

Software Defined Enterprise Passive Optical Network

Ahmed Amokrane^{*†} Jinho Hwang[†] Jin Xiao[†] Nikos Anerousis[†]

^{*}*LIP6 / Pierre and Marie Curie University, Paris, France*

[†]*IBM T.J. Watson Research Center, Yorktown Heights, NY, USA*

Abstract

In the last few years, changing infrastructure and business requirements are forcing enterprises to rethink their networks. Enterprises look to passive optical networks (PON) for increased network efficiency, flexibility, and cost reduction. At the same time, the emergence of Cloud and mobile in enterprise networks calls for dynamic network control and management following a centralized and software-defined paradigm. In this context, we propose a software-defined edge network (SDEN) design that operates on top of PON. SDEN leverages PON benefits while overcoming its lack of dynamic control. This paper is a work-in-progress focusing on enabling key flow control functions over PON: dynamic traffic steering, service dimensioning and realtime re-dimensioning. We also discuss how SDEN edge network can integrate with core SDN solutions to achieve end-to-end manageability. Through case experiment studies conducted on a live PON testbed deployment, we show the practical benefits and potentials that SDEN can offer to enterprise networks redesign.

1 Introduction

Enterprise networks today are under tremendous pressure to change. A recent study conducted by the Economist Intelligence Unit Research, sponsored by Juniper Networks, points out that over 50% of the businesses surveyed consider IT operations a core business enabler, and yet they find that their current IT infrastructure largely falls short of expectations in driving business growth [1]. The problem stems from changes in infrastructure and business requirements. Infrastructure wise, enterprise networks desire operational efficiency, management simplicity, green and cost effectiveness; Business requirement wise, enterprise networks need to cope with disruptive yet business vital technologies such as Cloud and mobile. In fact, Cloud and mobile are fundamentally changing the traffic characteristics of enterprise networks, and how best to manage them. It comes from the fact that network traffic pattern is far more dynamic and uncertain when enterprise embraces these new technologies in their midst. Cloud workload varies significantly with time and mobiles' traffic is migratory and volatile. Taking the stadium enterprise as an example, the traffic patterns exhibited during a game day is highly

dependent on the phase and condition of a game. Before the game starts, the majority of the traffic comes from the gate entrance; during game periods, they are concentrated at the seating areas. During half-time, the concentration shifts to the concourse; and after the game, they migrate to the parking lots. The transitory traffic volume and burst intensity is also highly related to changes in game states (e.g., a remarkable touchdown scored by the home team is likely to trigger a large surge in mobile traffic). Over the past few years, a vast majority of the stadiums that have hosted the final games of the Superbowl have experienced significant network congestion/failure on game days. Many campuses also experienced similar conditions on the eve of a major iOS or mobile game release.

In this context, Passive Optical Networks (PON) is a promising technology that contains the key ingredients in addressing the new infrastructure requirements. Originated from Fiber-to-the-Home (FTTH) sector, PON is becoming an attractive fiber-based LAN edge network solution for enterprises. Some of the key benefits PON brings to enterprise networks are: significant reduction in capex and opex, centralized control and management, high capacity, flexible deployment, and strong physical and communication security. IBM and its PON alliance partners have observed significant increase in the demand for enterprise PON since 2013. This is exemplified by a recent 10 million plus USD deal with Texas A&M Kyle Field stadium to renovate their network infrastructure. From an infrastructure perspective, much of the savings PON brings to enterprise is due to the replacement of distribution layer active equipments (i.e., network switches) with passive optical splitters, while relying on the core for flow control and traffic management. On the other hand, there is a need to redesign edge network control and management that is flexible, responsive, and adaptive to realtime conditions. We think Software Defined Networking (SDN) is a promising paradigm that satisfies the new enterprise network requirements.

In light of this, we are actively investigating, at IBM Research, the application of SDN on top of PON to achieve enterprise network efficiency and agility. In this paper, we discuss two challenges we have encountered and addressed: firstly, SDN is a layer 3 technology that does not have corresponding translation to layer 2 edge network, which is required for End-to-

End manageability; secondly, flow control and management in PON cannot be achieved by traditional edge switching due to PON's lack of distribution switching fabric. Innovation is therefore required to enable primary flow control functions such as flow steering, dimensioning and performance management. In this paper, we introduce software-defined edge network (SDEN), a software-defined enablement of enterprise PON that has the flexibility, responsiveness and realtime control capability to meet the new enterprise network requirements. As this paper reports on our work-in-progress, we focus on how flow management functions such as steering, dimensioning and re-dimensioning can be supported in PON, and discuss how SDEN interacts with SDN controller to achieve end-to-end manageability. For feasibility study, we implemented SDEN flow control functions in software and conducted case experiments on live PON testbed deployment.

The rest of this paper is organized as follows. Section 2 presents PON and SDN technologies and related works in the literature. Section 3 presents our SDEN framework, followed by the experiments in Section 4. Section 5 presents our vision in terms of End-to-End manageability. Finally, Section 6 concludes the paper.

2 Background and Related Work

2.1 Passive Optical Networks

A typical PON is a set of Optical Network Terminals (ONTs), passive splitters and the Optical Line Terminal (OLT). The ONTs connect edge devices into the PON network via Ethernet ports. Digital signals from edge devices are converted to optical signal in the ONT. The optical splitters split the light signal multiple ways to ONTs and transmit the multiplexed signal to the OLT. The OLT aggregates all optical signals from the ONTs and converts them back to digital for the core router. The OLT may support a range of built-in functionalities such as integrated Ethernet bridging, VLAN capability and security filtering. Compared with traditional copper network, PON replaces switches in the access and aggregation layers with splitters, and the traditional distribution layer is collapsed back to a few OLTs at the core. An OLT may support 8-72 fiber ports, with each port connecting a fiber cable to the splitter. The splitter can support different splitting ratios with 1-32 or less being the recommended ratio. Therefore each OLT port can potentially support 32 ONTs. Different ONT configurations are available ranging from 2 to 24 Ethernet ports. Enterprise PON uses the ITU-T Gigabit PON (GPON) standards [2, 3, 4]. We, therefore, use PON and GPON interchangeably in this paper.

Compared to traditional copper networks, PON has a

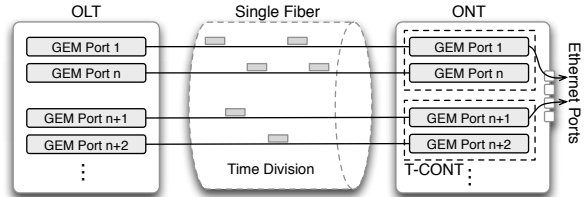


Figure 1: Traffic management in GPON networks

number of salient advantages. In fact, the optical fibers in PON can travel up to 20 Km from the core to the access, capable of delivering upstream to 1.2 Gbps up and 2.4 Gbps downstream to the port, in current generation, and the fiber is much lighter than copper cables. Moreover, PON eliminates active equipments in the distribution resulting in significant capex and opex savings (up to 40% and 60% respective savings compared to traditional enterprise copper networks [5]). Furthermore, PON offers much stronger security with enhanced data encryption and physical protection [6]. For more details about the enterprise PON technology and its benefits over traditional copper network, please refer to [5].

The entire PON network constitutes an Ethernet LAN. In fact, users' Ethernet frames are encapsulated in GTC Encapsulation Method (GEM) frames. Each GEM frame belongs to a GEM port. A GEM port represents a logical connection (channel) between an ONT and an OLT, with a class of service and a unique identifier. A typical architecture for traffic management in GPON is illustrated in Figure 1. A Transmission Container (T-CONT) is an ONT object representing a set of GEM ports that appear as a single entity for the purpose of upstream bandwidth assignment on the PON. In the upstream direction, bandwidth allocation for ONTs is done in a TDMA manner by the OLT, where each slot is allocated for a given T-CONT. More specifically, users' Ethernet frames are assigned N-VLAN tags (Network VLAN) and CoS (802.1p) values based on Physical Port of the ONT, Subscriber VLAN ID, 802.1p bits and/or DSCP, as defined by the ITU-T GPON standard. Then, each of these N-VLAN and CoS combination is mapped into a specific GEM port, and the QoS of the T-CONT to which the GEM port belongs applies to the frame for scheduling. In the downstream direction, traffic is transmitted in a TDM manner, where each ONT forwards the traffic to the appropriate GEM port.

2.2 Software Defined Networks

SDN has recently emerged as new norm for networks. In a nutshell, SDN relies on (i) decoupling the control plane from the data plane, (ii) logically centralized controller and (iii) a standard protocol, such as OpenFlow [7], for communication between the controller and the forward-

ing elements in the network. SDN has mainly been used in data center networks, with mainly an Ethernet copper-based switching fabric. As SDN offers flexibility, manageability and agility, number of proposals advocated to extend SDN for wireless networks, be it cellular [8] or WLANs wifi-based networks [9, 10], wireless mesh networks [11], campus copper-based networks [12]. Moreover, one active and interesting effort is to extend SDN to optical networks [13, 14, 15]. The objective is to ease management and flexibility that are often rigid and cumbersome. In enterprise networks, SDN helps to address the problems of flow control, network load balancing and performance management (quality assurance and congestion control), required by increasingly heterogenous, mobile and dynamic user traffic profiles.

On the other hand, optical networks are becoming an attractive solution as they offer higher capacity and reduced opex and capex. Logically, SDN should eventually be extended to incorporate PONs in the years to come. In fact, the ONF created The Optical Transport Working Group (OTWG) [13]. The OTWG will work towards identifying use cases, defining a target reference architecture for controlling Optical Transport Networks (OTNs) incorporating OpenFlow, and identifying and creating OpenFlow protocol extensions. Gringeri et al. [15] identified some of the key requirements, benefits and challenges of extending SDN concepts to OTNs. However, these works focused on OTNs, which are capable of active switching and use GMPLS for creating virtual circuits on top of the optical backbone, and did not address the challenging aspects of PONs. The first work to introduce SDN paradigm in PONs was proposed by Parol et al. [16]. In this work, authors proposed extensions to OpenFlow protocol, which consist mainly on mapping flows (as defined by the OpenFlow protocol) to GEM ports, in addition to pushing and popping VLAN tags from the packets. However, such proposal requires changes in the ONTs and OLTs to be implemented. Additional works, such as [17], which considers the specific requirements of an ISP GPON-based networks, have also proposed hints for integrating SDN in optical networks. However, the dynamic and mobility pattern of enterprise network traffic and the need for agility have not been addressed in this work.

3 Software Defined Edge Network

3.1 SDEN Framework

Figure 2 illustrates our proposed architecture for SDEN. SDEN defines a common interface through APIs between the controller and the PON nodes (OLT). The interface provides a standardizable and vendor neutral set of functionalities that a controller can use. On top of the

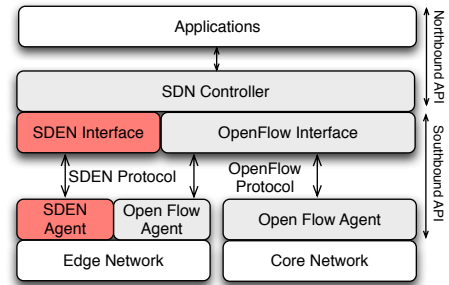


Figure 2: Software Defined Edge Network Extension with PON

controller, one can have different applications for network management and optimization. For instance, in this paper, we show how an application can perform dynamic traffic and capacity steering in PON.

To realize this design, we need to tackle two major challenges. First, there is no concept of flows in PON, and the existing PON nodes (i.e., OLT and ONT) do not talk or understand the OpenFlow protocol [16]. Therefore, we define in SDEN a mapping between a flow and a set of PON primitives (fiber connectivity, PON ports and service profiles). Moreover, PON can not perform flow control in the conventional sense as there is no active switching fabric. Instead, PON management is conducted centrally at the OLT by setting service profiles and defining PON ports attributes. Second, PON native management interfaces in the OLT allow a human administrator to perform these configurations manually (e.g., Tellabs PON provides a command line interface). To enable automatic real-time flow control in PON, A new SDEN module is appended to the SDN controller to call the APIs of the SDEN agent located in the edge network. More specifically, the SDEN agent translates high level flow control requests from a controller into a set of native PON configuration commands, and sends them to the OLT via the command line interface (CLI).

In this way, we are able to follow a software-defined paradigm to dynamically control and manage flows on the fly in response to changes in traffic characteristics induced by end user workload changes and mobility. To this end, there are three key functionalities the SDEN agent needs to support: online flow and capacity steering, service dimensioning, and realtime service re-dimensioning. In the following subsections, we discuss the motivation and how each one of these functionalities is enabled in PON.

3.2 Flow Control and Steering

One issue in enterprise and campus networks is to achieve agility. In fact, in most of the current deployments, capacity is statically allocated to parts of the net-

work. However, nowadays enterprise traffic patterns vary significantly with time and mobiles' traffic is migratory and volatile. As a matter of fact, it is of paramount importance to provide capacity in an on-demand fashion to parts of the network. Taking the stadium enterprise as an example, the capacity should be allocated to the different parts of the network depending on the phase and condition of the game. To achieve this, we propose a two-fold solution.

First, we advocate a crisscross deployment where multiple fibers run from different GPON cards and/or ports towards the same area. Note that other fibers from the same GPON cards and/or ports might run to different areas of the campus. For instance, a single WAP is connected to multiple ONTs, each of the ONTs is connected to a different GPON card/port. As such, we achieve an N:N mapping of the WAPs and the ONTs-OLTs. Note that this deployment allows traffic steering and bandwidth management in a dynamic manner.

Second, we propose a dynamic approach to direct the capacity in the PON network towards specific areas, depending on the traffic conditions in the areas. To do so, we propose a dynamic approach that (i) monitors the GPON links (from the ONT towards the OLT) and (ii) dynamically reroutes traffic through different GPON ports to provide more capacity to overloaded areas of the network. More specifically, we propose a simple yet efficient algorithm that uses predefined thresholds to dynamically perform capacity steering in the network. In a nutshell, if a link utilization (i.e., fiber at the GPON port) is above a threshold α , we redirect part of the traffic that runs through an ONT (one Ethernet port at the ONT) to use a different path, going through a different GPON port. At the same time, if the utilization of a link is below a threshold β , we consolidate the traffic through an underutilized GPON port. The objective is to achieve elasticity in resource utilization in the network in response to the traffic load.

It is worth noting that our proposed dynamic capacity steering can run as an application on top of the SDN controller. For ease of explanation, we do not include the controller in this paper and consider the application to be the only one running on the controller.

3.3 Dynamic Flow Dimensioning

As traffic in enterprise networks becomes heterogeneous, differentiating flows with priorities and requirement is another key functionality that PON should offer. More specifically, one should be able to offer guaranteed QoS for classes of traffic independently of the network state.

In our proposed framework, we provide guaranteed bandwidth for specific traffic flows based on the user devices. In fact, we can provide guaranteed committed

rates regardless of the traffic and state of the network. The idea is to isolate critical traffic from congestions that might occur in the network by providing classes of services, with strict priority. In an enterprise network, one can think of VOIP traffic as being critical and should not be affected by the other traffic flowing in the network.

3.4 Real-time Flow Re-dimensioning

In addition to initial flow dimensioning, PON should also allow for realtime flow re-dimensioning by dynamically adjusting the allocated resources. The aim of dynamic re-dimensioning is to free more resources for traffic with higher priority and efficiently use the resources in the network. For instance, in the absence of VOIP traffic, one can allocate the available bandwidth to data traffic. However, in the case of high VOIP traffic, data traffic is delayed to leave room for VOIP traffic.

Our proposed framework enables dynamic re-dimensioning of services on the fly and in realtime. More specifically, we propose to prioritize certain services over others by dynamically adjusting the allocated resources in the network to the different services based on their corresponding priorities. Similar to capacity steering, dynamic service re-dimensioning is performed thanks to resource and demand monitoring in the network.

4 Experiments

In this section, we demonstrate the effectiveness of our proposal, and illustrate the benefits of the three key enablers for SDEN: traffic and capacity steering, service dimensioning and dynamic service re-dimensioning. To do so, we conducted experiments on a real PON testbed deployed in our lab. Figure 3 illustrates the experimental setup. We used a Tellabs 1150 OLT with one single mounted GPON card. We run two fibers from two different GPON ports (GPON port 1 and 2) of the same GPON card. We attach to each of the two ports one ONT, Tellabs

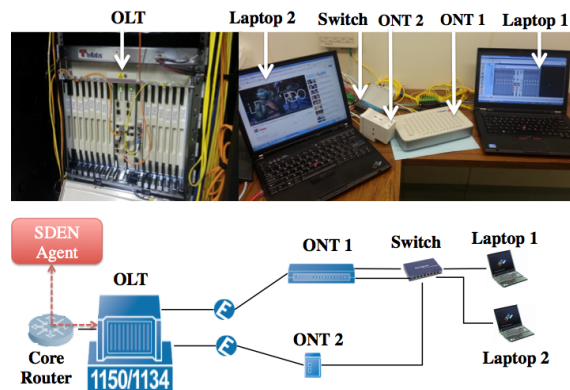
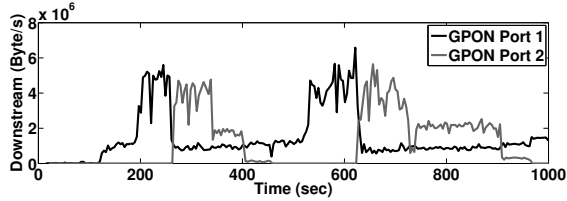
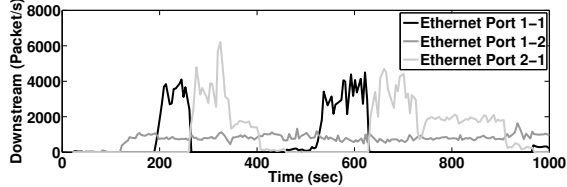


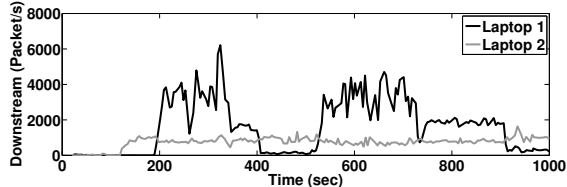
Figure 3: Experimental setup



(a) Downstream traffic at the GPON ports



(b) Downstream traffic at the Ethernet ports



(c) Downstream traffic at the laptops

Figure 4: Dynamic traffic and capacity steering depending on the load in the network

728 and 709, to GPON port 1 and 2, respectively. Note that we put a splitter between each ONT and the corresponding GPON port. Two laptops are connected to the same switch, and the switch is connected to ethernet port 1 and port 2 of ONT 1 (Ethernet port 1-1 and 1-2, respectively), and port 1 of ONT 2 (Ethernet port 2-1), as illustrated in Figure 3. From each laptop, traffic can flow either through Ethernet port 1-1, Ethernet port 1-2 or Ethernet port 2-1, using the L2 unmanaged switch. Note that the switch is used only to offer different paths for each laptop and has no active function. In the following, we present and discuss the experiments that showcase capacity steering, service dimensioning and dynamic service re-dimensioning.

Dynamic Capacity Steering: To illustrate the dynamic bandwidth management and capacity steering, we run the first set of experiments. We launch FTP downloading sessions from Laptop 1 and 2 and observe the traffic flows at the GPON ports, and the ethernet ports to which the switch is connected. Note the upper bound α and lower bound β are set to 4 Mbps and 1 Mbps respectively. The results are shown in Figure 4. We gradually increase the downloading rate in the two laptops over time. At the beginning, all traffic is routed through the GPON port 1 (see time ≤ 260 s). As traffic load increases and reaches the upper threshold α (see time 260-270 s), our algorithm switches the traffic of Laptop 1 to

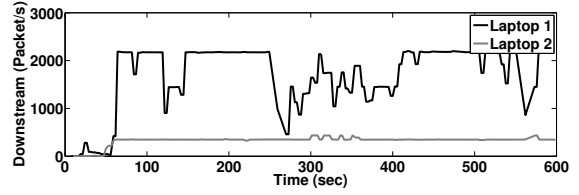


Figure 5: Service dimensioning and bandwidth guarantees

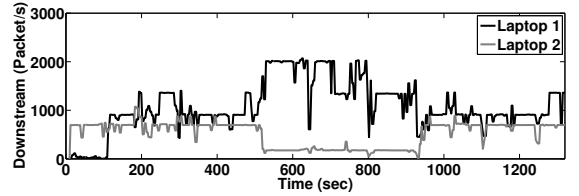


Figure 6: Dynamic service re-dimensioning

flow through GPON port 2 (see time 270-400 s). When one of the GPON ports is underutilized (see GPON port 2 at time 400-450 s), its traffic is consolidated with the traffic flowing through the second GPON port (GPON port 1, time 450-610 s).

Service and Flow Dimensioning: In the second set of experiments, we illustrate how our framework provides bandwidth guarantees and service dimensioning. To do so, we run two different services (Service 1 and Service 2) from the two laptops (Laptop 1 and Laptop 2, respectively). Service 2 is set to have a guaranteed bit rate. Service 1 is assumed to be of variable bit rate. Note that we use in this case the same ONT for both laptops, i.e., the two laptops use the same GPON port. We vary the traffic rate from Laptop 1 over time and monitor the bandwidth at Laptop 2. The results are plotted in Figure 5. The figure shows a guaranteed bandwidth for Laptop 2 even though Laptop 1 increases its traffic rate. However, note that the traffic at Laptop 1 is bounded to prevent it from disrupting the traffic of Laptop 2.

Dynamic Service Re-dimensioning: Let us now illustrate the dynamic service re-dimensioning. To do so, we run two different services (Service 1 and Service 2) from the two laptops (Laptop 1 and Laptop 2). Service 1 is set to have higher priority than Service 2. Similar to the previous experiment, the two laptops use the same GPON port. We vary the bandwidth demand from Laptop 1 over time, while keeping the traffic in Laptop 2 at a fixed rate. The results are plotted in Figure 6. From the figure, we note that the traffic in Laptop 2 is reduced as traffic in Laptop 1 increases. This is done in our framework by dynamically adjusting the committed rate of Service 2 to cope with the increase in traffic of Service 1 as they share the same network resources.

5 Ongoing Challenges: End-to-End Manageability with SDN

In our work, we presented dynamic flow management through steering, dimensioning and realtime re-dimensioning. A flow is defined by the Subscriber VLAN, 802.1p bits and/or DSCP, as defined in the GPON standard. Each flow is then mapped into a single GEM port, with a CoS. As such, a flow defined in our GPON is the aggregation of multiple flows, where each of them is a single flow defined in OpenFlow standard. One difficulty is that our current definition of a flow, using GPON standard, is coarse grained. In fact, to get a 1:1 mapping of our GPON flows and the flow definition of OpenFlow, we need a more detailed inspection of the Ethernet frames, by looking at L3 and L4 headers, at the ONT for upstream and OLT for downstream.

Currently, in our effort towards a full End-to-End SDN manageability, we embrace the incremental SDN deployment and integration in networks [18]. To do so, we define a mapping of the aggregate flows to integrate the SDEN into an SDN managed core network. The traffic steering, service dimensioning and dynamic service re-dimensioning offer a considerable agility and manageability in PON, without using active equipment at the edge, nor modification in the PON equipment design. In fact, this allows us to direct capacity dynamically based on the traffic in the network, assuring premium traffic and guaranteed QoS. However, defining fine grained flows and match the same definition of flows as OpenFlow, is one of our key challenges to address.

On the other hand, we are actively investigating the implementation of OpenFlow in SDEN. One promising track is to define a mapping of OpenFlow messages and commands into PON commands. Another option to build plugins for SDN controllers such as OpenDaylight, without going through OpenFlow, is also under investigation.

6 Conclusion

In this paper, we investigated the implementation of software defined control and management in enterprise PON. Our objective is to introduce agility, flexibility and dynamic adaptation to the PON network in order to meet the new enterprise business requirements, and to cope with traffic dynamicity introduced by Cloud and mobile. To this end, we proposed the SDEN framework, discussed how traffic steering, dimensioning and re-dimensioning can be achieved in PON. Through experiments conducted on a live PON testbed network, we demonstrated the practicality and potential benefits software-defined control brings to enterprise PON. Based on what we have learned and established in this

paper, we are actively investigating the integration of SDEN agent with SDN controller and the translation of OpenFlow protocol to PON.

References

- [1] Juniper Networks. Scaling the new service provider network at the edge. *Juniper Networks Whitepaper*, 2013.
- [2] ITU-T Gigabit capable Passive Optical Networks (GPON). General characteristics. *ITU-T G.984.1*, 2008.
- [3] ITU-T Gigabit capable Passive Optical Networks (GPON). Transmission convergence layer specification. *ITU-T G.984.3*, 2008.
- [4] ITU-T Gigabit capable Passive Optical Networks (GPON). Ont management and control interface specification. *ITU-T G.984.4*, 2008.
- [5] IBM Global Technology Services. Smarter networks with passive optical lans. *IBM Whitepaper*, 2013.
- [6] Jason Young. Comparative Risk Analysis Between GPON Optical LAN and Traditional LAN Technologies. *White Paper, SANS Institute InfoSec Reading Room*, October 2013.
- [7] Nick McKeown, Tom Anderson, Hari Balakrishnan, Guru Parulkar, Larry Peterson, Jennifer Rexford, Scott Shenker, and Jonathan Turner. Openflow: Enabling innovation in campus networks. *SIGCOMM Comput. Commun. Rev.*, 38(2), March 2008.
- [8] L.E. Li, Z.M. Mao, and J. Rexford. Toward software-defined cellular networks. In *European Workshop on Software Defined Networking (EWSN)*, pages 7–12, 2012.
- [9] Julius Schulz-Zander, Lalith Suresh, Nadi Sarrar, Anja Feldmann, Thomas Hühn, and Ruben Merz. Programmatic orchestration of wifi networks. In *2014 USENIX Annual Technical Conference (USENIX ATC 14)*, pages 347–358, Philadelphia, PA, June 2014. USENIX Association.
- [10] Lalith Suresh, Julius Schulz-Zander, Ruben Merz, Anja Feldmann, and Teresa Vazao. Towards programmable enterprise WLANS with Odin. In *Proceedings of the first workshop on Hot topics in software defined networks, HotSDN*, pages 115–120, 2012.
- [11] P. Dely, A. Kassler, and N. Bayer. OpenFlow for Wireless Mesh Networks. In *Proceedings of 20th International Conference on Computer Communications and Networks (ICCCN)*, pages 1–6, Aug 2011.
- [12] ONF. SDN in the Campus Environment. *ONF Solution Brief*, October 2013.
- [13] The ONF Optical Transport Working Group (OTWG), <http://bit.ly/1mRiYVE>.
- [14] P. Parol and M. Pawlowski. Towards networks of the future: Sdn paradigm introduction to pon networking for business applications. In *Federated Conference on Computer Science and Information Systems (FedCSIS)*, pages 829–836, Sept 2013.
- [15] S. Gringeri, N. Bitar, and T.J. Xia. Extending software defined network principles to include optical transport. *IEEE Communications Magazine*, 51(3):32–40, March 2013.
- [16] Pawel Parol and Michal Pawlowski. Towards networks of the future: Sdn paradigm introduction to pon networking for business applications. *Federated Conference on Computer Science and Information Systems (FedCSIS)*, 2013.
- [17] H. Woesner and D. Fritzsche. Sdn and openflow for converged access/aggregation networks. In *Optical Fiber Communication Conference and Exposition and the National Fiber Optic Engineers Conference (OFC/NFOEC)*, pages 1–3, March 2013.
- [18] Dan Levin, Marco Canini, Stefan Schmid, Fabian Schaffert, and Anja Feldmann. Panopticon: Reaping the benefits of incremental sdn deployment in enterprise networks. In *2014 USENIX Annual Technical Conference (USENIX ATC 14)*, pages 333–345, Philadelphia, PA, June 2014. USENIX Association.