

Kent Academic Repository

Full text document (pdf)

Citation for published version

Al-Hares, Mohamad Kenan and Assimakopoulos, Philippos and Muench, Daniel and Gomes, Nathan J. (2017) Traditional queuing regimes and time-aware shaping performance comparison in an Ethernet fronthaul network. In: 2017 19th International Conference on Transparent Optical Networks (ICTON). IEEE ISBN 9781538608586.

DOI

<https://doi.org/10.1109/ICTON.2017.8025067>

Link to record in KAR

<http://kar.kent.ac.uk/64306/>

Document Version

Author's Accepted Manuscript

Copyright & reuse

Content in the Kent Academic Repository is made available for research purposes. Unless otherwise stated all content is protected by copyright and in the absence of an open licence (eg Creative Commons), permissions for further reuse of content should be sought from the publisher, author or other copyright holder.

Versions of research

The version in the Kent Academic Repository may differ from the final published version.

Users are advised to check <http://kar.kent.ac.uk> for the status of the paper. **Users should always cite the published version of record.**

Enquiries

For any further enquiries regarding the licence status of this document, please contact:

researchsupport@kent.ac.uk

If you believe this document infringes copyright then please contact the KAR admin team with the take-down information provided at <http://kar.kent.ac.uk/contact.html>

Traditional Queuing Regimes and Time-aware Shaping Performance Comparison in an Ethernet Fronthaul Network

Mohamad Kenan Al-Hares¹, Philippos Assimakopoulos¹, Daniel Muench² and Nathan J. Gomes¹

¹Communications Research Group, University of Kent, Canterbury, UK

²ADVA Optical Networking SE, Munich, Germany

ABSTRACT

This paper compares the performance of traditional priority-based queuing regimes with a time-aware shaping scheduler in an Ethernet fronthaul. Different use-cases are considered for the high and low priority traffic generation, which are made to represent precision-time protocol traffic and traffic originating from different LTE functional subdivisions (function splits) respectively. It is shown that the relative performance characteristics of the three scheduling regimes depend strongly on the utilization of the fronthaul links where contention is taking place and on the traffic generation characteristics of the different traffic sources.

Keywords: Fronthaul, C-RAN, LTE, Ethernet, Background traffic.

1. INTRODUCTION

Traditional centralised (or cloud) radio access networks (C-RANs) will not be able to scale to the data rates envisioned for next generation mobile networks. This being a result of implementing a centralised LTE functional subdivision, traditionally between a baseband processing unit, located centrally, and a remote radio head that is located remotely [1]. Data exchanged over such a fronthaul network is on the form of In-phase and Quadrature (IQ) radio samples resulting in very high data rates that do not scale with the cell load. The implementation of additional functional subdivisions (or function splits) has been proposed as a means of solving the scalability problem and obtaining statistical multiplexing gains [2-4]. Standardisation bodies such as the 3GPP [5] as well as other interest groups [4] are in the process of identifying a number of candidate splits for the future mobile network. An additional potential driver for scalability and cost reductions is Ethernet technology [6, 7]. The ubiquity of Ethernet equipment and its ability to provide operational and structural convergence will play an important role in the roll out of new 5th generation (5G) services. However, envisioned 5G use cases will have very tight latency and latency variation constraints. To this extend a number of time sensitive networking profiles (TSN) are being considered. The IEEE802.1CM group is in the process of defining profiles for TSN in Ethernet fronthauling [8]. These include strict priority (SP), a more traditional priority-based scheduling regime, and pre-emption. However, both profiles are defined for the traditional IQ-based fronthaul. Implementation of new function splits will result in an increase in the number of traffic flows, originating from the data plane, the timing plane (e.g. precision-time protocol, *PTP*), and the control, and management (C&M) plane. Furthermore, these flows will have different characteristics. The advanced-type schedulers that are being considered for the Ethernet fronthaul, might suffer from increased end-to-end latencies while others might not be able to scale to the numbers of flows. It is interesting then to compare the performance of these scheduling regimes with more traditional ones. Defining use-cases where the latter can be used can lead to higher traffic management flexibility in future fronthaul implementations.

2. STRICT PRIORITY AND WEIGHTED ROUND ROBIN QUEUING REGIMES

This section introduces the three types of schedulers that will be compared in Section 4. Two are traditional Ethernet schedulers, weighted-round robin (WRR) and strict priority (SP) [9]. However the latter has gained interest recently as it is being considered by Ethernet fronthaul standardisation. The advanced scheduler is a time-aware shaper (TAS) implementation based on IEEE802.1Qbv [10]. SP and WRR are readily configured for use in OPNET as models already exist for them, while the TAS model is implemented in OPNET.

2.1 Strict Priority (SP)

With SP, queues transmit their packets in order of priority. As a result, a lower priority queue will not transmit until a higher priority queue has been completely serviced by the scheduler (i.e. the queue is now empty). SP is beneficial for high priority traffic but can lead to starvation of lower priority traffic flows.

2.2 Weighted Round Robin (WRR)

With a frame-based WRR scheduler, each queue is allowed to transmit a specific number of frames in every transmission round. This queuing regime is fairer than the SP regime. Each queue is assigned a weight, which corresponds to the number of frames that it can transmit before it must cede transmission to other queues. Increasing the number of transmitted streams will reduce the share of each stream in the link since each stream

will be allocated a smaller percentage in the link resources. Note that SP can be considered as a special case of WRR with a weight equal to zero for the lower priority stream(s).

2.3 Time-Aware Shaper (TAS)

The main concept of the TAS is to allow transmissions from port queues in network bridges and traffic sources based on explicit time scheduling. The total transmission time (i.e. encompassing all traffic sources) is divided into sections and subsections where individual queues are allowed to transmit. In essence, the port queues are “gated” according to the window section boundaries. A protected section (PS) and subsections (PSS) within (if applicable), are assigned to high priority streams while a best effort section (BES) and subsections (PSS) within (if applicable), to the lower priority streams. A generic layout of the different sections and subsections that make up the transmission window (TW) is shown in Fig. 1. Note that for this scheme to work a level of time synchronisation is implied. A guard period (GP), in essence an idle time where all transmissions cease, is used to ensure that best-effort traffic does not overrun into the protected section.

The implementation of TAS in OPNET is shown in Fig. 2. The window sections are implemented in the media-access control (MAC) layer of the Ethernet switch. A packet received in one of the input ports is inserted into a queue while its reception time is used to identify whether it is received within its section (or subsection). If it is, it is forwarded to the output port queue and transmitted over the link. If it is not, then the packet is dropped. On the sending station side, windows are also defined so that packets that are generated by the application layer are only sent if they are received by the station MAC layer within their allocated sections. If they are not, they are buffered and their transmission deferred for the next transmission window. The two traffic sources (traffic generator 1, *TG1* and traffic generator 2, *TG2*) are assigned different virtual-local area identifiers (VLAN IDs) which take effect in the switch (i.e. a port-based VLAN configuration). The first switch implements the scheduling and is followed by a second switch with a trunk link between them that allows both VLANs to pass through. The WRR and SP schedulers are applied in the output port queues of the first switch. The second switch is then followed by the receiving stations. All links are optical 1 GbE (gigabit-Ethernet).

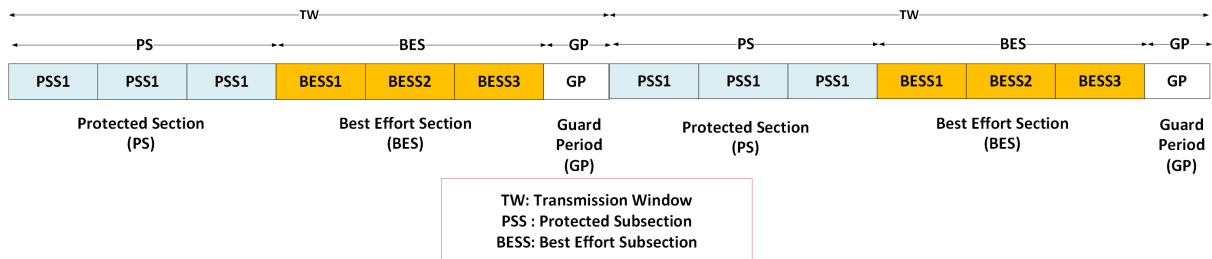


Fig.1. Generic window section and subsection plan based on IEEE 802.1Qbv.

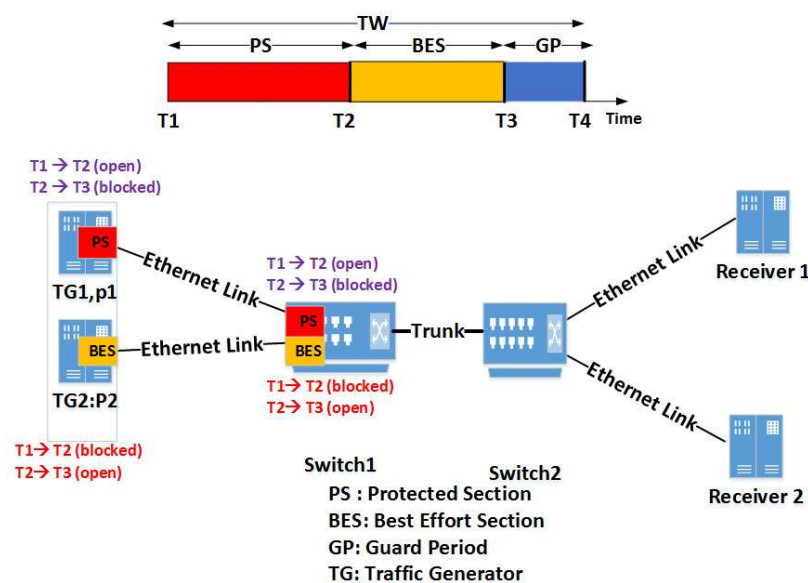


Fig.2. TAS implementation in OPNET.

3. MEASUREMENT RESULTS

For the results presented in this Section, the high priority traffic is a PTPv2-emulating stream generated by traffic generator TG1 and is composed of 32 timing messages per second per PTP slave (Receiver 1). The number of slave stations is 50 (these are modelled through the amount of traffic generated and the corresponding utilization in the trunk and not as separate receivers) while each PTP ‘sync’ message is formed as a 68 octet frame. Note that the amount of background traffic that shares the trunk link with the PTP traffic may not correspond to the same number of receiving stations. The background traffic is made to represent the traffic produced by a functional split at the LTE MAC/PHY interface. In such a split, transport blocks, each corresponding to a different user, are produced at the beginning of each LTE subframe (1 ms, corresponding to the scheduling resolution of LTE). The weights used for WRR are 8:1, 8:4, 8:6 and 8:8 with the higher weight corresponding to the high priority traffic. The performance comparisons presented here are based on the frame-delay variation (FDV) experienced by the PTP traffic. FDV can have detrimental effects in the PTP timestamping accuracy.

Fig. 3(a) shows the results for the background traffic configured with a constant frame size (1000 octets), constant burst size (10 frames per burst) and an inter-frame gap of 30 μ s. While for the results of Fig. 3(b), the background traffic is configured with a variable burst size following a uniform distribution (1-10 frames).

It can be seen that both WRR and SP schemes result in the same performance irrespective of the weights used for WRR. Note that only one WRR trace is shown here as the traces for all weights overlap. This is a result of a combination of two factors. The first is that the aggregate traffic data rate (trunk utilisation) is low. The second is that the high priority traffic is constant frame rate and thus when it contends with the low priority traffic only a single frame will be affected at any given time.

Furthermore, the worst-case performance for TAS in terms of both an average and peak FDV, is equivalent to the SP and WRR performance. But as the GP is increased, the peak and average FDV reduce consistently. Zero FDV is achieved for a GP that is equivalent to a full background frame serialisation delay (8 μ s). The step-like behaviour for the TAS results is an effect of the resizing of the BES in order to accommodate the GP (i.e. the TW remains constant). As the GP is increased, there is no change in FDV until the GP “eliminates” the frame from the burst that is closer (in time) to the GP boundary. This can be seen in Fig.3(a) where both the burst and frame sizes are kept constant, by observing that the step changes for the peak FDV occur at GP values that when added to the corresponding FDVs are approximately equal to one background frame serialisation. The same relation is not observed in Fig3(b) where the burst size is random, as the random generator seed values are changed from one simulation run to the next.




Fig.3. High priority traffic emulating PTP transmissions and background traffic with constant frame size and (a) constant burst size and (b) variable burst size.

4. CONCLUSIONS

A comparison between two traditional priority-based scheduling regimes, strict priority (SP) and weighted-round robin (WRR) with a time-aware shaper implementation based on IEEE802.1Qbv is presented. All three schedulers are modelled/configured in OPNET. The high priority traffic is made to represent a PTP traffic stream, while the low priority traffic a stream produced by an LTE MAC/PHY split. Frame delay variation results are presented and are used to show that the relative performance characteristics of the three scheduling regimes depend strongly on the utilisation of the links where contention is taking place and on the traffic generation characteristics of the individual traffic sources. Scheduling regimes like the ones presented here will play an important role in next-generation mobile network fronthauls in ensuring the timely delivery of traffic flows with high latency and latency variation constraints. The results presented here are more applicable for fronthaul networks with low levels of

aggregation or for the edge of such networks, towards the radio sites, where aggregation levels are expected to be low.

Acknowledgment

This work was carried out within the framework of the European Union's Horizon 2020 research and innovation programme under grant agreement No 644526 (iCIRRUS project) and EPSRCs "Towards an Intelligent Information Infrastructure (TI3)" programme (NIRVANA project). Mohamad Kenan Al-Hares acknowledges the funding through an EPSRC Doctoral Training Partnership (DTP) and Philippos Assimakopoulos by the NIRVANA project. Data used in this work is stored in Kent Academic Repository (<https://kar.kent.ac.uk/>).

REFERENCES

- [1] CPRI (Oct. 2015), "CPRI Specification V7.0, Interface Specification," [Online]. Available: <http://www.cpri.info/spec.html>.
- [2] U. Dötsch et al, "Quantitative analysis of split base station processing and determination of advantageous architectures for LTE," in Bell Labs Tech. Journal, vol. 18, no. 1, pp. 105-128, June 2013.
- [3] China Mobile et al, "White Paper of NGFI," ver. 1.0 (En) [Online]. Available: <http://labs.chinamobile.com/cran/>.
- [4] NGMN (Mar. 2015), "A deliverable by the NGMN alliance: Further study on critical C-RAN technologies," [Online]. Available: <https://www.ngmn.org/publications/technical.html>
- [5] 3GPP, "Study on New Radio Access Technology; Radio Access Architecture and Interfaces (Release 14)," 3GPP TR 38.801 V0.4.0, Aug. 2016 [Online]. Available: <http://www.3gpp.org/DynaReport/38-series.htm>
- [6] iCIRRUS (Jan. 2016), "D3.1: Verification of Ethernet as transport protocol for fronthaul / midhaul," [Online]. Available: <http://www.icirrus-5gnet.eu/category/deliverables/>
- [7] N.J. Gomes, V. Jungnickel, P. Chanclou, J.-P. Elbers, and P. Turnbull, "A flexible, Ethernet fronthaul for 5th generation mobile and beyond (Invited)," in Optical fiber Commun. Conf. (OFC), 2016, Anaheim, CA, 2016, paper W3C.1
- [8] "Time-Sensitive Networking for Fronthaul," IEEE Standard P802.1CM [Online]. Available: <http://www.ieee802.org/1/pages/802.1cm.html>
- [9] Cisco.(30 May 2006), *QoS Scheduling and Queueing on Catalyst 3550 Switches* [Online]. Available: <http://www.cisco.com/c/en/us/support/docs/lan-switching/lan-quality-of-service/24057-187.html> .
- [10] "Enhancements for Scheduled Traffic," IEEE standard 802.1Qbv [Online]. Available: <http://www.ieee802.org/1/pages/802.1bv.html>