Predictive Analytics in Hybrid Optical/Electrical DC Networks

(Invited Paper)

Zuqing Zhu, Wei Lu, Lipei Liang, Bingxin Kong

University of Science and Technology of China, Hefei, Anhui 230027, China, Email: zqzhu@ieee.org

Abstract: We explain how to leverage knowledge-defined networking (KDN) to realize automatic network control and management (NC&M) for orchestrating the IT and bandwidth resources in a hybrid optical/electrical datacenter network (HOE-DCN). **OCIS codes:** (060.1155) Software-defined optical networks; (060.4251) Networks, assignment and routing algorithms.

1. Introduction

Due to the raising of cloud computing and services, global datacenter (DC) traffic has been growing at a compound annual rate of 25% since 2016 and would reach 20.6 Zettabytes (1 Zettabytes = 10^{21} bytes) in 2021, of which 71% is within DCs [1]. Consequently, to adapt to such tremendous traffic growth, DC operators are facing a few difficult challenges when designing and operating their DC networks (DCNs) [2]. First of all, from the perspective of capital expenditure (CAPEX), a DC operator has to properly balance the tradeoff between its investment on network equipment/bandwidth and the margin left for service upgrades. Secondly, the operational expenditure (OPEX) of a DCN (*e.g.*, energy consumption) should be carefully considered too. For instance, when undergoing capacity upgrades, the power used in a DC that interconnects servers with hierarchical electronic switches would increase rapidly [3, 4]. Thirdly, it is crucial for a DCN to provide differentiated services to applications with various quality-of-service (QoS) requirements, while new applications with stringent QoS demands on bandwidth, latency, jitter, *etc.*, are emerging with an unprecedented rate [5]. Lastly but most importantly, the automation and agility of the DCN's network control and management (NC&M) requires continuous efforts from the academia and industry. Although software-defined networking (SDN) offers new opportunities to architect more programmable, adaptive and reliable NC&M for DCNs [6], how to dynamically orchestrate IT and bandwidth resources in a DCN effectively and respond timely to network status changes for ensuring various QoS requirements from applications still needs further exploration.

In this work, we explain our attempt to address the aforementioned challenges with the combination of hybrid optical/electrical DCNs (HOE-DCNs) and knowledge-defined networking (KDN). As shown in Fig. 1(a), an HOE-DCN tries to take the advantages of both electronic packet switching (EPS) and optical circuit switching (OCS) [3]. Specifically, in the HOE-DCN, the top-of-rack (ToR) switches are interconnected with both high-speed Ethernet switches and optical cross-connects (OXCs). Therefore, for the applications that require short setup latency and/or generate highly dynamic traffic, their inter-rack traffic can be forwarded with EPS, while the inter-rack traffic from bandwidthintensive applications that last relatively long can occupy a direct lightpath through the OXC(s). More importantly, OCS is much more energy efficient than EPS, and thus HOE-DCNs can not only handle various QoS requirements but also achieve energy saving. Although HOE-DCNs are promising due to the reasons mentioned above, their advantages can never be explored without an effective NC&M scheme. This is exactly the reason why we consider to leverage KDN [7, 8], which is essentially a combination of SDN and artificial intelligence (AI) as shown in Fig. 1(b). Specifically, with the centralized NC&M provided by SDN, a DC operator can obtain a precise global view on its DCN by collecting rich telemetry data proactively, and then, it can leverage AI-assisted data analytics to abstract knowledge from the data through deep learning (DL) and use the knowledge to reach intelligent decisions for automatic NC&M.



Fig. 1. (a) Architecture of the HOE-DCN considered in this work, ToR switch: top-of-rack switch, OXC: optical crossconnect, KD-O: knowledge-defined orchestrator, IT-C: IT resource controller, NET-C: network controller, (b) Principle of knowledge-defined networking (KDN), and (c) Control plane design for the HOE-DCN, NSD: network service database, TED: traffic engineering database, NAM: network abstraction module.

Even though the idea of KDN is quite straightforward, how to incorporate it in the NC&M of HOE-DCNs is still an open and challenging problem. The reasons are twofold. Firstly, to realize end-to-end service provisioning for applications, the KDN-based NC&M scheme needs to manage not only the network elements in both the EPS and OCS clouds but also the virtual machines (VMs) in the servers. In other words, it is required to efficiently orchestrate the IT and bandwidth resources in an HOE-DCN. Secondly, the telemetry data from network elements and servers would be in huge volume, while the useful information buried in it for intelligent and automatic NC&M (*i.e.*, the knowledge) is sparse. Hence, we need an effective AI-assisted method to extract the knowledge in a time-efficient manner. We design a knowledge-defined control plane system to address the two aforementioned challenges. As shown in Fig. 1(a), the control plane consists of three DL-based modules, *i.e.*, the knowledge-defined orchestrator (KD-O), IT resource controller (IT-C), and network controller (NET-C). They work coordinately to orchestrate the IT and bandwidth resources in the data plane based on predictive analytics. In the following, we will first elaborate on the design of the DL-based modules and how they cooperate to realize predictive analytics, then briefly describe the system implementation to prototype the HOE-DCN system in Fig. 1(a), and finally discuss the open questions in this area.

2. Predictive Analytics

Note that, the telemetry data from the devices in an HOE-DCN generally includes historical requests on VMs and traffic flows, distribution and status of VMs, in-service applications and their QoS parameters, configuration and status of intra- and inter-rack networks, status of traffic flows, etc. Apparently, if we plan to use one DL module to analyze all of them, the DL module would not only be prohibitively complex but also take unreasonably long time to be trained. Therefore, we decide to incorporate three DL modules to accomplish the task, and our detailed design of the control plane for an HOE-DCN is shown in Fig. 1(c). The three DL-based modules are organized in a hierarchical structure, where KD-O sits on top of IT-C and NET-C to orchestrate them. The IT resource monitor in IT-C collects telemetry data related to VMs proactively, and then stores it in the network service database (NSD) and also forwards the data to the DL-based IT demand predictor. The VM management module takes the instructions from KD-O to implement VM deployment, migration and scaling accordingly. Note that, there are also communication channels in between KD-O and the NSD and DL-based IT demand predictor, for IT-C to report current VM configuration and forecasted IT demands. The design of NET-C includes more modules. Here, the network abstraction module (NAM) collects both the configuration and status of the HOE-DCN and the status of traffic flows among VMs, stores this telemetry data in the traffic engineering database (TED), and also sends the data to the DL-based traffic predictor. The DL-based traffic predictor and TED report necessary information to KD-O, which will then instruct the flow provisioning module and the OXC reconfiguration module to adjust the network part of HOE-DCN for preparing for future traffic demands.

Note that, even though the knowledge extracted by IT-C and NET-C is important, it is still fragmented and thus cannot be utilized directly for orchestrating the IT and bandwidth resources in the HOE-DCN. This is also the reason why we refer to it as "low-level knowledge" in Fig. 1(c). Taking the low-level knowledge as input, the DL-based resource orchestrator in KD-O further abstracts high-level knowledge out of it. Here, the high-level knowledge refers to the matching degree between the HOE-DCN's configuration and the applications running in it, *i.e.*, the information that is exactly required to effectively orchestrate the IT and bandwidth resources. This actually mimics the knowledge regarding shapes and colors, and then leverage it to define objects for identifying the person's face or figure (*i.e.*, high-level knowledge). Moreover, the cooperation of IT-C, NET-C and KD-O follows the principle of predictive analytics in human behaviors, as explained in Fig. 2(a). Specifically, the control plane first uses IT-C and NET-C to predict future demands on IT and bandwidth resources, respectively (*i.e.*, forecasting based on memory) and then lets KD-O find the optimal configuration of the HOE-DCN according to the predictions (*i.e.*, decision making based on knowledge).

3. System Prototyping

In order to verify the effectiveness of our proposed HOE-DCN system, we prototype it in a real network testbed as shown in Fig. 2(b) and conduct experiments. Most of the functional modules in IT-C are implemented based on OpenStack, for VM monitoring and management [9], while the DL-based IT demand predictor is developed based on TensorFlow. Similarly, the modules in NET-C, which are related to network monitoring and management, are implemented based on ONOS [10], while the DL-based traffic predictor is also programmed based on TensorFlow. KD-O is also based on TensorFlow. We leverage the external application programming interfaces (APIs) provided by OpenStack, ONOS and TensorFlow to realize the communications and cooperation among IT-C, NET-C and KD-O. The data plane of the HOE-DCN prototype includes servers stacked in three racks. In each rack, the servers are connect to their ToR switch through 1GbE ports. The inter-rack EPS cloud is built with aggregation and core switches organized in a hierarchical structure. Here, all the ToR, aggregation and core switches support OpenFlow and can be



Fig. 2. (a) Principle of predictive analytics, and (b) Prototype of our proposed HOE-DCN.

managed with a centralized SDN controller. In the mean time, the ToR switches are also interconnected through a 24×24 OXC to emulate the inter-rack OCS cloud. In order to distinguish the capacities of the EPS and OCS inter-rack clouds, we let ToR switches connect to them with different type of ports, *i.e.*, the ToR switches use 1GbE ports to communicate with their aggregation switches while 10GbE optical ports are utilized to connect to the OXC [2].

In the experiments, we train the three DL modules using the real telemetry data collected in the HOE-DCN testbed. Specifically, we first use the DCT2GEN tool to generate the traffic matrixes among VMs, and then realize the traffic matrixes in the HOE-DCN testbed by running iPerf to pump the specified traffic volumes among the desired VMs. Meanwhile, KD-O instructs IT-C and NET-C to implement random configurations on the HOE-DCN and collect key performance metrics such as average data-transfer latency, average CPU usage, and number of active ports. Here, the number of active ports actually reflects the power consumption of the HOE-DCN system. Among the three DL modules, the DL-based traffic predictor and the DL-based resource orchestrator can be trained within a few hundred seconds, while the DL-based IT demand predictor takes around 15.5 hours to be trained [2] due to the complexity of the task. The trained DL modules are then put in the control plane to reach intelligent decisions for automatic NC&M. Promisingly, the experimental results reported in [2,8] suggested that KD-O can consolidate the traffic distribution to make intra-rack traffic be dominated, and thus both the average data-transfer latency and the number of active ports can be reduced. Therefore, unlike the traditional IT and bandwidth resource orchestrators that can only trade data-transfer latency for energy efficiency, our proposal can leverage predictive analytics to optimize both metrics simultaneously.

4. Discussion and Future Work

We explained our attempt to realize intelligent and automatic NC&M for orchestrating the IT and bandwidth resources in an HOE-DCNs. Specifically, by leveraging KDN, we designed the control plane to include three DL modules, *i.e.*, KD-O, IT-C and NET-C, and made them work coordinately based on predictive analytics.

Note that, there is still a relatively big gap between our prototype and a practical HOE-DCN. To be practical, we first need to address the scalability issue. A practical DCN could easily have hundreds of thousands of servers and switches, and thus its telemetry data would be in tremendous volume to make automatic NC&M extremely complicated. Hence, the scheme to scale KD-O, IT-C and NET-C efficiently is a must. Moreover, when the HOE-DCN scales up, the two-level model for knowledge abstraction needs to be extended and consider multiple levels for improved effectiveness.

References

- "Cisco Global Cloud Index: Forecast and Methodology, 2016-2021," Cisco [Online], Available: https://www.cisco.com/c/en/us/solutions/service-provider/visualnetworking-index-vni/index.html.
- [2] W. Lu et al., "AI-assisted knowledge-defined network orchestration for energy-efficient datacenter networks," *IEEE Commun. Mag.*, in Press, 2018.
- [3] N. Farrington *et al.*, "Helios: a hybrid electrical/optical switch architecture for modular data centers," ACM SIGCOMM Comput. Commun. Rev., vol. 40, pp. 339-350, Oct. 2010.
- [4] P. Lu et al., "Highly-efficient data migration and backup for Big Data applications in elastic optical inter-datacenter networks," IEEE Netw., vol. 29, pp. 36-42, Sept./Oct. 2015.
- [5] P. Lu et al., "Distributed online hybrid cloud management for profit-driven multimedia cloud computing," IEEE Trans. Multimedia, vol. 17, pp. 1297-1308, Aug. 2015.
- [6] S. Li et al., "Improving SDN scalability with protocol-oblivious source routing: A system-level study," IEEE Trans. Netw. Serv. Manag., vol. 15, pp. 275-288, Mar. 2018.
- [7] A. Mestres et al., "Knowledge-defined networking," ACM SIGCOMM Comput. Commun. Rev., vol. 47, pp. 2-10, Jul. 2017.
- [8] W. Lu *et al.*, "Leveraging predictive analytics to achieve knowledge-defined orchestration in a hybrid optical/electrical DC network: Collaborative forecasting and decision making," in *Proc. of OFC 2018*, paper Th3F.2, Mar. 2018.
- [9] B. Kong et al., "Demonstration of application-driven network slicing and orchestration in optical/packet domains: On-demand vDC expansion for Hadoop MapReduce optimization," Opt. Express, vol. 26, pp. 14066-14085, 2018.
- [10] Z. Zhu et al., "Build to tenants' requirements: On-demand application-driven vSD-EON slicing," J. Opt. Commun. Netw., vol. 10, pp. A206-A215, Feb. 2018.