

Kent Academic Repository

Full text document (pdf)

Citation for published version

Assimakopoulos, Philippos and Al-Hares, Mohamad Kenan and Hill, Simon and Abu-Amara, Ahmad and Gomes, Nathan J. (2016) Statistical distribution of packet inter-arrival rates in an Ethernet fronthaul. In: 2016 IEEE International Conference on Communications Workshops (ICC). IEEE pp. 140-144. ISBN 978-1-5090-0449-2.

DOI

<https://doi.org/10.1109/ICCW.2016.7503778>

Link to record in KAR

<http://kar.kent.ac.uk/57862/>

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>

Statistical Distribution of Packet Inter-Arrival Rates in an Ethernet Fronthaul

Traffic Protection Issues and Implications on Buffer Management

Philippos Assimakopoulos¹, Member, IEEE, Mohamad Kenan Al-Hares¹, Simon Hill², Ahmad Abu-Amara¹ and Nathan J. Gomes¹, Senior Member, IEEE

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

²Viavi Solutions UK Ltd, Newbury, Berkshire, UK

Abstract— This paper investigates the effects of background traffic streams in the packet inter-arrival rates of an LTE traffic stream, when these streams are transported over the same Ethernet fronthaul network. Contention of background traffic with LTE traffic can occur in a Cloud-RAN that is transporting traffic streams originating from constant bit-rate (CBR) sources such as the Common Public Radio Interface (CPRI) and from other non-CBR sources originating from different LTE physical layer functional subdivisions. Packet inter-arrival statistics are important in such a network, as they can be used to estimate and/or predict buffer sizes in receiving network nodes. Buffer management will also be important for traffic streams originating from functional splits (such as direct LTE MAC transport block transportation) where user plane data and control primitives have to be time aligned at the receiving node.

Index Terms— Fronthaul, C-RAN, 4G, 5G, LTE, VLAN, Ethernet, priority, background traffic

I. INTRODUCTION

The increased capacity demands [1] for current mobile network implementations (4G) and near-term future implementations (5G) will require the use of new mobile networking architectures as a means for meeting part of these demands [2]. A fronthaul architecture that employs Ethernet as the transportation technology can help reduce costs for operators, as Ethernet is an ubiquitous and potentially low cost technology. Additionally, Ethernet equipment used in the fronthaul can be re-used by the operator for backhaul links. Currently, C-RAN deployments use a fully centralised approach whereby Base station Baseband Units (BBUs), that handle all the digital processing, are kept in a central location. The analogue processing part (with some limited digital processing) is carried out in the remote locations by Remote Radio Heads (RRHs). Currently, the transportation between the BBUs and RRHs is done through the Common Public Radio Interface (CPRI) by transporting In-phase and Quadrature (I/Q) quantised radio samples. This centralised approach has a main drawback when it comes to networks with multiple antennas (e.g. Multiple-input and Multiple-output, MIMO), massive MIMO or carrier aggregation (up to 5 component carriers in LTE-A with a bandwidth of 100 MHz). The bandwidth requirements for these applications become prohibitive for

practical deployments and the situation becomes much worse when considering potential 5G data rates. An overview of data rate requirements based on the number of physical antennas for a fully centralised approach (I/Q transportation) is shown in Table I.

A method of reducing the data rate requirements over the fronthaul is through the implementation of different LTE physical layer functional subdivisions. Under such a regime part of the processing currently residing in the BBU, is moved into the RRH. This may lead to a more complicated RRH but this increase in complexity may not be so significant, since RRHs contain processing capabilities that are underutilised in current implementations. By moving part of the functionality in the RRH, reductions in data rate requirements become available. A number of subdivisions (or “splits”) are possible starting with a frequency domain one, where the split is located prior to the inverse-fast Fourier transform, (in the downlink), to one where all the physical layer functionality is placed in the RRH (LTE MAC transport block (TB) transportation) [3-5].

Through a functional split, additional techniques become more tractable. For example virtualization can be applied in the BBU pool for processing load balancing [2, 6, 7], while software defined networking (SDN) techniques can be used at layer 2, for traffic steering based on Quality-of-Service (QoS) and/or link utilisation primitives extracted through an SDN controller. But more importantly, for some split points, statistical multiplexing gains become a possibility, as the data rates over the fronthaul links will depend on the cell load (number of users).

Such split functionality is the central focus of two current projects, the NIRVANA [8] and iCIRRUS [9, 10], while both projects propose the use of Ethernet in the fronthaul as a standardized convergence layer.

Then, investigating the performance of a fronthaul network that will be transporting a mixture of traffic streams that can include generic I/Q or CPRI-type traffic streams as well as traffic streams originating from different physical layer splits, as well as backhaul traffic, becomes important.

To this extent in this paper, we present a detailed testbed set-up that employs “smart” probing techniques to sample

Table I. Data rates for LTE-A and 5G (est.) system bandwidths per RU sector for different no. of MIMO antennas (including massive MIMO implementations).

Channel BW/MHz	Sample rate /MHz ¹	Data rate (16 bpS) /Gbps ²					
		No. of antennas per sector at RU					
		2	4	8	16	64	128
20	30.72	1.97	3.93	7.86	15.72	62.91	125.82
40	61.44	3.93	7.86	15.73	31.46	125.82	251.65
60	92.16	5.9	11.8	23.59	47.184	188.73	377.47
80	122.9	7.9	15.72	31.46	62.91	251.64	503.3
100	153.6	9.8	19.68	39.36	78.72	314.9	629.8

¹Sample rate= IFFT_size/T_s

²Data rate= N x sample_rate x 2 x 16 bpS, (factor of 2 for I and Q, 16-bits per sample for N antennas)

(or capture) the traffic in the fronthaul. We then investigate the statistical distributions of packet inter-arrival rates of the LTE I/Q data traffic when it is transported over the same fronthaul links with generic Ethernet traffic. The latter is meant to represent traffic that would potentially be produced from different functional splits and/or backhaul traffic in a converged fronthaul architecture.

Section II shows an overview of pure layer-2 fronthaul, transporting different-type traffic streams and briefly discusses the issues regarding protection and buffering. In Section III we present an overview of the testbed used for the measurements, the results of which are presented in Section IV.

II. PROTECTION AND BUFFERING IN THE FRONTHAUL

Fig. 1 shows a C-RAN fronthaul architecture that combines fully centralised processing (this is shown by the BBU pool to RRH connections) and distributed processing (i.e. different physical layer functional splits). Under the latter the BBU is generalised to a Digital Unit (DU) and the RRH to a Remote Unit (RU). Data streams originating from the BBU and DU pools are “switched” to the respective destinations through VLAN ids. Multiple trunks are formed within the network that carry a number of VLANs. Depending on the data being transported through each trunk, different layer-2 priority regimes can be applied to offer more (or less) protection to some streams over others. The protection requirements will depend on what is being transported (i.e. on the split point).

Table II shows an overview of the protection requirements for a number of LTE channels and data that need to be transported by different functional splits. These channels and data will generally have different protection, data rate and latency requirements. For example, the PBCH may be transported only when there is a master information block change (every 40 ms). It is also implementation dependent whether this channel is transported through the split or whether it is generated at the RU through control primitives. Furthermore, dropping Ethernet frames that carry the PRACH will lead to increased delays in user access (for a number of users) and uplink resource grants. Regarding the transportation of TBs, protection requirements can be medium for the downlink but potentially high for the uplink, as UEs can send control channel data through the physical uplink shared channel. The handling of DMRS is implementation dependent. Although, the DMRS assists the DU in demodulating

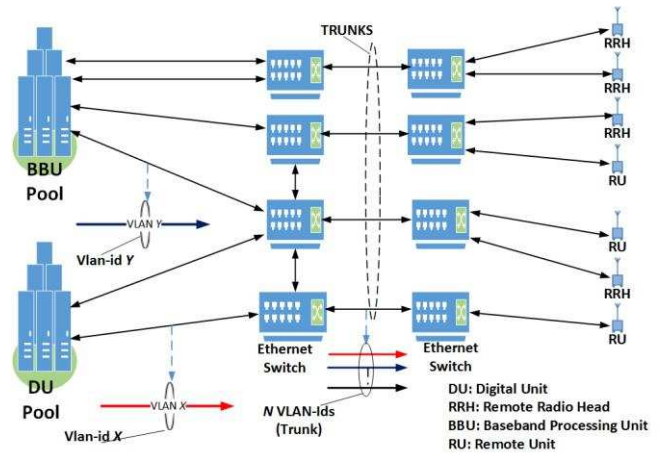


Fig. 1. A fronthaul architecture combining fully centralised functionality (BBU to RRH) and different physical layer function splits (DU to RU).

Table II. Protection requirements of LTE channels and data. DMRS=Demodulation Reference Signal, PBCH= Physical Broadcast Channel, PRACH=Physical Random Access Channel, SRS=Sounding Reference Signal, PDSCH=Physical Downlink Shared Channel.

Data transported	Protection requirement
DMRS	High
MAC control primitives	High
Transport blocks, DL (UL)	Medium/(High)
PBCH	High
PRACH	High
Radio “slice”-time domain	High
SRS	Low
PDSCH	Medium

user-plane and control-plane data on an individual user basis, if transported as a “block” per Transmission Time Interval (TTI) i.e. encapsulated in a single Ethernet frame, it will have implications for all user allocations in that TTI. The protection requirements for the PDSCH are also implementation dependent. If a number of user queues are encapsulated in a single Ethernet frame, implications can be more severe.

An additional issue when transporting data from different functional splits is the requirement for time alignment at the receiving node (for example, the alignment of control primitives and user-plane data for the radio resource mapping). Fig. 2 shows a conceptual example for time alignment, through buffering queues, between LTE MAC control primitives and user-plane data. Sequencing is used to synchronise the queues

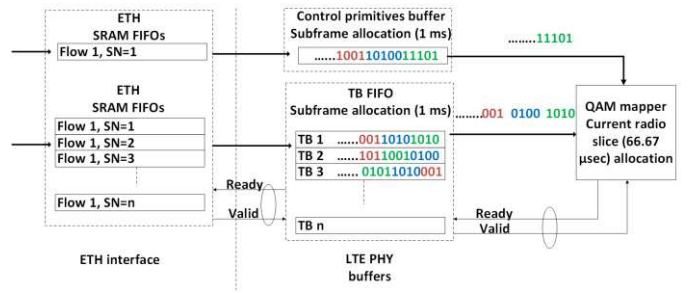


Fig. 2. Control primitives and TB data need to be time aligned at the receiver. SN=Sequence Number, ETH=Ethernet interface, SRAM=Static RAM, FIFO=First-in First-out.

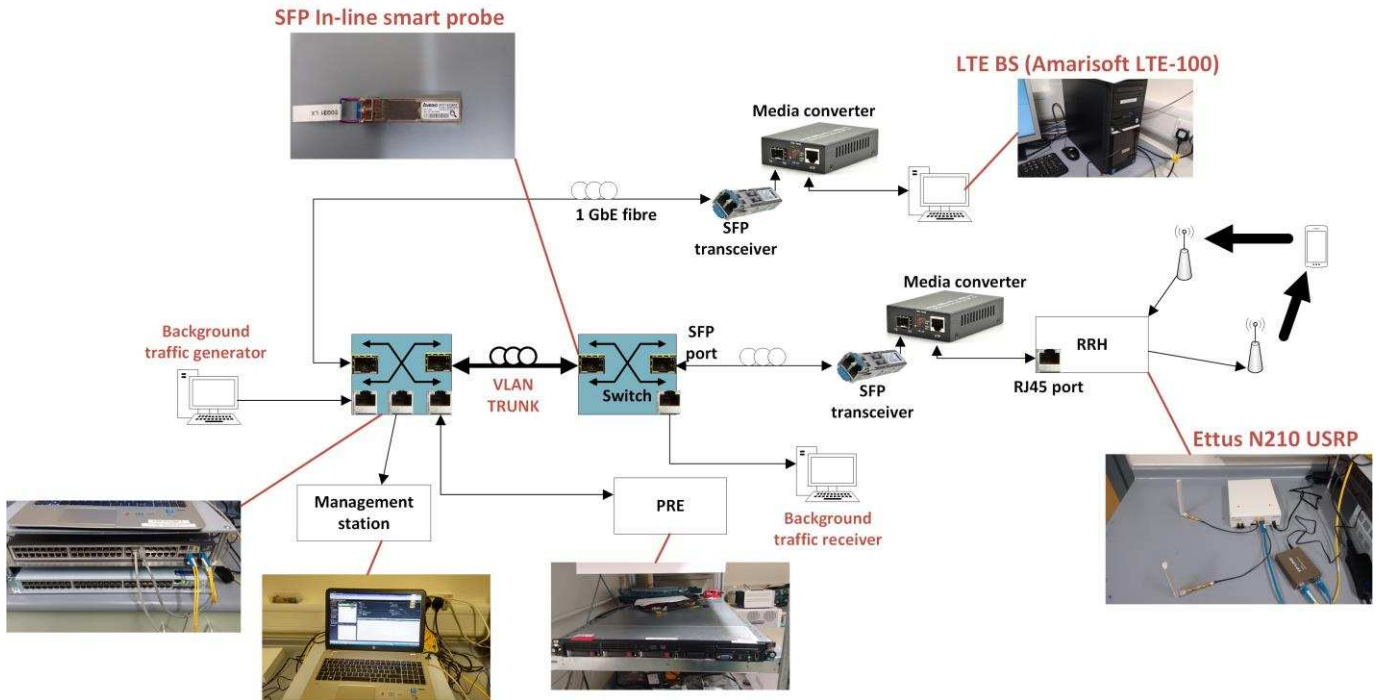


Fig. 3. Testbed used for the measurement procedure. PRE=Packet Routing Engine, GbE=Gigabit Ethernet, VLAN=Virtual Local Area Network, SFP=Small Form-Factor pluggable, RRH=Remote Radio Head, USRP=Universal Software Radio Peripheral. Arrows indicate the direction of traffic flow.

at the RU while the different data streams that are being transported are separated into flows. The buffer sizes need to be such, so that they can accommodate the encountered packet inter-arrival delays, otherwise overflows or underflows can occur and frames can be dropped. Whether a particular frame can be dropped without serious consequences will depend on what is being transported by the frame (see Table II for example). It is also expected, that in the fronthaul, data that has lower protection requirements may be ignored (essentially dropped) in cases where the frames transporting these data are excessively delayed. This can be beneficial as otherwise the buffer sizes at the receiving nodes would have to accommodate these worse case delays leading to excessive end-to-end latencies.

III. MEASUREMENT SET-UP

Fig. 3 shows the testbed used for the measurement procedure. A workstation runs an emulated LTE base station (Amari LTE-100) that produces I/Q samples corresponding to a 5 MHz channel bandwidth (sampling rate of 6.25 MHz). The samples are then inserted into the payload section of a UDP packet and transmitted over a pure layer 2 network. The network comprises of two 3COM-5500G Ethernet switches with standard 1000BASE-LX small form-factor pluggable (SFP) transceivers with LC connectors and Single Mode Fiber (SMF) patch-cords. The stream of packets containing the I/Q samples is received by an Ettus N210 RRH where, following Ethernet processing, the samples are de-quantised and sent to a digital-to-analog converter (DAC). Following the DAC they pass through a digital modulator for DC offset correction and are then up-converted to one of the LTE bands and amplified prior to transmission over the wireless channel. In the uplink

the reverse processes take place.

An additional workstation is used to generate background traffic of variable payload sizes and at variable data rates, using an open-source Linux-based traffic generator (Ostinato). The two streams of traffic are assigned two different VLAN ids (through the switch port i.e. a port-based VLAN configuration). The link between the two switches forms a Trunk that allows the pair of VLAN ids to pass through. As both VLAN ids will be transmitted through the same port there will be traffic contention. A first run of measurement is carried out without a priority implementation while a second run is carried out by applying VLAN priority using a weighted round robin queueing algorithm, whereby the LTE traffic is given a higher priority.

In both cases we sample the LTE traffic using a Viavi in-line Ethernet probe (smart probe). The probe comes in the form of a 1000BASE-LX SFP which under normal conditions operates as a standard transceiver. However, the probe contains Application-specific integrated circuit (ASIC) logic that allows it to capture the traffic going through it based on a user-defined filter definition. Additionally, the probe can be set to capture whole packets or just packet headers. In this experiment, a filter is used that instructs the probe logic to capture all packets containing in their headers the destination MAC address and destination UDP port of the RRH. Once the packets are captured, they are timestamped (using a propriety form of the Precision Time Protocol, PTP) and re-encapsulated with the normal network encapsulation, and with an additional Viavi proprietary header that contains the timestamp (in addition to other fields). The captured packets are re-injected into the network as Frame Result Packets (FRPs) and sent to a Packet Routing Engine (PRE), through which they are routed to a management station for further processing. Injection of the FRP

occurs between inter-packet gaps that are longer than the 12-octet time minimum (96 ns for 1 GbE). Once the FRPs are received in the management station, they are extracted using Wireshark for offline processing, which includes Matlab routines for estimating the statistics of the packet inter-arrival delays. The employed algorithm for obtaining the results is shown in Fig.4.

The switches used in these measurements operate in store-and-forward mode. This standard mode of operation is interesting for investigating cases when contention or, under certain conditions, a rate transitions occurs through a switch (a high-end carrier grade switch used in an Ethernet fronthaul would otherwise operate in cut-through mode). We are then interested to see how the queued packets from the two process queues (from the two traffic streams) are handled, with and without a priority implementation. To this extent, we do the following manipulation to the data during offline processing: We remove all packet inter-arrival delay values larger than 100 μs (link rate transmissions corresponds to an inter-arrival delay of approximately 32 μs). The large delay values are a result of the generation process in the LTE software base station and they can potentially “mask” the effects we want to measure. The total number of measurements that remain after extracting the high delay values are approximately 300000. The data rate of the LTE traffic was constant, for the duration of the experiment, at 200 Mbps, using jumbo frames with a length of 4000 bytes.

IV. MEASUREMENT RESULTS

Fig. 5 shows the Complementary Cumulative Distribution Functions (CCDFs) for the LTE traffic under different background traffic regimes. The blue trace represents the LTE traffic only (i.e. no background traffic transmitted). Fig. 5a shows the CCDF with background traffic packet size of 1500 bytes and different data rates, while Fig. 5b is the same result

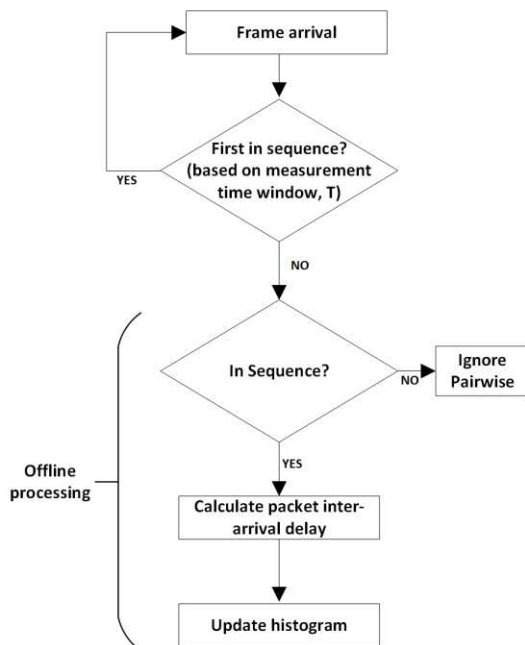


Fig. 4. Algorithmic view of the measurement procedure.

but with a packet size of 4000 bytes. Finally Fig. 5c shows the CCDF for a packet size of 4000 bytes, but with priority enabled. Note that the CCDFs are variable size step functions (discrete distributions).

There are clear differences between the individual subplots. For the 1500 bytes case there is some differentiation on the cumulative behavior between the different data rates. But for the 4000 byte case most of the values from the different data rates are grouped together indicating that for larger packet sizes, the probability of encountering larger delays, is higher.

With priority a clear improvement is seen as points are now distributed over longer delay values. This is more evident in Fig.6 by comparing the traces for the two packet sizes with and without a priority implementation. It is clear that with priority, points become more distributed (i.e. energy is removed from the smaller delay values and transferred into higher delay values). Fig.7 is a relative frequency count (histogram) of all measurements combined. This histogram is obtained by combing the packet inter-arrival delays for a number of packet sizes (512, 1500 and 4000) and a number of background traffic data rates (40, 100, 200 and 400 Mbps).

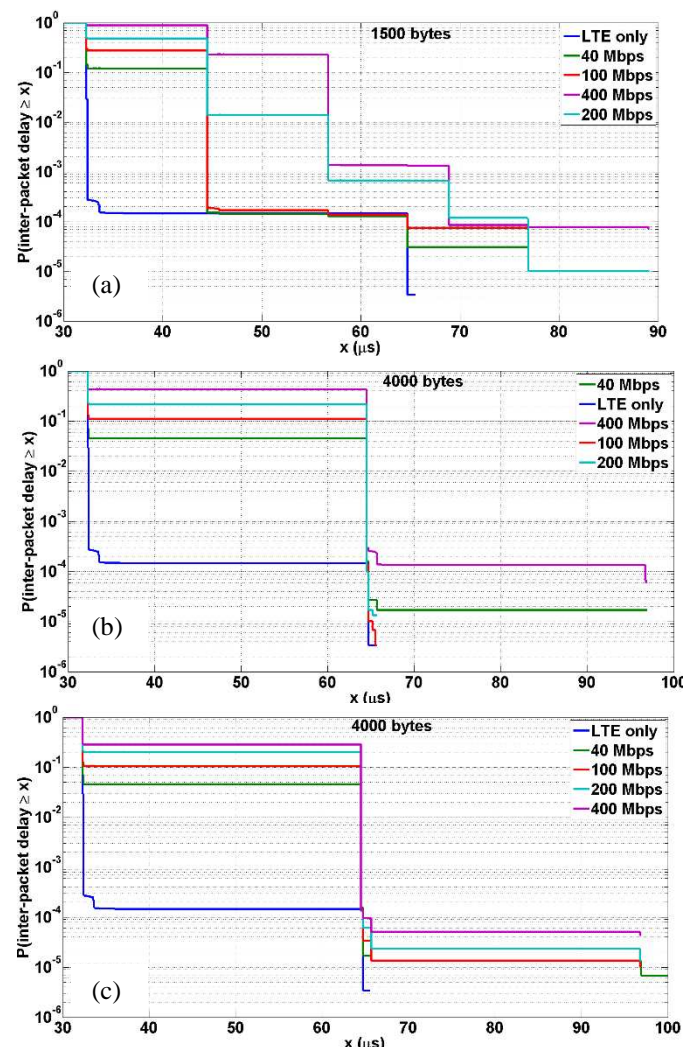


Fig. 5. CCDF of packet inter-arrival delays of the LTE traffic under different background traffic packet sizes and data rates. (a) 1500 bytes, (b) 4000 bytes, (c) 4000 bytes with priority enabled.

The plot shows that, by using a priority regime, energy from the larger delay bins is transferred into the smaller delay bins, indicating that values are more concentrated into the smaller delay values (i.e. mean and standard deviation are reduced).

The same effect is noticed on the combined CCDF in Fig.8. For buffer management, a measurement such as the one shown in this Figure is important as it can be used to inform an algorithm for resizing buffers according to statistical predictions. For example, in this case two values are annotated in the plot: Assuming buffers that can accommodate a packet inter-arrival delay of $68\mu\text{s}$, a frame drop will occur approximately once every 4000 frames without priority but once every 13000 frames with a priority implementation.

V. CONCLUSIONS

In this paper, we discuss some of the main issues (protection requirements, buffer management) in future fronthaul implementations, where different LTE physical layer functional

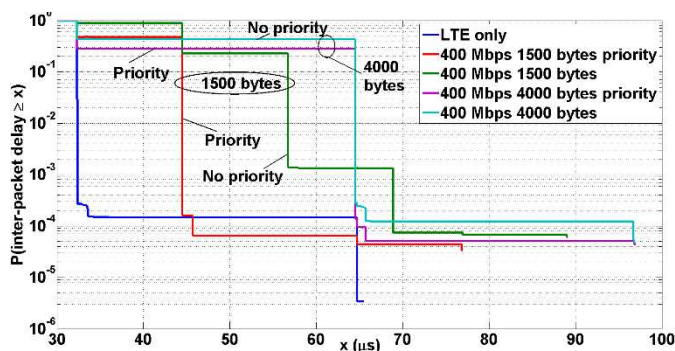


Fig. 6. Combining priority and non-priority results for 1500 and 4000 bytes packet sizes for a data rate of 400 Mbps.

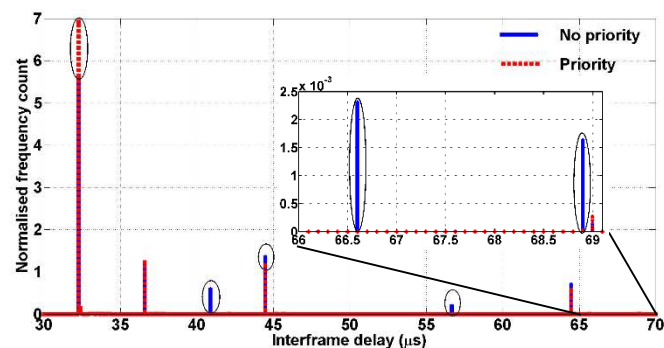


Fig. 7. Histogram plot for all measurements combined.

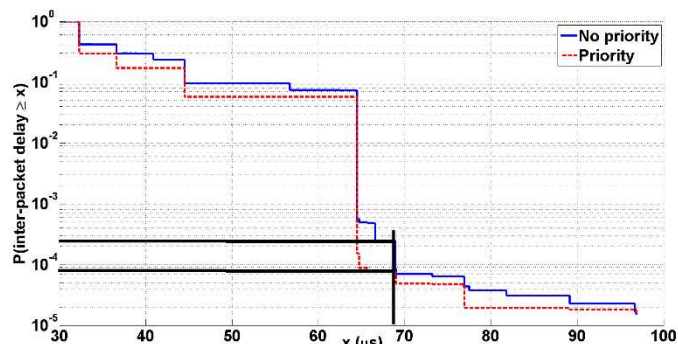


Fig. 8. CCDF plot for all measurements combined.

subdivisions are implemented. We have presented statistical distribution for packet inter-arrival delay measurements, from a testbed transporting emulated LTE traffic (I/Q transportation) and generic Ethernet traffic. The testbed employs smart probing techniques to filter the LTE traffic. Measurements were carried out for different background traffic packet sizes and data rates. The obtained distributions show how the statistics of the delays are affected by different background traffic packet sizes and data rates. Additionally we show how, through a priority regime, these distributions change. A combined PMF and CCDF for all measurement results (all data rates and packet sizes) is presented and the CCDF is used to show how such a measurement can potentially be used to adapt buffer sizes in a receiving node in the fronthaul (albeit without covering in this work the buffering algorithm that will be used following such a measurement). Buffer management, based on statistical distributions such as the ones presented here, is very important for future C-RAN fronthaul implementations transporting traffic from different LTE physical layer functional subdivisions as well as Constant-bit rate traffic (e.g. CPR).

ACKNOWLEDGMENTS

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

REFERENCES

- [1] Cisco (2015, Feb.), Visual Networking Index: Global Mobile Data Traffic Forecast Update 2014–2019 White Paper [Online]. Available: http://www.cisco.com/c/en/us/solutions/collateral/service-provider/visual-networking-index-vni/white_paper_c11-520862.pdf.
- [2] China Mobile (2011, Oct.), C-RAN: The Road Towards Green RAN (white paper) [Online]. Available: http://labs.chinamobile.com/cran/wp-content/uploads/Cran_white_paper6_v2_5_EN.pdf.
- [3] U. Dötsch et al, "Quantitative analysis of split base station processing and determination of advantageous architectures for LTE," in *Bell Labs Tech. J.*, vol. 18, no. 1, pp. 105-128, June 2013.
- [4] NGMN (2015, Mar.), A deliverable by the NGMN alliance: Further study on critical C-RAN technologies [Online]. Available: <https://www.ngmn.org/publications/technical.html>
- [5] China mobile et al (2015, Sept.), White Paper of NGFI (ver. 1.0 En) [Online]. Available: <http://labs.chinamobile.com/cran/>.
- [6] L. Zhao, M. Li, Y. Zaki, A. Timm-Giel and G. Gorg, "LTE Virtualization: from Theoretical Gain to Practical Solution," in *23rd Int. Teletraffic Congr. (ITC)*, San Francisco, CA, 2011, pp. 71-78.
- [7] W. Kiess et al, "Base station virtualization for OFDM air interfaces with strict isolation," in *IEEE Int. Conf. on Commun. Workshops (ICC)*, Budapest, Hungary, 2013, pp. 756-760.
- [8] <http://www.intelligent-nirvana.net/>
- [9] N. J. Gomes, P. Chanclou, P. Turnbull, A. Mageec, V. Jungnickel, "Fronthaul evolution: From CPRI to Ethernet," in *Optical Fiber Tech.*, vol. 26, par. A, pp. 50-58, Dec 2015.
- [10] iCIRRUS (2015, Jul.), iCIRRUS intelligent C-RAN architecture [Online]. Available: <http://www.icirrus-5gnet.eu/category/deliverables/>