

# Joint Dynamic Bandwidth and Queue Management for Upstream SLA-Oriented QoS in Multi-Tenant, Multi-Service PONs

Frank Slyne<sup>1</sup> and Marco Ruffini<sup>2</sup>

**Abstract**—We devise a PON upstream scheduler that jointly manages intra-ONT and inter-ONT capacity across VNOs. This provides balancing between throughput and delay, providing SLA-oriented QoS supporting low latency 5G applications even in challenging multi-tenant environments.

## I. INTRODUCTION

Passive Optical Networks (PONs) were developed as a bandwidth upgrade to Asymmetric Digital Subscriber Line (ADSL) for residential broadband services. Due to their cost advantage as an optical access network and their ability to provide statistical multiplexing, they are being considered also to support new 5G-oriented services in multi-tenant environments, where strict Quality of Service (QoS) must be enforced. However, while ITU-T standards define the mechanism by which bandwidth can be apportioned, they do not specify how performance objectives such as low latency and low jitter can be attained. Performance objectives, typically, are fulfilled not through careful traffic planning, but through the over-provision of capacity. When implemented across the network, this approach is wasteful of resources and does not support the implementation of traffic profiles, a cause of revenue assurance issues. In this paper, we propose an architecture that integrates Optical Network Terminations (ONT) queue management with the upstream DBA mechanism, to provide support for new services in multi-tenant PONs. This approach is similar to that proposed in [5] but for downstream traffic. The architecture supports efficient bandwidth redistribution across ONTs and Virtual Network Operators (VNOs), as well as differentiation in Service Level Agreement (SLA)-oriented traffic profiles.

## II. BACKGROUND

Several ITU-T standards for PON define the mechanism by which ONTs communicate buffer occupancy on a periodic per-frame basis, and how the Optical Line Termination (OLT) arbitrates upstream transmission of the ONT Traffic Container (T-CONT) buffers using Bandwidth Maps (BWmaps). T-CONTs are the logical channels used to carry traffic upstream from the ONT to the OLT. The standards do not define the algorithm used to calculate the BWmap, or how objectives such as the optimisation of upstream throughput and delay can be achieved. In a PON system, the amount of bandwidth allocated and the frequency at which it is calculated, affect both throughput and delay. As the service

interval increases, traffic throughput will also increase since less bandwidth is being consumed by overheads such as Dynamic Bandwidth Allocation (DBA) messages and guard times. Unfortunately, packet delay also increases due to the increased time between allocations. Typically there are between 4 and 16 T-CONTs per ONT, each defined by a bandwidth type (1) Fixed, (2) Assured, (3) Non-Assured, and (4) Best Effort. For example, in a 70% loaded Gigabit Passive Optical Network (GPON) system, the average delay for type 2 and type 3 traffic is 0.5 ms, and is 1 ms for type 4 [1]. The alternative family of IEEE PON standards, Ethernet Passive Optical Network (EPON), has evolved a number of approaches to balance the trade-off between throughput and delay and which could be adapted for use in GPON. These include bandwidth redistribution, SLA partitioning and inter-ONT/intra-ONT scheduling [2]. Bandwidth redistribution has already been used to reduce variations in GPON average delay, because service intervals tend to be defined for each T-CONT [3], and because fixed type T-CONTs are allocated every downstream frame while flexible type T-CONTs (such as assured and non-assured types) are allocated according to a variable service interval. Partitioning reduces idle time, by processing Dynamic Bandwidth Report Units (DBRUs) from one group of T-CONTs while bandwidth for another group of T-CONTs is being allocated. The excess bandwidth from one group can then be added to the available bandwidth for the other group. Typically, type one and type two T-CONTs are grouped together, and type three and four are grouped together. Intra-ONT scheduling refers to scheduling between queues on the same ONT. While ITU-T recognises the principle when it is applied to traffic of different priorities in the same T-CONT, it does not mandate that it occurs between T-CONTs. This is because individual T-CONTs can be managed directly by the DBA process located at the OLT. In this instance, the only requirement is for the ONT to classify the traffic and route it into the appropriate T-CONT type [4].

## III. UPSTREAM SCHEDULER

In order to balance the trade-off between throughput and delay, we have devised an upstream scheduler that implements all the approaches described above, that is, bandwidth redistribution, SLA partitioning and segregated inter-ONT/intra-ONT scheduling, with the addition of traffic prediction. In a typical SLA that defines QoS performance, Committed Information Rate (CIR) is the rate that the operator commits to deliver to the customer at any time, while Excess Information Rate (EIR) is the additional (maximum)

<sup>1</sup>F. Slyne is with CONNECT Centre, Trinity College Dublin fslyne at tcd.ie

<sup>2</sup>M. Ruffini is with CONNECT Centre, Trinity College Dublin mruffini at tcd.ie

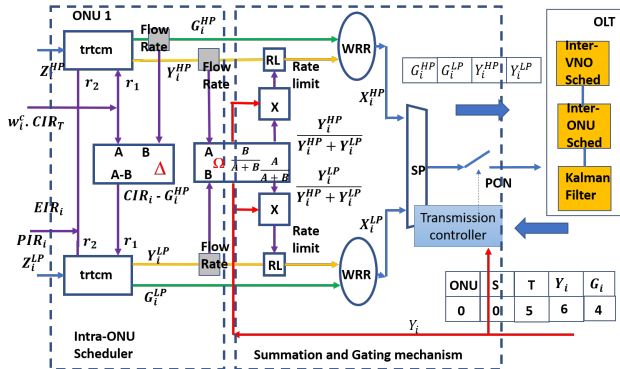


Fig. 1. Proposed Upstream Scheduler

rate that the operator will provide, if there is unused capacity from other customers, but without any assurance. The policies we reported in [5], were devised by the Metro Ethernet Forum (MEF) to define appropriate performance distribution in multi-tenant environments, assisting with characterisation of SLA partitioning. In particular, Policy 1 dictates that Contracted CIR should be fully provisioned to ONTs offered High Priority (HP) traffic; Policy 2 that traffic forwarded as HP should not exceed contracted CIR; Policy 3 states that any excess CIR should be availed of by an ONTs's Low Priority (LP) traffic; Policy 4 states that EIR should be apportioned to HP and LP traffic in proportion to the relative weighting of each ONTs within the VNO it belongs to; finally, Policy 5 says that any imbalance in HP or LP traffic in one VNO should not affect another VNO. The expected high priority  $X_i^{HP}$  traffic and the expected low priority  $X_i^{LP}$  traffic directed upstream towards the PON from the  $i^{th}$  ONT, respectively, are as the sum of HP traffic marked as green (assured) and yellow (non-assured) and the sum of LP traffic marked as green and yellow.

$$X_i^{HP} = G_i^{HP} + w_i^c \cdot EIR_T \cdot \frac{Y_i^{HP}}{Y_i^{HP} + Y_i^{LP}} \quad (1)$$

$$X_i^{LP} = G_i^{LP} + w_i^e \cdot EIR_T \cdot \frac{Y_i^{LP}}{Y_i^{HP} + Y_i^{LP}} \quad (2)$$

Where  $G_i^{HP}$ ,  $G_i^{LP}$ ,  $Y_i^{HP}$  and  $Y_i^{LP}$  are the HP and LP traffic marked as green and yellow. Any traffic marked as yellow contends for the total EIR ( $EIR_T$ ).  $w_i^c$  and  $w_i^e$  are defined as the weighting given to the packets from the  $i^{th}$  ONT and are calculated as a fraction of the configured CIR and EIR respectively for ONT  $i$  over the total CIR and EIR for ONTs. In[5], we instantiated the principles described by equations (1), (2) and (3), for downstream traffic, using standard components such as Weighted Round Robin (WRR) and Strict Priority (SP) schedulers and two rate three colour marker (trTCM) policers as defined by the IETF. In this paper we have elaborate a new mechanism, for tackling the integration of DBA and SLA-oriented queue management in our upstream scheduler, shown in Fig. 1. Logically, the scheduling algorithm is distributed between the ONT (queue management) and the OLT (DBA) and intra-ONT scheduling

occurs at each ONT, close to the source of each traffic stream. This allows traffic related to the same customer to be managed efficiently and in a responsive manner. Inter-ONT and Inter-VNO scheduling occurs at the OLT, and due to the time delay of exchanging information between the two domains (across a PON), operates in a slower cycle than the Intra-ONT scheduler. The Inter-ONT bandwidth reallocation algorithm executed at the OLT, with the associated pseudocode shown in Fig. 2), uses the weights  $w_i^c$  and  $w_i^e$  to iteratively reallocate EIR between ONTs.

**Algorithm 1** Calculate Intra-VNO dispersion of EIR

```

Input  $w, y, r, EIR$ 
 $w \leftarrow$  vector weights for all ONTs
 $y \leftarrow$  vector requests for all ONTs
 $r \leftarrow$  vector results for all ONTs
 $EIR \leftarrow$  total unallocated EIR bandwidth
for all  $ONT_i$  do
   $request \leftarrow$  yellow traffic needed by  $ONT_i$ 
   $offer \leftarrow$  portion of EIR allocation to  $ONT_i$ 
   $(take, surplus) \leftarrow trade(request, offer)$ 
   $take \leftarrow$  amount of bandwidth taken by  $ONT_i$ 
   $surplus \leftarrow$  excess bandwidth for distribution
  if  $ONT_i$  has received requested bandwidth then
    remove  $ONT_i$  from future calculations
     $redist$   $ONT_i$  weight  $w$  to other  $ONT$ 's in VNO $_i$ 
  end if
end for

```

Fig. 2. Inter-ONT Scheduling algorithm

The Inter-VNO algorithm operates in a similar iterative manner to reallocate any subsequent remaining bandwidth after the Inter-ONT algorithm has run. A Kalman Filter, based on Least Squares Estimation, is employed for the purposes of mediating between faster and slower cycle scheduling domains at the ONT and OLT, respectively, as well as performing traffic prediction. Standard 10 Gigabit PON (XG-PON) DBA message passing, that is, DBRu in the upstream and BWmap in the downstream, is used to communicate between the scheduler components at the ONT and the OLT.

## IV. RESULTS

We have executed extensive simulations to compare the behaviour of our upstream scheduler against a classic DBA, based on the GigaPON Access Network (GIANT) algorithm [1]. We configured 2 VNOs that share a common XG-PON upstream capacity of 2.488 Gbps, of which 70% is apportioned to CIR. An ONT may have one of two upstream service profiles, which is expressed as (CIR, EIR) tuple in Mb/s: Profile-1 (10,100) and Profile-2 (100,1000). In order to assess the isolation performance, we create an imbalance between VNOs, so that VNO "A" has 24 ONTs and VNO "B" has 8 ONTs, half with Profile-1 and half with Profile-2. Traffic is generated according to a Poisson distribution packet arrival rate, with averages varying  $\pm 20\%$  over the nominal profile rates. These are labelled as \_U (-20%) and \_O (+20%)

in our results in Fig. 3 and Fig. 4, which show the deviation, in pivot table form, of the schedulers' performance from the ideal scheduler, represented by Eq. (1). The independent

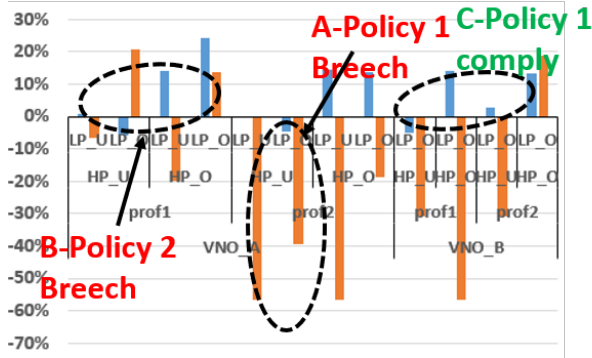


Fig. 3. HP Actual vs. expected per VNO




Fig. 4. LP Actual vs. expected per VNO

variables are the Operator (VNO\_A or VNO\_B), Profile (1 or 2), High Profile offered traffic (condition HP\_U or HP\_O) and Low Profile offered traffic (condition LP\_U or LP\_O). The graphs report the deviation from ideal of HP or LP traffic for 2 schedulers: our upstream scheduler (blue) and the classic DBA scheduler, based on GIANT (orange). The labelling in the x-axis reports the scenario under which the simulation was obtained: for example, the third group of bars was obtained considering VNO "A", profile-1, HP\_O and LP\_U. We see, in Fig. 3 that a classic DBA would be in breach of Policy 1 (label A) since CIR is not provisioned, while our proposed scheduler is within 10% of the ideal case and so is in compliance (label C). However, both our proposed mechanism and the classic DBA are in breach of Policy 2 (labelled B in the Figure) for the HP\_O case, although only by up to 20%. We see in the LP traffic graph, Fig. 3, that the classic DBA scheduler is in breach of Policy 4 (label F), whereas our scheduler allows excess CIR to be used for LP traffic (label E). Similarly, we can see that the classic DBA breaches Policy 4 (label D), as the transmitted LP traffic is determined by the amount of offered traffic rather than by the profile: when LP\_U changes to LP\_O (label D), the deviation swings widely from negative to positive. There is no such change in deviation for our scheduler which

instead shows compliance. . In Fig. 5 and 6, we look at traffic

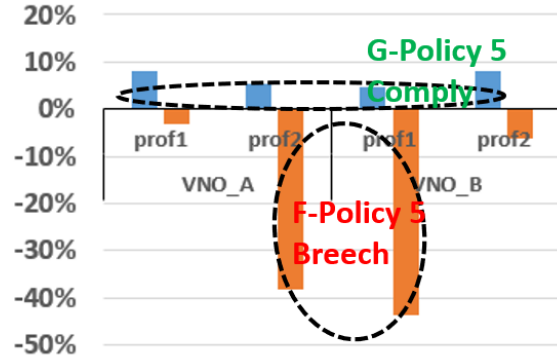


Fig. 5. HP Actual vs. Expected per VNO.

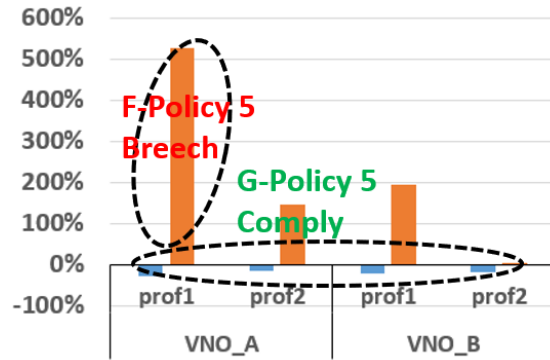


Fig. 6. LP Actual vs. Expected per VNO

separation between the VNOs using a pivot analysis of all ONTs grouping by VNO. For the classic DBA, VNO\_A's LP traffic benefits significantly to the detriment of VNO\_B's HP traffic. This shows that Policy 5 is breached (label F). On the other hand, our proposed scheduler shows close alignment between actual and expected HP and LP traffic for both VNOs, and hence compliance with Policy 4 and 5 (Label G).

In conclusion, we showed how our proposed joint scheduler can properly handle SLA-oriented QoS, which is necessary for PONs to support the next generation of multi-tenant and multi-service 5G and beyond networks.

#### ACKNOWLEDGEMENT

Financial support from SFI grants 14/IA/252 (O'SHARE) and 13/RC/2077 is gratefully acknowledged.

#### REFERENCES

- [1] J. Angelopoulos et al., "Efficient transport of packets with QoS in an FSAN-aligned GPON," IEEE Comm Mag, 2004
- [2] H. Alshaer et al., "An end-to-end QoS scheme for GPON access networks" IEEE GCC, 2011
- [3] M.S. Han et al., "Efficient dynamic bandwidth allocation for FSAN-compliant GPON", JOCN, 2008.
- [4] ITU-T, "Gigabit-capable PON (G-PON): ONT management and control interface specification," Rec G.984.4, 2008.
- [5] F.Slyne et al., "Single-stage scheduler for accurate QoS delivery in virtualised multi-tenant Passive Optical Networks", ECOC 2019.