

Next-Generation Optical Fronthaul in the iCirrus Project

Jörg-Peter Elbers¹, Jim Zou¹, Philippos Asimakopoulos², Nathan Gomes², Kai Habel³, Volker Jungnickel³, Gregor Linne⁴, Christoph Juchems⁴, Philippe Chanclou⁵, Patrik Ritoša⁶, Howard Thomas⁷

¹ADVA Optical Networking SE, Fraunhoferstr. 9a, 82152 Martinsried, Germany. ²University of Kent, Canterbury, Kent, England, CT2 7NZ, UK.

³Fraunhofer HHI, Einsteinufer 37, 10587 Berlin, Germany. ⁴IAF GmbH, Berliner Straße 52 J, 38104 Braunschweig, Germany.

⁵Orange Labs, 2 Avenue Pierre Marzin, 22300 Lannion, France. ⁶Telekom Slovenije, Cigaletova ulica 15, 1000 Ljubljana, Slovenia.

⁷Viavi Solutions, London Road, Newbury Berkshire, RG14 2PZ, UK.

jelbers@advaoptical.com

Abstract: We discuss next-generation fronthaul solutions for 5G and legacy radio access networks. Architectures, findings and experimental results from recent lab and field trial activities are reported.

OCIS codes: (060.2330) Fiber optics communications; (060.4250) Networks

1. Introduction

The 5th generation mobile network (5G) is much more than just a new radio (NR) interface [1]. Set out to serve diverse markets and industries such as automotive, manufacturing, energy, eHealth, and entertainment, it needs to deliver services with very different bandwidth, latency and reliability characteristics; enhanced Mobile Broad-Band (eMBB), ultra-Reliable Low-Latency Communications (uRLLC), and massive Machine-Type Communications (mMTC) have to be provided over a single network. Fulfillment of 5G performance metrics [2] such as 1000 times higher cell capacity, 100 times higher peak data rates, 10 times lower latency and 10 times better reliability compared to 4G¹ requires both technology and architecture innovation.

Network slicing, edge computing and network function virtualization (NFV) will turn the radio access network (RAN) into a programmable resource. A cloud-RAN (C-RAN) architecture boosts efficiency by centralized resource pooling and can be complemented by distributed networking, storage and compute functions as network or application performance dictates. One of the main C-RAN implementation challenges lies in the fronthaul network that connects central units with remote units. In this paper, we investigate a next-generation optical fronthaul based on Ethernet. We discuss architectures and report results of the Horizon2020 iCirrus project.

2. Architectural considerations

In today's RAN, optical fronthaul and backhaul are served by different technologies: Backhaul links use standard Ethernet technology. Fronthaul links rely on layer 1 protocols such as CPRI (Common Public Radio Interface) to transport digitized I-Q, control & management, and synchronization data. Originally designed for connections from the base of a cell tower to the antenna on top, CPRI imposes stringent requirements on bandwidth, latency and admissible time error. A more than 10-fold higher data rate (compared to the user data), a one-way latency of less than 75 μ s², a timing error of less than ± 8 ns, and a frequency error of less than 2 ppb are common specifications for CPRI transport [3]; fulfilling these in a C-RAN context requires dedicated fiber links, posing economic and operational challenges to service providers. Moving from LTE to 5G, it becomes unsustainable to simply follow the CPRI route for C-RAN, as the radio bandwidths would require transport bit-rates in excess of 100 Gb/s to each remote unit.

Next-generation fronthaul solutions offer convergence on the basis of a common Ethernet layer [4]. Using different transport profiles, front-, back- and mid-haul (between macro-cell and subtended small cell sites) traffic can be transported over the same (x-haul) network. This approach eliminates stovepipe implementations and allows the use of established Ethernet OAM (operation, administration & maintenance) tools. By introducing new functional splits [5], the transmission of baseband bits or frequency domain I-Q symbols instead of time domain waveforms leads to a reduction in transport bit-rates compared to CPRI. The fronthaul becomes dependent on the user traffic and statistical multiplexing in the transport network can be exploited while the benefits of central resource pooling and coordination can (at least partly) be retained.

Following current IEEE P1914.1 terminology [6], the fronthaul transport network can be split into a central unit (CU), distributed unit (DU) and remote unit (RU). Depending on deployment scenario, the DU can be combined with either CU or RU. Transport between these units is provided by the next-generation fronthaul interface (NGFI), which exists in two versions.

NGFI-II relates to the higher layer split (HLS) in 3GPP (F1 interface) [7] and describes the fronthaul segment connecting the CU to the DU. The agreed HLS (3GPP Option 2) is derived from the LTE dual connectivity (DC)

¹ 1 million devices or 10 Tb/s per km², 10 Gb/s peak rate, 99.999% availability, 1ms user plane latency.

² If support of coordinated multi-point transmission (CoMP) is not necessary, a relaxation to approx. 250 μ s is possible.

configuration. It separates the radio protocol stack such that the PDCP and higher layer functions are located in the CU, while RLC and lower layer functions are located in the DU. The HLS increases bit-rates by less than 10% compared to those in the backhaul and has comparable latency and timing requirements. Yet, it allows centralization of CPU-intensive PDCP tasks and central aggregation (e.g. of 5G and LTE traffic).

NGFI-I relates to the lower layer split (LLS) in 3GPP (future F2 interface) and describes the fronthaul segment connecting the DU to the RU. Multiple LLS options are under consideration in both 3GPP and the new eCPRI standard [8]. No agreement has been reached to converge on a single option yet. A split point inside the PHY (3GPP Option 7.2, eCPRI Split IId+Iu) allows the same distributed massive MIMO (multiple input, multiple output) or CoMP operation as CPRI, but pushes the resource management to the RU thereby offering a 5-10 times reduction in data rate. A MAC-PHY split (3GPP Option 6) further reduces the transport bandwidth reduces transport bit-rates to less than twice backhaul rates (depending on the amount of control traffic between DU and RU), but still allows joint transmission and centralized scheduling. The slightly higher split points among the LLS options offer greater statistical multiplexing as bit-rate requirements scale more directly with load, at the expense of complexity in the decoupling of inter-layer dependencies. Latency and timing requirements stay similar to CPRI for all LLS options as they are more determined by the radio technology (carrier aggregation [CA], transmitter [TX] diversity, MIMO, CoMP) than a particular LLS choice.

The iCirrus optical fronthaul architecture is based on a converged Ethernet network which provides 10GbE connections to the RUs and aggregates traffic from multiple antenna sites in linear, chain or ring topologies using 100 GbE trunk links. It can accommodate variable split points based on operator requirements, cell load and aggregation levels (flexible RAN concept). Legacy CPRI signals can be transported over the same network after Ethernet encapsulation (CPRIoEth) by the IEEE P1914.3 Radio-over-Ethernet (RoE) standard [9]. As low latency and timing error is crucial for network performance, frequency syntonisation and time synchronisation is necessary, which can be provided by the IEEE 15882v2 precision timing protocol (PTP) in conjunction with synchronous Ethernet (SyncE) (using e.g. the ITU-T G.8275.1 telecom profile).

3. Theoretical investigations

Based on the architecture described above, work in iCirrus focused on time-sensitive networking approaches meeting the latency and timing error requirements of next-generation x-haul networks. Maximum latency and time error values for the different splits are still under discussion in 3GPP and the eCPRI group. Table 1a summarizes maximum transport latency values used iCirrus [10] and in the most recent eCPRI specification [8], binned into three classes of service (CoS). Current requirements are based on LTE time-division duplex (TDD) data. An eCPRI update taking into account 5G NR information is expected in 07/2018. The fiber transmission delay (5 μ s/km) needs to be taken into account in addition. Tolerable time errors are listed in Table 1b. An additional margin on the eCPRI specifications is required if the slave clock for timing recovery is located in the radio equipment, not at the transport edge. The class C requirement is an absolute time error which should be fulfilled by all radio transport solutions. Classes A+ to B requirements apply to a local radio cluster in the CPRI or LLS case and denote relative errors.

CoS	Example use case	One-way latency (iCirrus / eCPRI)
High	CPRI, LLS user plane (fast)	75 μ s / 100 μ s
Medium	CPRI, LLS user plane (slow), C&M plane (fast)	440 μ s / 1ms
Low	HLS & backhaul user plane, C&M plane (slow)	60ms / 100ms

Class	Maximum time error TE (iCirrus / eCPRI)	Example use case
A+	10ns / 20ns (relative)	MIMO or TX diversity
A	30ns / 70ns (relative)	Intra-band contiguous CA
B	n/a / 200ns (relative)	Intra-band non-cont. CA Inter-band CA
C	1360ns / 1100ns (absolute)	3GPP LTE TDD

Tab. 1: a) Fronthaul latency requirements, b) maximum permissible timing error.

Mechanisms for time-sensitive Ethernet transport investigated in iCirrus were: Strict priority (SP) queuing (IEEE 802.3br) without and with frame preemption (IEEE P802.1Qbu) as also considered in IEEE P802.1CM (profiles A and B) [11], time-aware shaping (TAS, IEEE 802.1Qbv) and a deterministic latency approach called FUSION [12]. Compared to round-robin scheduling, SP queuing can reduce average frame delay variation, but cannot reduce peak delay variation or worst case latency which are the relevant parameters here [4]. Frame preemption (FP) can reduce the peak delay variation, as it effectively lowers the maximum transmission unit (MTU) of the lower priority traffic. SP with FP may be able to meet class C specifications, but will require additional means such as a play-out buffer with high precision timing to meet the more stringent classes A+ to B requirements [13]. Time-aware shaping can in principle eliminate any packet delay variation of high priority traffic [14] and therefore offers deterministic latency, but requires a very accurate, network-wide flow scheduling. FUSION maintains deterministic latency for high priority traffic by preserving the timing of its frames and filling low priority traffic into the inter-frame gaps thereby

eliminating the need for any central scheduling. Both, TAS and FUSION incur additional fixed latency, but this could be offset by smaller playout buffers in the radio equipment.

4. Experimental results

To verify the theoretical findings, different experimental building blocks were implemented and successfully tested in iCirrus [10]. A common basis for the hardware demonstrators was a flexible FPGA prototyping platform developed in the project (see Fig. 1a). The list summarizes the main experimental results:

- A CPRIoEth mapper was developed to transport CPRI streams over a converged Ethernet network (3GPP Option 8). Measurement results verified a back-to-back latency of $35\mu\text{s}$ at 2.5G and $10\mu\text{s}$ at 10G operation.
- A full LTE virtual RAN (vRAN) DU/RU setup with MAC-PHY LLS (3GPP Option 6) was developed and achieved $266\mu\text{s}$ round-trip delay at 0.4Gb/s fronthaul data rate. The software packet generation is a significant contributor to latency.
- A 60GHz DU and RU were prototyped implementing a LLS (3GPP Option 7) and operating at $<250\mu\text{s}$ round-trip latency with 1Gb/s throughput. A less bursty packet generation profile would reduce the larger latency values.
- A full LTE vRAN using a HLS (3GPP Option 2) was tested. No impact on the throughput was demonstrated up to 20ms Ethernet link latency.
- A deterministic Ethernet aggregator using the FUSION approach was implemented. Measured maximum one-way latency was $14.3\mu\text{s}$ at 16kB MTU size and the delay variation was less than 136ns at 10G.

In a final field trial set-up (Fig. 1b) the individual demonstrators were integrated at Telekom Slovenije (TS) facilities. CPRIoEth traffic, 10GbE best effort traffic, 60GHz radio, vRAN LTE signals and Femtocell traffic out of the TS core network were aggregated in the central office (CO) and successfully transmitted together with multicasted IEEE 1588 PTP traffic over Telekom Slovenje's metro network to the remote nodes (RN). An IEEE 1588 PTP slave signal was used to generate a 10MHz synchronization signal for the CPRIoEth demapper. Apart from the additional fiber latency, no additional performance degradations could be observed.

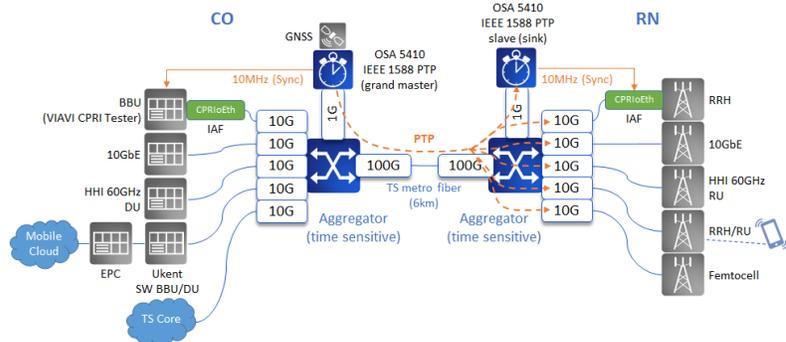
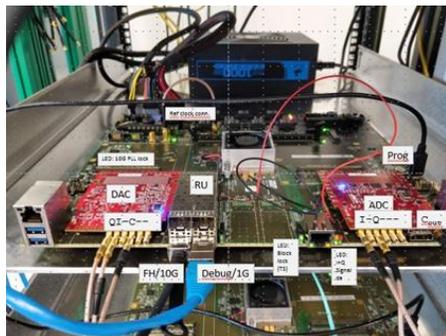


Fig. 1: a) iCirrus 5G FPGA prototyping platform in a 60 GHz RU implementation, b) iCirrus field trial set-up demonstrating multiple radio and split approaches over a converged time-sensitive Ethernet link (right).

5. Conclusions

Next-generation Ethernet-based fronthaul networks offer convergence of fronthaul, backhaul and midhaul traffic. We discussed and investigated key building blocks to meet the stringent capacity, latency and timing requirements.

6. Acknowledgments

The work leading to these results has received funding from the European Union's Horizon 2020 Research and Innovation Program under Grant Agreement n° 644526 (iCIRRUS). The authors want to thank Raimena Veisllari, Jan Petter Braute, Steinar Bjørnstad and Mickaël Fontaine from Transpacket for their support.

- [1] 5G Infrastructure Association, <https://5g-ppp.eu/wp-content/uploads/2018/01/5G-PPP-5G-Architecture-White-Paper-Jan-2018-v2.0.pdf>
- [2] 5G Infrastructure Association, <http://5g-ppp.eu/wp-content/uploads/2015/02/5G-Vision-Brochure-v1.pdf>
- [3] Nokia, https://onestore.nokia.com/asset/192728/Nokia_Evolution_to_Centralized_RAN_with_Mobile_Fronthaul_White_Paper_EN.pdf
- [4] N. J. Gomes, et al., "Boosting 5G through Ethernet", IEEE Vehicular Technology Magazine, doi: 10.1109/MVT.2017.2782358
- [5] 3GPP, TR38.801, "Study on new radio access technologies: Radio access architecture and interfaces", v14.0.0
- [6] IEEE P1914.1, "Standard for Packet-based Fronthaul Transport Networks", <http://sites.ieee.org/sagroups-1914/p1914-1/>
- [7] 3GPP, TS 38.470, "NG-RAN; F1 general aspects and principles", v15.0.0
- [8] CPRI group, "eCPRI Specification V1.0" and "Requirements for the eCPRI Transport Network V1.1", <http://www.cpri.info/>
- [9] IEEE P1914.3, "Standard for Radio Over Ethernet Encapsulations and Mappings", <http://sites.ieee.org/sagroups-1914/p1914-3/>
- [10] iCirrus, Deliverable D5.4, "Validation test results and analysis evaluation", <http://www.icirrus-5gnet.eu/category/deliverables/>
- [11] IEEE P802.1CM, "Time-Sensitive Networking for Fronthaul", <https://1.ieee802.org/tsn/802-1cm/>
- [12] R. Veisllari, et al., "Experimental demonstration of 100Gb/s optical packet network for mobile fronthaul with load-independent ultra-low latency", ECOC 2017
- [13] J. Farkas, et al., "P802.1CM simulation results for profiles A & B", <http://ieee802.org/1/files/public/docs2016/cm-farkas-profiles-A-and-B-0316-v01.pdf>
- [14] M. K. Al-Hares, et al., "Modeling time-aware shaping in an Ethernet fronthaul", IEEE GLOBECOM 2017.