

Kent Academic Repository

Full text document (pdf)

Citation for published version

Assimakopoulos, Philippos and Birring, Gurtej S. and Al-Hares, M. Kenan and Gomes, Nathan J. (2017) Ethernet-based fronthauling for cloud-radio access networks. In: 2017 19th International Conference on Transparent Optical Networks (ICTON). IEEE ISBN 9781538608586.

DOI

<https://doi.org/10.1109/ICTON.2017.8025034>

Link to record in KAR

<http://kar.kent.ac.uk/64307/>

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>

Ethernet-Based Fronthauling for Cloud-Radio Access Networks

Philippos Assimakopulos, Gurtej S. Birring, M. Kenan Al-Hares, Nathan J. Gomes
Communications Research Group, University of Kent, Canterbury, UK

ABSTRACT

Standardization bodies such as IEEE and 3GPP, as well as other interest groups, are in the process of defining and standardizing different functional subdivisions within mobile network base stations, primarily to reduce the data rate requirements imposed on the transport architecture by 4th and 5th generation mobile systems. Ethernet is considered the leading candidate for the transport architecture as it brings benefits from structural and operational convergence in what is now termed the “xhaul” (x signifying *front*, *mid* or *back*). This paper reports on the performance of different functional subdivisions over a bridged Ethernet network and presents some of the main issues that occur when internetworking at the Ethernet transport level. In particular, it examines Priority-based Scheduling within an Ethernet fronthaul, which is also a subject of active standardization by IEEE.

Keywords: Functional split, LTE, Ethernet, fronthaul, cloud-radio access network, time-sensitive networking

1. INTRODUCTION

Next generation cloud radio-access networks (C-RANs) must scale to much higher data rates with the advent of 5th generation (5G) bandwidths and new applications [1]. This requires a move away from traditional schemes based on the common-public radio interface (CPRI) where digitised In-phase and Quadrature (IQ) radio samples are transported from baseband processing units (BBUs) to remote radio heads (RRH) [2]. New functional subdivisions are proposed, whereby more processing is moved from BBU towards the air interface [3, 4]. As a result, BBUs and RRHs are more generalised (in terms of processing) into digital units (DUs) and remote units (RUs). These units potentially offer flexibility in processing, consistent with current moves towards network function virtualisation (NFV), where different RAN users can be allocated different processing “slices”. Such multi-user/multi-operator schemes will benefit from the use of standardised carrier-grade Ethernet equipment, and its use in the fronthaul has been proposed as a means of achieving structural and operational convergence [5, 6]. However, both functional splitting and Ethernet come with significant challenges. For the former, a number of splits are possible and choosing one split that can meet all performance targets is not possible. Both 3GPP [7] and IEEE [8] groups are in the process of developing a consensus on the choice of at least a few candidate splits. Ethernet’s main challenge is the lack of synchronisation features, which are necessary for carrier-grade operation. Precision-time protocol (PTP) and Synchronous Ethernet (SyncE) offer potential solutions [1]. The combination of functional splitting and Ethernet means that new traffic mechanisms become available due to statistical multiplexing gains in aggregation nodes (Ethernet switches). However, this means that fronthaul links must be provisioned for the timely delivery of fronthaul traffic. To this extent, IEEE 802.1 CM [9] is defining/adapting time-sensitive networking (TSN) profiles for fronthauling. Currently focusing on CPRI flows, similar profiles will be required for functional split traffic and any in-line timing protocol (e.g. PTP). To test these TSN profiles with real functional split traffic, hardware testbeds and simulation models will be required. In this paper, we show experimental results from an Ethernet-based testbed transporting functional split traffic and present simulation results for an Ethernet fronthaul employing a simplified time-aware shaper (TAS) and frame bursting to eliminate contention-induced delay and delay variation.

2. FUNCTIONAL SPLIT EXPERIMENTATION

A reference scenario, shown in Fig. 1(a), is used as the basis for the testbed presented in this section and the simulation set-up in Section 3. The fronthaul is formed by a pool of DUs which are connected, through Ethernet, to RUs or RRHs. These distributed RAN entities can perform different split functionalities, with some using centralised processing (IQ radio transportation, DU to RRH pair) while others use a new split, e.g. at the LTE MAC/PHY interface (DU to RU pair). The DUs perform part of the LTE protocol layer processing and are themselves fed with data from higher LTE layers, coming from RAN nodes closer to the core, through an x -haul. Timing flows exist over the network and provide timing references to the different RAN entities. In this case, an example use of a PTP boundary clock (BC) is shown. Fig. 1(b) shows a testbed based on this reference scenario. The testbed generates and transports MAC/PHY split data alongside background traffic, which is used to emulate either a centralised split or an additional MAC/PHY split. The LTE functionality in the DU, RU and user equipment (UE) runs in a software emulation environment based on the open source OpenAirInterface (OAI) software libraries (see the OpenAirInterface software alliance [13]) and specifically on the ‘OAI5G’ source code. A *Fronthaul Interface Library (FIL)* is used to encapsulate data exchanges between functional split entities and to provide a useable abstraction (mapping functions) to the new functionality. The DU performs all eNodeB processing down to and including the MAC layer. It then generates a number of flows and packetizes them into Ethernet frames. The resulting packet-types include, downlink shared channel (DL-SCH), downlink control

information (DCI), also used for MAC/PHY primitive exchange, system information (SI) and random access response (RAR). The flows are transported over an Ethernet network and received by the RU, which performs all LTE physical layer (PHY) processing. The radio frequency (RF) processing and the air interface are simulated. For simplicity, in the uplink a single packet type is used to aggregate all transmissions. Internet protocol (IP)-level data is injected directly into the fronthaul at the PDCP layer (i.e. there is no mobile core). The network comprises Ethernet switches that form a trunk through which the different flows can contend. The MAC/PHY split and background traffic flow are assigned to different virtual-local area network identifiers (IDs).

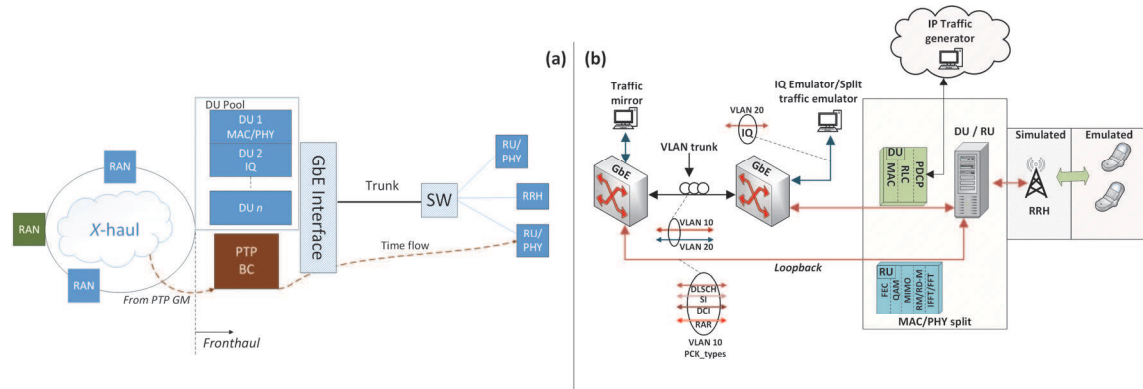


Figure 1. (a) Reference architecture and (b) Measurement set-up. SW: Ethernet switch; PTP BC: Precision-time protocol boundary clock; GM: Grandmaster; RRH: Remote radio head.

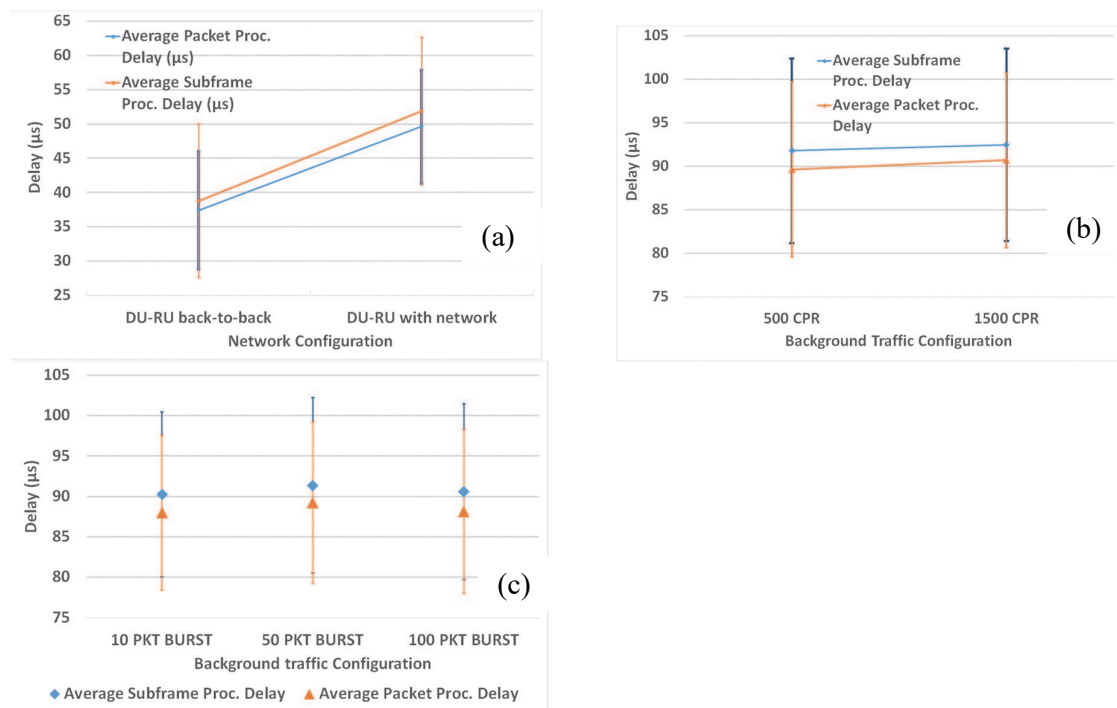


Figure 2. Average per-packet and per-LTE subframe processing delay. (a) Baseline case without background traffic, (b) With constant packet-rate (IQ/CPRI-type) background traffic and (c) With bursty background traffic.

Figure 2 shows measurement results for processing delay over the fronthaul with different background traffic configurations. For all cases, three UEs are connected to the LTE network while the DLSCH size per UE is approximately 340 octets long. The measured processing delays (per-subframe and per-packet) include both the *FIL* processing and networking-induced delays. Figure 2(a) shows results for the baseline case without background traffic. The per-packet and per-subframe delays are close in value. This is due to the fact that while a packet is being processed or is waiting in the DU transmit buffer (prior to being pulled out by the network interface card), the next packet to follow is also being processed. As a result, the packets that constitute the subframe are bursted out. The difference in average delay between the back-to-back case and the case with switching is approximately 13 μ s and this difference can be broken down to the two extra serialisation delays, from the switches, and two

fabric delays. Figure 2(b) shows results with two different background traffic configurations. Both are for constant packet rate traffic with a data rate of 400 Mbps with either 500 octet or 1500 octet frames. The data rate is equivalent to a generic IQ transport for an LTE bandwidth of 10 MHz and sample widths of 16 bpS (bits-per-sample). The measured delay is slightly higher for the 1500-octet case but the per-packet and subframe-delays converge. This can be attributed to cases where queuing delays (due to contention) in the switch absorb some of the time difference between consecutive packets from the same subframe. Figure 2(c) shows results with bursty background traffic. Three different burst sizes 10, 50 and 100 are shown, all with a packet size of 500 octets. The delays are similar with those from the previous case, and similar for the different burst sizes (as the data rate is the same, a higher burst size means that a burst will occur less often). For all results, there is significant delay variation (shown in the form of superimposed error bars) mainly due to variations in processing delay. Therefore, the encountered variation has a strong implementation dependency. The RAN entities are running on off-the-shelf generic processing platforms. Therefore, these variations in delay would be significantly reduced with hardware offloading techniques (for example by moving packetization/de-packetization functions into an FPGA board). However, with the advent of softwarisation in the RAN, where hardware abstraction techniques are expected to be employed, the resulting timing instabilities will have to be taken into account in the design of a fronthaul.

3. TIME-SENSITIVE NETWORKING

Previous work has investigated priority based scheduling in the fronthaul for IQ-based traffic using weighted round-robin (WRR) [10] and strict priority (SP) [11] algorithms. Additional techniques of interest include frame pre-emption (IEEE 802.1Qbu) [12] and TAS (IEEE802.1Qbv) [13]. For the latter, simulation results with CPRI [14] and traffic flows emulating functional split and PTP traffic [15] have been presented. Beyond priority-based scheduling, in [16] it was proposed (specifically for IQ transport) that frame bursting can be employed for trunk links, whereby uncontended “windows” are obtained for different flows; this is promising for transporting MAC/PHY split traffic. Figure 3(b) shows an architecture, where two DU/RU pairs contend for trunk access.

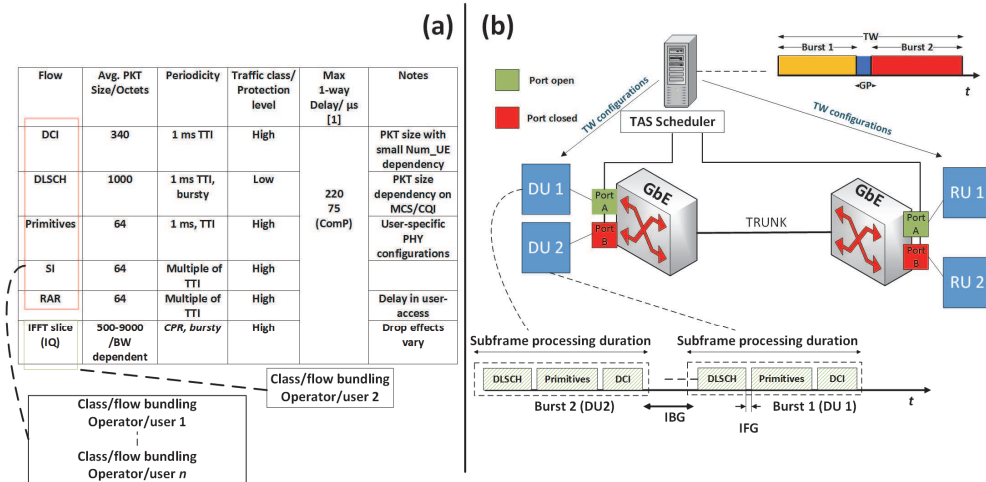


Figure 3. (a) Traffic flow information and (b) Simulation set-up. IFG: Interframe gap, IBG: Inter-burst gap; TTI: Transmission time interval; CPR: Constant packet rate.

By assigning different time window (TW) sections to each pair, traffic can pass through the trunk uncontended. A simplified version of a TAS based on IEEE 802.1Qbv (Fig. 3(b)) guarantees that the two bursts do not contend, for example due to timing instabilities, by employing a guard period (GP) where no end-node is allowed to transmit. The different flows that make up the MAC/PHY split traffic are shown in Fig. 3(a). Figure 4 shows simulation results for the set-up of Fig. 3. Without TAS the two flows contend, resulting in an average delay higher than the baseline case (single traffic flow). The level of contention is quite low, therefore this case is representative of bursty flows that can overrun due to timing instability. With TAS, contention is removed and the delay reduces to the baseline value and delay variation is completely eliminated.

4. CONCLUSIONS

Measurement results from a testbed transporting LTE MAC/PHY split data over an Ethernet network have been presented. The results show the fronthaul-induced processing delays and how these are affected by different types of traffic flows. Considering transmission and processing timing instabilities that will result from future implementations where RAN functionalities are distributed (split) in generic processing platforms with hardware abstraction/virtualisation techniques, a time-aware shaper design is presented that combined with frame bursting can result in a fronthaul where contention-induced delays and resulting delay variation are completely removed.

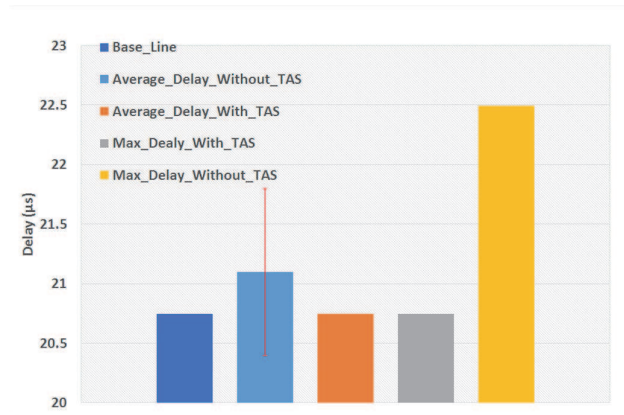


Figure 4. Simulation results with a TAS and traffic flows emulating bursty MAC/PHY split traffic

Acknowledgment

This work was carried out within the framework of the EU's H2020 research and innovation programme under grant agreement No 644526 (iCIRRUS) and EPSRC's TI3 programme (NIRVANA project). M K Al-Hares acknowledges funding through an EPSRC 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] iCIRRUS (Jul. 2016), "D3.2: Preliminary Fronthaul Architecture Proposal," [Online]. Available: <http://www.icirrus-5gnet.eu/category/deliverables/>
- [2] CPRI (Oct. 2015), "CPRI Specification V7.0, Interface Specification," [Online]. Available: <http://www.cpri.info/spec.html>
- [3] 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.
- [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] 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
- [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] 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>
- [8] "Next Generation Fronthaul Interface" IEEE 1914 Working Group [Online]. Available: <http://sites.ieee.org/sagroups-1914>
- [9] "Time-Sensitive Networking for Fronthaul," IEEE Standard P802.1CM [Online]. Available: <http://www.ieee802.org/1/pages/802.1cm.html>
- [10] P. Assimakopoulos, M. K. Al-Hares, S. Hill, A. Abu-Amara and N. J. Gomes, "Statistical Distribution of Packet Inter-Arrival Rates in an Ethernet Fronthaul," in *IEEE Int. Conf. on Commun. Workshops (ICC)*, Kuala Lumpur, Malaysia, 2016, pp. 140-144
- [11] M. K. Al-Hares, P. Assimakopoulos, S. Hill, and N. J. Gomes "The Effect of Different Queuing Regime On a Switched Ethernet Fronthaul," in *IEEE Int. Conf. on Transp. Optic.Net. (ICTON)*, Trento, Italy, 2016, pp. 1-4
- [12] "Frame Preemption," IEEE standard 802.1Qbu [Online]. Available: <http://www.ieee802.org/1/pages/802.1bu.html>.
- [13] "Enhancements for Scheduled Traffic," IEEE standard 802.1Qbv [Online]. Available: <http://www.ieee802.org/1/pages/802.1bv.html>
- [14] T. Wan and P. Ashwood-Smith, "A Performance Study of CPRI over Ethernet with IEEE 802.1Qbu and 802.1Qbv Enhancements," in *Global Commun. Conf. (GLOBECOM)*, San Diego, CA, 2015, pp. 1-6
- [15] M. K. Al-Hares, P. Assimakopoulos, D. Muench and N. J. Gomes, "Scheduling in an Ethernet Fronthaul Network," in *European Conf. on Networks and Commun. (EUCNC)*, Oulu, Finland, 2017. Accepted for publication
- [16] P. Assimakopoulos, M.K. Al-Hares, and N.J. Gomes, "Switched Ethernet fronthaul architecture for cloud-radio access network", in *OSA/IEEE J. Optical Commun. and Netw.*, vol: 8, no. 12, pp. B135-B146, Dec. 2016