

Single-stage scheduler for accurate QoS delivery in virtualised multi-tenant Passive Optical Networks

Frank Slyne⁽¹⁾, Bruno Cornaglia⁽²⁾, Marco Boselli⁽²⁾, Marco Ruffini⁽¹⁾

⁽¹⁾ CONNECT Research Centre, Trinity College Dublin, {fslyne,ruffinm}@tcd.ie,

⁽²⁾ Vodafone Group Technology, {bruno.cornaglia,marco.boselli}@vodafone.com

Abstract We propose a single-stage downstream scheduler assuring accurate Quality of Service delivery across Virtual Network Operators in a multi-tenant PON. We show performance close to an ideal scheduler in the ability to control QoS and compare performance to that of standard schedulers.

Introduction

Passive Optical Networks (PONs) are an economically efficient means of providing high capacity bandwidth across dense urban geographical areas and are increasingly attractive for delivering non-retail services such as Long Term Evolution (LTE)^{1,2}, reducing, for example, the cost of small cell deployment. PON sharing across multiple Virtual Network Operators (VNOs) is also becoming increasingly important to increase competition without replicating costly hardware infrastructure. In a multi-tenant scenario each VNO should have the ability to assign its portion of assured and best effort traffic to its customers, without affecting or being affected by other VNOs. Given that LTE, when implemented as Cloud Radio Access Networks (C-RAN) has a low tolerance to packet jitter and latency, the VNO needs granular slicing and resource control³. Achieving separation between VNOs in a PON environment is complicated by the fact that upstream and downstream transmissions operate in different manners. Granular resource control was proposed in⁴ for upstream traffic and recently standardised⁶, while downstream operations are typically based on proprietary vendors implementations. Moving central office infrastructure towards virtualisation and slicing⁷ however opens up new possibilities for the design and implementation of novel schedulers⁸. In current vendor implemen-

by chaining multiple stages (Fig. 1), in order to obtain finer granularity of control. For example, the scheduler in the left hand side of Fig. 1, which we call PON level scheduler, works on a single stage, providing only the ability to differentiate QoS at the service level, with no further differentiation across the entire PON. If the operator wants to also police bandwidth for individual Optical Network Units (ONUs), it needs to implement a second stage (shown in the right hand side of the Figure). This additional stage however increases the system complexity and latency (since it uses additional queues and scheduling algorithms). In multi-tenant environments, the traditional vendor approach to providing VNO-level QoS isolation is to add an additional, third stage. In this way, unused capacity can preferentially be redistributed within a VNO slice, before being shared with other VNOs. However this further increases complexity and latency. In this paper, we design, from first principles, a novel scheduler that is capable of operating in a multi-tenant PON environment, but only using a single stage.

Ideal Scheduler Behaviour

The Ideal scheduler should use standard components such as Weighted Round Robin (WRR) and Strict Priority (SP) schedulers and two rate three colour marker (trTCM) policier, as defined by the IETF⁹. It should satisfy the following policies, which we have drafted, based on the MEF Service Level Specification⁵ for Ethernet Virtual Circuit (EVC):

Policy 1: contracted Committed Information Rate (CIR) should be fully provisioned to ONUs offered High Priority (HP) traffic.

Policy 2: traffic forwarded as HP should not exceed contracted CIR.

Policy 3: any excess CIR should be availed of by an ONUs's Low Priority (LP) traffic.

Policy 4: Excess Information Rate (EIR) should

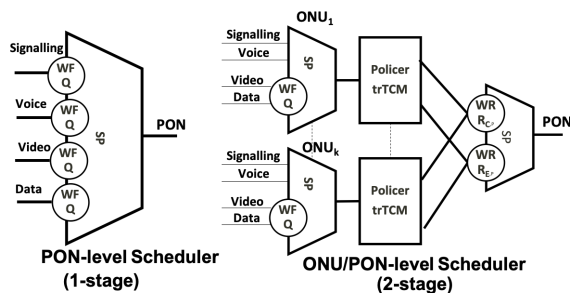


Fig. 1: Classic PON downstream schedulers
tations, downstream PON schedulers are created

be apportioned to HP and LP traffic in proportion to the relative weighting of each ONUs within the VNO it belongs to.

Policy 5: Any imbalance in HP or LP traffic in one VNO should not affect another VNO.

Here, we formulate the behaviour of an ideal scheduler, which satisfies all the above policies, and against which current and proposed architectures may be compared. To do this, we calculate the expected HP and LP traffic, X_i^{HP} and X_i^{LP} , directed from the PON towards the i^{th} ONU, given offered HP (Z_i^{HP}) and LP (Z_i^{LP}) traffic. The high priority traffic G_i^{HP} (i.e., coloured as green) should not exceed the CIR allocated to that ONU. In the event that the HP traffic exceeds the CIR, the excess is remarked as LP and recoloured as yellow (Y_i^{HP}).

$$G_i^{HP} = \min(Z_i^{HP}, CIR_i) \quad (1)$$

$$Y_i^{HP} = Z_i^{HP} - G_i^{HP} \quad (2)$$

In addition, any excess CIR may be used to recolour part of the LP traffic as green (while the remaining LP traffic is coloured as yellow).

$$G_i^{LP} = CIR_i - G_i^{HP} \quad (3)$$

$$Y_i^{LP} = Z_i^{LP} - G_i^{LP} \quad (4)$$

The expected HP traffic directed from the PON towards the i^{th} ONU (i.e., X_i^{HP}) is the sum of HP traffic marked as green and HP traffic remarked as yellow, contending for the total EIR (EIR_T). With w_i^c and w_i^e as defined in Eq. (7), X_i^{HP} is:

$$X_i^{HP} = G_i^{HP} + w_i^e \cdot EIR_T \cdot \frac{Y_i^{HP}}{Y_i^{HP} + Y_i^{LP}} \quad (5)$$

Similarly, the expected LP traffic directed from the PON towards the i^{th} ONU (i.e., X_i^{LP}) is the sum of LP traffic marked as green and LP traffic marked as yellow, also contending for the total EIR (EIR_T).

$$X_i^{LP} = G_i^{LP} + w_i^e \cdot EIR_T \cdot \frac{Y_i^{LP}}{Y_i^{HP} + Y_i^{LP}} \quad (6)$$

w_i^c and w_i^e are the weighting given to the packets from the i^{th} ONU

$$w_i^c = \frac{CIR_i}{\sum_{j=1}^k CIR_j}, w_i^e = \frac{EIR_i}{\sum_{j=1}^k EIR_j} \quad (7)$$

calculated as a fraction of the configured CIR and EIR respectively for ONU i over the CIR and EIR for all (k) ONUs.

Proposed multi-tenant single-stage scheduler

In this section, we show our development of a single-stage scheduler, summarised in Fig. 2,

aiming to replicate the ideal behaviour described by equations (5) and (6), using only standard WRRs, SPs and trTCMs components. The first action is to separate the offered HP traffic Z_i^{HP} into G_i^{HP} and Y_i^{HP} flows using a trTCM functional block with the CIR and Peak Information Rate (PIR) set to the proportion of overall CIR and PIR allocated to that ONU. While the Green traffic G_i^{HP} flows unchanged, the Yellow traffic Y_i^{HP} is shaped using a Rate Limiter. A functional block Ω calculates the amount of rate limiting as a product of the fixed fraction of the total EIR allocated to the PON (across all ONUs and VNOs) and the dynamic proportion of the HP marked as Yellow out of the total traffic marked as Yellow for that ONU. A WRR combines the G_i^{HP} and rate shaped Y_i^{HP} .

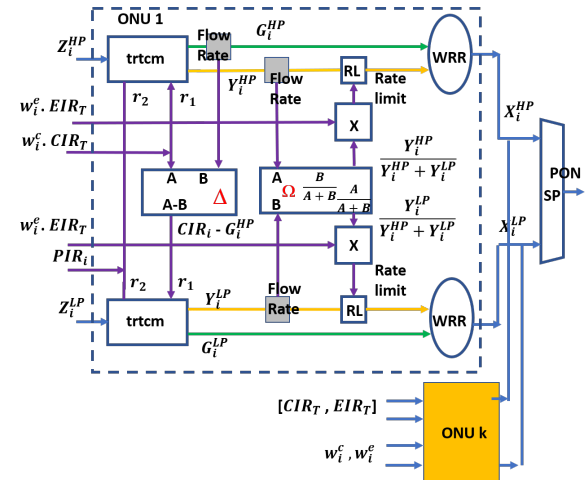


Fig. 2: Proposed Single Stage Scheduler

The second action is to separate the offered LP traffic Z_i^{LP} into G_i^{LP} and Y_i^{LP} flows using a second trTCM functional block with the trTCM CIR and PIR set to the proportion of overall excess CIR and PIR allocated to that ONU. Excess CIR is any CIR unallocated to the HP for that ONU, and is calculated in realtime by functional block Δ . While the Green traffic G_i^{LP} flows unchanged, the Yellow traffic Y_i^{LP} is shaped using a Rate Limiter. The same functional block Ω calculates the amount of rate limiting as a product of the fixed fraction of the total EIR allocated to the PON (across all ONUs and VNOs) and the dynamic proportion of the LP marked as Yellow out of the total traffic marked as Yellow for that ONU. A second WRR combines the G_i^{LP} and rate shaped Y_i^{LP} . The schedulers for all ONUs are aggregated at the egress to the PON using a SP queue.

Simulation and Results

We have carried out extensive simulations comparing the behaviour of our proposed single-stage

scheduler, with the reference ONU/PON level and PON level architectures shown in Fig. 1. Simulations were carried out using the well-known Python SimPy¹⁰ discrete event simulator, assuming 2 VNOs sharing a common GPON downstream capacity of 2.488 Gbps, of which 70% is allocated to CIR. In our simulations, an ONU may have one of two service profiles, expressed as the tuple (CIR, EIR) in Mb/s: Profile-1 (10,100) and Profile-2 (100,1000). In order to assess fairness between VNOs, we configure VNO "A" to have 24 ONUs and VNO "B" to have 8 ONUs, half on Profile-1 and half on 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 labeled in our results in Fig. 3 and 4 as $_U$ (-20%) and $_O$ (+20%).

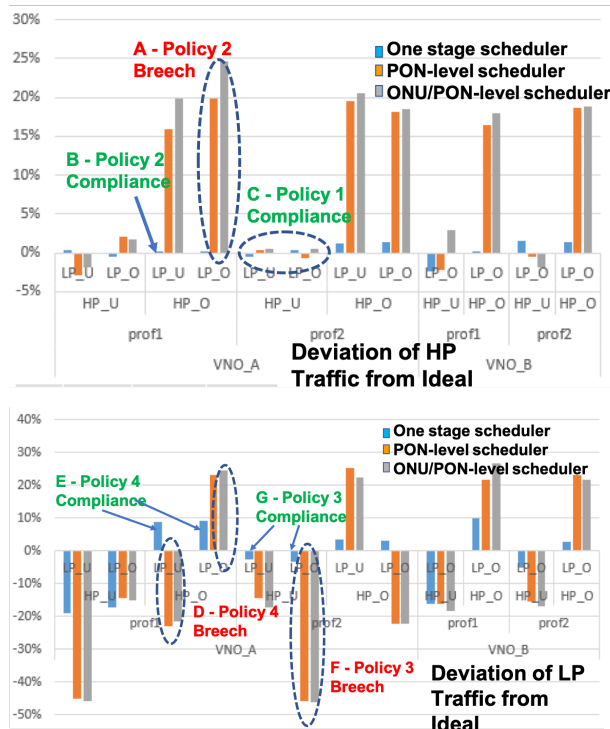


Fig. 3: Deviation of HP (top) and LP (bottom) traffic from ideal

The results of our simulations are reported in pivot table form in Fig. 3, showing the deviation of the schedulers' performance from the ideal scheduler represented by Eq. (5) and (6). The independent 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 (top graph) or LP (lower graph) traffic for 3 schedulers: our novel single-stage scheduler (blue), the PON level scheduler (orange), and the ONU/PON level scheduler (grey). The labeling in the x-axis report the scenario under which the simulation was ob-

tained: for example, the third group of bars was obtained considering VNO "A", profile-1, HP_O and LP_U. We see from the top graph that all 3 schedulers honor Policy 1 (condition HP_U), as the deviation from ideal is low for all schedulers (labelled C in the Figure). However, both standard schedulers are in breach of Policy 2 (labelled A in the Figure) for the HP_O case, unlike our proposed scheduler (label B). Looking at the lower graph of Fig. 3, we can see how both standard schedulers breach Policy 3 (label F), unlike our scheduler, which allows excess CIR to be used for LP traffic (label G). Similarly, we can see that both standard schedulers breach 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

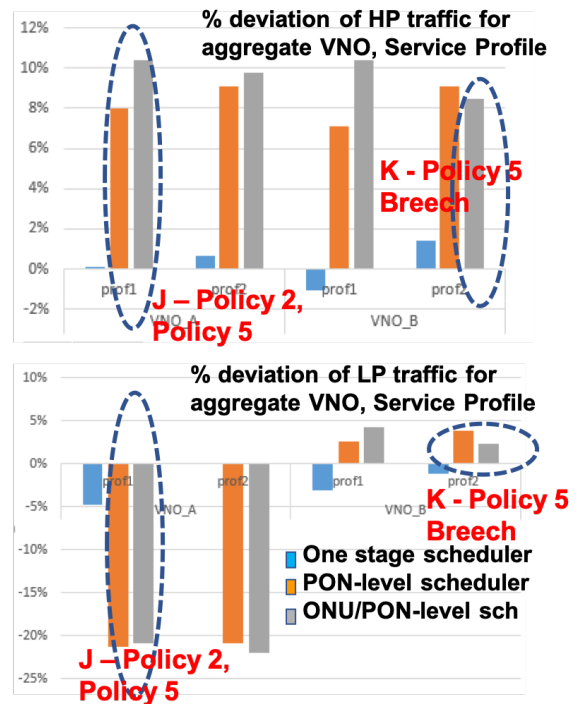


Fig. 4: Performance of multi-tenant VNO HP and LP traffic Fig. 4 we have aggregated all ONUs from VNO A and all ONUs from VNO B. The results show that both standard schedulers affect VNO_A's LP traffic significantly by lack of compliance to Policy 2 (label J). This unfairness is not shared by VNO_B, which implies that it occurs at the expense of the VNO_A, and thus is a breach of Policy 5 (label K). Our schedulers shows instead minimal deviation from ideal.

Acknowledgements

Financial support from SFI grants 14/IA/2527 and 13/RC/2077 is gratefully acknowledged.

References

- [1] S. Zhou, X. Liu, F. Effenberger and J. Chao, Low-latency high-efficiency mobile fronthaul with TDM-PON (mobile-PON). *Journal of Optical Communications and Networking (JOCN)*, Vol. 10 No. 1, Jan. 2018.
- [2] S. Bidkar, J. Galaro and T. Pfeiffer, First Demonstration of an Ultra-Low-Latency Fronthaul Transport Over a Commercial TDM-PON Platform. *Optical Fibre Communications conference (OFC)*, paper Tu2K.3, March 2018.
- [3] Tashiro et al., A Novel DBA Scheme for TDM-PON Based Mobile Fronthaul. *OFC 2014*, paper Tu3F.3.
- [4] A. Elrasad, N. Afraz and M. Ruffini, Virtual Dynamic Bandwidth Allocation Enabling True PON Multi-Tenancy. *Optical Fibre Communications conference (OFC)*, paper M3I.3, March 2017
- [5] Technical Specification 10.2. Section 7.11. Metro Ethernet Forum. 2009 http://www.mef.net/PDF_Documents/technical-specifications/MEF10.2.pdf
- [6] BBF TR-402 technical report, "Functional Model for PON Abstraction Interface", October 2018.
- [7] M. Ruffini, F. Slyne. Moving the Network to the Cloud: the Cloud Central Office Revolution and its Implications for the Optical Layer. *IEEE/OSA Journal of Lightwave Technology*, Vol. 37, No 7, April 2019.
- [8] K. Asaka, H. Ujikawa, H. Uzawa, H. Nakamura, J. Kani, A. Otaka and J. Terada. Disaggregation of Time-Critical Applications in Flexible Access System Architecture. *Journal of Optical Communications and Networking (JOCN)*, Vol. 1, No 1, January 2019.
- [9] Internet Engineering Task Force standards. <https://www.ietf.org/standards/>
- [10] N. Matloff. Introduction to discrete-event simulation and the simpy language, Feb., 2008. Available at: <http://heather.cs.ucdavis.edu/~matloff/156/PLN/DESIntro.pdf>