Realizing AI-Assisted Multi-Layer Restoration in a Software-Defined IP-over-EON with Deep Learning: An Experimental Study

Siqi Liu[1], Baojia Li[1], Zuqing Zhu[1]

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

Abstract: By using deep learning, we experimentally demonstrate AI-assisted multi-layer restoration in an IP-over-EON, which recovers affected traffic timely with congestion-avoidance rerouting.

OCIS codes: (060.1155) All-optical networks; (060.4251) Networks, assignment and routing algorithms.

1. Introduction

Elastic optical networks (EONs) can achieve agile spectrum management in the optical layer to support upper-layer applications adaptively [1, 2]. The ever-growing IP services in the Internet make a rational combination of IP and EON technologies an inevitable trend [3]. However, in an IP-over-EON, service interruptions can be caused by failures in both the EON and IP layers. More importantly, a recent analysis on the failures in Google's wide-area networks suggested that failures in the IP layer (*e.g.*, router outages) actually happened much more frequently than those in the optical layer [4]. Hence, to ensure a survivable IP-over-EONs, effective cross-layer orchestration is needed to realize multi-layer restoration (MLR) for addressing various failure events [5,6]. This can be achieved by leveraging software-defined networking (SDN) [5]. Specifically, when a failure happens, the centralized network control and management (NC&M) can respond quickly and reroute the affected traffic in both the IP and EON layers.

Previous studies on MLR in IP-over-optical networks [5, 6] usually tried to reroute the affected flows based on the current network status. This, however, might not always be effective since IP traffic is usually highly dynamic and bursty. Therefore, even though the current network status allows us to groom the affected flows with existing ones and transfer them in established lightpaths, congestions may happen in the future due to the sudden increase of certain flows' data-rates. Such a congestion can cause cascaded failures [4], which not only fails the MLR but also complicates the failure handling. Hence, it would be desirable that the centralized NC&M in software-defined IP-over-EONs (SD-IPoEONs) could be more intelligent, such that it can analyze and predict the traffic fluctuation on each established lightpath to achieve artificial intelligence (AI) assisted MLR, for avoiding future congestions. Nevertheless, to the best of our knowledge, this type of SD-IPoEON system has not been studied or demonstrated before.

In this work, we design and demonstrate an SD-IPoEON system that can realize AI-assisted MLR by leveraging the deep learning (DL) based on a long short-term memory neural network (LSTM-NN) [7]. The control plane (CP) of the SD-IPoEON is developed based on ONOS [8] to realize universal control of both the OpenFlow switches (OF-SWs) in the IP layer and the OpenFlow-enabled bandwidth-variable wavelength-selective switches (BV-WSS') in the EON, for cross-layer orchestration. To realize traffic analysis and prediction, we use TensorFlow to implement a DL module in the CP. When a network failure happens, the CP achieves AI-assisted MLR by letting ONOS calculate the most effective rerouting scheme for affected flows based on the traffic prediction provided by the DL module. Then, the MLR scheme is implemented in the data plane (DP). By sending real traffic through its DP, our experimental demonstrations verify that the proposed system can recover affected flows timely during network failures, and our AI-assisted MLR achieves congestion-avoidance traffic rerouting based on the predication from the DL module. In all, our proposal reveals a promising use-case of knowledge-defined networking (KDN) [9] in optical networks.

2. Network Architecture and System Design

Fig. 1 shows the architecture of our SD-IPoEON system. The IP layer in the DP is based on OF-SWs, which are equipped with optical ports for 10 Gigabit Ethernet (10GbE) connections. The optical ports are interconnected with each other through the BV-WSS' in the EON. On each OF-SW, there are several locally-attached servers, which generate/receive dynamic IP traffic. Hence, as indicated in Fig. 1, two servers can talk with each other through either an end-to-end lightpath that connects a pair of 10GbE ports on their local OF-SWs directly or multiple lightpaths that invoke O/E/O conversion(s) and traffic de-/re-grooming at intermediate OF-SW(s). The actual communication scheme is determined and implemented by the ONOS-based CP, which consists of an ONOS module and a DL module. The OpenFlow controller (OF-C) in the ONOS communicates with both the OF-SWs and BV-WSS' to realize cross-layer orchestration. Specifically, the OF-C can insert, update and remove flow-tables in the OF-SWs to achieve traffic de-/re-grooming, and meanwhile, it also communicates with the OpenFlow agents (OF-AGs) on the BV-WSS' based on the OpenFlow with optical transport protocol extensions (OF w/ OTPE) to realize lightpath setup/removal.

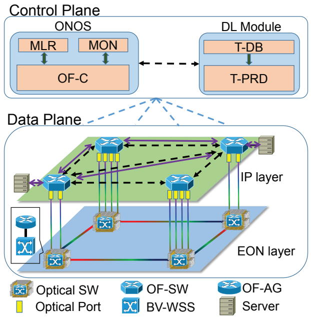

Fig. 1. Network architecture of SD-IPoEON.

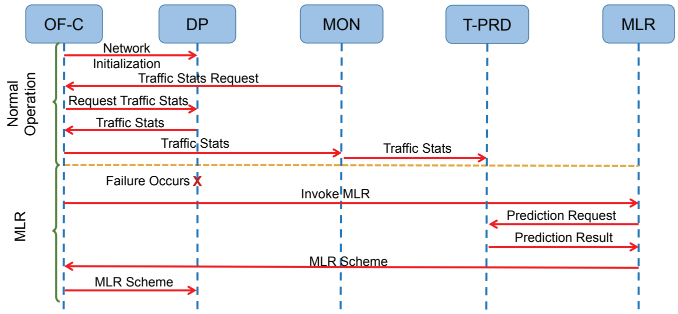

Fig. 2. Operation procedure.

To ensure network survivability, we implement a monitor (MON) in the ONOS to observe the working status of each network element in the DP as well as collect traffic statistics on established lightpaths periodically. If a network failure is detected, the ONOS would use the MLR submodule to calculate the rerouting schemes for affected flows. Note that, for AI-assisted MLR, the MLR utilizes the knowledge on traffic prediction provided by the DL. Here, the traffic statistics collected by the MON are forwarded to the DL consistently, where they get stored in the traffic database (T-DB) and are used to train the traffic predictor (T-PRD) based on an LSTM-NN. Specifically, the T-PRD can take historical traffic statistics as input and predict future traffic fluctuations accurately after proper training. This is achieved with our design of the LSTM-NN, which, as Fig. 3 shows, consists of three neural network layers, *i.e.*, the input, hidden and output layers. The input layer takes in $n+1$ historical traffic samples, the output layer buffers m predicted traffic samples, while the hidden layer is an LSTM layer that can be trained to map input to output. Moreover, to improve the prediction accuracy, we use real-time recurrent learning (RTRL) to update the LSTM-NN.

The operation procedure of our SD-IPoEON system is depicted in Fig. 2. During network initialization, the OF-C populates flow-tables on each network element in the DP to make sure that the traffic among the servers can be delivered smoothly. Then, the MON acquires network status from the OF-C periodically. For the traffic statistics on the lightpaths, we collect the data-rate of each optical port on the OF-SWs. The MON also sends the traffic statistics to the DL, which uses them to train the T-PRD. When a network failure (*e.g.*, an OF-SW outage) happens, the OF-C would invoke the MLR immediately for failure recovery. The MLR, in turn, asks for the traffic prediction of related lightpaths from the T-PRD, and then uses an MLR algorithm to calculate the rerouting schemes for the affected flows. Here, for the MLR algorithm, we adapt the AG-NE reported in [6] and add forecast-based congestion-avoidance in it. Specifically, if the affected flows could be groomed with the existing ones on established lightpaths and would not cause congestions in the future according to the traffic prediction, we would just reroute them in the IP layer. Otherwise, new lightpath(s) would be set up in the EON to provide extra transmission capacity for the failure recovery.

3. Experiment Demonstration

We implement the proposed SD-IPoEON system in a practical network testbed. The ONOS and DL modules in the CP are realized on commodity servers, while for the DP, the OF-SWs are Pica-8 switches equipped with optical 10GbE ports (*i.e.*, operating in the C-band), and the EON consists of four Finisar 1×9 BV-WSS'. Each BV-WSS covers the wavelength range within $[1528.43, 1566.88]$ nm and supports a bandwidth allocation granularity of 12.5 GHz. Its operation is controlled by an OF-AG, which is also implemented on a commodity server. To fully demonstrate the effectiveness of our SD-IPoEON system, we send traffic through the DP and conduct AI-assisted MLR experiments with live and dynamic traffic on. The topologies of and traffic in the IP and EON layers are shown in Fig. 4. Specifically, we have five lightpaths (*i.e.*, $\{LP1, \cdots, LP5\}$) to carry the five flows (*i.e.*, $\{F1, \cdots, F5\}$) in the IP layer. To emulate the case in a real network, the traffic fluctuation on each lightpath are set according to the traces collected by ISPs [10].

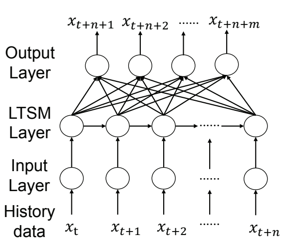

Fig. 3. Proposed LSTM-NN.

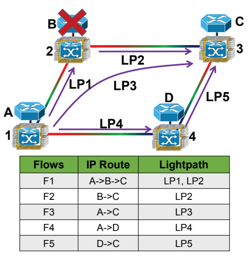

Flows	IP Route	Lightpath
F1	A->B->C	LP1, LP2
F2	B->C	LP2
F3	A->C	LP3
F4	A->D	LP4
F5	D->C	LP5

Fig. 4. Experimental Setup.

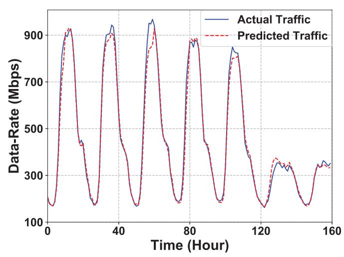

Fig. 5. Prediction results.

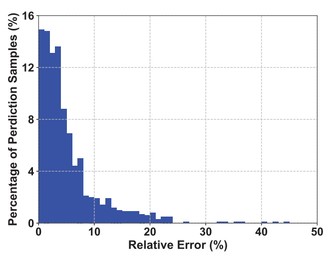

Fig. 6. Relative error of prediction.

In the experiments, we divide the traffic samples in each trace into two sets, *i.e.*, the training and testing sets. The training set is purely used to train the T-PRD, while the testing set is used in the experimental demonstration of AI-assisted MLR. The T-PRD's performance on traffic prediction is shown in Figs. 5 and 6. We can see that the predicted traffic matches with the real traffic well in Fig. 5, and the relative error distribution in Fig. 6 suggests that the average relative error is 4.28% and 82% of predicted results have a relative error less than 10%. We run iPerf on the servers to generate dynamic traffics for $\{F1, \cdots, F5\}$ according to the traces. Note that, we can hardly saturate the capacity of optical 10GbE ports on the OF-SWs in the experiments, because the servers are connected with the OF-SWs through 1GbE ports and we only have a limited number of servers. Hence, to emulate a real network in which congestions can happen, we scale down the traces with a peak throughput of 1 Gbps and also limit the capacity of each optical 10GbE port as 1 Gbps. Then, the MON in the CP collects the average data-rate of each lightpath every 10 seconds to sample its traffic fluctuation. During normal operation, the T-DB buffers 30 historical traffic samples for each lightpath. When a network failure happens, the T-PRD will predict 10 future traffic samples for each lightpath to assist the MLR.

Then, we emulate an OF-SW outage in the system by shutting down *OF-SW B* manually. Note that, when *OF-SW B* is down, only the traffic that experiences de-/re-grooming on it will be affected (*i.e.*, *F1* and *F2*), while the traffic that only uses its underlying BV-WSS will not be affected (*i.e.*, *F3*) since *BV-WSS 2* is intact. Here, *F2* cannot be recovered before *OF-SW B* is fixed, because it uses a local server on *OF-SW B* as the source. Hence, the CP detects the failure and determines that *F1* needs to be rerouted. Fig. 7 shows the predicted traffic fluctuations on lightpaths *LP3* and *LP4* if we groom *F1* on one of them in the MLR, which indicates that even though *LP3* or *LP4* has enough capacity to carry *F1* at the time of recovery, congestions will happen in the future. Therefore, the AI-assisted MLR decides to set up a new transparent lightpath between *OF-SWs A* and *C* to restore *F1*. This can be verified by the spectrum results in Fig. 8, which are measured at the input of *BV-WSS 3*. We observe that before the OF-SW outage, the spectrum of *LP2* can be seen in Fig. 8(a), while after the outage, *LP2* is gone in Fig. 8(b) and the MLR sets up a new lightpath to reroute *F1* to *BV-WSS 3* directly. Fig. 9(a) shows the receiving ratio of *F1*, which indicates that the AI-assisted MLR recovers its traffic with 1.94 seconds after the OF-SW outage, and the receiving ratio would not decrease after the MLR. For comparison, we disable the DL module in the CP, only use the ONOS for the MLR, and redo the experiment. This time, the CP decides to groom *F1* onto *LP3* in the MLR. The receiving ratio of *F1* is shown in Fig. 9(b), which shows that future congestion on *LP3* will cause severe packet losses on *F1*, even though grooming it on an existing lightpath induces much shorter traffic interruption. Specifically, we measure the packet loss ratio of *F1* for its whole lifetime and found that the results are 0.5% and 37.5% for with and without AI assistance, respectively. In all, the experimental results verify that our proposed system can recover affected flows timely with congestion-avoidance rerouting.

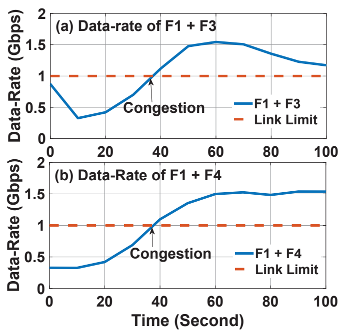

Fig. 7. Predicted traffic fluctuation.

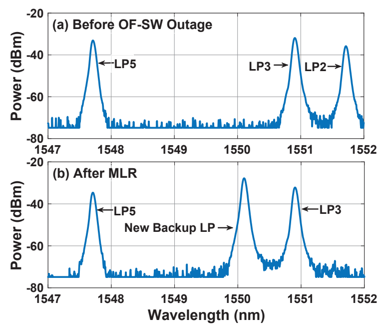

Fig. 8. Spectra measured on input of *BV-WSS 3*.

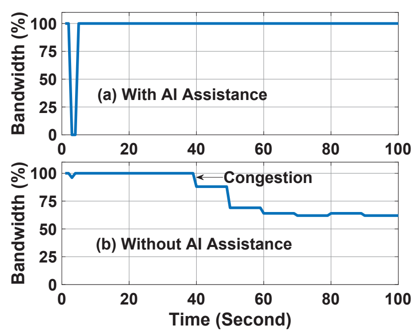

Fig. 9. Receiving ratio of *F1*.

4. Summary

We designed and experimentally demonstrated an SD-IPoEON system, which realizes AI-assisted MLR with DL and achieves congestion-avoidance traffic rerouting based on the precise traffic predication from the DL.

References

[1] O. Gerstel *et al.*, *IEEE Commun. Mag.*, vol. 50, no. 2, pp. s12-s20, Feb. 2012.
[2] Z. Zhu *et al.*, *J. Lightw. Technol.*, vol. 31, no. 1, pp. 15-22, Jan. 2013.
[3] O. Gerstel *et al.*, *IEEE Commun. Mag.*, vol. 52, no. 1, pp. 44-51, Jan. 2014.
[4] R. Govindan *et al.*, *in Proc. of ACM SIGCOMM 2016*, pp. 58-72, Aug. 2016.
[5] I. Maor *et al.*, *in Proc. of ECOC 2016*, paper Tu.1.B.3, Sept. 2016.
[6] S. Liu, *et al.*, *in Proc. of OFC 2017*, paper W3I.5, Mar. 2017.
[7] S. Hochreiter *et al.*, *Neural Comput.*, vol. 9, no. 8, pp. 1753-1780, Nov. 2016.
[8] ONOS. [Online]. Available: http://onosproject.org/
[9] A. Mestres *et al.*, *ACM SIGCOMM Comput. Commun. Rev.*, vol. 47, no. 3, pp. 4-10, Jul. 2017.
[10] [Online]. Available: https://data.is/TSDLdemo