



Kent Academic Repository

Gomes, Nathan J., Assimakopoulos, Philippos, Chanclou, Philippe, Habel, Kai, Jungnickel, Volker, Ritosa, Patrik, Zou, Jim, Elbers, Jörg-Peter, Thomas, Howard, Linne, Gregor and and others (2019) *Testbed Verification of New Fronthaul Technology for 5G Systems*. In: 2019 IEEE International Conference on Communications Workshops (ICC Workshops). 2019 IEEE International Conference on Communications Workshops. . IEEE ISBN 978-1-72812-373-8.

Downloaded from

<https://kar.kent.ac.uk/73129/> The University of Kent's Academic Repository KAR

The version of record is available from

<https://doi.org/10.1109/ICCW.2019.8756785>

This document version

Author's Accepted Manuscript

DOI for this version

Licence for this version

UNSPECIFIED

Additional information

Versions of research works

Versions of Record

If this version is the version of record, it is the same as the published version available on the publisher's web site. Cite as the published version.

Author Accepted Manuscripts

If this document is identified as the Author Accepted Manuscript it is the version after peer review but before type setting, copy editing or publisher branding. Cite as Surname, Initial. (Year) 'Title of article'. To be published in *Title of Journal*, Volume and issue numbers [peer-reviewed accepted version]. Available at: DOI or URL (Accessed: date).

Enquiries

If you have questions about this document contact ResearchSupport@kent.ac.uk. Please include the URL of the record in KAR. If you believe that your, or a third party's rights have been compromised through this document please see our [Take Down policy](https://www.kent.ac.uk/guides/kar-the-kent-academic-repository#policies) (available from <https://www.kent.ac.uk/guides/kar-the-kent-academic-repository#policies>).

Testbed Verification of New Fronthaul Technology for 5G Systems

Nathan J. Gomes¹, *Senior Member, IEEE*, Philippos Assimakopoulos¹, *Member, IEEE*, Philippe Chanclou², Kai Habel³, Volker Jungnickel³, *Member, IEEE*, Patrik Ritoša⁴, Jim Zou⁵, *Member, IEEE*, Jörg-Peter Elbers⁵, *Member, IEEE*, Howard Thomas⁶, Gregor Linne⁷, Christoph Juchems⁷

¹Communications Research Group, University of Kent, Canterbury, Kent, CT2 7NT, UK

²Orange Labs, 2 Avenue Pierre Marzin, 22300 Lannion, France

³Fraunhofer HHI, Einsteinufer 37, 10587 Berlin, Germany

⁴Telekom Slovenje, Cigaletova ulica 15, 1000 Ljubljana, Slovenia

⁵ADVA Optical Networking SE, Fraunhoferstr. 9a, 82152 Martinsried, Germany

⁶Viavi Solutions, London Road, Newbury Berkshire, RG14 2PZ, UK

⁷IAF GmbH, Berliner Straße 52 J, 38104 Braunschweig, Germany

Abstract—The fronthaul for 5th generation mobile systems (and beyond) has evolved with new splits for the radio access network functions defined, and the transport for these split interfaces having very different requirements. Testing of the transport for such split interfaces is reported, and it is shown that an Ethernet fronthaul transport network, which is capable of bringing efficiency gains through statistical multiplexing, can meet stringent latency and latency variation requirements, assuming buffering and playout of the radio waveforms and that timing/synchronization signals are prioritized. An aggregation technique for a 100 Gb/s Ethernet trunk which provides for such timing signals is demonstrated. Real-time monitoring of the Ethernet fronthaul for software-defined networking control and performance optimization is also shown.

Keywords—5th generation mobile systems, evolved fronthaul, radio over Ethernet, Open Radio Access Network

I. INTRODUCTION

The fronthaul of 5th generation mobile (5G) systems and beyond will be very different to that used in previous generations [1]. The reasons for this arose mainly due to the realization that following the legacy approach of sampled time-domain waveform transmission, such as has been used in industry “standards” such as the Common Public Radio Interface (CPRI) or Open Base Station Architecture Initiative (OBSAI), for the much larger bandwidths and higher numbers of antenna “layers” in multi-antenna techniques, would result in an infeasible bit-rate explosion [2]. There was, nevertheless, a need to retain the centralization offered by a fronthaul, for centralized- or cloud-radio access networks (C-RANs) [3], particularly in view of the increased interest in the use of network function virtualization (NFV) in 5G. Most solutions to the fronthaul bit-rate problem proposed moving some functions away from the central unit (CU) or base station baseband unit (BBU) pool, into the remote radio unit (RRU) [1], [4]-[6]. As more functions are moved to the remote location, the bit-rate required for the

new fronthaul reduces, first through only sending frequency domain symbols, then by sending symbols prior to pre-coding for the antenna layers, then data for individual users, etc. The bit-rate tends to more-and-more approach the backhaul, user data rates as functions are distributed to remote units [4]. Further efficiency savings for bit-rate reductions are possible as the opportunities for statistical multiplexing gains increase (assuming packet-mode transport, such as Ethernet or IP, is used over the fronthaul) [1]. Latency requirements between the CU and remote units also become more relaxed, especially once the functions for the hybrid automatic repeat request (HARQ) loop are moved fully to the remote radio unit [5]. However, there are disadvantages, too, in moving functions away from the CU. Apart, from the general loss of centralization advantages for NFV, distributed beamforming such as joint transmission/reception with coordinated multipoint (CoMP), centralized coordination of interference cancellation, all become difficult or impossible to achieve as the functions required for them are distributed to remote units [6].

Now, it is generally accepted that there are two types of new functional split for the Radio Access Network (RAN) [7]. A higher-layer split, which is outside of the HARQ timing loop, with relatively relaxed latency requirements on the transport over this fronthaul segment. 3GPP has defined this interface as that between a CU and distributed unit (DU) [8]. The split point defined, between the Packet Data Convergence Protocol (PDCP) and Radio Link Control (RLC) layers of the RAN, is that used for “dual-connectivity” (micro-base station connection to a macro-base station) and is of interest as it may enable connection of user plane, local compute and storage functions (referred to as multi-access or mobile edge computing, MEC).

The lower-layer split (somewhere in the RAN protocol stack Physical (PHY) layer) was not specified by 3GPP, who define it as the interface between the DU (or CU/DU) and RRU [8]. The CPRI corporation have defined an Ethernet-based fronthaul transport specification, eCPRI, for such an intra-PHY layer RAN split [9], and groups such as Open RAN, now part of TIP [10], have continued work on specifications for such split points. The split point may also be variable, depending on particular scenarios, implementations and, even, service requirements. As stated

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). Philippos Assimakopoulos acknowledges funding by the EPSRCs “Towards an Intelligent Information Infrastructure (TI3)” program (NIRVANA project).

previously, timing constraints are stringent, demanding low latency and low latency variation from the fronthaul transport between DU and RRU. Clearly statistical multiplexing gains, which can reduce aggregate bit-rates, will also lead to increased delay and delay variation due to port contention in Ethernet switches and aggregators. Some loosening of delay requirements is possible through buffering of received packets and subsequent “payout” if there is strict and known timing throughout the RAN [11]. Such timing can be provided by protocols such as Precision Timing Protocol (PTP, IEEE1588) [12], but the packets transporting PTP signals must themselves suffer from minimal delay variation.

Further requirements for the new fronthaul stem from the need to enable end-to-end network slicing for 5G and beyond systems. Separation of physical and virtual network functions within the RAN will play a role. The use of an IP/Ethernet-based fronthaul should much more easily enable the use of software-defined networking (SDN) techniques already being developed for transport networks [13]. Then, for orchestration and control of the virtualized RAN and fronthaul, performance monitoring is required [14]. Measurements should be sent to the orchestration/control function to ensure slice service requirements are met.

In this paper, testbed measurements on new, Ethernet-based fronthaul technologies for mobile systems are presented. In Section II, results for a higher-layer split option are presented. Here, the Ethernet transport is also overlaid onto a passive optical network (PON) to demonstrate convergence possibilities with fixed access networks [15]. In Section III, an integrated testbed demonstration is presented. The transport of two lower-layer splits is demonstrated: one for a “5G-like” waveform (high bit-rate for wide-bandwidth, millimeter-wave wireless signal), the second for a 4G signal, but with the complete RAN protocol stack operational, through the use of an Open Air Interface (OAI) software base station. In addition, to show convergence with legacy systems, these signals are aggregated with the fronthaul transport of a CPRI signal and with the backhaul transport to/from a femtocell. The aggregator enables highly accurate transfer of a PTP signal through the use of a gap-filling method for traffic prioritization [16]. Section IV of this paper presents measurements on the fronthaul using hardware probing, and example use of such measurements by SDN-controlled layer 2 switches, as a step towards a software-defined, virtualized RAN. A summary is presented in Section V.

II. HIGHER-LAYER SPLIT TESTBED

The higher-layer split demonstration shown in Fig.1 consisted of Ethernet point-to-point (PtP) transport in which Ethernet impairments could be injected, and Ethernet over GPON passive optical network point-to-multipoint (PtMP) transport, illustrating convergence possibilities with fixed access deployments. Global Positioning System (GPS) timing enabled synchronization and accurate delay measurements. A 15 MHz bandwidth 2x2 MIMO 3GPP LTE (4G) signal was used in the demonstration. The top-of-rack (ToR) switch enabled the CU component, the virtualized base station baseband unit (vBBU) and its associated management and core network, to communicate with intelligent remote radio heads (iRRHs), in Fig.1, but effectively DUs in 3GPP terminology over either the PtP or

PtMP links. The key performance parameters measured in this demonstration were the fronthaul latency (round-trip delay) and fronthaul data rate. For the former, a round-trip delay of between 5 – 10 ms was measured for the PtP case. The delay for the PON case was 1 ms greater, caused by the dynamic bandwidth allocation (request/grant) mechanism; this additional delay could be avoided by using fixed bandwidth allocation. The fronthaul data rate was measured to be 120 Mb/s, around 20% higher than the user data rate, but around 8 times less than the CPRI bit-rate required. Of course, the peak user data rate was one to two orders of magnitude less than would be expected for 5G. For 5G rates, efficient use of the fronthaul bandwidth will be more important, which would make statistical multiplexing gains more attractive, and also require the use of some form of dynamic bandwidth allocation in a PON to exploit them.

III. CONVERGED FRONTHAUL TESTBED WITH AGGREGATOR

The converged, aggregated Ethernet fronthaul demonstration considered two implementations of a lower layer split (one to confirm the split operation through complete user equipment to core network connection and operation, the second to confirm transport for “5G-like” high data rate waveforms), an LTE/4G femtocell backhaul link and a CPRI link (these latter two links representing legacy systems) all transported over the same fronthaul. The overall system testbed is shown in Fig.2. The key component in this evolved fronthaul was a high-speed Ethernet aggregator which multiplexed all systems over a 100 Gb/s fronthaul carried over 6km of metro fiber in Telekom Slovenije’s network. Also carried through the aggregator was a PTP synchronisation signal, obtained via a grand master clock locked to a GPS time reference, and background Ethernet traffic between a traffic source and sink. As the key element to the aggregation, and enabling the time synchronization while aggregating, the high-speed aggregator will be examined first.

A. High-speed, gap-filling Ethernet aggregator

The aggregator used is based on the Transpacket FUSION® gap-filling technology [16]. The high-priority Ethernet packets are buffered for a fixed period before being forwarded onto the appropriate output. Lower priority packets are checked and placed in the available gaps between the high-priority Ethernet packets, thus suffering variable delays. In the testbed implementation, the PTP traffic was given the high priority (and was used as the timing for the CPRI transmission). The clock synchronization at the RRH with the BBU clock showed no measurable difference when transported through the fronthaul, better than the 2 parts per billion requirement. This was carried out with all traffic sources at maximum rates and some 40% load on the 100 Gb/s Ethernet trunk link.

B. LTE software base station MAC-PHY split

A LTE software base station functional split at the MAC-PHY interface was implemented [17]. The split functions used an interface library for mapping protocol data units, control primitives etc. into Ethernet packets, generally following specifications defined by the IEEE 1914.1 Next Generation Fronthaul Interface (NGFI) group [18]. In the testbed demonstration, the objective was to show continuous operation of the split (continuous connection of a UE with

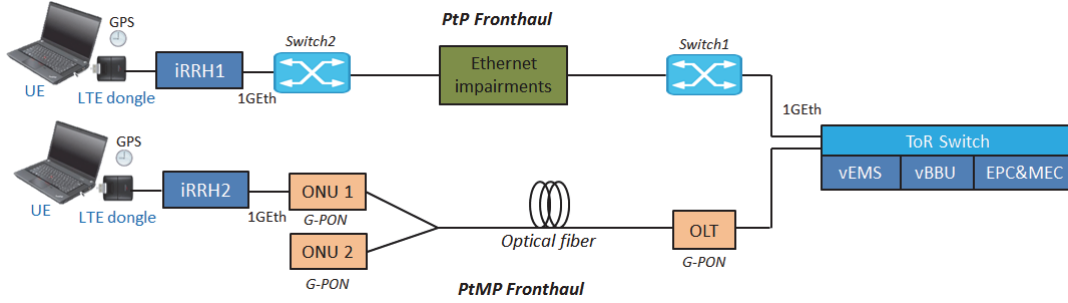


Fig.1. Higher-layer split demonstration.

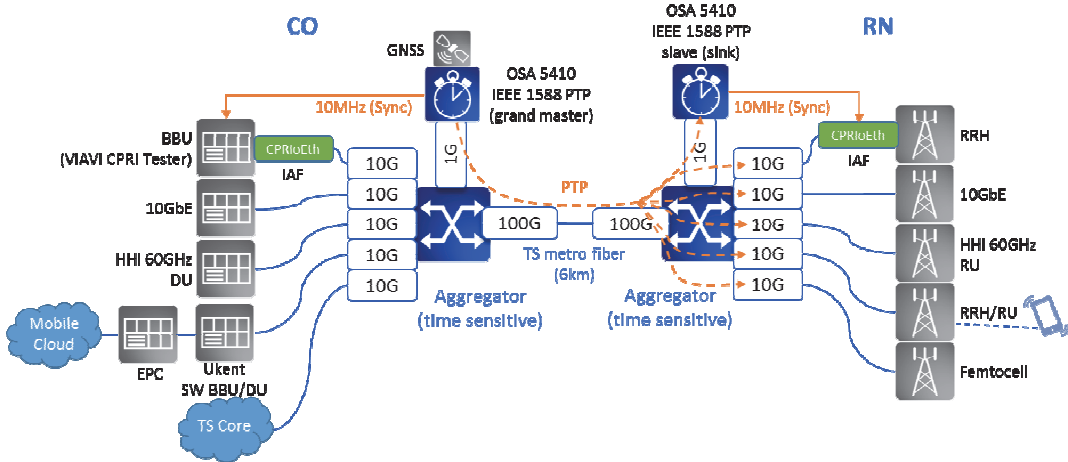


Fig.2. Converged, aggregated Ethernet fronthaul demonstration.

the software evolved packet core, EPC) over the aggregated fronthaul, with port contention occurring for this fronthaul transport with the packets from the other sources.

The average delay for the MAC-PHY split fronthaul (scaled for 10 Gb/s) was estimated to be 156 μ s (round-trip), which is close to the requirements specified for functions such as CoMP. In fact, a large part of this delay was due to the software generation of the data and frames at the end stations and could easily be reduced using hardware acceleration technologies.

C. High-speed Upper-PHY layer split

Hardware-based OFDM transmitter and receiver pairs were developed with split functionality at the upper-PHY layer. The Field-Programmable Gate Array (FPGA) hardware-based OFDM boards were used to interface to 60-GHz wireless modules. On the data input/output side of the FPGA boards Ethernet mapping/demapping functions, generally following the definitions specified by the IEEE 1914.3 Radio-over-Ethernet group were implemented. For the testbed demonstrator with the aggregated fronthaul, the split functions were placed either side of the aggregator/dis-aggregator pair connecting to the 100 Gb/s Ethernet trunk link, with the DU and RU modules connecting to these via 10 Gb/s Ethernet links, as shown in Fig.3. The 60GHz radio was used in loopback mode for testing. The key performance indicator measured for the high-speed Upper-PHY layer split was the overhead-induced data rate (the excess compared to the user/backhaul rates). A flow control

mechanism was implemented in the setup to guarantee wireless 60 GHz transmission at a bit-rate of 2.5 Gb/s. This meant that the measured rates over the Ethernet links adjusted in order to fulfil this transmitted bit-rate condition, and showed dependency on the Ethernet frame size, as Ethernet headers/trailers are stripped off before the processing for the wireless signal transmission (in the forward path). Fronthaul throughputs of between 2.3 and 2.5 Gb/s were measured, with the Internet Protocol packet payload throughput ranging between approximately 1.5 and 2.2 Gb/s, depending on Ethernet frame size. The overhead was of the same order as that measured for the MAC-PHY split. The multi-Gb/s throughput and loopback operation at 60 GHz, were clear indicators that 5G-like waveforms were being experimentally demonstrated.

Delay and delay variation across the aggregated fronthaul were also measured. The delay variation (measured as Ethernet packet jitter) showed some dependency on frame size, but was always less than 7 μ s. A comparison with measurements without the aggregator indicated that less than 2 μ s of this jitter was due to the aggregation. This level of jitter could be recovered by appropriate buffering and a time synchronization signal, e.g. using PTP, to enable precise playout of the waveform. The delay measurement was compromised by the buffering for the flow control that enabled the continuous 2.5 Gb/s wireless bit-rate. Again, comparison with the measurements without the aggregator confirmed only very small contributions from the aggregator and aggregated Ethernet transport.

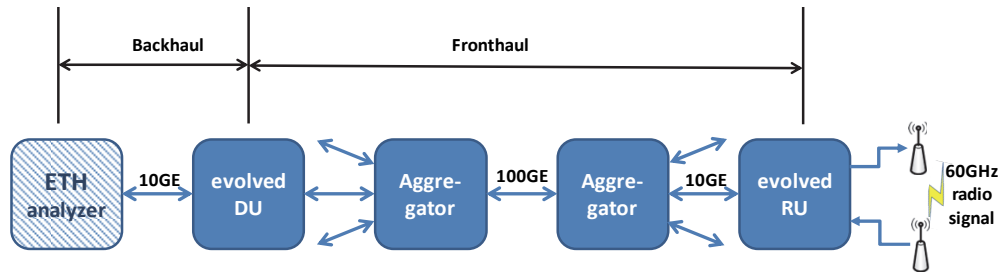


Fig.3. High-speed fronthaul for 60GHz wireless.

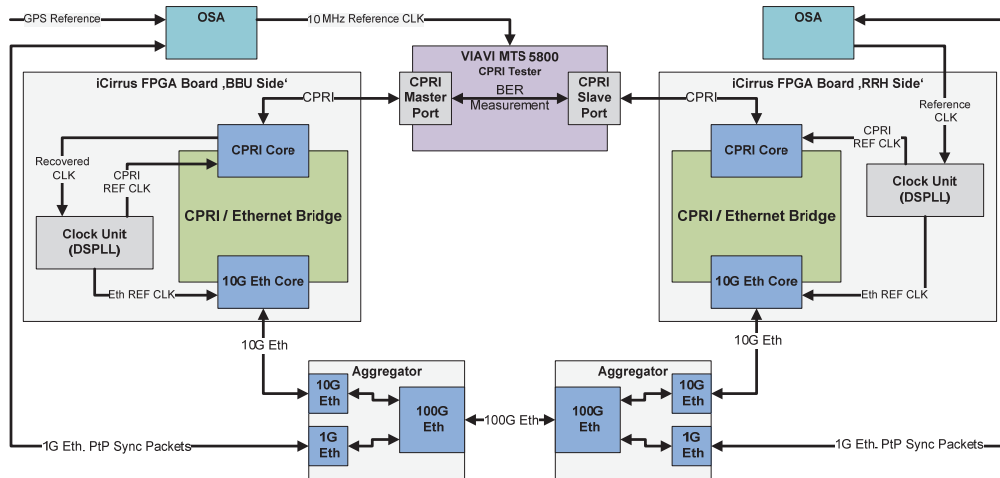


Fig.4. CPRI over Ethernet fronthaul.

D. Legacy system transport

Two forms of legacy system were transported over the aggregated Ethernet fronthaul. The first, as shown in Fig.1, was a backhaul connection between a 4G/LTE femtocell and Telekom Slovenije’s core network. Continuous operation of the femtocell was verified, with no measureable errors. As the femtocell operated with no intercell coordination (only for handovers), the delay requirement of 60 ms was easily met (a round-trip delay of 30 ms was measured, only 70 μ s of which was contributed by the aggregator).

The second legacy system transported was a CPRI fronthaul, generated and received by Viavi Solutions CPRI testers. The CPRI signals are mapped into Ethernet frames in an FPGA with the CPRI/Ethernet bridging function generally following specifications defined by the IEEE 1914.3 Radio over Ethernet group. The CPRI tester timing was derived from the same master clock source as used and distributed by the PTP packets, sent through the aggregated 100 Gb/s link. The Ethernet functions in the FPGAs, and the synchronization for CPRI and Ethernet functions on the FPGA derived their clock from the CPRI streams. The aggregator/dis-aggregators were Synchronous Ethernet compliant, so all timing was derived from the same clock reference and PTP signals, as shown in Fig. 4. As the CPRI packets contend with other packets in the aggregators for the 100 Gb/s link, variable packet delay is generally experienced. This means that the continuous CPRI stream is no longer continuous when retrieved from the CPRI over Ethernet packets. In order to maintain operation with these

variable fronthaul delays, a receive first-in first-out (FIFO) buffer is implemented. Payout from the buffer only commences after the buffer has filled to a pre-defined level; the waveform is “played out” according to timing information retrieved from the CPRI over Ethernet packets, but the pre-filling of the buffer allows jitter in the frame arrivals to be absorbed. The larger the fill level of the FIFO buffer, the more jitter can be absorbed, but the buffer also causes a fixed delay. The tests found that a 10 μ s buffer latency was sufficient to enable error-free operation for both 2.45 Gb/s and 9.83 Gb/s CPRI streams over the aggregated fronthaul with 6km 100 Gb/s trunk link. The overall round-trip delay in these cases was 55 μ s and 30 μ s, respectively, considerably lower than requirements for CoMP (approximately, 150 μ s). In addition to the measurements of CPRI stream latency, the frequency and time synchronization were measured. The former was found to within the 2 parts per billion requirement. The timing synchronization was 30ns on average, just within requirement, but was frequently greater than this. However, it was found that the buffering of the packets allowed this to be absorbed, and the CPRI over Ethernet fronthaul was able to function without any problem.

IV. KPI MEASUREMENT AND SDN

A testbed was also set up to demonstrate real-time measurements of the fronthaul traffic using in-line hardware probes, the collection of the measurements and their analysis in an “intelligence server”, which effectively acts as a simple control/orchestration engine for the fronthaul, and the use of software-defined networking (SDN) at the Ethernet layer, as

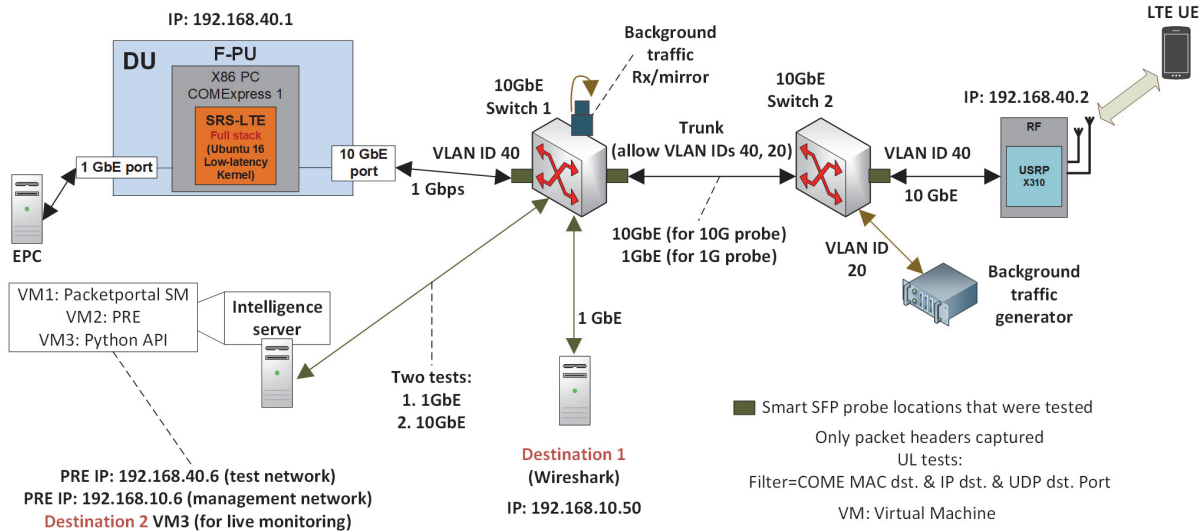


Fig. 5. Testbed for in-line probe measurements with “intelligence server”.

shown in Fig.5 [18]. The in-line, hardware probes are small form-factor pluggable (SFP) modules which replace standard 1 Gb/s or 10 Gb/s Ethernet SFPs for connection to Ethernet ports for optical fiber links. The probes contain FPGA hardware which can filter Ethernet packets based on header fields, send frame result packets (FRPs) to a monitoring/management station (the “intelligence server” in this case), and can timestamp these FRPs. The FRPs are sent when gaps in transmission are detected, resulting in minimal disturbance to the transported traffic on the network. At the intelligence server, the FRPs are analysed such that, for example, frame delay and delay variation, and interframe delay and delay variation can be calculated. Throughput through Ethernet ports can also be monitored.

In Fig. 6, example captures from the real-time outputs calculated by the intelligence server are presented. In this case, the measurements were carried out for sampled time-domain waveform (CPRI-like) traffic, hence the usually constant interframe delay. The used bandwidth of the link is similarly constant in this case, the spikes resulting from just the occasional irregularity in frame sizes (which simultaneously affect the interframe delay). In the testbed, such outputs are also used to inform the SDN-enabled switch (Switch 2) in the network. For example, we have steered traffic from a background source (sending packets to the traffic mirror in Fig.6) away from the trunk link used by the mobile fronthaul uplink when the load exceeds a certain value.

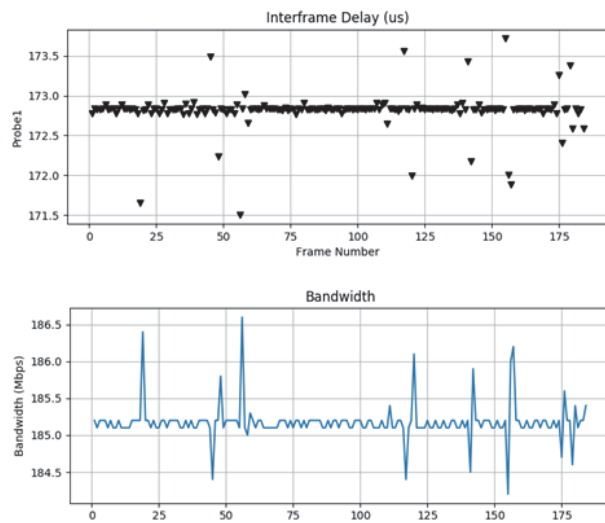


Fig. 6. Example real-time analysis: interframe delay (top), throughput (bottom).

V. CONCLUSION

The fronthaul for 5G is evolving with new RAN functional splits between a CU and DU, and DU and RRU, with the latter having a yet-to-be defined, and possibly varying interface. For many splits, statistical multiplexing is possible as the bit-rate will be dependent on the load and number of spatial streams being used at any given time, and packet-based transport will enable such efficiency gains to be made. However, statistical multiplexing also leads to port contention in an Ethernet or IP-based fronthaul, and meeting the strict latency and latency variation requirements of the lower-layer RAN splits (between DU and RRU) becomes problematic. Meeting these latency and latency variation requirements is the most demanding aspect for a packet-based fronthaul.

Techniques for meeting the requirements for both higher- and lower-layer split points have been demonstrated. These show that the statistical multiplexing can be used beneficially, while latency and latency variation is met using specific techniques: use of a timing protocol (PTP) to allow receiving nodes to employ buffers to absorb packet delay variations and play out radio waveforms with precise timing;

use of a time-sensitive networking protocol to guarantee minimal latency variation for the PTP packets; use of probing and fronthaul measurements to assure performance, employing SDN to take corrective actions, as necessary. A testbed was also demonstrated which successfully aggregated fronthaul traffic from new 5G-like waveforms, and LTE traffic (both for new RAN lower-layer new functional splits), legacy CPRI traffic, backhaul traffic and the timing protocol traffic.

ACKNOWLEDGMENT

The authors would like to thank Telekom Slovenije for the use of their test network facilities, and acknowledge useful discussions with other collaborators in the iCIRRUS project.

REFERENCES

- [1] N.J. Gomes, P. Chanclou, P. Turnbull, A. Magee and V. Jungnickel, "Fronthaul evolution: from CPRI to Ethernet," *Optical Fiber Technology*, vol. 26A, pp. 50-58, Dec. 2015
- [2] U. Dötsch, M. Doll, H.-P. Mayer, F. Schaich, J. Segel, and P. Sehier, "Quantitative analysis of split base station processing and determination of advantageous architectures for LTE," *Bell Labs Techn. J.*, vol. 18, pp. 105-128, 2013
- [3] China Mobile Research Institute, "C-RAN: the road towards green RAN", White Paper, V3.0, Dec. 2013
- [4] P. Rost, C.J. Bernardos, A. De Domenico, M. Di Girolamo, M. Lalam, A. Maeder, D. Sabella and D. Wübben, "Cloud technologies for flexible 5G radio access networks", *IEEE Commun. Mag.*, vol. 52, pp. 68-76, May 2014
- [5] iCIRRUS (Jul 2016), D3.2 Preliminary Fronthaul Architecture Proposal, [Online]. <http://www.icirrus5gnet.eu/category/deliverables/>.
- [6] K. Miyamoto, S. Kuwano, T. Shimizu, j. Terada and A. Otaka, "Performance evaluation of Ethernet-based mobile fronthaul and wireless CoMP in split-PHY processing," *J. Opt. Commun. Netw.*, vol. 9, pp. A46-A54, 2017
- [7] N.J. Gomes, P. Sehier, H. Thomas, P. Chanclou, B. Li, D. Münch, P. Assimakopoulos, S. Dixit, and V. Jungnickel, "Boosting 5G through Ethernet: How evolved fronthaul can take next-generation mobile to the next level," *IEEE Vehicular Technol. Mag.*, vol. 13, pp. 74-84, March 2018
- [8] 3GPP, "Study on new radio access technologies: Radio access architecture and interfaces," 3rd Generation Partnership Project, Sophia Antipolis, France, Tech. Rep. TR 38.801, v14.0.0, 2017
- [9] Common Public Radio Interface (2017). Industry leaders release new CPRI specification for 5G. [Online]. Available: <http://www.cpri.info/press.html>
- [10] Telecom Infra Project, OpenRAN project group, see: <https://telecominfraproject.com/openran/>
- [11] J.-P. Elbers, J. Zou, P. Assimakopoulos, N. Gomes, K. Habel, V. Jungnickel, G. Linne, C. Juchems, P. Chanclou, P. Ritoša and H. Thomas, "Next-generation optical fronthaul in the iCirrus project," *Optical Fiber Communications Conference, OFC 2018, San Diego, CA, USA, March 2018*
- [12] IEEE Precision Clock Synchronization Protocol for Networked Measurement and Control System, IEEE Std 1588-2008
- [13] S. Gringeri, N. Bitar, T.J. Xia, "Extending software defined network principles to include optical transport," *IEEE Commun. Mag.*, vol. 51, pp.32-40, March 2013
- [14] P. Assimakopoulos, J. Zou, K. Habel, J.-P. Elbers, V. Jungnickel, and N.J. Gomes, "A converged evolved Ethernet fronthaul for the 5G era," *IEEE J. Selected Areas Commun.*, vol. 36, pp. 2528-2537, November 2018
- [15] Z. Tayq, L. Anet Neto, B. Le Guyader, A. De Lannoy, M. Chouaref, C. Aupetit-Berthelemot, M. Nelamangala Anjanappa, Si Nguyen, K. Chowdhury, P. Chanclou "Real Time Demonstration of the Transport of Ethernet Fronthaul based on vRAN in Optical Access Networks", Th3A.2, OFC 2017, 19-23 March, 2017, USA
- [16] R. Veisllari, S. Bjornstad and J.P. Braute, "Experimental demonstration of 100 Gb/s optical network transport and aggregation for Ethernet fronthaul with low and bounded delay," *Optical Fiber Communications Conference, OFC 2018, San Diego, CA, USA, March 2018*
- [17] G.S. Biring, P. Assimakopoulos, and N.J. Gomes, "An Ethernet-based fronthaul implementation with MAC/PHY split LTE processing", *IEEE Global Communications Conference, GLOBECOM 2017, Singapore, December 2017*
- [18] iCIRRUS (Dec 2016), D3.3 SLA and SON Concept for iCIRRUS, [Online]. <http://www.icirrus5gnet.eu/category/deliverables>