

Kent Academic Repository

Full text document (pdf)

Citation for published version

Assimakopoulos, Philippos and Al-Hares, Mohamad Kenan and Gomes, Nathan J. (2016) Switched Ethernet Fronthaul Architecture for Cloud-Radio Access Networks. *Journal of Optical Communications and Networking*, 8 (12). B135-B146. ISSN 1943-0620.

DOI

<https://doi.org/10.1364/JOCN.8.00B135>

Link to record in KAR

<http://kar.kent.ac.uk/60771/>

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>

Switched Ethernet Fronthaul Architecture for Cloud-Radio Access Networks

Philippos Assimakopoulos, *Member, IEEE*, Mohamad Kenan Al-Hares, *Student Member, IEEE*, and Nathan J. Gomes, *Senior Member, IEEE*

Abstract— A fronthaul design for current and future mobile networks based on the transport of sampled radio signals from/to base station baseband processing units (BBUs) to/from remote radio heads (RRHs), is presented. The design is a pure-Ethernet switched architecture that uses virtual local area network (VLAN) identifiers for the RRHs and flow identifiers for the antenna ports, and is compatible with current standardization definitions. A comprehensive analysis for the limits of the Ethernet fronthaul in terms of the total number of antennas that can be supported is carried out, based on the latency imposed by the Ethernet network. The analysis assumes the transportation of control and management (C&M) and timing information (based on the precision-time protocol, PTP) but is valid for other types of background traffic (for example, that generated by the implementation of different long-term evolution (LTE) functional subdivisions, in a fronthaul with mixed processing). A low-cost testbed using “smart SFP” in-line probes is presented and used to obtain measurements from an Ethernet fronthaul, transporting mixed traffic. The measurements show how background traffic affects hybrid-automatic repeat request (HARQ) retransmissions, and are used to validate the analysis. The effects of contention of PTP packets is discussed and a simple solution to overcome the effects of contention is proposed.

Index Terms— C-RAN, Ethernet, fronthaul, LTE, 4G, 5G, radio-over-Ethernet, PTP

I. INTRODUCTION

NEW mobile network architectures are essential for meeting the capacity demands of current and future mobile systems (4G and 5G). These capacity demands are fueled by the increasing number of mobile devices and the increased number of applications that require ever-higher data rates [1]. The cloud-radio access network (C-RAN) is

an extension of the classical RAN implementation and can offer multiple advantages [2]. An example of a C-RAN implementation is shown in Fig.1. Here, a pool of baseband processing units (BBUs) are connected to remote radio heads (RRHs) through high speed links that transport the digitized radio signals. The main RAN complexity is placed in the central location while the RRHs are implemented at low cost and complexity. Operational, management and energy costs for the operator are reduced while the scalability of the infrastructure is improved, and the spectral efficiency of the operator’s network increased [2].

Digital implementation is likely to be the choice for the future (as opposed to analog transportation), leveraging current standardization efforts [3-7] and the potential use of off-the-shelf equipment. The use of Ethernet equipment can also pave the way for structural and operational convergence, where the infrastructure is shared by backhaul and midhaul links (note that in the context of this work, the term midhaul is used to refer to the transport between a small cell and its controlling macro-base station) and, possibly, fixed-access links. Current fronthaul implementations focus on-existing standards, mainly the common public radio interface (CPRI) [3], while current standardization efforts (e.g., the IEEE P1914.3) investigate encapsulating CPRI by Ethernet [6]. In addition, the IEEE P802.1CM work group is looking to overlay a synchronisation system on top of an Ethernet bridged network transporting CPRI-type streams, by defining standard time-sensitive networking (TSN) profiles (this includes the choice of existing TSN profiles and features and the definition of new ones if needed) [7].

The implementation of different functional subdivisions (or “split processing”) as a means of reducing the data rate requirements and obtaining statistical multiplexing gains is also currently being investigated (see for example the work by the IEEE1914 working group [8]). However, due to prior investment and backward compatibility issues, the fully centralized approach (that of digitized In-phase and Quadrature (I/Q) radio transportation) will remain in use, perhaps in fronthaul networks that offer a mixture of split and centralized processing, targeted at the specific needs of operators, for example, in multioperator scenarios.

In this paper, we investigate the use of Ethernet in the fronthaul: Ethernet equipment can offer economies of scale

P. Assimakopoulos, M. K. Al-Hares and N. J. Gomes are with the Communications Research Group, University of Kent, Canterbury, UK, (p.assimakopoulos@kent.ac.uk).

(it is practically ubiquitous), while Ethernet switching equipment can allow for a direct integration of virtualization/cloudification techniques and traffic load balancing. Additionally, “carrier Ethernet” offers fully standardized operations, administration and maintenance (OAM) for telecommunications [9-11].




Fig.1. Example of a C-RAN implementation.

The main challenge in using Ethernet in the fronthaul comes from the lack of frequency and phase/time synchronization (both are required for carrier grade operation) in standard Ethernet networking equipment. However, there are existing technologies that can either directly or through some modifications, be used in an Ethernet-based fronthaul. These include the use of synchronous Ethernet (SyncE) [12-14] and the precision-time protocol (PTP) [15].

In this work, an investigation into the limits of an Ethernet architecture for the fronthaul is carried out. While it is assumed that the use of SyncE can provide the necessary frequency synchronization, phase/time synchronization which arises due to latency constraints in the Ethernet fronthaul is a more significant challenge and is the focus of this investigation. The existence of background traffic (with focus on PTP and C&M) is also taken into account but the analysis can be used for any type of background traffic (including, for example, that generated by other functional split implementations). A low-cost testbed for monitoring key performance indicators (KPIs) in a mixed-traffic Ethernet fronthaul using “smart small factor pluggable (SFP)” in-line Ethernet probes is presented and measurement results obtained are used to validate the analysis.

The paper is organized as follows: in Section II, the bandwidth requirements and the manner of insertion of long-term evolution (LTE) radio frames into Ethernet frames are presented. Section III presents the Ethernet fronthaul design concept and the analysis for the timing constraints of the architecture and the number of physical antenna ports that can be supported by the design. Section IV presents measurement results that show the effects of violating these timing constraints and how contention in the network can affect the PTP performance. Finally, the conclusions are presented in Section V.

II. BANDWIDTH REQUIREMENTS FOR 4G/5G

Transporting the LTE radio frame over a fronthaul requires that the sampled output of the IFFT is quantized

prior to it being framed for transportation as shown in Fig. 2. Table I shows the resulting bit-rates for different choices of sample lengths (bits per sample) and different channel bandwidths. As bit-rate scales with the signal bandwidth, the transport of high bandwidth signals through this type of fronthaul becomes challenging. For example, although 5G systems are not yet standardized, it is possible that such systems will have channel bandwidths of the order of 1 GHz. For a sampling rate at the Nyquist limit and 16 bits per sample, the expected bit-rate would be approximately 32 Gbps, already a challenge for current Ethernet technologies. The situation becomes even more challenging by considering that these data rates are for a single physical antenna port. Multiple antenna systems, such as multiple-input and multiple-output (MIMO), would require this value to be multiplied by the number of antennas. Then, the required data-rate $D_{I/Q}$, which will need to be accommodated by the fronthaul per RRH is given by

$$D_{I/Q} = 2N_A f_s b f, \quad (1)$$

where N_A is the number of physical antenna ports, b is the number of bits per sample, f_s is the sampling rate for a 20 MHz channel, f is a carrier aggregation factor normalized to a 20 MHz channel, and the factor of two is for the In-phase and Quadrature components. As shown in Table II, very large aggregate bit-rates result.

The total baud rate B_r , per RRH, has to include the transportation infrastructure framing, control and management (C&M), precision-time protocol (PTP) and encoding overheads and is given by

$$B_r = (D_{I/Q} O_{E,I/Q} + D_{C\&M} O_{E,C\&M} + D_{PTP}) C, \quad (2)$$

where $O_{E,I/Q}$ and $O_{E,C\&M}$ are the encapsulation overheads for the I/Q data and C&M channels respectively, represented as a proportion of the data traffic, $D_{C\&M}$ is the C&M channel data rate, D_{PTP} is the PTP data rate over the fronthaul segment and C is the line coding overhead (either 10/8 for 8B/10B or 66/64 for 64B/66B). The encapsulation overhead will vary depending on the size of the Ethernet frame used, and whether jumbo frames are supported over the network. Note that jumbo frames are not standardized by IEEE or other standardization bodies, which can be an issue when considering vendor interoperability within a fronthaul network, however they are allowed by at least some specifications such as by the Metro Ethernet Forum (MEF) and most networking device vendors support jumbo frame sizes in excess of 9000 octets.

Assuming maximum transmission unit (MTU) sizes in the range of 500 to 9000 octets the overhead can vary from 10.2% down to 0.56%, respectively.

Inserting the sampled signals into an Ethernet frame involves a mapping of the quantized outputs of the IFFT into the maximum transmission unit (MTU) portion of an Ethernet frame, as shown in Fig. 3, in effect, “slicing up” the radio frame in each resource element column (frequency axis).

TABLE I
TOTAL DATA RATE REQUIRED AFTER SAMPLING OF IN-PHASE AND QUADRATURE COMPONENTS

Channel BW (MHz)	Sample rate (MHz)	Data rate (Gbps)		
		20 (bpS)	16 (bpS)	8 (bpS)
20	30.72	1.229	0.983	0.492
40	61.44	2.458	1.966	0.983
60	92.16	3.686	2.949	1.475
80	122.88	4.915	3.932	1.966
100	153.6	6.144	4.92	2.458
5G ¹	1000	40	32	16

¹Expected for 5G and assuming a bandwidth of 1 GHz and sampling at the Nyquist rate theoretical limit.

TABLE II
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)	No. of antennas per sector at RU				
		4	8	16	64	128
		Data rate (16 bpS) (Gbps)				
20	30.72	3.93	7.86	15.73	62.91	125.8
40	61.44	7.86	15.72	31.47	125.82	251.7
60	92.16	11.8	23.59	47.18	188.74	377.5
80	122.88	15.73	31.45	62.91	251.7	503.3
100	153.6	19.68	39.36	78.72	314.9	629.8
5G	1000	128	256	512	2048	4096

Each “slice” shown in Fig. 2 will have a bandwidth dependent on IFFT size and a time duration of T_s . As an example, a 20 MHz LTE signal with 16-bit quantization is assumed here. (Note that although only 1200 data subcarriers are shown in the slice, the actual time domain signal is oversampled due to the inclusion of null subcarriers resulting in an IFFT size of 2048).

The encapsulation overheads include the L2 headers (including virtual local area network (VLAN) identifier (ID), following the IEEE802.1Q standard [16]) and the radio-over-Ethernet (RoE) header, where a size of 8 octets is assumed (see [6] for the latest RoE packet formats).

The number of bits per slice b_s , is given by

$$b_s = 2bfN_{IFFT}, \quad (3)$$

where N_{IFFT} is the maximum IFFT size defined in LTE (i.e. 2048). The number of Ethernet frames per slice F_s , is given by

$$F_s = \lceil b_s / MTU \rceil, \quad (4)$$

where MTU is the MTU size of the frame. The result of the mapping of slices into Ethernet frames is shown in Table III for $b=16$. With Ethernet jumbo frames, using a 20 MHz radio bandwidth, one slice fits into one jumbo frame. A 100

MHz radio slice requires 5 jumbo frames per slice.

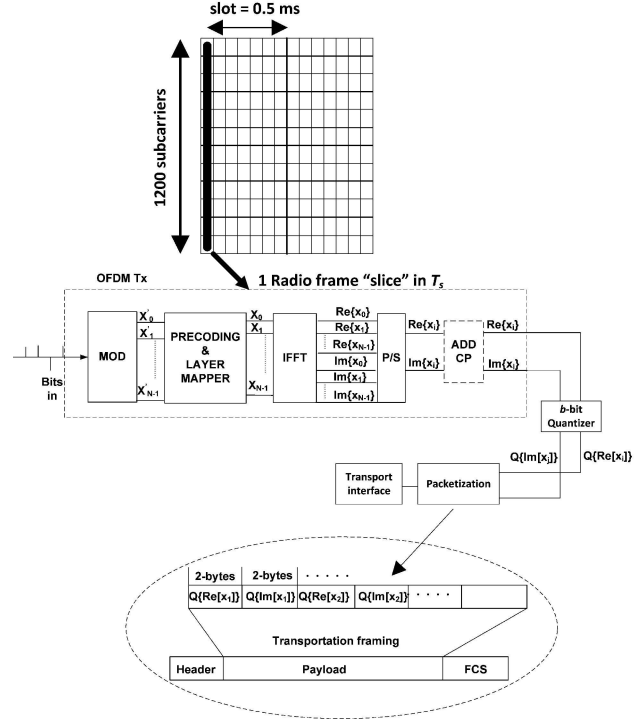


Fig. 2. “Slicing” of a radio frame. A 2048 IFFT (20 MHz) signal is assumed here with a sample size, b , of 16 bits per I and Q sample. MOD=modulation, IFFT=inverse fast Fourier transform, CP=cyclic prefix, P/S=parallel-to-serial, FCS= frame check sequence, OFDM=orthogonal frequency-division multiplexing.




Fig. 3. Insertion of I/Q data with into Ethernet frames. Note that the full MTU size is used. The MAC header includes source and destination addresses, ETH type and VLAN ID fields. IFG=interframe gap.

Chan nel BW (MHz)	Octet s (16- bit) per slice	Jumb o frame s per slice	Stand ard frame s per slice	No. of jumbo frames per radio frame	No. of standard frames per radio frame
20	8192	1	6	140	840
40	16384	2	11	280	1540
60	24576	3	17	420	2380
80	32768	4	22	560	3080
100	40960	5	28	700	3920

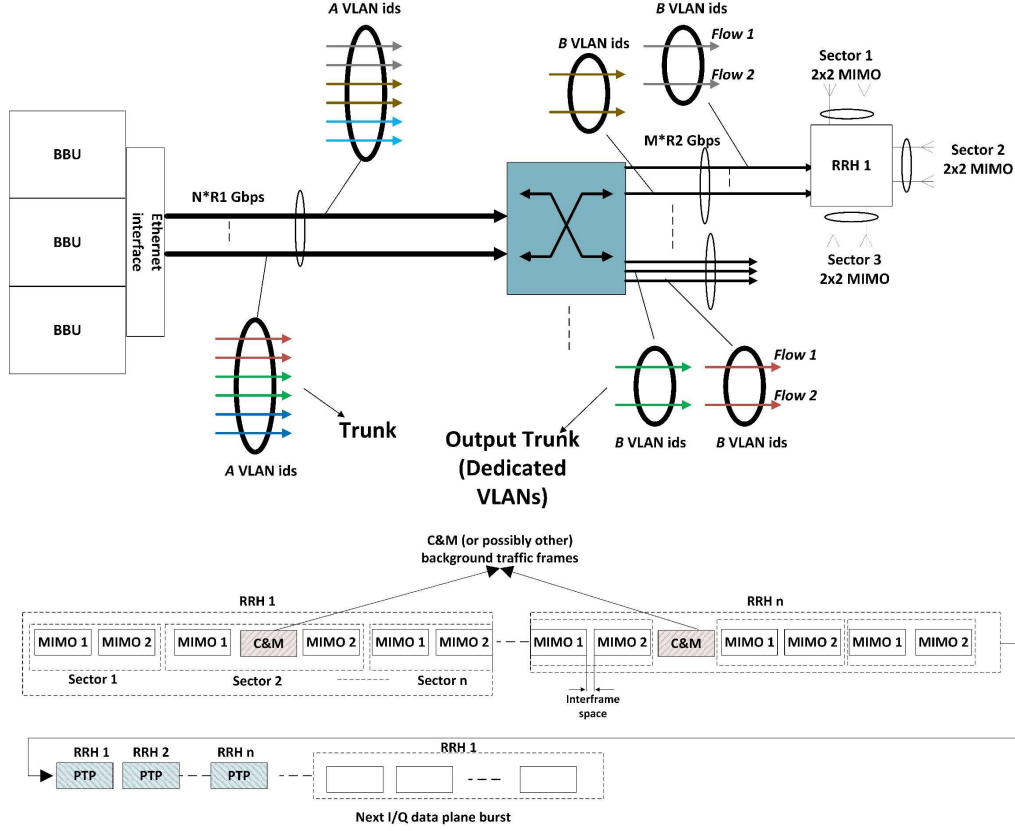


Fig. 4: Fronthaul architecture with *VLAN trunking* for addressing of the RRHs/sectors. An example RRH with three sectors and 2x2 MIMO is shown with $B=2$. Two RoE flows are generated for each sector each one addressing a separate antenna (top). One possible addressing scheme for the antennas (round-robin) is shown here. The I/Q data plane is produced as a burst which then allows for PTP traffic to be transported without contention (bottom). Note that the Ethernet frame sizes are not shown to scale.

III. FRONTHAUL ARCHITECTURE AND ETHERNET LATENCY

Fig. 4 (top) shows the fronthaul architecture, which illustrates the concept of *VLAN trunking*. The different RRHs and/or sectors within RRHs are addressed through VLAN IDs, with the use of dedicated VLANs, while individual antennas are addressed through flow IDs. The first trunk is set in the R1 link while the second trunk is set on the R2 link where generally the condition $A > B$ will hold, with A being the number of antenna ports supported by a single R1 link and B the number of antenna ports supported by a single R2 link.

Note that although the flows shown in Fig. 4 (top) are for the downlink direction the same flows will exist in the uplink direction as the architecture is reciprocal. The round-robin scheme shown in Fig. 4 (bottom) is only an example but fundamentally there is no need for a specific addressing scheme in the downlink or uplink directions provided that certain time limits are met. The study and analysis of these time limits is the focus of this Section.

At both ends of the fronthaul, a convergence layer is used (interfacing the LTE physical layer with the Ethernet layers). Here, it is assumed that the convergence layer is based on RoE [6]. Thus, the RoE mapper will encapsulate the radio samples into a RoE packet and then the Ethernet process will encapsulate the RoE packet according to Fig. 3.

A detailed view of these processes is shown in Fig. 5. The software-defined networking (SDN)-type controller entity is used to estimate uncontended windows for PTP transmissions (an issue that will be discussed in Section IV), but can also be used for purposes such as traffic steering and protection. Additionally, a self-optimized network (SON) entity can be used to obtain long-term performance measures and can instruct the SDN controller to adapt fronthaul operation based on those measures. Both entities obtain inputs from a probing system implemented through the use of “smart SFP” Ethernet probes [17, 18]. VLAN IDs are set during initial fronthaul configuration (the switch can be configured remotely) but antenna ports in sectors and/or RRHs that are addressed by each BBU in the pool can be dynamically re-assigned during fronthaul operation by configuring VLAN IDs (through the SDN controller).

Fig. 6 shows an overview of the main delay components on a switched Ethernet network. The serialization delay T_{se} , is given by

$$T_{se} = P / R, \quad (5)$$

where P (bits) is the frame size and R (bits/s) is the supported interface speed. Serialization in the switch occurs when the switch operates in store-and-forward mode (this is

the standard operating mode for an Ethernet bridge, as specified in IEEE802.1D [19]).

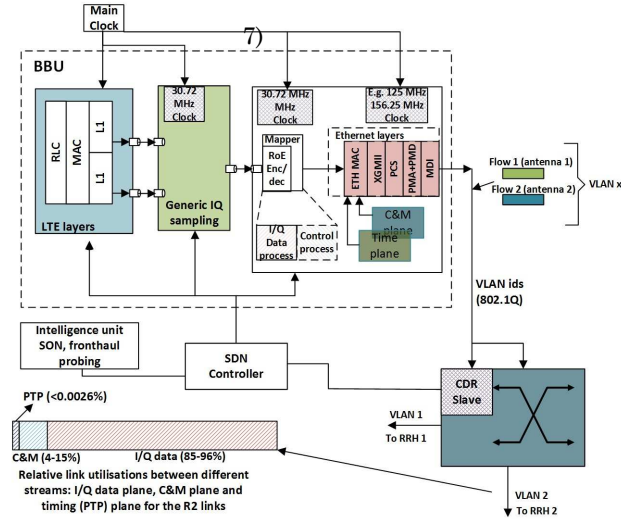


Fig. 5 Detailed view of the BBU side processing (DL direction). For the Ethernet layer definitions, see the IEEE802.3-2015 [20]. CD=clock and data recovery, SDN=software-defined networking.

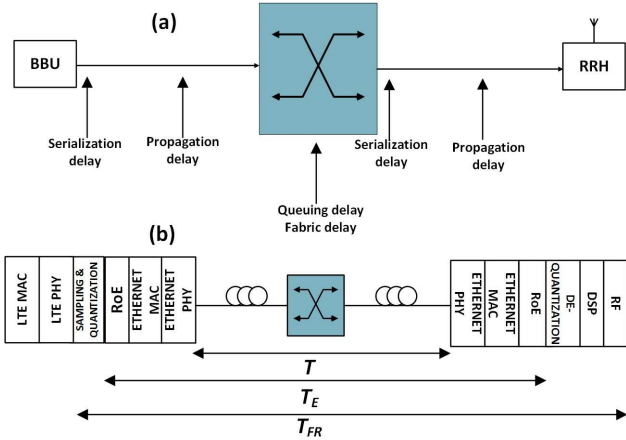


Fig. 6. Main delay parameters in a switched Ethernet fronthaul. Note that T is based on the delay definitions shown in (a) i.e. including serialization delays. Higher LTE protocol layers are not shown.

On the other hand, a cut-through switch operates in a store-and-forward mode only when multiple frames need to be output from the same port, such that buffering has occurred, and/or under certain conditions when there is a rate transition between an input port and the output port. Fig. 7 shows the serialization delays for different frame sizes and Ethernet link rates.

The propagation delay T_p , in fiber is approximately

$$T_p = d / (c / n), \quad (6)$$

where n is the refractive index of silica (≈ 1.45), c is the speed of light in vacuum and d is the distance from BBU to RRH.

The total physical layer end-to-end latency T_{s-n-f} ,

neglecting the interframe space, for a frame in a store-and-forward switching regime, is given by

$$T_{s-n-f} = (N + 1)T_{se} + T_p + N(T_f + T_q), \quad (7)$$

where N is the number of switches in the path, T_f is the fabric delay for each switch and T_q is the queuing delay in each switch. Equation (7) treats the architecture as a single networking entity and thus considers first-bit-in to first-bit-out (FIFO) delays.




Fig. 7. Serialization delay for different frame sizes and different Ethernet link rates.

The total Ethernet delay shown in Fig. 6, T_E , includes the processing delay of the Ethernet and RoE layers, while the total fronthaul delay T_{FR} , further includes the processing following the LTE physical layer in the BBU and the processing following the Ethernet layers in the RRH. Note that the second serialization delay in Fig.6 (a) at the output port of the switch is valid only in store-and-forward operation. The total end-to-end delay in a cut-through switching regime (queuing is neglected in this case), T_{c-t} is given by

$$T_{c-t} = T_{se} + N(T_c + T_f) + T_p, \quad (8)$$

where T_c represents the serialization delay introduced by the switch having to store and read the fields up to and including the VLAN ID field. The delay introduced by a single switch is now reduced from $T_{se} + T_f$ in store-and-forward operation to $T_c + T_f$ in cut-through operation.

The total end-to-end latency in the fronthaul needs to be constrained such that LTE protocol timings are not violated: namely, HARQ acknowledgements and physical random access channel (PRACH) procedures. The former has the much more stringent requirements. The processing in the BBU (from reception of uplink transmission to generation of ACK) is generally assumed to be complete after 2.75 ms (a typical vendor specification [21]). The HARQ ACK message, corresponding to an uplink (UL) transmission, needs to be received after three subframes (a subframe is a sub-division of the LTE radio frame corresponding to a time duration of 1 ms and is also known as the transmission time interval, TTI) following the subframe in which that uplink transmission occurred.

Therefore, the additional round-trip time (RTT) allowed for the fronthaul will be in the order of 250 μ s. A similar RTT budget is available if a negative-ACK (NACK) is transmitted in the downlink (DL). The user equipment (UE) needs to resend its data using the same resources (a nonadaptive retransmission) three subframes after the subframe that the downlink transmission occurred in. Based on the latency definitions in Fig. 6, the maximum allowed RTT and the corresponding available processing delay budget for the fronthaul are shown in Fig. 8.

The queuing delay will depend on the switching regime and on the A to B ratio, with certain choices of values for this ratio resulting in queuing at the switch.

There are two important aspects regarding the delay in the fronthaul and the transportation of LTE radio frames. The first is the total end-to-end latency. This needs to be constrained as previously discussed, so that LTE protocol timings, mainly HARQ processes, are not violated. The second has to do with the physical layer LTE radio frame timings: the average time duration from reception of one slice to the next for the same antenna port, needs to be within T_s . This average inter-slice arrival time $T_{<slice>}$, must meet

$$T_{<slice>} = F_s T_{<frame>} \leq T_s, \quad (9)$$

where $T_{<frame>}$ is the inter-frame (Ethernet) arrival time. The time window over which the average is taken in this case is equal to the scheduling resolution of LTE (i.e. the 1 ms subframe or transmission time interval, TTI).

This requirement means that the separation in time between received slices in each antenna port in the RRH must be within the fundamental slice duration T_s , of the transported LTE signals.




Fig. 8. RTT fronthaul delay and fronthaul processing delay budget based on HARQ timings for different numbers of Ethernet switches (N) and switching regimes. Assuming: $T_f = 5 \mu$ sec, 10 GbE operation with VLAN addressing (refer to (7), (8) and Fig. 6). S-n-f=store-and-forward.

If this requirement is not met, then the subframe timings of the transmitted LTE signals will be violated.

Higher bandwidth signals will require more frames to be received (taken into account in (9) by F_s) in order to reassemble the whole slice and thus will have a shorter inter-arrival time requirement for the Ethernet frames.

This can be compensated with the use of higher Ethernet link rates, offering lower serialization delays, but will also depend on the number of served antenna ports per each R1 link. Additionally, the requirement becomes more stringent with increased overheads in the Ethernet layer. These include the encoding overhead and the addition of the Ethernet and RoE headers (see Fig. 3). As an example, for a 20 MHz radio frame (i.e. for $f=1$), a jumbo frame will contain the whole slice (see Table III) and thus the inter-slice and inter-frame arrival delays will be equal. But, for a 100 MHz radio frame (i.e. for $f=5$), the inter-frame arrival time will need to be smaller by a factor of five as the antenna port will need to receive five jumbo frames to reassemble the whole slice.

Most importantly, the inter-frame delay in (9) will also include any background traffic transmissions (see Fig. 4 (bottom)), the effect of which will be dealt with in Section IV. In general, for an architecture as in Fig. 4, with arbitrary numbers of MIMO antennas, sectors and RRHs, the requirement that must be met is given by

$$R \geq B_r, \quad (10)$$

where R is the Ethernet link rate. However, (10) can be expressed in terms of the inter-slice timing requirement in (9) according to

$$F_s A T_1 + T_{1,I} \leq T_s, \quad (11)$$

where A , is the number of antenna ports serviced by a single R1 link, T_1 is the frame serialization delay at the input of the R1 links and $T_{1,I}$ is the frame serialization delay for the intermittent traffic (includes C&M and PTP traffic).

The penalty for an architecture that does not meet the requirement in (11), will be the onset of HARQ retransmissions. But it is important to note that these retransmissions will be a result of exceeding the allowed inter-slice delay and not a result of exceeding the RTT budgets shown in Fig. 8. As the LTE subframe timings need to be maintained over-the-air, the RRH will need to schedule and transmit radio frame slices even if due to increased latency in the fronthaul, it has not yet received all the radio slice samples through the Ethernet interface. How the RRH does this will be implementation dependent but one example is through the insertion of dummy data (e.g. nulls) in the place of the missing samples. Experimental results that show this effect will be presented in Section IV.B.

Equation (11) is valid provided that

$$(T_2 + T_{2,I}) / (T_1 + T_{1,I}) < A / B, \quad (12)$$

is always true. This reflects the fact that the R2 link can represent a bottleneck in the system. If (12) is not met, then by the time a frame has been serialized out of the R2 port in the switch, the next frame to follow is already being queued for transmission. This is not a sustainable condition as with time it can lead to the formation of an infinite queue (or in practical conditions, to dropped frames in the switch due to

buffer overflows). Note that (12) assumes that the link rates from the switch to the RRH(s) are the same. However, it would be straight forward to modify this expression to include different rates. One example would be to use the worst case delay (corresponding to the lowest rate from Switch to RRH) which would then replace the T_2 delays in (12). Note also that in such a case although the rates will be different, they will still be related (i.e. integer multiples).

The condition in (12) is important for another reason: if this condition is met, it is implied that other than having an effect on the RTT, the use of cut-through switching (when this is possible) will not affect in any way the inter-frame (and inter-slice) delay. Note that cut-through switching requires very specific conditions to be met, namely the guarantee of no contention and the avoidance of buffering due to rate transitions. Therefore, although cut-through switching is treated in this work for completeness, it is considered as a rather special case. The experimental results that will be presented on Section IV will concentrate only on store-and-forward-operation.

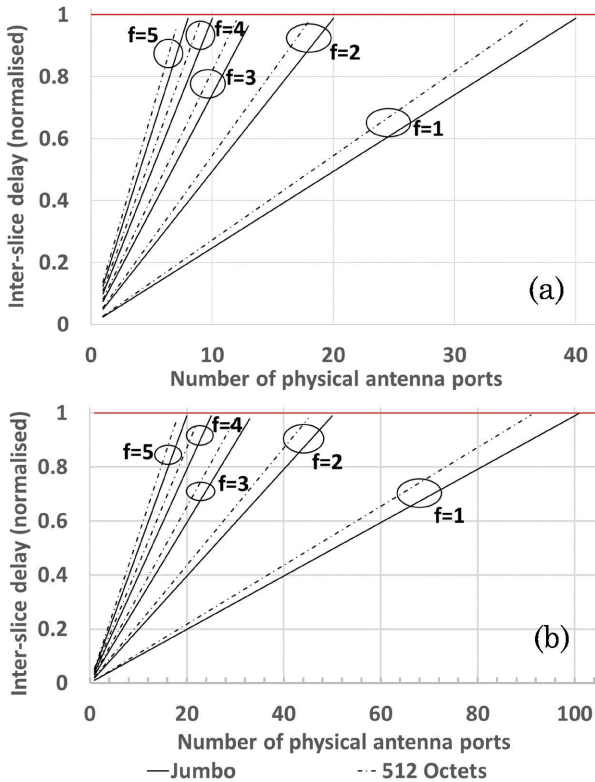


Fig. 9. Inter-slice delay versus number of supported physical antennas for different carrier aggregation factors and for two different lengths for the Ethernet frames transporting LTE. (a) R1 links are 40 GbE and (b) R1 links are 100 GbE.

Furthermore, the architecture and corresponding predictions presented here are valid for both cut-through and store-and-forward switching but for the later, only provided that condition (12) is met. Note also that these predictions are valid for both uplink and downlink directions as the architecture is reciprocal. Using (11), the limits of the architecture, in terms of the

number of physical antenna ports that can be supported, can be calculated. The result of this is shown in Fig. 9 for two different frame lengths (jumbo and 512 Octets).

Note that with the smaller frame length, the traces move towards the left, i.e. smaller numbers of antennas, can be supported. This is a result of reduced overhead efficiency when using smaller frames. Table IV summarizes the maximum number of antennas that can be supported.

TABLE IV
NUMBER OF ANTENNA PORTS THAT CAN BE SUPPORTED BY THE ARCHITECTURE FOR DIFFERENT ETHERNET LINK RATES

f	LTE frame size/octets	1	2	3	4	5
Max Antenna ports		R1=40 GbE				
	Jumbo	40	20	13	10	8
	512	36	18	12	9	7
		R1=100 GbE				
	Jumbo	101	50	33	25	20
	512	91	45	30	22	18

Fig. 10 shows the allowed background traffic burst sizes within a subframe duration for different background traffic and LTE-carrying Ethernet frame lengths for a 20 MHz signal ($f=1$). Again, the effect of improved overhead efficiency with smaller length of LTE-carrying frames is evident. In terms of the background traffic frame length, obviously the burst size reduces for larger frame length but what is not evident from the graph is that the actual throughput of the background traffic improves with larger frame lengths due to the improved overhead efficiency.




Fig. 10. Allowed background traffic burst sizes for different background traffic and LTE-carrying Ethernet frame lengths, for $f=1$. Assuming that the R1 links are 100 GbE.

IV. EXPERIMENTAL RESULTS

The testbed used for the measurement results that are presented in this Section is shown in Fig. 11. LTE signals are generated by a software base station (Amari LTE-100) in the form of time-domain I/Q samples. The samples are quantized and inserted into user datagram protocol (UDP) packets (similar to the process described in Fig. 2, albeit

with more overheads). The corresponding Ethernet frames are transported through the fronthaul and received by a RRH (Ettus N210 USRP) where they are de-packetised, de-quantised and up-converted to RF and transmitted over-the-air. The fronthaul comprises of two 1 GbE switches operating in store-and-forward mode, with a trunk between them where contention can take place with background traffic generated by a Viavi hardware-based traffic generator (note the choice of a store-and-forward switching here, as we examine cases where contention arises, there is no obvious advantage in trying to use the cut-through mode of operation.).

The traffic streams (LTE and background) are logically separated by VLAN IDs. All links are formed by standard 1000BASE-LX Small Form-factor Pluggable (SFP) transceivers with LC connectors and Single Mode Fiber (SMF) patch-cords. The LTE traffic is captured using a number of Viavi in-line Ethernet probe that come in the form of a 1000BASE-LX SFP. A filter is applied that instructs the probe logic to capture all packet headers containing the destination MAC address and destination UDP port of the RRH. Once captured, the headers are timestamped (using a propriety form of the Precision Time Protocol, PTP) and re-encapsulated (with the discovered network encapsulation), and with an additional Viavi proprietary header that includes the timestamp (in addition to other metadata fields). The captured packet headers are re-injected into the network as frame result packets (FRPs) and sent to a packet routing engine (PRE) which routes them to a management station for further processing. Additional fields include the SFP probe ID and FRP injection number, both of which are used as inputs to the in-house algorithm, implemented in Matlab.

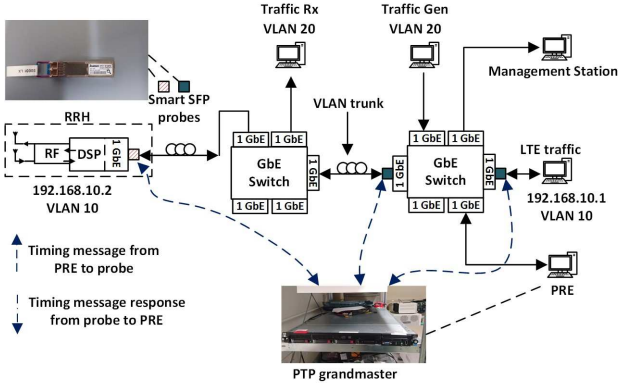


Fig. 11. The testbed used for the measurements. PRE=packet routing engine, RF=radio frequency, DSP=digital signal processing.

A flow chart for the algorithm is shown in Fig. 12. Results are obtained from the management station through a Wireshark capture of the FRP frames (step a). Within the FRP frames, sending and receiving IP addresses correspond to the eNodeB and RRH (and not to the management station and PRE). As there is no guarantee that the PRE will route FRPs to the management station in a specific order, the timestamps need to be ordered according to which probe they came from (step b). Additionally, the range and values of the injection numbers across the two

probes may not be the same. Therefore, these have to be normalized to start at the same value (Step c). The algorithm looks for missing injection numbers, for example, due to dropped packets or captures that were not injected back in-line due to congestion, and removes the equivalent injection number for the other probe (step d). Then, transit delays are calculated by subtracting timestamps corresponding to the equivalent injection numbers of the two probes (step e). At this stage, additional KPIs can be extracted, including frame-delay variation and inter-frame delays (step f) [17, 18]. The final step is the calculation of the corresponding statistics (step g).

Note the correspondence of the testbed with the architecture of Fig. 4. The first switch in the testbed corresponds to the Ethernet interface in the architecture, the trunk between the switches to the R1 links and the second switch to the switch of Fig. 4. Thus, the testbed can be used to verify the predictions of (11) with some limited modification in the calculations carried out for the results of Figures 9 and 10 regarding the overhead (addition of UDP and IP headers and non-inclusion of RoE headers).

Regarding the impact of packetisation and de-packetization, these processes are fundamentally implementation dependent. For this work we ensure that any latency introduced by this processing is not such that the RTT requirements for correct HARQ operation are violated. Furthermore, the results that will be presented in Subsection B investigate the inter-slice and inter-frame delays. As the packetisation and de-packetisation delays will be approximately the same for all packets, these processes will not have an effect on the inter-slice or inter-frame delays.




Fig. 12. Flow-chart of in-house algorithm for obtaining KPIs from the testbed presented in Fig. 11.

A. PTP delay asymmetry issues

For proper operation of PTP, there should be no delay asymmetry between the PTP grandmaster and the slave clocks. Delay asymmetry will result from networking design (e.g. variability between uplink and downlink networking segments) and from different levels of contention between DL and UL paths.

Whatever the source, the result will be an error in the timestamping equal to one-half of the delay asymmetry. This error will remain constant until the next timing update takes place.

Fig. 13 shows three different estimates for the delay

through the fronthaul obtained by timestamping through PTP. Note that all estimates are relatively stable, but there is a large offset between them. Two of the estimates are wrong (indicated as stable PTP errors), while the actual fronthaul delay (approximately 37.7 μ s) is more accurately estimated during a time period in which contention does not take place.

One method of overcoming the contention issue, is to make use of bursting and transmit the PTP packets during “silent” periods in-between the bursts.

Note that bursting is not a necessary method to gain an uncontended time period. For a transmission scheme where the I/Q data traffic is transported at a constant frame rate instead, a scheduler could be used to transmit the PTP packets during periods between LTE-carrying frames.




Fig. 13. Measured delay estimates for the fronthaul obtained through PTP timestamping. Two of the traces indicate erroneous values for the fronthaul delay (a stable error), while one is close to the actual value, which is approximately 37.7 μ s.

But with a bursting implementation, a scheduler that transmits the PTP packets after the I/Q burst has been transmitted, is a possibly lower-complexity implementation. This fact is indicated in Fig. 14 where an example of bursting to 40 antennas is shown (these can be divided amongst RRHs and sectors). The uncontended time window for PTP transmission is indicated on the graph, normalized to the subframe time. If the scheduler knows the number of served antennas, the Ethernet link rate, and the length of the Ethernet frames used, then it is straightforward to estimate when this window will occur and its size.




Fig. 14. By allowing bursting of the I/Q carrying frames, a “time window” for PTP traffic becomes available, shown here normalized to the slice duration. The example value shown here is for the DL direction for a 20 MHz bandwidth ($f=1$) and 100 GbE R1 link.

B. HARQ Retransmissions Results

The predictions shown in Fig. 10 for the background traffic allowance, do not take into account the switch scheduler operation. A scheduler in an Ethernet switch makes decisions based on whole frames (unless some form of pre-emption is implemented). It will generally attempt to balance the traffic from different input ports that need to be scheduled over the same output port. As this is done for whole frames, the prediction from (11) needs to be rounded based on the number of background frames the scheduler will insert into the output queue.

An example of a scheduler attempting to balance the traffic load over a trunk port is shown in Fig. 15. The y-axis here is the inter-frame delay of the LTE-carrying frames and is normalized to the serialization delay of the frame, which in this case has a length of 2000 octets. The background traffic frame length is 512 octets. Obviously, without background traffic, the inter-frame delay will be equal to one (normalized). But when there is background traffic, the scheduler in the switch will balance the load by inserting into the output trunk port queue four 512 octet frames for every LTE-carrying Ethernet frame. Therefore, the inter-frame delay will increase by a factor of two (approximately). This will continue until the end of the burst where the normalized inter-frame delay will fall back to one, as shown.




Fig. 15. Measurement results of inter-frame delays. Circle indicates the point where background traffic frames (approximately 4 frames in this result) are inserted into the switch queue in-between LTE-carrying frames. Note, that the delay is normalized to the serialization time of the LTE-carrying frame.

By rearranging and expanding (11), the maximum inter-frame delay (given as the number of background traffic Ethernet frames $N_{f,I}$), before HARQ retransmissions are triggered (i.e. on the onset of HARQ retransmissions), will be given by

$$N_{f,I} \leq \text{round} \left(\frac{R(T_s - F_s A T_1)}{P_I F_s}, k \right), \quad (13)$$

where P_I is the length of the background traffic Ethernet frame, and k is a factor that takes into account the scheduler operation and thus determines how the rounding is carried out. For this specific scheduler it is given by

$$k = \lceil P_{LTE} / P_I \rceil, \quad (14)$$

where P_{LTE} is the length of the LTE-carrying Ethernet frame.

The agreement between the theoretical predictions from (13) and measurement results is shown in Fig. 16 (note that here, the plotted results from (13) are normalized to the LTE-carrying Ethernet frame length). But, the improvement expected by the use of larger frame lengths for the LTE traffic (due to improved overhead efficiency) is not seen, as it is being “masked” by the rounding function in (13). Some of the values in Fig. 16 will cause higher delays than that allowed by (11). As the scheduler can only schedule whole frames for transmission, this is unavoidable and it means that for certain background frame sizes, there is no smooth transition into delays that are longer than that allowed by (11), and at which HARQ retransmissions begin to take place. This is also the reason that the traces do not seem to follow a clearer trend. Obviously, if the occurrence of excess delay is more frequent, the number of retransmissions will increase. Furthermore, the effect will be stronger for longer background traffic frame lengths. This is shown in Fig. 17 in terms of LTE transport block (TB) retransmissions (as a percentage of total TBs transmitted in each subframe).




Fig. 16. Comparison of measured and theoretical results for the inter-frame delay on the onset of HARQ retransmissions, for different LTE-carrying and background traffic Ethernet frame lengths.




Fig. 17. Transport block retransmissions versus background traffic Ethernet frame size for the same burst size and a bit rate of 50 Mbps. The LTE-carrying frame size is 2000 octets.

For this result, the burst size of the background traffic was kept constant while the data rate was set to 50 Mbps. Note

that the increase of retransmissions with frame size is generally linear in nature.

Figure 18 shows complementary cumulative distribution functions (CCDFs) for the results of Fig. 17. Note that smaller frame sizes with a non-fractional relation to the LTE-carrying frame size can lead to longer delays. As was discussed for the results in Fig. 16, the scheduler will attempt to balance the traffic load but this is problematic for non-fractional frame size ratios. For example, the 1500 octet trace will lead to larger delays than the 2000 octet trace when the scheduler inserts two background traffic frames for each LTE-carrying frame (which for these results is 2000 octets long). But fundamentally, these longer delays do not occur as often so as to induce more retransmissions (as shown in the results of Fig. 17). This fact is indicated by the circle annotation in the Figure around the 30 μ s value, that shows the values in the delay statistics that are responsible for the increased re-transmissions with increased background traffic frame size. These results depend on the scheduler operation. For example, a packet-based round-robin scheduler with equal weight queues (like the one presented in [17, 18]), will result on a similar amount of retransmission for the 2000 octet frame size case but smaller number of retransmissions for the non-fractional frame sizes of 1500, 1700 and 1850 octets.




Fig. 18. CCDF plots for the results of Fig. 17. The circle annotation indicates the group of values that are responsible for the increased retransmissions with larger background traffic frame size. Retx=retransmissions.

It is also important to note that the results of Fig. 18 indicate that when monitoring the KPI performance of the fronthaul, a moving average filter with a window size set properly (to match the averaging window in (9)) will have to be used. Otherwise the statistics obtained will not provide an accurate estimate regarding the onset and amount of HARQ retransmissions.

V. CONCLUSION

A comprehensive analysis of the latency constraints imposed by the use of Ethernet in an I/Q-based C-RAN fronthaul, transporting mixed traffic, has been carried out. It was shown that for larger 4G bandwidths (in the order of 100 MHz) or those envisaged for 5G, the only viable transportation options will require the use of very high Ethernet link rates (40 GbE and possibly 100 GbE), otherwise the number of supported antennas will be significantly limited. Design options for meeting subframe timing requirements by controlling the inter-frame delay in the Ethernet fronthaul were presented. For PTP

requirements, it was proposed that a simple solution making use of frame bursting could be used. A low cost testbed and in-house algorithm for obtaining KPIs were presented. The testbed makes use of “smart SFP” Ethernet probes to obtain the KPI measurements. The measurements results show the effects of contention of the PTP packets and validate the latency analysis by showing the onset of HARQ retransmissions. It can be concluded that the Ethernet switch scheduler operation would be a key component of the fronthaul and has to be considered when making capacity predictions. Furthermore, the accuracy of performance predictions that are based on the statistics of the fronthaul KPIs will depend on whether the scheduler operation and the amount of buffering at the end nodes are properly taken into account.

ACKNOWLEDGMENT

This work was carried out within the framework of the European Union’s Horizon 2020 research and innovation program under grant agreement No 644526 (iCIRRUS project) and EPSRCs “Towards an Intelligent Information Infrastructure (TI3)” program (NIRVANA project). Mohamad Kenan Al-Hares acknowledges funding through an EPSRC Doctoral Training Partnership (DTP), and Philippos Assimakopoulos by the NIRVANA project. The authors would like to thank Simon Hill (Viavi Solutions) and Peter Turnbull (ADVA Optical Networking) for interesting discussions. Data used in this work is stored in Kent Academic Repository (<https://kar.kent.ac.uk/>).

REFERENCES

- [1] A. Zakrzewska, S. Ruepp and M. S. Berger, “Towards converged 5G mobile networks-challenges and current trends,” in *Proc. ITU Kaleidoscope Academic Conf.: Living in a converged world - Impossible without standards?*, St. Petersburg, Russia, 2014, pp. 39-45.
- [2] China Mobile (2013, December), *C-RAN: The Road Towards Green RAN (white paper)* [Online]. Available: <http://labs.chinamobile.com/cran/2014/06/16/c-ran-white-paper-3-0/>
- [3] CPRI (2014, July), *CPRI Specification V6.1, Interface Specification* [Online]. Available: <http://www.cpri.info/spec.html>
- [4] ETSI (2014, October), *Open Radio equipment Interface (ORI); ORI interface Specification; Part 1: Low Layers (Release 4)* [Online]. Available: <http://www.etsi.org/technologies-clusters/technologies/ori>
- [5] Metro Ethernet Forum (2014, May), *MEF 12.2 Carrier Ethernet Network Architecture Framework Part 2: Ethernet Services Layer* [Online]. Available: <https://www.mef.net/carrier-ethernet/technical-specifications>
- [6] IEEE P1914.3, *Standard for Radio Over Ethernet Encapsulations and Mappings* [Online]. Available: <http://sites.ieee.org/sagroups-1914/p1914-3/>
- [7] IEEE P802.1CM, *Time-Sensitive Networking for Fronthaul* [Online]. Available: <http://www.ieee802.org/1/pages/802.1cm.html>
- [8] IEEE 1914 WG, *Next Generation Fronthaul Interface* [Online]. Available: <http://sites.ieee.org/sagroups-1914>
- [9] IEEE 802.3ah (2004, September), *Part 3: Carrier Sense Multiple Access with Collision Detection (CSMA/CD) Access Method and Physical Layer Specifications Amendment: Media Access Control Parameters, Physical Layers, and Management Parameters for Subscriber Access Networks* [Online]. Available: http://www.ieee802.org/21/doctree/2006_Meeting_Docs/2006-11_meeting_docs/802.3ah-2004.pdf
- [10] IEEE 802.3ag, *Connectivity Fault Management* [Online]. Available: <http://www.ieee802.org/3/ag/index.html>
- [11] Metro Ethernet Forum (2006, January), *MEF 16: Ethernet Local Management Interface (E-LMI)* [Online]. Available: <https://www.mef.net/carrier-ethernet/technical-specifications>
- [12] ITU (2013, August), *G.8261 Timing and synchronization aspects in packet networks* [Online]. Available: <https://www.itu.int/rec/T-REC-G.8261>
- [13] ITU (2015, January), *G.8262 Timing characteristics of a synchronous Ethernet equipment slave clock* [Online]. Available: <https://www.itu.int/rec/T-REC-G.8262>
- [14] ITU (2014, May), *G.8264 Distribution of timing information through packet networks* [Online]. Available: <http://www.itu.int/rec/T-REC-G.8264-201405-I>
- [15] ITU (2014, July), *G.8275.1 Precision time protocol telecom profile for phase/time synchronization with full timing support from the network* [Online]. Available: <https://www.itu.int/rec/T-REC-G.8275.1/en>
- [16] IEEE (2014, December), *802.1Q-2014-Standard for Local and metropolitan area networks—Bridges and Bridged Networks* [Online]. Available: <http://standards.ieee.org/about/get/802/802.1.html>
- [17] 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.
- [18] M. K. Al-Hares, P. Assimakopoulos, S. Hill and N. J. Gomes, “The Effect of Different Queuing Regimes on a Switched Ethernet Fronthaul,” in *Proc. Int. Conf. on Transparent Optical Networks (ICTON)*, Trento, Italy, 2016, pp. 1-4.
- [19] IEEE (2004, June), *802.1D-Standard for Local and metropolitan area networks-Media Access Control (MAC) Bridges* [Online]. Available: <http://standards.ieee.org/about/get/802/802.1.html>
- [20] IEEE (2016, March), *802.3-2015-Standard for Ethernet* [Online]. Available: <http://standards.ieee.org/about/get/802/802.3.html>
- [21] H. J. Son and S. M. Shin (2014, April), *Fronthaul size: calculation of maximum distance between RRH and BBU* [Online]. Available: http://www.netmanias.com/en/?m=view&id=blog_mb1&no=6276

Philippos Assimakopoulos received the B.Eng. degree in electronic engineering from the University of Bath, Bath, U.K., in 2003, and the M.Sc. degree in broadband and mobile communication networks and Ph.D. degree in electronic engineering from the University of Kent, Canterbury, U.K., in 2007 and 2012, respectively. He is currently with the Communications Research Group, University of Kent. His research interests include low-cost microwave RoF networks for indoor and outdoor applications and the design of cloud-radio access network for 4G and 5G applications.

Mohamad Kenan Al-Hares received the B.Eng. degree in computer and communication engineering from the Arab International University, Daraa, Syria and the M.Sc. degree in Network Computing from Coventry University, Coventry, U.K. He is currently with the Communications Research Group, University of Kent. His research interests include cloud radio access network and Ethernet fronthaul development for the future mobile network.

Nathan J. Gomes (M’92–SM’06) received the B.Sc. degree from the University of Sussex, Sussex, U.K., in 1984, and the Ph.D. degree from University College London, London, U.K., in 1988, both in electronic engineering. From 1988 to 1989, he held a Royal Society European Exchange Fellowship with ENST, Paris, France. Since late 1989, he has been with the University of Kent,

Canterbury, U.K., where he is currently Professor of Optical Fibre Communications. His current research interests include fiber-wireless access, and RoF systems and networks. Professor Gomes was a TPC Co-chair for the MWP 2014 IEEE Topical Meeting in Sapporo and TPC Chair for IEEE International Conference on Communications, ICC 2015 in London.