunately, this issue can be resolved by enabling a receiver timeout counter when a data frame is received. If the count terminates and there is no data frame to send, the receiver will send an ACK. The sender also adds a emitter timeout counter, if the counter ends without receiving confirmation, the sender assumes packet loss, and sends the frame again [43].

B. In-Network and Centralized FD Schemes

Although centralized and in-network schemes use the same components (i.e., collector and classifier), they differ in the implementation such that the centralized scheme is implemented in a CU (e.g., SDN controller) whereas the in-network scheme is implemented as a kernel module in the hypervisor/virtual-switches. While the centralized scheme is used to classify the flows of unmodifiable hosts, the in-network scheme is used to classify the flows of modifiable hosts. Due to its position in the hypervisor/virtual-switch, the in-network is faster than the centralized scheme in terms of detection speed. Although the administrative overhead of the in-network scheme is lighter than that of the in-host scheme, it needs to be installed in every hypervisor/virtual-switch, and the DC is expected to have various versions and brands. On the other hand, the centralized scheme is simpler where the network administrator needs only to install it in the CU, and it will configure the switches with the needed configurations. In what follows, we provide the implementation details of LightFDG in both centralized and in-network cases.

1) *In-network Scheme*: As shown in Fig. 3, the in-network LightFDG scheme is designed as module installed on the kernels of hypervisor of edge servers or in the virtual switches. The kernels of these systems provide a degree of control on exchanging packets, i.e., access to their packet headers

Fig. 4: Flowchart of the centralized FD scheme.

and unencrypted payload. For EF detection, LightFDG exploits two main components: collector and classifier. The collector is responsible for learning general flow information from the messages of the low-frequency phase. Moreover, it has a flow-information table to store the flow information (e.g., source/destination IP/MAC addresses and TCP sequence numbers). On the other hand, the classifier is accountable for detecting EFs from headers of the high-frequency phase messages.

2) *Centralized Scheme*: The network switch/router configurations can be performed physically or remotely: In the former, portable machines are connected to console port on the switch/router to perform the configuration steps, which is repeated for every switch/router in the network. In the latter, configurations are remotely written by using telnet/ssh terminal over physical links established between every switch/router and the CU after the necessary remote configuration activation commands are physically entered into the switches/routers. Alternatively, LightFDG employs a programmable interface (i.e., OpenFlow) that remotely communicates with ESs and configure them to capture the low and high frequency phase packets. Unlike the aforementioned methods, OpenFlow protocol has more abstracted and advanced commands that facilitate the configuration of network devices. Taking the dynamic nature of FGF into account, OpenFlow protocol is also a key solution to create virtual topologies and reconfigure it based on changes driven by FGF policy.

As demonstrated in Fig. 4, the centralized FD scheme configures ES to capture the TCP first phase packets (i.e., SYN+ACK and RST/FIN) and ACKs within the second phase packets. When these packets are present in a flow, the ESs send a *flow-notification* to collector of the CU which reads and writes the flow information and ISNs in a flow-information table. To accelerate the classification process, the flow-information table has been divided into multiple small tables each of which is indexed by a subnet ID. When the collector receives a captured message, it needs to figure out which table this message belongs to. Since every OpenFlow header has the ID of the source switch and the flow information, it can

Fig. 5: The LightFDG steers the ACK and non-ACK packets to separate virtual interfaces on every ES.

use the source IP or the data-plane ID to achieve this task⁴. To determine how many bytes have been successfully received, the flow classifier compares the sequence number of captured ACKs, t_c , with their correspondent ISNs in the subnet flow-information table. When the number of bytes exceeds the EF threshold value, the flow is then considered as an EF.

VI. MINIMIZING THE NETWORK OVERHEAD

The propagation of low and high frequency phase packets through the DCN may introduce a queuing delay and processing overhead on the network. Before proceeding to the proposed overhead reduction methods, it is important to demonstrate the feasibility of the fundamental idea by investigating two kinds of datasets; DCN [3] and enterprise [44] datasets. In Table I, the collections are ordered such that the sequence number is a suffix added to the name of the datasets. Table I shows that normal LightFDG without overhead minimization techniques (i.e., capturing only the low and high frequency packets) delivers a low overhead in comparison with capturing all the packets. Although LightFDG is required to collect much less packets than the total number of packets in DCN datasets, enterprise dataset may cause a significant overhead since almost all packets are required to be captured at some cases. Accordingly, following overhead reduction methods are proposed:

A. Sample Only ACKs Packets

In practice, there are two methods to capture and notify the classifier for every ACK message: 1) Configuring the ESs with particular OpenFlow rules to mirror every ACK message to the CU and 2) Using sFlow to sample every packet. Although the former method has a high accuracy, mirroring every ACK message to the collector introduces high communication and processing overhead. On the other hand, the random packet sampling has a low overhead and at the same time low accuracy. Therefore, the number of sampled packets is similar to the number of sampling from every TCP packet. To get the best of both methods, their advantages are combined as shown in Fig. 5 where sampling is only done for ACK packages by steering them to a separate virtual interface. In this way, the unnecessary overhead caused by sampling data (non-ACK) packets are eliminated. These interfaces could be

⁴In this work, we prefer to use the data-plane ID in the OpenFlow messages.

TABLE I:
The impact of LightFDG on real packet traces.

Dataset(ID)	# Packets	# Flows	Normal LightFDG	Ratio
Inter. 1	5700526	662	2313433	=2.5 ×
Inter. 2	2261261	3512	2174151	=1.04 ×
Inter. 3	84574	172	82817	=1.02 ×
Univ. 2.0	11917501	589	6333	=1881.8 ×
Univ. 2.1	11917527	775	14636	=814.3 ×
Univ. 2.3	11915859	727	7787	=1530 ×
Univ. 2.4	11915613	844	10528	=1131.8 ×

a high priority queue or a virtual interface associated with every interface in the ESs. Here, we especially focus on the ESs to reduce the overhead of capturing multiple packets of the same flow during a short period.

B. Stop Useless Notifications

The useless notifications are the messages that lost their value for the flow classifier such as the ACK messages of the already classified flows or notifications regarding flows that can be classified from their packet headers. Accordingly, the proposed FD mechanism relieves the overhead in two ways:

- 1) It is widely known that some portion (>30%) of flows can be pre-classified from application types (e.g., FTP, Syslog, and NTP) [45]–[47]. Indeed, LightFDG is capable of pre-classification based on application types because it already captures low-frequency phase packets. At the beginning of each flow, LightFDG can immediately and accurately identify the flow class of these applications from TCP port numbers [45]. This naturally eliminates the need for capturing high-frequency phase packets and resulting overhead.
- 2) Furthermore, the CU is triggered to reconfigure the ESs to stop capturing its ACK messages when the classifier detects an EF. In OpenFlow, for instance, the controller is responsible to perform such actions and install the necessary flow-entries with higher priority on the ESs. In this case, the classifier will not receive the notification messages about the classified flows, which naturally reduces the overhead.

VII. PERFORMANCE EVALUATIONS

In this section, we first evaluate the performance of the proposed FD mechanisms and show how they can improve the overall DCN throughput and reduce flow completion times (FCTs). To demonstrate the light overhead of LightFDG in a real DCN environment, LightFDG is directly run on the same datasets with the minimization methods explained in Sec. VI. Thereby, the LightFDG performance is influenced by the network configurations (e.g., QoS settings) and the link loads of the datasets. The following subsections show the classification characteristics of LightFDG during different traffic loads and communication scenarios. Finally, the virtues of LightFDG in improving throughput and FCT performance are illustrated in comparison with other schemes.

Fig. 6: Accuracy comparison of FD schemes in different traffic configurations during the mix scenario.

A. Network Setup

Evaluations are conducted using Mininet emulator [48] and POX [49] controller on machine with $16 \times (2.5\text{GHz-Intel Xeon CPU E5-2680v3})$ processors and 128 GiB memory. A leaf-spine topology is created with eight spines and leaves, each with 40 hosts. Iperf is used to generate flows whose arrivals follow exponential distribution. The Iperf permits specifying the flow size and communicating units. The hosts in Mininet are divided into two separate sets (clients and servers) and Iperf is configured to loop these two sets till all the members are visited and the launched flows are completed. The clients and the servers are always selected from different racks, and the number of MFs and EFs are determined according to the required percentage (e.g., 10%, 15%, and 20%). The flow ID of flows are directly collected from the Mininet and stored in another file to compare them with the output results from the classifier. Our evaluation is based on network instances (one-shot), where a new instance is run after ensuring that all flows of the ongoing instance are complete. Linux monitoring features are used to monitor the processing of evaluation components such as Iperf PID of every running flow.

Although the CU and the network management servers of a real-life DC have high computational power to process large number of flows, such processing capacity can not be offered in a contained lab environment. To mimic the traffic load of real DCs, rather than randomly generating thousands of flows, different traffic loads are steered into a specific link where our classification and the competitor solutions are executed. In this respect, the traffic between the first and second subnet is steered on a single path to make sure that the number of flows on that path is as needed, i.e., 100, 500, and 1K. During the flow mix evaluations, the percentage of MFs is set to 90% of the total flows. Unless stated explicitly otherwise, the size of MFs and EFs are 100KB and 128MB, respectively. EF detection threshold is 1MB. For comparison purposes, we consider OpenSample-TCP and sFlow as our benchmark detection schemes and their sampling-rate is configured to 1-to-1000 as suggested in [29], [30], [50]. For the sampling rate of the “Sample Only ACK Packets” method proposed in Section VI-A, there is no clear and direct way of deciding on the sampling rate. Thus, different sampling rates are evaluated

from 1-to-50 to 1-to-1000, and 1-to-100 is selected as it gives a desirable performance in terms of accuracy, speed and overhead. For the sake of a fair comparison with the other detection schemes, presented results do not consider any pre-classifiable flows to evaluate the performance of LightFDG at the worst-case. Nonetheless, it is worth to point out that overall LightFDG performance may be further improved with the increasing portion of pre-classifiable flows. At the extreme case of having all flows are pre-classifiable, the LightFDG will not suffer from classification delay and accuracy.

B. Detection Accuracy

The FD mechanisms may have true-negatives (i.e., an EF is considered to be an MF) and false-positives (i.e., considering an MF as an EF). In particular, the percentage of true-negative incidents is critical as keeping the EF on the path of MFs can cause a sever queuing delay for delay-sensitive flows. To evaluate the detection accuracy, both of the FD schemes is investigated under different flow scenarios; pure MFs, pure EFs, as well as a mix of MFs and EFs. Fig. 6 illustrates the accuracy levels of different schemes when 100, 500, and 1000 flows are exchanged. Both of the FD schemes obviously achieve a 100% detection accuracy, i.e., the percentage of true-negative incidents is 0%, in all scenarios. The high accuracy of the proposed FD mechanisms is mainly because of the fact that we only need to subtract the sequence number of the ACK messages from the initial sequence number. Since existing solutions’ accuracy varies for different threshold values, we compare our methods with the best case performance of the existing solutions. Fig. 6 clearly shows that sFlow provides quite a low accuracy in comparison with LightFDG. However, OpenSample-TCP obtains around 90% accuracy for a varying number of flows. Therefore, we consider only the OpenSample-TCP for comparison in the remainder of the evaluation.

Fig. 7 depicts the performance comparison of various schemes for 10, 100, and 500 flows in pure MF and EF scenarios⁵. Similar to the mixed scenarios, LightFDG schemes

⁵During the performance evaluations, Mininet is observed to crash in 1000 EF case. Since such errors have serious impacts on the evaluation accuracy, we focus on the comparisons between 10, 100 and 500 flows.

Fig. 7: Accuracy comparison of FD schemes in different traffic configurations during the pure scenario.

achieved a high level of accuracy in both pure MF and pure EF cases. Also, OpenSample-TCP attained 100% in pure MF evaluation, and 90%-98% in EF experiment. The accuracy increases with the increase in traffic load due to the long transmission delay of EFs during high congestion. In order to examine both schemes in a short transmission delay and high traffic loads, flow size is reduced to 64MB. Performing the 10, 100, and 500 experiments on the OpenSample-TCP algorithm, a high percentage of true-negative incidents is observed; 40%, 35.3%, and 19% respectively, which in turn is corresponding to 60%, 64.7%, and 81% accuracy. On the same evaluation basis, the proposed FD solutions achieved 0% of true-negative incidents.

At this point, it is necessary to discuss the reasons behind the low performance of sampling based methods, which is summarized as follows:

- In sFlow, packet sampling randomly captures one out of every S packets and forward it to a collector for analysis, where S is the sampling rate. Accordingly, the collector estimates the flow size by multiplying S with the size of the collected packets from the same flow. However, the calculated flow size is practically less than the actual one especially in bursty traffic scenarios where a large number of packets arrives in a time less than the sampling period. Therefore, detection speed of the sampling methods heavily depends on the flow size estimation as well as the sampling and flow arrival rates
- A single TCP packet has partial insight about a flow class. To make accurate and timely decisions, random sampling methods are required to collect multiple samples from the same flow as soon as possible, which is hard to ensure if flows arrive and complete much faster than the sampling rate.
- Random sampling has a low efficiency since the captured samples could be from the already classified flows or a flow completed before it has been classified.

Although these drawbacks can be partially mitigated by sampling at higher rates, resulting network overhead can exhaust the processing power of switches. Accordingly, the 1-to-1000 sampling rate is widely selected in literature [29], [30], [50] to strike a good trade-off between the detection performance and overhead. In Fig. 6, sFlow performance is enhanced

by increasing the sampling rate from 100 to 500. However, increasing the sampling rate is always not a solution since increasing from 500 to 1000 did not deliver a significant improvement.

On the other hand, OpenSample-TCP is built on top of sFlow to provide a better performance by 1) using TCP packet headers to enhance the efficiency of sFlow and 2) setting a shorter control loop period (100 ms) than that of sFlow (1-5 s). However, OpenSample-TCP still needs high sampling rates as stated in [29] to learn accurate and up-to-date statistics about the link utilizations. Notice that OpenSample-TCP requires several packets from the same flow to measure the size by comparing the EF threshold with difference of the first and last captured packets. Therefore, both OpenSample-TCP and sFlow show low accuracy in case of small EF sizes (e.g., 12.8 MB and 32 MB [8], [51]) because it is not possible to collect enough number of samples during the EF flow lifetime.

To this end, a high accuracy is intuitively expected from the proposed solutions since both centralized and in-network schemes are configured to capture almost every ACK. To this end, detection speed must be questioned for such high accuracy performance. In addition to its decentralized nature, the location of the in-network scheme in the pipeline of the packet paths makes it highly accurate and fast. Nonetheless, centralized scheme speed is enhanced by 1) steering the ACK packets to a separate virtual interface, 2) classifying the flows of some applications based on the main flow information, and 3) stop capturing ACK messages of the flows which are already classified.

C. Detection Speed

In order to investigate the detection speeds, we examine different scenarios of network loads; 10, 100, 200, or 600 flows. Notice that these are not the total number of flows in the evaluated DCN, they are rather the number of flows that have been steered into a single link where the collector is implemented. The average speeds of detection during the different traffic loads and configurations are illustrated in Fig. 8a. According to the load on the path, which range from 2.2 to 580 milliseconds, the centralized scheme is $223.67 \times$ to $3 \times$ faster than OpenSample-TCP. Such performance degradation is because of that EF detection is slower than the

Fig. 8: Detection speed comparisons under various link loads.

Fig. 9: Speed and accuracy comparisons under varying flow sizes.

measured link utilization. That is, OpenSample-TCP [29] can measure the link utilization in 100 milliseconds while it needs about 1 second to detect EFs in this evaluation. The comparison between centralized and in-network scheme is depicted in Fig. 8b where in-network scheme provides a faster performance especially if link loads increase. This is expected since having large number of flows cause more overhead in the CU.

A varying flow size setup is also investigated as follows: Every MF (EF) randomly selects its size from a range of 10KB to 1MB (10MB to 128MB), where threshold is set to 1MB. Under 500 and 1K traffic loads, we consider 1-to-500 and 1-to-100 sampling rates for OpenSample-TCP and sFlow. Since the in-network scheme is already shown to perform better, we narrow our focus on the comparison with the centralized scheme. For 1-to-500 sampling rate, Fig. 9a shows that the centralized scheme still achieves 100% accuracy while OpenSample-TCP at 500 load reaches about 95%, which reduces slightly at 1K load. The Fig. 9b demonstrates the superiority of the centralized scheme's speed compared to OpenSample-TCP at 500 load. On the other hand, setting sampling rate to 1-to-100 results in a chain of CU failures and 100% CPU utilization, hence, a comparison between LightFDG and OpenSample schemes was not possible. Although sFlow did not encounter such technical problems since it uses a different collector, both 1-to-100 and 1-to-500 scenarios detected more than 50% of the flows as EF. That is, more than 55% of the MFs are considered EF. Hence, we omitted the sFlow results and content ourselves with the comparison of the centralized and OpenSample-TCP schemes.

D. Detection Overhead

In this section, we measure the FD overhead on the CU and the network. Even though the previous network setup

Fig. 10: The processing overhead during the first 120 sec.

and topologies are still valid, the CU is installed in a virtual machine with four processors and 10GB memory. To collect the statistics (e.g., CPU and memory utilization) of the CU, we use a `psutil` python library [52], configure `top` Linux tool to collect the device statistics for every second, and store them on a particular file. For all evaluations, the CU needs to set up the network and prepare the OvS switches with necessary configurations as well as the sFlow rules. This start-up phase consumes about 100% of a single CPU processing capacity and installs about 7KB of flow-entries. The communication and computation overhead of the in-network scheme is lower than that of the centralized scheme. However, the computation of in-network scheme is at the servers, while the CU is responsible for executing the centralized algorithm. It is also observed that the computation overhead of the in-network FD scheme is very small.

Since the centralized FD configures sFlow in the virtual interfaces with a fast sampling rate, the communication and the processing overhead of the centralized FD is expected to be worse than the others. The OpenSample-TCP continues the collection of samples for all of the flows regardless whether the flow has been classified or not. Thereby, the communication overhead and the processing capacity have not been improved after the classification. Hence, the total communication overhead has increased with the increase in the size of the transmitted file. On the contrary, centralized FD has the overhead minimization techniques that capture only the ACK packets of the unclassified flows. This is justified when the flow size is increased. We found that the communication overhead of the centralized FD is about 70KB while it is 63KB for OpenSample-TCP. For larger flow sizes (e.g., 256MB), the communication overhead of the centralized FD interestingly remains close to 70KB while it increases dramatically for OpenSample-TCP.

Thanks to its simplicity and faster detection, the computational overhead of the centralized FD is $5\times$ lower than that of the OpenSample-TCP, and close to that of sFlow (1-to-10k). The results are illustrated in Fig. 10 where statistics are collected during the transmission of 1k mix-flows. When we run the proposed overhead minimization algorithm on the enterprise and DCN datasets, the overhead of ACK messages reduced to minimal values compared to the ACK arrival-rate

TABLE II:
The arrival-rate of LightFDG ACK messages.

Dataset(ID)	Dataset ACK/Sec	LightFDG ACK/Sec	Ratio	Dataset(ID)	Dataset ACK/Sec	LightFDG ACK/Sec	Ratio
Inter. 1	3854.2	62.6	61.6x	Univ. 2.0	5.1	3.33	1.5x
Inter. 2	3615	62.5	57.84x	Univ. 2.1	11.9	0.798	15.1x
Inter. 3	137.5	4.15	33.13x	Univ. 2.3	6.892	0.19	36.3x

(a) MF throughput for different algorithms.

(b) EF throughput for different algorithms.

(c) MF-FCT for different algorithms.

Fig. 11: Throughput and FCT performance comparisons during different traffic scenarios.

of the datasets themselves. The results are shown in Table. II, ranging from $1.5\times$ up to $61.6\times$. Moreover, the collected results from `top` show that the centralized LightFDG and other solutions have a small impact on the memory ($\leq 1\%$).

E. Network Throughput Results

Throughout this subsection, the size of MFs and EFs are set to 50KB and 128MB, respectively. Since the links in Mininet are limited by the processing capacity of the host machine, we configured the FSO-links with 10 Gbps and the cabled links with 1 Gbps which means FSO links are ten times faster than wired DCN links. The WDM-FSO links can be thought as parallel independent channels, however, they are coupled by the total available lightpath power intensity. Accordingly, in this evaluation part each FSO link consists of 4 wavelengths, which are realized as virtual links in the simulation. Since the emulator is limited in DCN size, optical channel gains are not distinguishably different due to similar link distances and thus assumed to be identical without loss of generality.

We use MapReduce to mimic workloads of real DCNs whose shuffle-phase communication pattern has k servers from every rack communicate with another k servers in a different rack. For instance, the hosts in $\mathcal{R}_i$ have been divided into two sets and every set has k servers, e.g., 4 MF and one EF. Each of k server is communicating with k servers of rack j , $j \neq i$, different than other k servers of the same rack. For accurate results, we evaluated the routing algorithms with various percentage of EF: 10%, 15%, and 20%. We set $k = 20$ in all evaluation scenarios, that is, the number of MFs in the 10% scenario is 20, and the number of EFs is one.

For the network throughput results, we consider the following cases:

- 1) *Equal-Cost-Multi-Path routing (ECMP)* [53] is a widely used DCN routing method which uses the packet header

information, such as the IP/MAC addresses and TCP port numbers, as a key for a hash function. The outgoing path is the output hash value modulo the number of outgoing paths. This strategy splits the flows among available paths. Since the header information for an individual flow is the same during the session, the packets of the same flow are always forwarded via the same path.

- 2) *ECMP-FSO* is an ECMP routing algorithm supported by the FSO technology. In this routing method, the link capacity is equally divided between the wavelengths, that is, the capacity of every wavelength is fixed to 2.5 Gbps. Each flow was assigned to a single wavelength. However, when flows are more than the available number of lightpaths, the packets of the waiting flows are enqueued until a lightpath is available for transmission. Since wavelengths capacities are fixed and uniform, this case can be regarded as a DCN with fiber connections.
- 3) *FG-FSO* refers to the proposed FGF with a priori knowledge of flow classifications as in [17], [18].
- 4) *LightFDG* is the proposed 3-step optical FG algorithm employing the distributed FD mechanism.

Notice that the first three cases have the perfect and a priori knowledge of flow classifications. By considering the first three cases, we compare the virtues of re-configurable flexibility of WDM-FSO based DCNs and bandwidth utilization performance of FG with wired DCNs, i.e., first two cases. On the other hand, the last case shows the impact of in-network LightFDG scheme on the proposed FG policy

Fig. 11 compares the throughput and FCT performance of above cases under various work load. Apparently, LightFDG provides much superior performance than the ECMP-FSO by grooming and forwarding each class separately over lightpaths provisioned based on the traffic conditions and QoS demands.

Notice that ECMP does not groom traffic and distinguish packets of different flow types while forwarding them over multiple paths for the sake of load balancing. Nonetheless, both MFs and EFs experience performance degradation. In particular, routing MFs along with EFs has a more significant impact on MF throughput (Fig. 11.a) and FCT (Fig. 11.c).

Since the time constraint is 1ms, the bandwidth demand for 5%, 10%, and 15% scenarios are 7.6Gbps, 7.2Gbps, and 6.8Gbps, respectively. As we initially treat all flows as MFs, EF detection delay has no impact on MF performance as shown in Fig. 11.a. That is why LightFDG achieves the same throughput of FG-FSO for MFs. On the other hand, Fig. 11.b shows that the impact of time elapsed for detecting EFs has a distinguishable impact on throughput albeit the fast and accurate FD performance of the in-network scheme. Referring to the detection speed comparisons in previous subsections, detection delays caused by the CU overhead would further deteriorate the system performance. In this regard, OpenSample TCP is expected to severely deteriorate the overall network throughput because of the poor detection speed and accuracy. Finally, Fig. 11.c demonstrates how the LightFDG and FG-FSO achieved about 98% of the flow completion demand of MFs, while about 70% of the MFs in ECMP-FSO and exceeded the time constraint and about 88% of the MFs in ECMP exceeded the time constraint (i.e., complete after the 1 ms time constraint).

VIII. CONCLUSIONS

LightFDG is an integrated approach to flow detection and grooming in optical wireless data center networks (DCNs). The LightFDG optically grooms flows of each class into R2R flows, which are separately forwarded over lightpaths of separate virtual topologies dedicated to each flow class. Since high speed and accurate flow-detection mechanisms are necessary to prevent EFs from congesting the MF lightpaths, LightFDG leverage TCP behaviors to develop a generic FD scheme that is suitable for in-network and centralized implementation. Numerical results show that classification accuracy and speed have a significant impact on the FGF performance. Performance evaluations also exhibits that proposed FD mechanisms outperforms sampling based methods in speed, accuracy, and overhead.

REFERENCES

- [1] A. Al-Ghadhban, A. Celik, B. Shihada, and M.-S. Alouini, "LightFD: A lightweight flow detection mechanism for traffic grooming in optical wireless DCNs," in *Proc. IEEE Global Commun. Conf. (GLOBECOM)*, Dec. 2019, pp. 1–6.
- [2] A. S. Hamza, J. S. Deogun, and D. R. Alexander, "Wireless communication in data centers: A survey," *IEEE Communications Surveys Tutorials*, vol. 18, no. 3, pp. 1572–1595, Third Quarter 2016.
- [3] T. Benson, A. Anand, A. Akella, and M. Zhang, "Understanding data center traffic characteristics," *ACM SIGCOMM Computer Communication Review*, vol. 40, no. 1, pp. 92–99, 2010.
- [4] S. Kandula, S. Sengupta, A. Greenberg, P. Patel, and R. Chaiken, "The nature of data center traffic: measurements & analysis," in *Proceedings of the 9th ACM SIGCOMM conference on Internet measurement conference*. ACM, 2009, pp. 202–208.
- [5] A. Roy, H. Zeng, J. Bagga, G. Porter, and A. C. Snoeren, "Inside the social network's (datacenter) network," *SIGCOMM Comput. Commun. Rev.*, vol. 45, no. 4, pp. 123–137, Aug. 2015.
- [6] M. Ghobadi, R. Mahajan, A. Phanishayee, N. Devanur, J. Kulkarni, G. Ranade, P.-A. Blanche, H. Rastegarfar, M. Glick, and D. Kilper, "Projector: Agile reconfigurable data center interconnect," in *Proceedings of the 2016 ACM SIGCOMM Conference*. ACM, 2016, pp. 216–229.
- [7] M. Al-Fares, A. Loukissas, and A. Vahdat, "A scalable, commodity data center network architecture," *SIGCOMM Comput. Commun. Rev.*, vol. 38, no. 4, pp. 63–74, Aug. 2008.
- [8] A. Greenberg, J. R. Hamilton, N. Jain, S. Kandula, C. Kim, P. Lahiri, D. A. Maltz, P. Patel, and S. Sengupta, "V12: A scalable and flexible data center network," in *Proceedings of the ACM SIGCOMM 2009 Conference on Data Communication*, ser. SIGCOMM '09, 2009, pp. 51–62.
- [9] A. Celik, B. Shihada, and M.-S. Alouini, "Wireless data center networks: Advances, challenges, and opportunities," *arXiv preprint arXiv:1811.11717*, 2018. [Online]. Available: <https://arxiv.org/abs/1811.11717>
- [10] X. Zhou, Z. Zhang, Y. Zhu, Y. Li, S. Kumar, A. Vahdat, B. Y. Zhao, and H. Zheng, "Mirror mirror on the ceiling: Flexible wireless links for data centers," *ACM SIGCOMM Computer Communication Review*, vol. 42, no. 4, pp. 443–454, 2012.
- [11] A. S. Hamza, S. Yadav, S. Ketan, J. S. Deogun, and D. R. Alexander, "OWCell: Optical wireless cellular data center network architecture," *IEEE International Conference on Communications (ICC)*, pp. 1–6, May 2017.
- [12] A. S. Hamza, J. S. Deogun, and D. R. Alexander, "Classification framework for free space optical communication links and systems," *IEEE Communications Surveys Tutorials*, pp. 1–1, 2018.
- [13] E. Ciaramella, Y. Arimoto, G. Contestabile, M. Presi, A. D'Errico, V. Guarino, and M. Matsumoto, "1.28 terabit/s (32x40 gbit/s) wdm transmission system for free space optical communications," *IEEE J. Sel. Areas Commun.*, vol. 27, no. 9, pp. 1639–1645, Dec. 2009.
- [14] A. Celik, B. Shihada, and M.-S. Alouini, "Optical wireless data center networks: potentials, limitations, and prospects," in *Broadband Access Communication Technologies XIII*, vol. 10945, International Society for Optics and Photonics. SPIE, 2019, pp. 102 – 108.
- [15] A. S. Hamza, "Recent advances in the design of optical wireless data center networks," *SPIE OPTO (PW190)*, 2019.
- [16] H. Zhang, J. Zhang, W. Bai, K. Chen, and M. Chowdhury, "Resilient datacenter load balancing in the wild," in *Proceedings of the Conference of the ACM Special Interest Group on Data Communication*, ser. SIGCOMM '17, 2017, pp. 253–266.
- [17] A. Celik, A. Al-Ghadhban, B. Shihada, and M.-S. Alouini, "Design and provisioning of optical wireless data center networks: A traffic grooming approach," in *Proc. IEEE Wireless Commun. Netw. Conf. (WCNC)*, Apr. 2018, pp. 1–6.
- [18] A. Celik, A. Al-Ghadhban, B. Shihada, and M. Alouini, "Design and provision of traffic grooming for optical wireless data center networks," *IEEE Trans. Commun.*, vol. 67, no. 3, pp. 2245–2259, Mar. 2019.
- [19] N. Katta, M. Hira, A. Ghag, C. Kim, I. Keslassy, and J. Rexford, "Clove: How i learned to stop worrying about the core and love the edge," in *Proceedings of the 15th ACM Workshop on Hot Topics in Networks*, ser. HotNets '16, 2016, pp. 155–161.
- [20] sFlow ECMP load balancing. [Online]. Available: <http://blog.sflow.com/2014/03/performance-optimizing-hybrid-openflow.html>
- [21] S. Huang and R. Dutta, "Dynamic traffic grooming: the changing role of traffic grooming," *IEEE Commun. Surveys Tuts.*, vol. 9, no. 1, pp. 32–50, First 2007.
- [22] G. C. Sankaran and K. M. Sivalingam, "Scheduling in data center networks with optical traffic grooming," in *proc. IEEE 3rd Intl.Conf. Cloud Netw. (CloudNet)*, Oct. 2014, pp. 179–184.
- [23] —, "Optical traffic grooming-based data center networks: Node architecture and comparison," *IEEE J. Sel. Areas Commun.*, vol. 34, no. 5, pp. 1618–1630, May 2016.
- [24] A. Curtis, W. Kim, and P. Yalagandula, "Mahout: Low-overhead data-center traffic management using end-host-based elephant detection," in *INFOCOM, 2011 Proceedings IEEE*, April 2011, pp. 1629–1637.
- [25] V. Sivaraman, S. Narayana, O. Rottenstreich, S. Muthukrishnan, and J. Rexford, "Heavy-hitter detection entirely in the data plane," in *Proceedings of the Symposium on SDN Research*, ser. SOSR '17, 2017, pp. 164–176.
- [26] J. C. Mogul and P. Congdon, "Hey, you darned counters!: get off my asic!" in *Proceedings of the first workshop on Hot topics in software defined networks*. ACM, 2012, pp. 25–30.
- [27] M. Yu, L. Jose, and R. Miao, "Software defined traffic measurement with opensketch," in *NSDI*, vol. 13, 2013, pp. 29–42.

- [28] M. Al-Fares, S. Radhakrishnan, B. Raghavan, N. Huang, and A. Vahdat, "Hedera: Dynamic flow scheduling for data center networks," in *NSDI*, vol. 10, 2010, pp. 19–19.
- [29] J. Suh, T. T. Kwon, C. Dixon, W. Felter, and J. Carter, "Opensample: A low-latency, sampling-based measurement platform for commodity sdn," in *Proceedings of the 2014 IEEE 34th International Conference on Distributed Computing Systems*, ser. ICDCS '14, 2014, pp. 228–237.
- [30] J. Rasley, B. Stephens, C. Dixon, E. Rozner, W. Felter, K. Agarwal, J. Carter, and R. Fonseca, "Planck: Millisecond-scale monitoring and control for commodity networks," in *Proceedings of the 2014 ACM Conference on SIGCOMM*, ser. SIGCOMM '14, 2014, pp. 407–418.
- [31] Y. Zhu, N. Kang, J. Cao, A. Greenberg, G. Lu, R. Mahajan, D. Maltz, L. Yuan, M. Zhang, B. Y. Zhao *et al.*, "Packet-level telemetry in large datacenter networks," in *ACM SIGCOMM Computer Communication Review*, vol. 45, no. 4. ACM, 2015, pp. 479–491.
- [32] Y. Li, R. Miao, C. Kim, and M. Yu, "Flowradar: A better netflow for data centers," in *NSDI*, 2016, pp. 311–324.
- [33] C. Lin, C. Chen, J. Chang, and Y. H. Chu, "Elephant flow detection in datacenters using openflow-based hierarchical statistics pulling," in *2014 IEEE Global Communications Conference*, Dec 2014, pp. 2264–2269.
- [34] A. AlGhadhban and B. Shihada, "Flight: A fast and lightweight elephant-flow detection mechanism," in *2018 IEEE 38th Inter. Conf. on Dist. Comp. Syst. (ICDCS)*, July 2018, pp. 1537–1538.
- [35] C. Chaintoutis, B. Shariati, A. Bogris, P. V. Dijk, C. G. H. Roeloffzen, J. Bourderionnet, I. Tomkos, and D. Syvridis, "Free space intradecenter interconnects based on 2d optical beam steering enabled by photonic integrated circuits," *Photonics*, vol. 5, no. 3, 2018.
- [36] F. Zhang, D. Zhang, and S. Pan, "Fast and wide-range optical beam steering with ultralow side lobes by applying an optimized multi-circular optical phased array," *Appl. Opt.*, vol. 57, no. 18, pp. 4977–4984, Jun. 2018.
- [37] A. Lapidith, S. M. Moser, and M. A. Wigger, "On the capacity of free-space optical intensity channels," *IEEE Trans. Inf. Theory*, vol. 55, no. 10, pp. 4449–4461, Oct 2009.
- [38] A. Chaaban, Z. Rezeki, and M. S. Alouini, "Fundamental limits of parallel optical wireless channels: Capacity results and outage formulation," *IEEE Trans. Commun.*, vol. 65, no. 1, pp. 296–311, Jan. 2017.
- [39] K. Zhu and B. Mukherjee, "Traffic grooming in an optical wdm mesh network," *IEEE J. Sel. Areas Commun.*, vol. 20, no. 1, pp. 122–133, Jan 2002.
- [40] J. Y. Yen, "Finding the k shortest loopless paths in a network," *Management Science*, vol. 17, no. 11, pp. 712–716, 1971.
- [41] T. Truong-Huu, M. Gurusamy, and V. Girisagar, "Dynamic embedding of workflow requests for bandwidth efficiency in data centers," *Computer Networks*, vol. 108, pp. 184 – 198, 2016.
- [42] A. Dixit, B. Lannoo, D. Colle, M. Pickavet, and P. Demeester, "Wavelength switched hybrid tdma/wdm (twdm) pon: A flexible next-generation optical access solution," in *14th International Conference on Transparent Optical Networks (ICTON)*, July 2012, pp. 1–5.
- [43] A. S. Tanenbaum and D. Wetherall, *Computer networks*. Prentice hall, 2011.
- [44] R. Pang, M. Allman, M. Bennett, J. Lee, V. Paxson, and B. Tierney, "A first look at modern enterprise traffic," in *Proceedings of the 5th ACM SIGCOMM Conference on Internet Measurement*, ser. IMC '05, 2005, pp. 2–2.
- [45] T. T. Nguyen and G. Armitage, "A survey of techniques for internet traffic classification using machine learning," *IEEE communications surveys & tutorials*, vol. 10, no. 4, pp. 56–76, 2008.
- [46] A. Madhukar and C. Williamson, "A longitudinal study of p2p traffic classification," in *14th IEEE International Symposium on Modeling, Analysis, and Simulation*. Citeseer, 2006, pp. 179–188.
- [47] T. Karagiannis, K. Papagiannaki, and M. Faloutsos, "Blink: multilevel traffic classification in the dark," in *ACM SIGCOMM computer communication review*, vol. 35, no. 4. ACM, 2005, pp. 229–240.
- [48] B. Lantz, B. Heller, and N. McKeown, "A network in a laptop: Rapid prototyping for software-defined networks," in *Proceedings of the 9th ACM SIGCOMM Workshop on Hot Topics in Networks*, ser. Hotnets-IX, 2010, pp. 19:1–19:6.
- [49] POX. [Online]. Available: <http://www.noxrepo.org/pox/about-pox/>.
- [50] A. R. Curtis and *et. al.*, "Devoflow: Scaling flow management for high-performance networks," *ACM SIGCOMM Computer Communication Review*, vol. 41, no. 4, pp. 254–265, 2011.
- [51] M. Alizadeh, A. Greenberg, D. A. Maltz, J. Padhye, P. Patel, B. Prabhakar, S. Sengupta, and M. Sridharan, "Data center tcp (dctcp)," in *ACM SIGCOMM computer communication review*, vol. 40, no. 4. ACM, 2010, pp. 63–74.
- [52] psutil cross-platform process and system monitoring module for Python. [Online]. Available: <http://code.google.com/p/psutil/>.
- [53] C. Hopps, "Analysis of an equal-cost multi-path algorithm," *RFC 2992, Internet Engineering Task Force*, 2000.

Amer AlGhadhban (S'18-M'19) is an assistant professor with the Electrical Engineering Department in the College of Engineering at the University of Hai'i. He obtained B.S. and M. Sc. degrees in Computer Engineering from King Fahd University of Petroleum and Minerals (KFUPM) in (2012). He received his Ph.D. in Electrical Engineering from King Abdullah University of Science and Technology (KAUST) in (2019). He was a Cisco Academy instructor and during that he earned multiple professional certificates: Cisco Certified Network Associate (CCNA), Cisco wireless, SANS- GIAC Certified Firewall Analyst (GCFW-Gold), Certified Ethical Hacker (CEH), Security Certified Network Engineer (SCNP), and SANS local mentor. His research interests span over multiple fields in networked systems and security.

Abdulkadir Celik (S'14-M'16-SM'19) received a B.S. degree in Electrical-Electronics Engineering from Selcuk University, Konya, Turkey in 2009. He received M.S. degrees in Electrical Engineering in 2013 and in Computer Engineering in 2015, and a Ph.D. degree in co-majors of Electrical Engineering and Computer Engineering in 2016 from Iowa State University, Ames, IA, USA. He is currently a senior post-doctoral fellow in Communication Theory Laboratory at KAUST. His research interests include of underwater things, and flying networks.

Basem Shihada (SM'12) is an associate & founding professor in the Computer, Electrical and Mathematical Sciences & Engineering (CEMSE) Division at King Abdullah University of Science and Technology (KAUST). He obtained his PhD in Computer Science from University of Waterloo. In 2009, he was appointed as visiting faculty in the Department of Computer Science, Stanford University. In 2012, he was elevated to the rank of Senior Member of IEEE. His current research covers a range of topics in energy and resource allocation in wired and wireless networks, software defined networking, internet of things, data networks, network security, and cloud/fog computing.

Mohamed-Slim Alouini (S'94-M'98-SM'03-F'09) was born in Tunis, Tunisia. He received the Ph.D. degree in Electrical Engineering from the California Institute of Technology (Caltech), Pasadena, CA, USA, in 1998. He served as a faculty member in the University of Minnesota, Minneapolis, MN, USA, then in the Texas A&M University at Qatar, Education City, Doha, Qatar before joining King Abdullah University of Science and Technology (KAUST), Thuwal, Makkah Province, Saudi Arabia as a Professor of Electrical Engineering in 2009. His current research interests include the modeling, design, and performance analysis of wireless communication systems.