Dynamic Capacity Management and Traffic Steering in Enterprise Passive Optical Networks

Ahmed Amokrane*, Jin Xiao†, Jinho Hwang†, Nikos Anerousis†
*LIP6 / UPMC - University of Paris 6; 4 Place Jussieu, 75005 Paris, France
†IBM T.J. Watson Research Center, Yorktown Heights, NY, USA
ahmed.amokrane@lip6.fr, {jinoaix, jinho, nikos}@us.ibm.com

***Abstract*—In the last few years, changing infrastructure and business requirements are forcing enterprises to rethink their networks. Enterprises look for network infrastructures that increase network efficiency, flexibility, and cost reduction. At the same time, the emergence of Cloud and mobile in enterprise networks has introduced tremendous variability in enterprise traffic patterns at the edge. This highly mobile and dynamic traffic presents a need for dynamic capacity management and adaptive traffic steering and appeals for new infrastructures and management solutions. In this context, passive optical networks (PON) have gained attention in the last few years as a promising solution for enterprise networks, as it can offer efficiency, security, and cost reduction. However, network management in PON is not yet automated and needs humain intervention. As such, capabilities for dynamic and adaptive PON are necessary.**

In this paper, we present a joint solution for PON capacity management both in deployment and in operation, as to maximize peak load tolerance by dynamically allocating capacity to fit varying and migratory traffic loads. To this end, we developed the novel approaches of capacity pool based deployment and dynamic traffic steering in PON. Compared with traditional edge network design, our approach significantly reduces the need for capacity over-provisioning. Compared with generic PON networks, our approach enables dynamic traffic steering through software-defined control. We implemented our design on a production grade PON testbed, and the results demonstrate the feasibility and flexibility of our approach.

I. INTRODUCTION

Enterprise networks today are under tremendous pressure to change. A recent study by the Economist Intelligence Unit Research, sponsored by Juniper Networks, reported 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]. One of the main problems faced by enterprise infrastructures today is that they are engineered for specific workload distributions and this rigidity cannot cope with disruptive yet business vital technologies such as Cloud and mobile. In fact, with increased reliance on Cloud and mobile technologies, enterprise network traffic patterns today are far more dynamic and nomadic. Moreover, Cloud workloads vary significantly over time and mobiles' traffic is migratory and volatile. Taking the stadium enterprise as an example, the traffic patterns exhibited during a game day are highly dependent on the phase and condition of the game. Before the game starts, the majority of the traffic comes from the gates entrance; during the game periods, it is concentrated at the seating areas. During the half-time, the concentration shifts to the concourse; and after the game, it migrates to the parking lots. The transitory traffic volume and burst intensity are also highly related to changes in the game states (e.g., a remarkable touchdown scored by the home team is likely to trigger a large surge in mobile traffic). To handle these traffic load shifts, traditional network designs use over-provisioning in different areas of the network to handle the peak traffic load. This often results in costly static solutions that inevitably fail as workloads evolve or unusual events occur. As such, a more dynamic and flexible network design is needed in which capacity can be dynamically allocated and steered based on ever-changing traffic load and distribution.

In this context, Passive Optical Networks (PON) is a promising technology that contains the key ingredients required to address the new infrastructure requirements. As such, PON is becoming an attractive fiber-based LAN edge network solution for enterprises. Some of the key benefits that 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. However, nowadays PON products are still not automated and need human intervention for management. Consequently, the new enterprise network requirements appeal for automated and smart network management through software-enabled control.

In our work towards enabling software-defined control of PON at IBM Research, we present in this paper a novel approach for dynamic capacity management of PON that achieves dynamic capacity allocation depending on the traffic load in the network. More specifically, we first formulate the solution as a joint PON deployment and online traffic steering problem in PON enterprise network. As the formulated problem is $\mathcal{NP}$-Hard, we propose a heuristic algorithm that achieves capacity steering in response to traffic surges and mobility in the network. This work is a follow-up of our previous work [2]. More specifically, we showed in [2] the feasibility of achieving programmability in PON networks. In this paper, we present an approach for dynamic capacity steering in PON that can handle variable traffic patterns in real-time. More specifically, the contributions of this paper are as follows:

- **Capacity pool based deployment of PON**: we explore the idea of deploying PON networks based on capacity pools. We define a capacity pool as a set of network capacity resources (e.g., ports and fibers) shareable across different areas covered by the network. Analogous to systems design, one can imagine a capacity pool as a pool of CPUs dedicated to a set of end users, and the association of CPUs to users is determined at runtime depending on their workload requirements. In this fashion, network resource over-provisioning is reduced and local peak surges can be tolerated at a higher level. Moreover, as capacity pools are realized through a crisscross design of the network physical connectivity, the network is more resilient as an area has multiple disjoint paths between the end devices and the core.

- **Dynamic traffic steering in response to surges in traffic in the network**: We propose a reactive algorithm for traffic steering in the PON network to accommodate surges in traffic in some parts of the network, analogous to runtime association of CPUs to end users based on their workload requirements. The traffic steering utilizes the multi-path connectivity in a capacity pool (both fibers and Optical Network Terminals) to make use of the shared capacity in the network.
- **SDN integration for PON management in an Enterprise environment**: we propose an architecture that integrates GPON management with an SDN controller. This integration is a precursor towards a full end-to-end management in enterprise networks that includes the core network.
- **Real deployment in a PON testbed**: we deployed our proposed solution in a real PON testbed in our lab at IBM T.J. Watson Research Center, New York. The deployment includes a central Java code module for dynamic capacity and traffic steering in the network.

We further note that theoretically, it is possible to have a single capacity pool for the entire enterprise for maximum flexibility in capacity management. However, this design will require a very large number of fibers and ports to form a complete graph of crisscrossed physical connectivity, a costly and complex solution. Therefore, an interesting and nontrivial problem we resolve in this paper is in determining how many capacity pools are needed and which areas each capacity pool should cover, depending on expected mobility patterns and load density distribution in a specific enterprise environment.

The remainder of this paper is organized as follows. Section II presents PON and SDN technologies and related works in the literature. Section III overviews the software-defined edge network architecture. Section IV presents the mathematical problem formulation of the problem of PON deployment and dynamic traffic steering, followed by the our proposed heuristic algorithms in Section V and experimental results in Section VI. Finally, Section VII concludes the paper.

II. BACKGROUND AND RELATED WORK

A. *Passive Optical Networks*

A Passive Optical Network (PON) consists of a set of Optical Network Terminals (ONTs), passive splitters and the Optical Line Terminal (OLT). The ONTs connect edge devices to the PON network via Ethernet ports, also called User Network Interfaces (UNI). Digital signals from edge devices are converted to optical signals 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 networks, PON replaces switches in the access and aggregation layers with splitters, and the traditional distribution layer is collapsed back to a few OLTs. 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-capable PON (GPON) standards [3]–[5]. We therefore use PON and GPON interchangeably in this paper.

Compared to traditional copper networks, PON has a number of salient advantages. The optical fibers in PON can travel up to 20 Km from the core to the access, capable of delivering 1.2 Gbps upstream and 2.4 Gbps downstream to the port in current generation, and the fiber is much lighter than copper cables. Moreover, PON eliminates active equipment in the distribution resulting in significant capex and opex savings (up to 40% and 60% respective savings compared to traditional enterprise copper networks [6]). Furthermore, PON offers much stronger security with enhanced data encryption and physical protection [7]. More details on the enterprise PON technology and its benefits over traditional copper networks can be found in [6].

The entire PON network constitutes an Ethernet LAN. Between the ONT and the OLT, Ethernet frames are encapsulated in GPON Encapsulation Method (GEM) frames, which are encapsulated in GPON 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 the Physical Port of the ONT, the subscriber VLAN ID, the 802.1p bits and/or DSCP, as defined by the ITU-T GPON standard. Then, each of these N-VLAN and CoS combinations 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

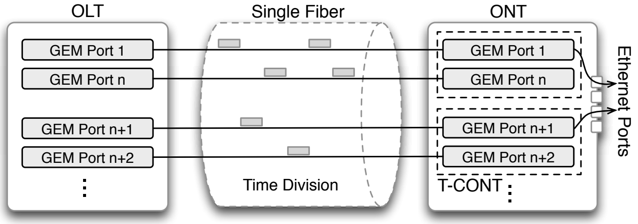

Fig. 1. Traffic management in GPON networks

the traffic to the appropriate GEM port.

B. 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 [8], for communication between the controller and the forwarding elements in the network. SDN has first been proposed for data center networks, with mainly an Ethernet copper-based switching fabric. As SDN offers flexibility, manageability and agility, a number of proposals advocated to extend SDN for wireless networks such as cellular networks [9], WLANs wifi-based networks [10], [11], wireless mesh networks [12] and campus copper-based networks [13]. Moreover, one active and interesting effort is to extend SDN to optical networks [14]–[16]. 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 Open Networking Foundation (ONF) created The Optical Transport Working Group (OTWG) [14]. 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. [16] 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. [17]. In this work, authors proposed extensions to OpenFlow protocol, which consist mainly on mapping flows (as defined by the OpenFlow standard) 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 [18], 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. PON has also been studied in the context of intra-data center networks. For instance, Gu et al. [19] proposed an SDN oriented architecture for data center networks to replace the switches at the edge and aggregation layers. More specifically, they propose to add SDN capabilities to the deployed ONT and introduce a Top of Rack (ToR) device to each wavelength allocation for the different ONTs. Moreover, network coding is used to further increase the network capacity. Such a design enables flexibility in the wavelength allocation, which can be leveraged to achieve high bandwidth utilization in the fibers. However, their proposed architecture introduces changes to the current PON networks on the market. In our case, we propose an architecture that takes into account the current PON implementations of the standard without introducing any changes. As our proposal is oriented towards deployment with nowadays available PON products, it is designed as a plug and play solution that can be easily integrated to any PON. In our previous work in [2], we presented the design of Software Defined Edge Network (SDEN) that brings agility and programmability to GPON networks. In this paper, we extend this design by presenting an approach for dynamic traffic steering in response to traffic load in the network. To do so, we reuse the key enabling points defined in [2] which are related to traffic steering, service dimensioning and realtime re-dimensioning.

III. SOFTWARE DEFINED EDGE NETWORK AND THE SUPPORTING MECHANISMS

A. Software Defined Edge Network (SDEN) Architecture

In our design and implementation of PON enterprise networks, we proposed SDEN in [2]. In this paper, we build our approach for dynamic traffic steering on top of the SDEN framework. The SDEN architecture is illustrated in Figure 2. It defines a common interface through APIs between the controller and the PON nodes (OLT). As PON does not support OpenFlow currently, our defined interface is used as a bidirectional high level interface. More specifically, the interface provides a standardizable and vendor neutral set of functionalities that a controller can use. On top of the controller, one can have different applications for network management and optimization. For instance, in this paper, we implemented our application that performs dynamic traffic and capacity steering in PON. As such, the network statistics and status are pulled from the application through the SDEN Agent API. Similarly, the decision on traffic routing and capacity steering computed by the application are passed also through the SDEN Agent to the OLT. However, as PON is a Layer 2 network, the traffic contra and steering is performed at a coarse-grained level by aggregating multiple Layer 3 and above flows.

B. Traffic Steering Enabling Mechanisms

In PON, each OLT is connected to ONTs through a passive splitter. Each ONT has a set of Ethernet ports, or User Network Interfaces (UNI). The OLT management entity enables the

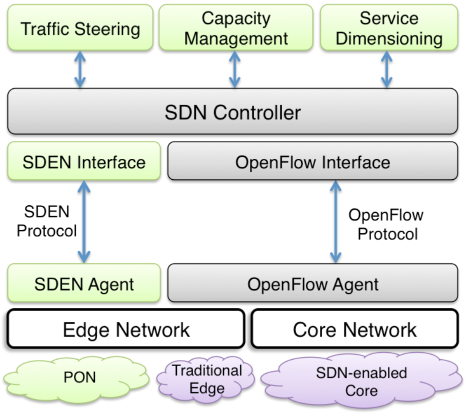

Fig. 2. Software Defined Edge Network Extension with PON

Ethernet ports to allow traffic to flow from the UNI to the core network, and achieves user traffic mapping into VLANs managed in the PON. More specifically, a Network Access Control (NAC) profile is assigned to each Ethernet port. The NAC profile defines among others, a service profile associated with it. Each service profile describes how traffic is tagged (associated VLAN tag), the committed and peak rates, and the traffic priority. Multiple service profiles can be associated with a single Ethernet port. In this case, user VLAN tags are used to identify the traffic belong to a specific service profile. For instance, an IP phone can tag the traffic and the regular laptops do not tag traffic. As such, the tagged traffic is mapped to the QoS guaranteed traffic in the PON, whereas the untagged traffic is mapped to a best effort service.

In our implementation, we use NAC-to-Ethernet port assignment to dynamically steer traffic. As PON is an L2 network, switching a user equipment from one Ethernet port to another by changing the NACs does not affect L3 (i.e., IP address) or high layer network constructs. Moreover, we enable/disable Ethernet ports to block or allow traffic to enter the PON.

IV. PROBLEM FORMULATION

In this section, we present the mathematical formulation of the problem of joint deployment and traffic steering of an enterprise network. We also present the specific case of traffic steering only, where the PON network is already deployed.

A. Joint Deployment and Traffic Steering in PON Enterprise networks

In this subsection we present the model for deploying the PON network to enable the creation of capacity pools and allow for dynamic traffic steering. We formulate the problem as an ILP where:

GIVEN

- The physical location of the different areas
- The estimation of the peak bandwidth demand in each area over time

FIND

- The number of needed GPON ports and ONTs
- The physical wiring of the ONTs to GPON ports
- Direct enough capacity to each area to route its peak traffic demand

TABLE I
TABLE OF NOTATIONS

Notation	Meaning
A	Set of areas
N	Set of ONTs
P	Set of GPON ports
T	Set of time slots
$D_{r,t}$	Traffic demand in area $r \in A$ during time slot t
$u_{r,j,t}$	Defines whether ONT j is used to route the traffic of Area r during time slot t
$x_{j,k}$	Defines whether ONT j is attached to GPON port k
$d_{k,r}$	Defines whether GPON port k can be assigned to area r
y_j	Defines whether an ONT j is deployed or not
$b_{j,r}$	Defines whether an ONT j is deployed in area r
z_k	Defines whether a GPON port k is used or not
C_k	The bandwidth capacity of GPON port k
S_k	The splitting ratio of GPON port k (1,2,8,16,32,64)
$v_{k,t}$	Defines whether GPON port k is used during time slot t

We assume the network is divided into areas. Such a division into areas of density to ensure coverage are one of the design recommendations of vendors. For instance, Cisco proposes deployment guidelines for WLAN to accommodate conference venues [20]. Research proposals also suggested dividing the network into areas [21]. Note that our work is not intended to design the deployment inside each area, and our focus is designed the PON backbone network for traffic routing in the network. For some examples and strategies for deploying wifi APs in each area, please refer to documents such as [20], or existing tools for planning such as WiROI [22]. Note that the problem of WAP-to-ONT association is solved by using L2 unmanaged switches. In fact, the number of WAPs per area can easily be wired through switches to offer the possibility to dynamically route the traffic and choose which ONT ports to use. For simplicity of the problem definition, it is not included in this paper. Again, the physical wiring of the switches to ONTs and WAPs to ONT is out of scope of this paper. However, we note that the deployment guarantees enough capacity for such wiring.

In the following, we present the mathematical formulation of the problem. For ease of understanding, we provide in Table I the list of used notations in the problem formulation.

The network is divided into areas A. We assume that time is divided into time slots, and during each time slot $t \in T$, the traffic load of each area is fixed. At the end of each time slot t, the traffic demand in each area might change. As this change might be due to density shift or additional users joining the network, we define the traffic demand in area $r \in A$ as $D_{r,t}$.

The network is composed of a set of n ONTs denoted by N, and a set of m GPON ports denoted by P. Each of the ONTs is connected to one and only one GPON port $k \in P$ through a splitter, where S_k is the splitting ratio of the splitter connected to GPON port k, and C_k is the bandwidth capacity of GPON port k.

The problem in the deployment phase is to find the minimum number of ONTs and GPON ports that are necessary to deploy and the association of ONT-GPON port in such a way to guarantee enough capacity for the peak demand in each area. The capacity steering consists of assigning GPON ports to areas in such a way to accommodate the traffic demand from each area during each time slot t.

We define also two decision variables:

$$u_{r,j,t} = \begin{cases} 1 & \text{If ONT } j \text{ is used to route the traffic} \\ & \text{of Area } r \text{ during time slot } t \\ 0 & \text{Otherwise.} \end{cases} \tag{1}$$

$$x_{j,k} = \begin{cases} 1 & \text{If ONT } j \text{ is attached to GPON port } k \\ 0 & \text{Otherwise.} \end{cases} \tag{2}$$

The objective is to minimize the number of deployed ONTs and used GPON ports while guaranteeing the traffic is drained. We define whether an ONT j is deployed or not by the following variable y_i:

$$y_j = \begin{cases} 1 & \text{If } \sum_{t \in T} \sum_{r \in A} u_{r,j,t} \geq 1 \\ 0 & \text{Otherwise.} \end{cases} \tag{3}$$

We define also the variable $b_{j,r}$ which determines whether an ONT is deployed in an area $r \in A$. It is defined as follows:

$$b_{j,r} = \begin{cases} 1 & \text{If } \sum_{t \in T} u_{r,j,t} \geq 1 \\ 0 & \text{Otherwise.} \end{cases} \tag{4}$$

We define $d_{k,r}$ that determines whether GPON port k can be assigned to area r, through a deployed ONT that is attached to GPON port k.

$$d_{k,r} = \begin{cases} 1 & \text{If } \sum_{j \in N} b_{j,r} \times x_{j,k} \geq 1 \\ 0 & \text{Otherwise.} \end{cases} \tag{5}$$

Similarly, we define the variable z_k that determine whether a GPON port $k \in P$ is used or not as follows:

$$z_k = \begin{cases} 1 & \text{If } \sum_{j \in N} x_{j,k} \times y_j \geq 1 \\ 0 & \text{Otherwise.} \end{cases} \tag{6}$$

The objective function is defined as follows:

$$Minimize\Big(\beta_{pon} \sum_{k \in P} z_k + \beta_{ont} \sum_{j \in N} y_j\Big) \tag{7}$$

where β_{pon} and β_{ont} are the costs of deploying a PON port and an ONT, respectively.

This problem is subject to the following constraints:

- An ONT is deployed in one and only one area:
$$\sum_{r \in A} \sum_{j \in N} b_{j,r} \leq 1 \tag{8}$$
- The number of ONTs that can be assigned to a single GPON port are limited by the splitting ratio of the GPON network:
$$\sum_{j \in N} x_{j,k} \leq S_k, \forall k \in P \tag{9}$$
where S_k is the maximum number of ONTs that might be connected to the GPON port k.
- For an Area $r \in A$ that has traffic to transmit during time slot t, there are enough paths towards GPON ports
$$\sum_{k \in P} d_{k,r} \times C_k \geq D_{r,t}, \forall t \in T \tag{10}$$
where C_k is the capacity of GPON port k.
- The capacities of the GPON ports should not be exceeded:
$$\sum_{r \in A} \sum_{j \in N} u_{r,j,t} \times D_{r,t} \times x_{j,k} \leq C_k, \forall t \in T \tag{11}$$

Note that this deployment problem is solved only at the deployment stage. It would be reasonable to use any solver or brute force exhaustive search to solve it. However, the traffic and capacity steering is dynamic and is achieved in response to the traffic pattern change in the network. Therefore, a non-time consuming solution is needed, which we present in the next section.

B. Traffic Steering in PON Enterprise networks

The joint deployment and traffic string presented in section IV-A is solved at the deployment phase. During operations, as the traffic patterns might change, the traffic should be steering accordingly, given the fixed deployment topology. We present in this section the modifications on the formulation in section IV-A to address traffic steering only.

As the deployment is already done, the variables $x_{j,k}$, $b_{j,r}$ and $d_{k,r}$ become input variables to the ILP. The only decision variable to determine is $u_{r,j,t}$, which defines which ONT is used to route the traffic from each area. Consequently, we modify the objective function de take into account the operational costs. In this case, our objective is to minimize the number of used ONTs and GPON ports when routing the traffic. The objective function is defined as follows:

$$Minimize\Big(\alpha_{ont} \times (\sum_{t \in T} \sum_{i \in A} u_{r,j,t}) + \alpha_{pon} \times (\sum_{k \in P} \sum_{t \in T} v_{k,t})\Big) \tag{12}$$

where α_{pon} and α_{not} are the costs of operating a PON port and an ONT during one time slot t, respectively, and $v_{k,t}$ is a binary variable that determines whether PON port k is used during a time slot t, and defined as:

$$v_{k,t} = \begin{cases} 1 & \text{If } \sum_{r \in A} \sum_{j \in N} u_{r,j,t} \times x_{j,k} \geq 1 \\ 0 & \text{Otherwise.} \end{cases} \tag{13}$$

The problem formulated in this section is at least a multi commodity flow problem (MCF). In fact, the formulation in section IV-B is an MCF. As a result, under large network deployments, finding the optimal combination of paths allocation is computationally expensive.

V. DYNAMIC TRAFFIC STEERING IN CAPACITY POOL ENABLED PON

The traffic in the network is dynamic and subject to end users' mobility. As the network experiences traffic load shift across areas, capacity should be dynamically directed to the overloaded areas. Moreover, the network should be reactive to the load shift. In addition to capacity steering, our approach integrates service dimensioning. In fact, under heavy traffic demands, low priority services are allocated less resources in the network. On the other hand, more resources are allocated to services with higher priority. From the implementation point of

Algorithm 1 Continuous Network Monitoring

```
while $true$ do
    $underUtilized \leftarrow \{p \in P, p.utilization < \gamma_{low}\}$
    $overUtilized \leftarrow \{p \in P, p.utilization > \gamma_{up}\}$
    $other \leftarrow \{p \in P, p \notin underUtilized \cup overUtilized\}$
    if $overUtilized.size > 0$ or  then
        Reconfigure the network by allocating the GPON ports to areas using Algorithm 2
    end if
end while
```

Algorithm 2 Greedy Dynamic Capacity Steering

```
1:  IN: $A$ The set of areas with traffic demands to route, each area has premium and best effort traffic demands
2:  IN: Network deployed topology (GPON ports: $P$, ONTs : $N$)
3:  OUT: Ethernet port, ONT and GPON port allocation to each area
4:  while $A \neq \phi$ and $\exists\ p \in P, p.residual > 0$ do
5:      $d \leftarrow 0$
6:      if $\exists a \in A, a.demand.premium > 0$ then
7:          $a \leftarrow$ any area in $A$ with $a.demand.premium > 0$
8:          $d \leftarrow a.demand.premium$
9:      else
10:         $a \leftarrow$ any area in $A$ with $a.demand.besteffort > 0$
11:         $d \leftarrow a.demand.besteffort$
12:     end if
13:     $p \leftarrow$ a port in $P$ with $p.residual > 0$, preferably $p$ is already in use (consolidation) and in a best fit manner
14:     Open enough available Ethernet ports in area $a$ connected to GPON port $p$, to drain the traffic demand $d$
15:     $done \leftarrow$ total routed demand out of $d$ by opening available Ethernet ports
16:     $p.residual \leftarrow p.residual - done$
17:     if $done == d$ then
18:         $A.remove(a)$
19:     else
20:         if $a.demand.premium > 0$ then
21:             $a.demand.premium \leftarrow a.demand.premium - done$
22:         else
23:             $a.demand.besteffort \leftarrow a.demand.besteffort - done$
24:         end if
25:     end if
26: end while
```

view, SDEN relies on Ethernet port allocation in each chosen ONT to serve an area. More specifically, a NAC profile is assigned to each Ethernet port. The NAC profile defines a service profile associated with it. Each service profile defines by itself how traffic is tagged (associated VLAN tag), the committed and peak rates and the traffic priority. For service dimensioning, we use the committed rates of services to adjust the amount of allocated resources in the network.

As the problem defined in the previous section is $\mathcal{NP}$-Hard, it becomes time prohibitive to solve. More specifically, in case of large networks and high traffic demand variability, the ILP should be solved very often, which can result in high computation overhead and slow response time to network changes. Consequently, a fast algorithm is needed. To do so, we propose a heuristic approach to find a feasible solution to allocated capacity to different areas in a reasonable time. First, we perform a continuous monitoring of the GPON ports utilization and traffic demands in areas. We define two thresholds for the GPON ports utilization: γ_{low} and γ_{up}, which determine the lower and upper bounds for under-utilized and over-utilized GPON ports. As such, a GPON port is considered under-utilized if its utilization is over γ_{up}, and under-utilized if its utilization is bellow γ_{low}. The algorithm for continuous monitoring is provided in Algorithm 1.

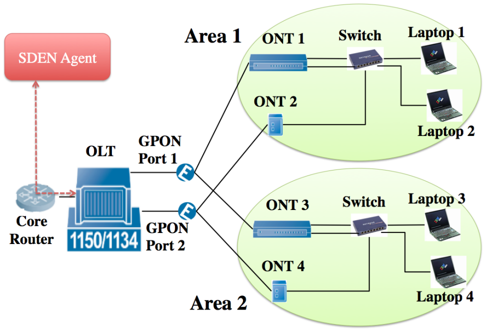

Fig. 3. Experimental setup

Then, we propose an algorithm for capacity allocation given in Algorithm 2. Each area has traffic demand. We assume there are two classes of services: premium and best effort. We incrementally open Ethernet ports in areas to route traffic by allocating available capacities in the GPON ports. More specifically, unused Ethernet ports are opened in the ONTs deployed in these areas. Note that we allocated the capacity by taking into account the capacity of the Ethernet and GPON ports. Moreover, we prioritize the premium services over the best effort ones (see lines 6-12). We also allocated the GPON ports to the areas in a best fit manner. In other words, for each area with traffic demand to route, we choose an already opened GPON port that has enough residual capacity (see line 13). This will allow for traffic consolidation and open reduction as provided in objective function in Eq. 12.

Note that the advantage of this approach is that it completely transparent from an end user perspective. In fact, the entire PON is L2, and the IP address of the device remains the same while moving across areas.

VI. PERFORMANCE EVALUATION

In this section, we present the results of our experiments in a real testbed deployed at IBM T.J. Watson, NY. More specifically, we study the viability of our dynamic capacity steering in managing mobile and dynamic traffic loads. To this end, a network topology is created as described below.

For our test scenario, we created two areas, Area 1 and Area 2 under the same capacity pool. A crisscrossed physical deployment allows sharing two GPON ports (GPON port 1 and GPON port 2) between the two areas. Two ONTs are deployed in each area, and each of these two is connected to a different GPON port. We use two laptops in each area to generate the traffic. The testing setup is illustrated in Figure 3. In what follows we show how our approach achieves dynamic capacity steering to allocate capacity to areas depending on the traffic load.

We developed a software service in Java that monitors the link utilization in the network. It runs Algorithm 1 and Algorithm 2. We set the threshold γ_{low} and γ_{up} to 2 $MBps$ and 12 $MBps$, respectively. Moreover, we scaled down the capacity of the GPON ports to 16 $MBps$ in order to generate traffic overload conditions in our test environment. The traffic

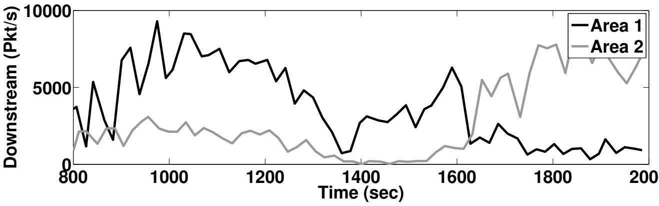

(a) Downstream traffic at the areas

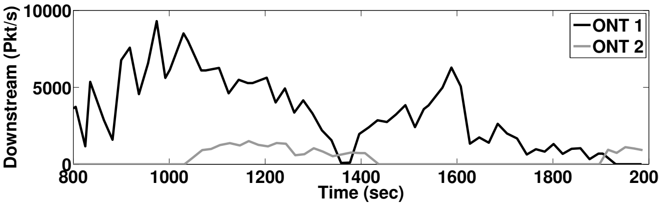

(b) Downstream traffic at the ONTs in Area 1

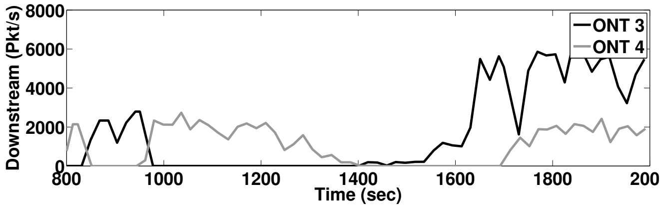

(c) Downstream traffic at the ONTs in Area 2

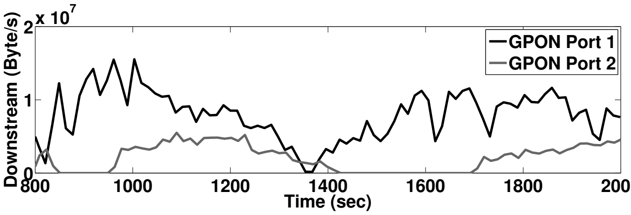

(d) Downstream traffic at the GPON ports

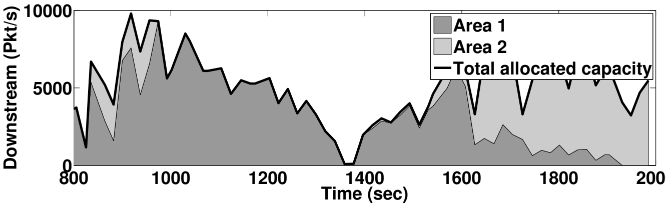

(e) Capacity allocation of GPON port 1 to the different areas

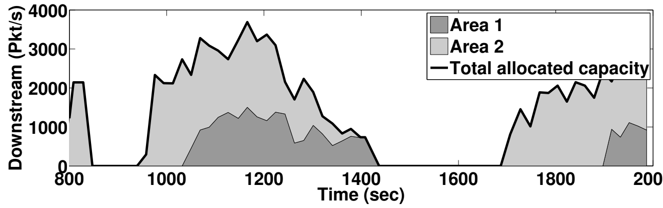

(f) Capacity allocation of GPON port 2 to the different areas

Fig. 4. Dynamic traffic and capacity steering depending on the network load

load generated by the laptops in the two areas varies over time as shown in Figure 4(a). We show also the traffic routing through the two ONTs in each area in Figures 4(b) and 4(c). The traffic load on the two GPON ports is shown in Figure 4(d), and the capacity allocation of the GPON ports to the different areas is shown in Figures 4(e) and 4(f). Note that we plot the interval of time we are interested in ($t \in [800, 2000]$), where we observed traffic variations. As we can observe, GPON port 1 is the primary carrier while GPON port 2 is the secondary carrier. When the traffic load on GPON port 1 threatens to overload, GPON port 2 is active to offload some of the traffic. Furthermore, we observe that Area 1 starts with heavy traffic load but peters out over time, while Area 2's load increases over time. It describes real life scenarios (e.g., stadium during a game) where high density crowd migrates over time. More specifically, for $t \in [800, 840]$, both GPON port 1 and GPON port 2 are used. However, these two ports experience under-utilization. As such, traffic is consolidated to GPON port 1 (see $t \in [840, 960]$). At $t = 960$, Area 1 experiences surge in traffic. In response, traffic of Area 2 is re-routed through GPON port 2. As such, GPON port 1 serves Area 1 exclusively and GPON port 2 serves Area 2. As the traffic surge continues in Area 1, more capacity in the GPON ports is needed. As such, at $t = 1030$, part of the traffic of Area 1 is routed through GPON port 1 and another part is routed through GPON port 2. In other words, more capacity is directed to Area 1 to cope with the surge in traffic. Again, traffic is consolidated to GPON port 1 only around $t = 1400$. At $t = 1620$, traffic starts to shift from Area 1 to Area 2. This triggers capacity allocation to Area 2. As such, both GPON ports are used to route the traffic of Area 2 (see $t \geq 1680$).

From the capex saving, note that a static PON over-provisioning scenario for this example would require 4 GPON ports (i.e., 2 GPON ports per area). Consequently, the dynamic capacity steering capability cuts the opex by around 50%. In comparison to copper networks, the resulting saving are around 80% given the results provided in [6].

On the other hand, our approach consolidates GPON cards when resource utilization is low in order to reduce operational costs such as energy. The fact that shared capacity can be dynamically allocated across areas when needed achieves both high local peak tolerance as well as high resource multiplex gain. As we have shown in our experiments, this design is effective in managing migratory and volatile traffic load at the edge (e.g., WAPs connected to ONTs rather than laptops). Our current implementation and deployment is quite limited in scale, and therefore we do not yet have results from large field deployment. It is an aspect we are actively working on.

VII. CONCLUSION

In this paper, we investigated the problem of PON deployment and dynamic capacity steering in enterprise networks. We propose a framework that introduces programmability and dynamic adaptation for GPON networks in response to traffic shifts in the network. We furthermore propose an ILP formulation for the problem of deployment and dynamic traffic steering and capacity allocation. To alleviate the time complexity of solving the ILP in large networks, we propose a heuristic algorithm that achieves traffic steering on existing GPON deployments. Through experimental study on a real testbed, we demonstrate the capacity steering using GPON port allocation to areas in the network depending on their traffic load. This shows capex reduction compared to static over-provisioning.

REFERENCES

[1] Juniper Networks. Scaling the new service provider network at the edge. *Juniper Networks Whitepaper*, 2013.

[2] Amokrane Ahmed, Hwang Jinho, Xiao Jin, and Anerousis Nikos. Software Defined Enterprise Passive Optical Networ. In *1st International Workshop on Management of SDN and NFV Systems, (ManSDN/NFV 2014)*, Rio de Janeiro, Brazil, November 2014.

[3] ITU-T Gigabit capable Passive Optical Networks (GPON). General characteristics. *ITU-T G.984.1*, 2008.

[4] ITU-T Gigabit capable Passive Optical Networks (GPON). Transmission convergence layer specification. *ITU-T G.984.3*, 2008.

[5] ITU-T Gigabit capable Passive Optical Networks (GPON). Ont management and control interface specification. *ITU-T G.984.4*, 2008.

[6] IBM Global Technology Services. Smarter networks with passive optical lans. *IBM Whitepaper*, 2013.

[7] Jason Young. Comparative Risk Analysis Between GPON Optical LAN and Traditional LAN Technologies. *White Paper, SANS Institute InfoSec Reading Room*, October 2013.

[8] 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.

[9] L.E. Li, Z.M. Mao, and J. Rexford. Toward software-defined cellular networks. In *European Workshop on Software Defined Networking (EWSDN)*, pages 7–12, 2012.

[10] 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.

[11] 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.

[12] 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.

[13] ONF. SDN in the Campus Environment. *ONF Solution Brief*, October 2013.

[14] The ONF Optical Transport Working Group (OTWG), http://bit.ly/1mRiYVE.

[15] 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.

[16] 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.

[17] 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.

[18] 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.

[19] Rentao Gu and Yuefeng Ji and Pei Wei and Shizong Zhang. Software defined flexible and efficient passive optical networks for intra-datacenter communications. *Optical Switching and Networking*, 14, Part 3(0):289 – 302, 2014.

[20] Florwick Jim, Whiteaker Jim, Cuelar Amrod Alan, and Woodhams Jake. Wireless lan design guide for high density client environments in higher education. *Cisco Design Guide*, November 2013.

[21] Jeffrey Erman and K.K. Ramakrishnan. Understanding the Super-sized Traffic of the Super Bowl. In *Proceedings of the 2013 Conference on Internet Measurement Conference*, IMC '13, pages 353–360, 2013.

[22] WiROI. WiROI Venue Tool, http://www.wireless2020.com/WiROIVenue/. 2013.