Leveraging Optics for Network Function Virtualization in Hybrid Data Centers

Tamal Das*, Aniruddha Kushwaha*, Ashwin Gumaste* and Mohan Gurusamy#
* Indian Institute of Technology, Bombay, India
# National University of Singapore, Singapore
{tamaldas, aniruddha}@cse.iitb.ac.in, ashwing@ieee.org, gmohan@nus.edu.sg

***Abstract*—Network Function Virtualization (NFV) has emerged as a hot topic for both industry and academia. NFV offers a radically new way to design and operate networks, by abstracting *physical network functions* (PNFs) to *virtual NFs* (VNFs). This disruptive innovation opens up a wide area of research, as well as introduces new challenges and opportunities – particularly in provisioning VNF forwarding graphs (or *network service chains*), and the resulting VNF placement issue. While forwarding graphs are often provisioned in the packet domain for fine-grained control over the respective traffic, we argue that doing so leads to lower efficiency; instead, provisioning forwarding graphs using optical transport proves to be far more efficient in intra-*datacenter* (DC) scenarios. While optical service chaining for NFV has already been proposed, we emphasize the use of optical bus architectures for the same. We present an architecture conducive for intra-DC NFV orchestration that can easily be extended to inter-DC scenarios. We deploy switchless optical bus architectures in both the frontplane and backplane of the DC. Our design particularly relies on readily available optical components, and scales easily. We validate our model using extensive simulations. Our results suggest that use of optical transport to provision VNF forwarding graphs can result in significant performance enhancement over packet-based electrical switch provisioning, in terms of packet drops and latency.**

I. INTRODUCTION

Network Function Virtualization (NFV) along with Software Defined networking (SDN) is considered as the game changers for next generation carrier-networks. While SDN will make the forwarding plane programmable, reduce the cost by including whiteboxes and bring generic agility into the network, NFV will allow the use of virtualization technologies to make complex network functions that existed in hardware to be placed in software. NFV, in some sense, facilitates the commoditization of networks by breaking service chains into network functions that can further be implemented on standard COTS platforms – IT-grade servers. The impact of virtualization is immense – NFV reduces CapEx and OpEx and facilitates better service velocity in terms of provisioning, upgrading, enabling elasticity to service chains. This promise of extreme cost-effectiveness and ability to softwarize the network is what has made NFV a popular research direction, not just in the academic community but also with providers – as evidenced by the ETSI initiative [1]. The NFV framework is undergoing severe consideration across vendors, providers, developers and academia. From a service provider standpoint, the question remains as to which are the best parts of a network to inculcate NFV. Given that at its core, the smallest indivisible part of NFV is the *virtual network function* or VNF – that exists as a standalone software chunk that can be moved around on VMs – the best position for placing a VNF is then the service provider data-center (DC). The thought of placing VNFs in provider data-centers is not new – it was first explored by the CORD project [2], where VNFs are placed in mini-data-centers that are formed by replacing traditional Central Office architecture with a bunch of servers and corresponding switches. While putting VNFs in the CO is a good idea for minimizing equipment churn towards the edge of network, another stress point is at the core of the network, where there is sizable need for network functions as well as storage of data. This is the reason why providers have data-centers in the core of the network, from where they can launch service chains as well as store data. Such a situation warrants that NFV technologies coexist with conventional data-technologies and, moreover, such an arrangement be integrated with the rest of the provider network.

The data-center, hence, becomes a key position in the network where we want to store, process, transport, work-upon data chunks. An ideal backbone data-center must be able to support huge amounts of data and network functions. Scalability is, hence, a key virtue that a DC must possess. Significant amount of research is available on DC design from a pure scalability perspective. We, in this paper, consider DC design from the perspective of both scalability and NFV provisioning. We require a DC to be able to scale to a large number of nodes that support both data storage as well as VNF storage.

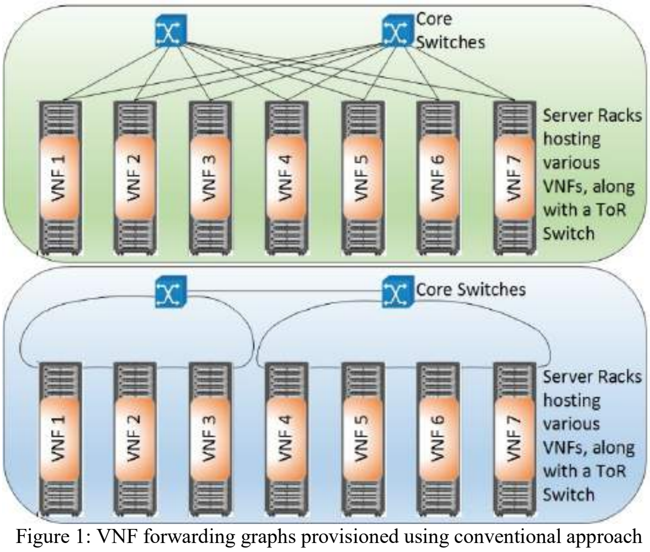

Figure 1: VNF forwarding graphs provisioned using conventional approach (top) and our approach (bottom).

978-3-903176-07-2 © 2018 IFIP

To this end, we proposed in [3] the DOSE architecture that facilitates the creation of a million node DC using optics in both its front plane and backplane. While HELIOS, OSA, WaveCube, FISSION, etc. [4-7] do use optics as an interconnection paradigm between top of the rack (TOR) switches, we go one step further – we use optics in the front plane as well; i.e. to connect servers to each other. Our concept assumes contemporary optics – that is easily available and mature. We do not rely on fast-moving optical devices, instead deploy an interesting architecture that is primarily based on the concept of broadcast and select of data across multiple fiber rings and wavelengths. The resulting DC is then ideally suited to house VNFs in terms of scalability, responsiveness, growth of services and churn in the network.

Our DOSE DC consists of sectors as a fundamental communication unit. A sector could have one-or-more TOR switches that are connected to servers. Sectors are interconnected using fiber rings. We could have as many fiber rings as one would want – subject to only port size of the cascade of sector-fiber switches and OSNR limitations. A sector is allocated a batch of wavelengths on which it perpetually can transmit into any one of the fiber rings. While a sector transmits into just one of the rings – it can receive data from all of the rings. The ring-to-sector interconnection is passive, implying a bus formation that is formed by the use of power-splitter, a combiner (coupler) as the interconnection element between the fiber ring and the sector. In addition to the use of optics in the backplane to connect sectors, we also use optics in the front plane to connect servers. Servers to communicate within a sector may use the TOR electronic switch or may use all-optical wavelength buses for communication. Fig. 1 illustrates the difference between provisioning a network service chain using the conventional spine-and-leaf architecture (top) vs. our proposed bus architecture (bottom). While the service chain needs to visit the core switch between any two VNFs, the same is not true for the bus architecture. The use of front plane optics is a way of saving on wiring as well as reducing latency for communication.

This paper is organized as follows. Section II discusses the related literature, while section III details our DC architecture. We present our simulation framework and results in section IV. Section V contains some concluding remarks.

II. RELATED WORK

In this section, we present some of the related literature in the context of this paper, and highlight our contributions.

Several approaches have been reported in the literature, addressing the VNF placement and deployment problem in the context of NFV. An instrumentation and analytics framework is presented in [8], which shows that use of embedded instrumentation provide opportunities for providers to fine-tune their NFV deployments from both the technical and economic perspectives. A micro-service-based NFV orchestrator TeNOR is presented in [9] that focuses on: (i) automated deployment and configuration of services composed of virtualized functions, and, (ii) management and optimization of networking and IT resources for VNF hosting. Authors in [10] proposed forecast-assisted service chain deployment algorithm that includes the prediction of future VNF requirements. Possibility of minimizing the expensive optical/electronic/optical conversions for NFV chaining in packet/optical datacenters by using on-demand placement of vNFs is identified in [11]. A model for NFV placement is presented in [12] which considers the utilization of links and servers to minimize the maximum utilization over all links and switches. [13] proposed a hybrid architecture (optical/electrical) suited for NFV.

We now summarize some of the leading DC architectures. Several data center architectures have been proposed in recent years. The fat-tree [14] data center architecture proposed a hierarchy of three layers of electrical switches – core, aggregate and edge switches and is the most commonly deployed variant. In the DCell [15] architecture, a server is interconnected with other servers as well as a mini-switch. Servers communicate either through their connection to the mini-switch or through their connections to other servers. In c-Through [16], optical paths between top-of-the-rack (ToRs) switches are shared based on inter-rack traffic demands, while, ToRs are also interconnected with dedicated packet-switched paths. Helios [4] also uses a topology manager to measure traffic and estimate demand of the ToRs, based on which it computes the optimal topology for circuit-switched paths. Architectures like OSA [5], WaveCube [6] use reconfigurable optical devices to create optical circuits at runtime. Reconfiguration delay for these optical devices is a bottleneck. Delay in order of 10ms is too high for latency-intensive or smaller granularity flows.

In FISSION [7], optical backplane consists of number of fiber rings which are divided into sectors. Each sector can receive from all the fiber rings but can transmit only to a single ring. Each sector consists of servers inter-connected using switches in clos architecture. DOSE [3, 17] is an extension of the FISSION, where fiber rings are used to interconnect servers within the sectors, thus leading to both optical backplane as well as optical frontplanes.

Our work is inspired from the DOSE architecture [3], which we apply in this paper in the context of provisioning VNF forwarding graphs. While we largely restrict our discussion in this paper to intra-DC scenarios, it is important to note that our approach can easily be extended to inter-DC scenarios as well.

III. SYSTEM ARCHITECTURE

In this section, we detail our datacenter architecture (see Fig. 2) where we deploy the optics in both frontplane and backplane to dynamically provision VNF forwarding graphs of services.

The fiber ring-based DC optical backplane comprises one or more fiber rings. Fiber based backplane support optical buses in a ring configuration. Multiple number of sectors are connected in each fiber ring. Each sector in a fiber is allocated a fix set of wavelengths to transmit in a specific fiber ring. At the receiver of a sector, each ring drops a composite WDM signal constitutes of all the wavelengths of the fiber. This configuration allows a sector to transmit in a single fiber while allowing a sector to receive from all the fibers.

Each sector consists of wavelength selective switches (WSSs) to split the composite WDM optical signal from the backplane fibers to its constituent wavelengths. WSSs are preconfigured to drop only select wavelengths to each sector (being restricted by the EOS port count). These dropped wavelengths are processed by Electro-Optical Switch (EOS) for the service requests. Based on the request of service chain rule and load, EOS forwards the service to one of the k frontplane fiber rings, which are at the k ports of the EOS. Each frontplane fiber ring has m interconnection points, which are the interface points for the ToRs. Each frontplane deploys a unidirectional wavelength bus shared amongst all racks. An interconnection point consists of two couplers (one each for adding/dropping wavelength to/from frontplane) separated by an optical switch as shown in Fig. 2. Being an optical bus-based frontplane, when a ToR/EOS transmits in it, all the downstream ToRs/EOS receive the data. In this scenario all the unintended recipient discards the data based on an electronic match at its receiver. On arriving at a server, the packet is processed and forwarded via the frontplane to the next VNF based on the service's forwarding graph. A rack may host one or more VNFs. Based on VNF forwarding graph, if a service is for the EOS, it is forwarded to the frontplane, and after all of a service's VNFs in that sector are processed, it is thereafter sent to the backplane.

There are different wavelength assignment schemes for the ToR switches in the frontplane fiber ring, based on the number of wavelengths:

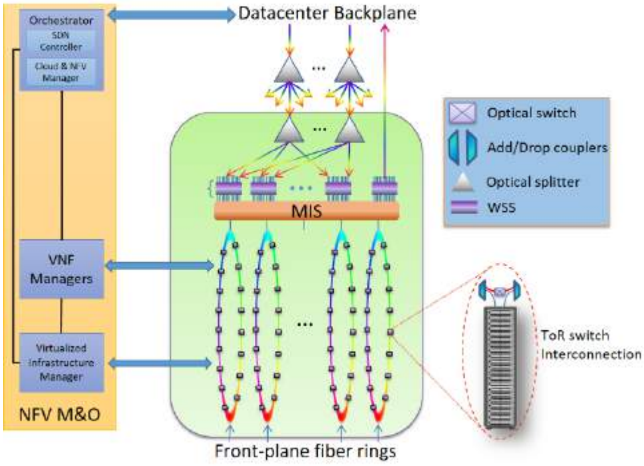

Figure 2: Proposed DC Frontplane architecture.

(a) Single Wavelength: In this wavelength assignment scheme, the frontplane fiber ring has only a single wavelength, where all ToRs are allocated with same wavelength to send and receive traffic. This single wavelength is time-shared between all ToRs and a token-based grant is used to avoid any simultaneous transmission of two or more ToRs/EOS. The arbitration is done using an out-of-band control channel ahead of time. Using a dedicated control channel for token allocation makes the architecture simple and helps in efficiently utilizing the data channel. This scheme also helps in reducing the load on the EOS as number of times a packet visit EOS will be less than the cardinality of its VNF forwarding graph.

(b) Multiple Wavelengths: This wavelength assignment scheme can be further divided into two sub-schemes:

Number of wavelengths = Number of racks: In this sub-scheme, there are a total of m wavelengths in a frontplane, and consequently each ToR is assigned a dedicated wavelength to send and receive traffic. In this case, if a VNF resides on some ToR, then the EOS will use a dedicated wavelength to send the packet to the respective ToR, and after processing the packet, the ToR will forward the packet to EOS for its next VNF processing. Since, each ToR has a dedicated wavelength, this scheme does not require any out-of-band control channel. But, the load on the EOS also increases as the service needs to visit the EOS after each VNF processing.

Number of wavelengths < Number of racks: In this sub-scheme, there are $< m$ wavelengths in a frontplane, which are time-shared to send/receive traffic. Because of time-sharing, each ToR is equipped with a tunable laser. Similar to the single wavelength scheme, an out-of-band control channel is used to arbitrate the pool of wavelengths between the ToRs and the EOS. This scheme also helps in reducing the load on the EOS as the ToR belonging to next service chain rule in the downstream can be directly reached without visiting the EOS.

Our architecture assumes an SDN-based central controller for service provisioning, which interfaces with the VNF manager for the instantiation and management of the VNFs on servers (see Fig. 2). The SDN controller populates the service chain rules and gather statistics to/from the EOS and ToRs. It shares the flow statistics and new service request information with the VNF manager. Based on the load and service request, VNF manager instantiates the VNFs on the server and shares this information with SDN controller for service provisioning and resource management.

IV. Optimization Model for Backplane Wavelength Assignement

In this section, we formulate an optimization model to deduce the backplane wavelength assignment. The goal is to connect most pair of sectors across the backplane using the minimum wavelengths. Our list of parameters and decision variables part are listed in Table 1 and Table 2, respectively.

Table 1: List of Parameters

Parameter	Meaning
$\boldsymbol{W_i}$	Wavelength of type i
$\boldsymbol{F_j}$	Fiber j
$\boldsymbol{S_k}$	Sector k
$\boldsymbol{\alpha}$	Wavelength Multiplicity in the backplane
$\boldsymbol{\beta}$	Contention factor at a sector's drop ports
$\boldsymbol{W}$	Number of wavelengths per backplane fiber
$\boldsymbol{n}$	Number of sectors per backplane fiber
$\boldsymbol{\gamma_p}$	Number of add ports at sector S_p

Table 2: List of Decision Variables

Variable	Meaning
λ_{ij}^{pq}	Wavelength of type W_i in fiber F_j from sector S_p to sector S_q.

The objective of our optimization model is to ensure maximum connectivity between every pair of sectors across the backplane, i.e.,

$$\max \sum_{i,j,p,q(\neq p)} \lambda_{ij}^{pq}$$

subject to the following constraints.

Each wavelength can be used at most once across a sector's add ports. For instance, a sector cannot transmit on the same wavelength on different fibers in the backplane. Moreover, a wavelength added by an ingress sector in the backplane may be dropped at a maximum of α sectors. Thus, each of the λ_{ij}^{pq} originating from ingress sector S_p on wavelength W_i can connect to atmost α egress sectors, i.e.,

$$\forall i, p, \sum_{j,q(\neq p)} \lambda_{ij}^{pq} \leq \alpha$$

An ingress sector can only transmit on wavelengths from a single fiber in the backplane. This results from the physical constraint that a sector's ADD WSS can be connected to only a single fiber, and consequently a sector's add wavelengths cannot be added across multiple fibers in the backplane. Thus, for a given ingress sector, there exists a unique fiber in the backplane on which it transmits, i.e.,

$$\forall p, \exists! j: \sum_{i,q(\neq p)} \lambda_{ij}^{pq} > 0$$

Such a uniqueness constraint can be handled by LP solvers using a Special Ordered Set (SOS) of type One.

Each wavelength in the backplane can be used by atmost one sector. This eliminates the case of multiple sectors transmitting on the same wavelength in the same backplane fiber. Thus, there exists a unique sector which transmits on a particular wavelength in a backplane fiber, i.e.,

$$\forall i, j, \exists! p: \sum_{q(\neq p)} \lambda_{ij}^{pq} \neq 0$$

Such a uniqueness constraint can be handled by LP solvers using a Special Ordered Set (SOS) of type One.

The number of distinct wavelengths added from a sector is bounded by the number of add ports at the sector. Let us first define few auxiliary variables to formulate this constraint.

$$\forall i, p: d_1^{ip} = \sum_{j,q(\neq p)} \lambda_{ij}^{pq} \text{ and } d_2^{ip} = \begin{cases} 0, \text{if } d_1^{ip} = 0 \\ 1, \text{otherwise} \end{cases}$$

Here, d_1^{ip} denotes the cardinality of the set of egress sectors receiving from sector S_p on wavelength W_i via the backplane, whereas d_2^{ip} is a binary variable which determines whether wavelength W_i is used by sector S_p to transmit in the backplane. The stated constraint can then be formulated in terms of these auxiliary variables as:

$$\forall p: \sum_{i} d_2^{ip} \leq \gamma_p$$

Each backplane fiber has W wavelengths, each of which can potentially be dropped at α sectors. Thus, a fiber can drop at up to αW port across n sectors, i.e. $\frac{\alpha W}{n}$ ports per sector. In addition, an egress sector is configured to receive at most β wavelengths of the same type from the backplane, i.e. it can receive up to $\frac{\alpha W}{n\beta}$ distinct wavelengths. Thus,

$$\forall q, j, \sum_{i,q(\neq p)} \lambda_{ij}^{pq} \leq \frac{\alpha W}{n\beta}$$

Given the contention factor, at most β backplane fibers can drop the same wavelength at an egress sector, i.e.,

$$\forall q, i, \sum_{j,p(\neq q)} \lambda_{ij}^{pq} \leq \beta$$

Each wavelength in a backplane fiber is dropped at upto α sectors.

$$\forall i, j, \sum_{p,q(\neq p)} \lambda_{ij}^{pq} \leq \alpha$$

Each sector has at least one drop port and at least one add port connected to the backplane.

$$\forall s: \sum_{i,j,p(\neq s)} \lambda_{ij}^{ps} \geq 1 \text{ and } \sum_{i,j,q(\neq s)} \lambda_{ij}^{sq} \geq 1$$

Atmost one wavelength connects every pair of sectors in the backplane.

$$\forall p, q(\neq p): \sum_{i,j} \lambda_{ij}^{pq} \leq 1$$

To compute the backplane wavelength assignment for a million server DOSE datacenter, the above formulation takes ~20 seconds on an Intel Quadcore i7 CPU@3.5GHz with 16GB RAM.

V. SIMULATION AND RESULTS

In this section, we evaluate our proposed DC design using a Python-based discrete event simulation, and discuss the observed results. We simulate a DC with the well-known fat-tree architecture [13] and compare its performance (primarily in terms of metrics such as latency and packet drops) with that of our proposed DC architecture.

Simulation Model: For a given number of servers, we generate the corresponding fat-tree DC. We assume one of the PODs interface with the Datacenter Inteconnection Point (DCIP), and is thus the source/sink of all DC traffic. Although edge, aggregate and core switches in a fat-tree DC are considered the same, for the sake of comparison, we consider all server-edge switch links at 1Gbps, all edge switch-aggregate switch links at 10Gbps and all aggregate switch-core switch links at 100Gbps. As traffic enters the DC via the DCIP, it visits various servers in succession depending on its VNF forwarding graphs, and on completion, exits the DC through the DCIP. We assume each server to host a single VNF. In the rest of this section, we refer to this case as the "FatTree" scenario.

Similarly, we also generate a DC network with our proposed architecture, comprising of sectors, each of which hosts a bunch of frontplanes, which in turn consist of a bunch of racks, while the sectors are interconnected via backplanes. Here too, we assume one sector to interface with the Internet (via DCIP), and is thus the source/sink of all DC traffic. We assume each frontplane to host all VNFs, one per server rack. As traffic enters the DC via DCIP, it visits the least-loaded frontplane in the DC and on completion, exits the DC via the DCIP. An EOS has three port types, namely, (a) backplane ports (to receive traffic from the backplane), (b) add ports (to send traffic to the backplane, and, (c) frontplane ports (each hosting a unique frontplane). In the backplane, each wavelength drops traffic at two sectors (we term this a "*wavelength multiplicity*" of 2). If two sectors are not directly connected (i.e. via a single-hop) via a backplane wavelength, we consider multi-hopping routing to route traffic between them. The server-ToR switch links are assumed at 1Gbps, the frontplane rings are assumed at 10Gbps, and the backplane rings are assumed to be at 100Gbps. In the rest of this section, we refer to this case as the "DOSE" scenario.

In both scenarios, to service a particular VNF requirement of a network service chain, of the many servers hosting the

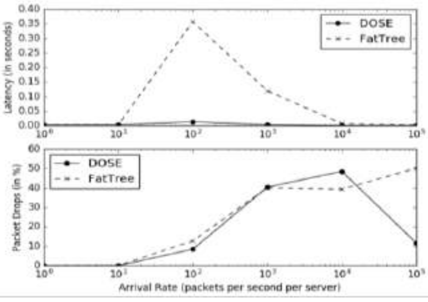

Fig. 3. Effect of Load on DOSE and FatTree DCs

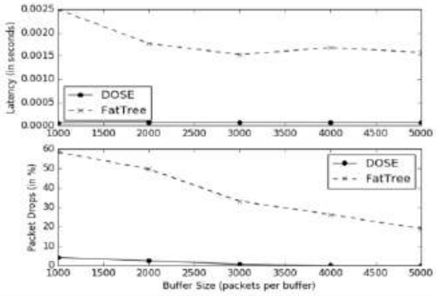

Fig. 4. Effect of buffer size on DOSE and FatTree DCs

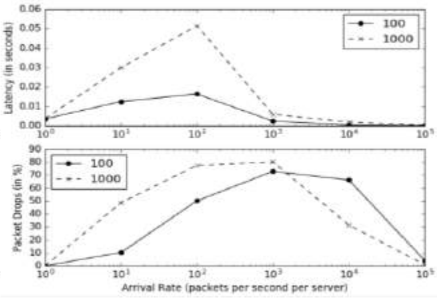

Fig. 5. Effect of network size on DOSE DC

required VNF, we choose the one with the least loaded path. Every switch/server port is assumed to have two buffers – one each for sending and receiving. All ports are assumed to be bidirectional. For a fair comparison between the two DC networks, we ensure the same traffic footprint in both scenarios. Each DC network is generated with 100 servers. Although for a given number of servers a unique topology is possible for the fat-tree architecture, the same is not true for our proposed architecture. In this evaluation, we consider an arbitrary topology for the DOSE DC, while deriving an optimal topology remains our future work. For a given number of VNFs, we generate the same number of services, each with a random service chain of varying lengths. All services are assumed to have the same priority level. Packet sizes are generated from an exponential distribution with a mean of 250 bytes, while packet arrivals are assumed to be Poisson distributed. A number of packet generators are placed on the Internet-facing side of the DCIPs to pump traffic composed of various services into the DC. A lookup delay of 300 nanoseconds and an average processing latency of 200 microseconds per packet is assumed. The port buffer capacities are considered proportional to their port rates, starting at 256kB for 1Gbps (~1000 packets per buffer) and so on. Leveraging the optical bus-based backplane in the DOSE DC, we consider each wavelength to be dropped to 2 sectors. We term this as "*wavelength multiplicity*". A wavelength multiplicity of 1 would imply a point-to-point connection (lightpath). To eliminate statistical errors, all results are averaged over 5 distinct traffic patterns.

Effect of Load: Fig. 3 contrasts the effect of load on the two DC architectures in terms of average end-to-end latency (in seconds, top) and average packet drops (in %, bottom). Both metrics are computed across all services. We vary the packet arrival rate per server from 1 to 100,000 packets per second. Note that the abscissa is plotted in log-scale.

The latency gradually increases from low to medium loads and drops thereafter. The decrease in latency from medium to high loads might seem counter-intuitive at first, but can be reconciled when observed in sync with the corresponding packet drops. At medium to high loads, packet drops significantly increase, and consequently lesser service chains are fully served. As a result, packets are either promptly dropped (resulting in higher packet drops), or promptly served (resulting in lower latencies). The benefit of DOSE over FatTree architecture is most pronounced at medium loads. Thus, in terms of latency, both packet and optical scenarios perform similarly at low and high loads, while benefit of optical backplane and frontplane is most pronounced at medium loads.

The average packet drops (or blocking probability) increases from low to high loads for both DC architectures, though the difference between the two is not much pronounced. In conclusion, while optics help bring down the latency, it does not improve the blocking probability as much.

Effect of Buffer Size: Fig. 4 plots the impact of buffer size on the two DC architectures. We vary the buffer size from a 1000 packets to 5000 packets in steps of 1000, and note the observed effect on the two performance metrics. These plots consider an arrival rate of 100,000 packets/second per server.

Both latency and packet drops decrease with rise in buffer size, the former only slightly while the latter considerably. This can be explained as follows. Larger the buffer, more packets can be stored, resulting in an increase in the observed end-to-end latency. An increase in buffer size essentially means more packets are accommodated, and in turn lesser packets dropped. The impact of buffer size is more pronounced for fat-tree architecture than for DOSE. This is attributed to the fact that the scope for betterment in latency/packet is rather low in case of a DOSE DC.

Effect of Network Size: Fig. 5 plots the effect of network size for a DOSE DC. The optical bus architecture employed in the DOSE DC significantly reduce the simulation run time, as compared to the fat-tree architecture; so much so that simulating a fat-tree network over servers becomes infeasible. Hence, the effect of network size could only be studied for the DOSE DC. We consider a 100 and 1000 node DOSE DC network, and vary the packet arrival rate per server from 1 to 100,000 packets per second, and plot the observed affect. Note that the abscissa is plotted in log-scale.

A larger topology leads to longer paths resulting in higher latencies. At low loads, larger topologies seem to lower the packet drops due to the larger cumulative buffer capacity across the network, although the buffer per server/switch remains the same. However, at medium to high loads, the packet drops tend to increase with larger topologies.

Effect of Service Chain Length: Fig. 6 plots the effect of service chain length on the two DC architectures. We vary the service chain length from 1 to 10, and note the observed effect on the two performance metrics. These plots consider an arrival rate of 1,000 packets/second per server. We generated a mix of 10 services each with a service chain length from 1 to 10.

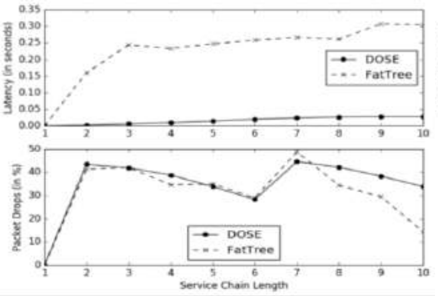

Fig. 6. Effect of service chain length on DOSE and FatTree DCs

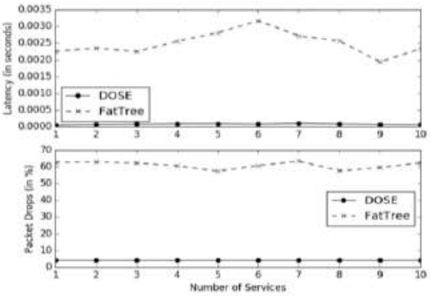

Fig. 7. Effect of number of services on DOSE and FatTree DCs

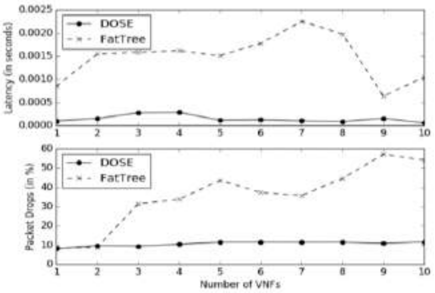

Fig. 8. Effect of number of VNFs on DOSE and FatTree DCs

With increasing service chain length, the latency increases in both the DC architectures, as longer service chains lead to longer service latencies. However, the latency for the DOSE DC is significantly and consistently less than that of a FatTree DC. With increasing service chain length, the packet drops largely increases, with no tangible improvement offered by a DOSE DC over a FatTree DC.

Effect of Number of Services: Fig. 7 plots the effect of varying number of services provisioned using the two DC architectures. We vary the number of services from 1 to 10, and plot the observed effect on the two-performance metrics. These plots consider an arrival rate of 100,000 packets/second per server.

The latency as well as the packet drops remain largely agnostic to the number of services in a DOSE DC, while a fat-tree DC seems to be slightly impacted. Thus, DOSE significantly outperforms fat-tree over a wide range of services.

Effect of Number of VNFs: Fig. 8 plots the effect of varying number of VNFs in the DC considering both architectures. We vary the number of VNFs from 1 to 10, and for a given number of VNFs, we generate as many services with varied service chain lengths. These plots consider an arrival rate of 100,000 packets/second per server.

Growing number of VNFs increases both the latency as well as the packet drops in case of a fat-tree DC, while a DOSE DC is hardly affected by the same. The performance of a DOSE DC is again better than that of a fat-tree DC across a varied range of VNFs.

VI. CONCLUSION

In this paper, we proposed a novel approach to provision network service chains for intra-DC scenarios. Our architecture heavily relies on optics, and deploys a switchless optical bus design both in the frontplane as well as in the backplane. Compared to the case with packet-based provisioning of network service chains, our architecture offers higher bandwidth due to use of optical fibers, as well as traffic-agnostic, as only the ports and not the links needs to be upgraded from time to time, unlike the packet-based scenario. We validated our model using extensive simulations, and compared our design with packet-based provisioning in terms of relevant metrics such as packet drops, latency as well as effect of design parameters such as buffer size, service chain length, topology size, etc. We observe that optics can play a significant role in improving the provisioning the VNF forwarding graphs for NFV.

REFERENCES

[1] ETSI NFVISG Whitepaper, “Network Functions Virtualization (NFV) – Network Operator Perspectives on Industry Progress,” SDN and OpenFlow World Congress 2013, Frankfurt, Germany.

[2] CORD Project [Online]. Available: https://opencord.org

[3] Gumaste, A. Kushwaha and T. Das, “DOSE: Double optics single electronics data-center using a switchless optical frontplane and backplane,” IEEE International Conference on Communications (ICC), 2016, Malaysia.

[4] N. Farrington et al., "Helios: a hybrid electrical/optical switch architecture for modular data centers." ACM SIGCOMM Computer Communication Review, Vol. 40, No. 4, pp. 339-350, 2010.

[5] K. Chen et al., “OSA: an optical switching architecture for data center networks with unprecedented flexibility,” IEEE/ACM Transactions on Networking, Vol. 22, No. 2, pp. 498-511, 2014.

[6] K. Chen *et al.*, “WaveCube: A scalable, fault-tolerant, high performance optical datacenter architecture,” in *Proc. IEEE Conf. Comput. Commun.*, Hong Kong, 2015, pp. 1903–1911.

[7] A. Gumaste et al., “On the Unprecedented Scalability of the FISSION (Flexible Interconnection of Scalable Systems Integrated using Optical Networks) Datacenter,” To appear in IEEE Journal of Lightwave Technology, 2016.

[8] P. Veitch, M. J. McGrath and V. Bayon, "An instrumentation and analytics framework for optimal and robust NFV deployment," in IEEE Communications Magazine, vol. 53, no. 2, pp. 126-133, Feb. 2015.

[9] J. F. Riera et al., "TeNOR: Steps towards an orchestration platform for multi-PoP NFV deployment," 2016 IEEE NetSoft Conference and Workshops (NetSoft), pp. 243-250, Seoul, 2016,

[10] Sun, Quanying, et al. "Forecast-Assisted NFV Service Chain Deployment based on Affiliation-Aware vNF Placement."

[11] Xia, Ming, et al. "Network function placement for NFV chaining in packet/optical datacenters." Journal of Lightwave Technology, vol 33, no. 8, pp 1565-1570, 2015.

[12] A. Mohammadkhan et al., “Virtual function placement and traffic steering in flexible and dynamic software defined networks,” ser. LANMAN 2015. IEEE, pp. 1–6, 2015.

[13] Xia, Ming, et al. "Optical service chaining for network function virtualization." IEEE Communications Magazine 53.4 (2015): 152-158.

[14] M. Al-Fares, A. Loukissas and A. Vahdat, “A scalable, commodity data center network architecture,” ACM SIGCOMM Computer Communication Review, Vol. 38, No. 4, 2008.

[15] C. Guo et al., “Dcell: a scalable and fault-tolerant network structure for data centers,” ACM SIGCOMM Computer Communication Review, Vol. 38, No. 4, 2008.

[16] G. Wang et al., “c-Through: Part-time optics in data centers,” ACM SIGCOMM Computer Communication Review, Vol. 40, No. 4, 2010.

[17] A. Kushwaha, T. Das and A. Gumaste, “Does it Make Sense to put Optics in Both the Front and Backplane of a Large Data-Center?,” Optical Fiber Communication Conference (OFC), Los Angeles, California, March 2017.