

Kent Academic Repository

Full text document (pdf)

Citation for published version

Gomes, Nathan J. and Sehier, Philippe and Thomas, Howard and Chanclou, Philippe and Li, Bomin and Munch, Daniel and Assimakopoulos, Philippos and Dixit, Sudhir and Jungnickel, Volker (2018) Boosting 5G Through Ethernet: How Evolved Fronthaul Can Take Next-Generation Mobile to the Next Level. IEEE Vehicular Technology Magazine, 13 (1). pp. 74-84. ISSN 1556-6072.

DOI

<https://doi.org/10.1109/MVT.2017.2782358>

Link to record in KAR

<http://kar.kent.ac.uk/66347/>

Document Version

Author's Accepted Manuscript

Copyright & reuse

Content in the Kent Academic Repository is made available for research purposes. Unless otherwise stated all content is protected by copyright and in the absence of an open licence (eg Creative Commons), permissions for further reuse of content should be sought from the publisher, author or other copyright holder.

Versions of research

The version in the Kent Academic Repository may differ from the final published version.

Users are advised to check <http://kar.kent.ac.uk> for the status of the paper. **Users should always cite the published version of record.**

Enquiries

For any further enquiries regarding the licence status of this document, please contact:

researchsupport@kent.ac.uk

If you believe this document infringes copyright then please contact the KAR admin team with the take-down information provided at <http://kar.kent.ac.uk/contact.html>

Ethernet-Based Evolved Fronthaul for Next-Generation Mobile Networks

Nathan J. Gomes, Philippe Sehier, Howard Thomas, Philippe Chanclou, Bomin Li, Daniel Münch, Philippos Assimakopoulos, Sudhir Dixit and Volker Jungnickel

Abstract:

Current approaches to the fronthaul for centralized- or Cloud-Radio Access Networks (C-RANs) need to be revised to meet the requirements of next-generation mobile networks. There are two major challenges: first, fronthaul signals need to be transported over public fixed access networks, such as passive optical networks (PONs), typically sharing them with other services; second, higher data rates must be catered for due to larger radio bandwidths and greater use of multi-antenna techniques, such as massive MIMO. Using Ethernet as a new transport protocol for the fronthaul allows statistical multiplexing and enables convergence between fixed and mobile services. This new approach more easily benefits from common developments being made for service level agreements, functional virtualization and software-defined networking. Higher data rates will be supported by the move to new, and possibly flexible, functional split points inside the radio access network (RAN) protocol stack of the processing located in the central and distributed units, as is being investigated by a number of bodies. However, there are technical challenges with regard to latency and packet delay variation. This article summarizes the benefits of an Ethernet-based fronthaul for the next generation of mobile networks, its main challenges and how these may be overcome.

1. Introduction

The 5th generation of mobile networks (5G) targets the rollout of new, customized and highly differentiated services, together with the associated business models for different vertical markets, having rather diverse sets of requirements. The key idea is to support these multiple services in parallel as “tenants” in a single network infrastructure, something also denoted as *network slicing*. Multiple such slices, one for each service, can be implemented in parallel as virtual networks in the same physical transport network. In this way, next generation mobile networks will achieve the required scalability, flexibility, CAPEX reduction, openness and portability [1].

In the radio access network (RAN), there is a need to support critical new technologies, such as small cells and the use of new spectrum with higher bandwidth in the mm-wave region [2]. New techniques such as massive multiple-input multiple-output (MIMO) and coordinated multipoint (CoMP) and new inter-cell interference management functions enable a higher spectral efficiency [3]. In general, these techniques require a higher degree of coordination in the next generation of mobile networks, and the so-called centralized Cloud-Radio Access network (C-RAN) is a favored approach to reach these goals while keeping complexity, energy consumption and costs low [4].

In a C-RAN, the functionality of a base station is split into a baseband unit (BBU) and remote radio heads (RRHs), with the transport between them denoted as the fronthaul (as opposed to the backhaul link between the RAN and the core network of the mobile operator). Until now, C-RANs have re-used equipment from traditional RANs with separated BBUs and RRHs, and stacked the BBUs at a central location. Sampled radio waveforms are transported over this fronthaul. There is a potential “pooling gain” through flexible interconnection of the stacked BBUs and the distributed RRHs. For instance, the same BBUs could be used for RRHs deployed in industrial and home areas during work and leisure times, respectively. Accordingly, the BBUs have been increasingly considered a pooled processing entity to which an increased amount of network virtualization can be applied, leading to the notion of the virtualized RAN (vRAN). The vRAN will enable new RAN techniques such as CoMP/MIMO and intercell interference management, that are applied on a per-user basis, to be incorporated into the end-to-end network slicing.

Increasingly centralized processing in operators’ networks will result in the fronthaul being part of a public network infrastructure, typically shared with other services, including fixed access, and open to other operators in some countries, due to telecommunications deregulation. The evolution of 4th generation (4G) and advent of 5G networks, with ever higher data rates, realized with higher bandwidths and more antennas, has led to the common understanding that the transport of sampled radio waveforms between the central site and the RRHs will no longer be feasible, as it would require extremely high data rates over this shared network infrastructure. To avoid sampled radio waveform transport, new split points for the partition of the RAN protocol functions between the central unit (CU) and distributed units (DUs) are now widely discussed [5]-[8].

Ethernet is a prime candidate for this evolved fronthaul, due to its flexibility and ubiquity/cost-effectiveness [7]. It allows sharing of the network infrastructure through standardized virtualization techniques and, through its packet-switched operation, the realization of statistical multiplexing and aforementioned pooling gains.

In this article, we report on the current state of standardization towards an Ethernet-based next-generation fronthaul interface within IEEE 1914. While the use of Ethernet in the fronthaul for a software-defined RAN or vRAN has been considered in previous work such as [9], we extend the analysis of the benefits provided by Ethernet of through its provision of not only transport but also standardized network control and management that can be employed for network optimization. The overview of the most interesting functional splits and their benefits and requirements is also given in the context of 3GPP standardization [6]. Further, in addition to highlighting fronthaul timing and synchronization requirements, an overview of techniques that can enable meeting such requirements in Ethernet networks is presented. Finally, as bandwidth requirements will continue to remain significant, we provide an overview of the optical fiber technologies, such as passive optical networks (PONs), which will enable the next generation fronthaul.

2. Consideration of Ethernet in the IEEE 1914 working group

The IEEE 1914 working group [7], Next Generation Fronthaul Interface (NGFI), founded in 2016, has been motivated by the flexibility of Ethernet. There are two ongoing projects, P1914.1 and P1914.3.

The P1914.1 project focuses on defining the architecture for the fronthaul transport networks while the P1914.3 project specifies the packetization of radio traffic over Ethernet.

The scope of the P1914.1 project is to specify: 1) An Ethernet-based architecture for the transport of mobile fronthaul traffic, including user data traffic, and management and control plane traffic; 2) Requirements and definitions for fronthaul networks, including data rates, timing / synchronization, and quality of service. It focuses on specifications from the fronthaul transport networks' perspective. The definition of new functional splits between CU and DUs is out of the scope of IEEE P1914.1; rather, proposed splits in other standardization groups, such as various options in the 3rd Generation Partnership Project (3GPP) [6], see section 3, Small Cell Forum (SCF) [5] and eCPRI in the CPRI group [8] will be considered.

Three general service classes are considered as the baselines for the fronthaul transport network requirements: Control & Management, Transport Network Control & Management (related to fronthaul transport, e.g. control of delay), and Data plane. Figure 1 illustrates the future fronthaul architecture: in the uplink direction, data flows from DUs to CU (which can comprise the BBU pools) are encapsulated into Ethernet frames before being transported. Radio over Ethernet (RoE) specifies how to encapsulate CPRI frames from 2G/3G/4G RRHs into Ethernet frames. Data flows from 5G DUs and enterprise services are encapsulated into Ethernet format based on example NGFI 1 and NGFI 2 specifications. All service classes will be specified for each NGFI, but the requirements differ. The CU supports processing of Ethernet frames in all of these formats. In the downlink direction, the CU sends Ethernet frames in the formats required by the destinations. RoE frames are de-encapsulated into CPRI frames before being transmitted to 2G/3G/4G RRHs. Ethernet frames sent to 5G DUs and enterprise services are de-encapsulated based on NGFI 1 and NGFI 2 specifications.

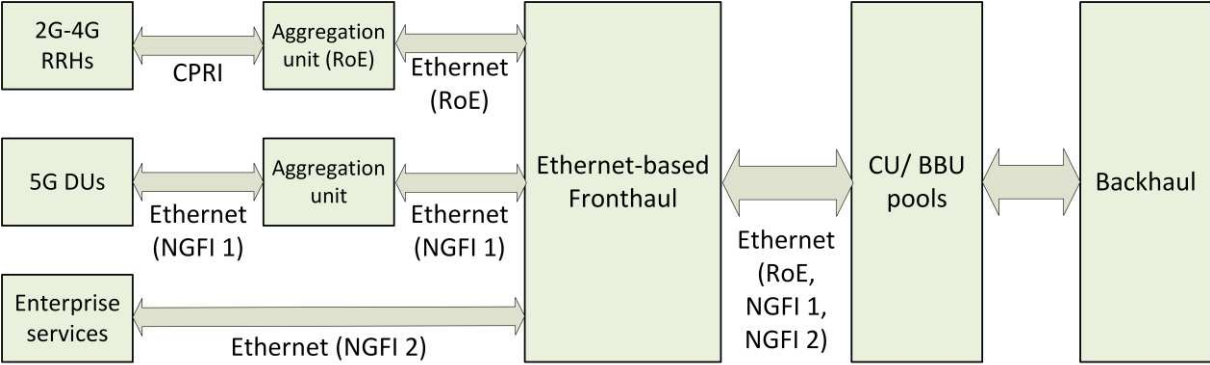


Figure 1 The flexible Future Fronthaul Network

The IEEE P1914.3 project is at its final review stage and aims for a first release by the end of 2017. By using a common header format for both data and control packets, RoE level sequencing, synchronization and multiplexing is supported. Two mappers are defined in the standard to support transporting existing radio transport protocols over Ethernet: structure-agnostic mapper and structure-aware mapper. The structure-agnostic mapper has minimal knowledge of the framing protocol it transports while the structure-aware mapper breaks a CPRI stream into antenna-carrier and control-data component streams to enable more efficient transportation and switching.

Packetization of in-phase and quadrature (I/Q) samples in both time domain and frequency domain will be defined in the standard. Future amendments can add support to radio data of other formats.

3. Split functionality in the RAN

Alternative function splits in next generation RAN architectures have gained significant interest in vendor ecosystems, as well as in 3GPP and other fora. A number of possible split points have been identified by 3GPP [6], see Fig.2, and analyzed. Variants have been identified for some split points.

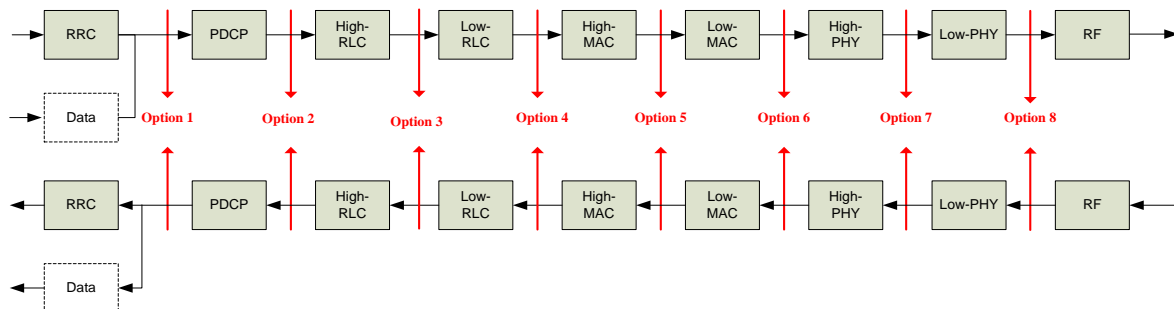


Figure 2: Functional split between central and remote units (reproduced with permission from [6])

The main criteria for analyzing these split points are: (1) data rate; (2) uniformity of the data rate depending on cell load (e.g. data pattern burstiness); (3) techno-economic scalability of the throughput (e.g. per number of antennas or layers?); (4) latency and jitter requirements; (5) synchronization requirements; (6) flexibility to support advanced features; (7) support of error detection/correction; (8) implementation complexity (e.g., required buffer sizes for link layer control); (9) compatibility with packet-optical transport networks.

The split points can be classified in two main categories:

- Low Layer Splits (LLS): those requiring a very low latency transport, typically below 250 μ s as the split is within the real time functions of the RAN. These are splits 4 to 8 in Fig.2.
- High Layer Split (HLS): those having less stringent requirements on latency, and therefore compatible with most existing transport networks. These are typically splits between the real and non-real time functions of the RAN, splits 1 to 3 in Fig.2 (including the traditional backhaul profile).

There is now a consensus in 3GPP on the main characteristics and merits of these split points and related transport profiles, summarized in Table 1.

	1	2	3-2	3-1	5	6	7-2	7-1	8
Baseline available	No	Yes (DC)	No						Yes (CPRI)
Traffic Aggregation	No	Yes							
ARQ location	Remote unit			Central unit					
Resource pooling in CU	Lowest	Increasing →						Highest	
Latency requirement	Loose				FFS	Tight			
Transport NW Peak BW requirement	N/A	Lowest	Increasing →					Highest	
	No UP req.	baseband bits					IQ (f)	IQ (t)	
	-	Scales with traffic load					Scales with antennas		
Multi-cell/freq. coordination	multiple schedulers (independent per remote unit)				centralized scheduler (can be common per central unit)				
UL Adv. Rx	FFS						Yes		

Table 1: Summary of characteristics of different central-remote unit split options (simplified from [6]). Note (definition of sub-options): Option 3.1: L-RLC: segmentation and concatenation functions, and High RLC: ARQ and re-ordering functions; Option 3.2: L-RLC: TX RLC and H-RLC: RX RLC; Option 7.1: Low-PHY: FFT (UL), iFFT (DL); Option 7.2: Low-PHY: option 7.1 plus, resource de-mapping and pre-filtering

3GPP is moving towards specifying only 2 split points, one LLS and one HLS, considering that no more splits are needed to fulfill all deployment and use case requirements. Although there is not yet a complete consensus at the time of writing of this article, it seems that 3GPP will focus its efforts for transport profiles on option 2 for the HLS, and on sub-options of option 7 for the LLS, as new additions to the traditional backhaul and IQ fronthaul (CPRI, option 8). An overview on implications of the different split options is given in Table 1.

HLSs have less stringent latency requirements, and their throughput needs are just 10-20% above those of traditional backhaul. This makes the transport profile of a HLS compatible with most existing packet-optical transport networks. Most or all functions above the split points can be virtualized and when located in the first or second aggregation point in the fixed access network, they can offer large-scale pooling (cloudification) gains. A HLS, being less tied to 5G New Radio (NR) waveforms, is inherently more future-proof with respect to L1 evolution than a LLS option. The downside is the limitation on supporting advanced cooperative features, and hence limited performance of the radio link, and the greater specialization (complexity) of the DUs.

The main objective of the LLSs is to exploit the entirety of the radio information of any given DU at the CU so as to improve radio performance (taking into account the severe interference-limited conditions in a cellular RAN). LLSs make it possible to use CoMP and distributed massive MIMO, as well as other advanced receiver techniques. LLSs are from this respect the most future-proof split option. The main drawbacks of LLSs are their stringent latency requirements and the higher bit-rates in the cases of options 7 (in particular for uplink) and 8 (for both link directions). Typical deployment scenarios of LLS profiles would be fiber-rich access environments where suitable transport can be established, most likely in rather local footprints to form cooperating clusters of DUs.

Moreover, there are techno-economic challenges for LLS option 7.1 as requirements scale with the number of antennas, making it unattractive, especially for massive MIMO. Option 7.2 introduces the possibility of doing part of the MIMO processing in the DU, at the cell premises [6] or directly in the neighborhood. As resource de-mapping can be performed at DUs only allocated resources are transmitted on the interface, offering traffic aggregation advantages when several cells are multiplexed.

5G networks will typically comprise several network aggregation points: for example, there may be a DU, an edge cloud and a central cloud. DUs may be limited to RF functions (conventional RRH), or include all or part of the RAN stack, or be a formed from a combination, e.g. in the case of multi-RAT sites. The general principle is illustrated in Fig. 3. The placement of RAN functions clearly depends on (i) transport network capabilities, (ii) service requirements (access to edge or centralized services, necessary response times) and (iii) load and availability.

Several split points could ideally be used at the same time in the same network. One intuitive example is of CoMP, requiring a LLS, which is significant for users at the cell edge, whereas single-cell processing may be sufficient for users near a DU, and a HLS would be adequate. By using a mix of split points, transport capacity could be minimized while network performance is maximized.

However, the dynamic re-allocation of functions between physically separated aggregation points is complex and currently considered only as a possible evolution. Similarly, mixing several processing levels on a UE or service flow basis requires the simultaneous support of several split points and transport network profiles. Thus, dynamic reconfiguration is not in the current scope of 3GPP specifications, but rather within other standardization bodies, such as ETSI NFV.

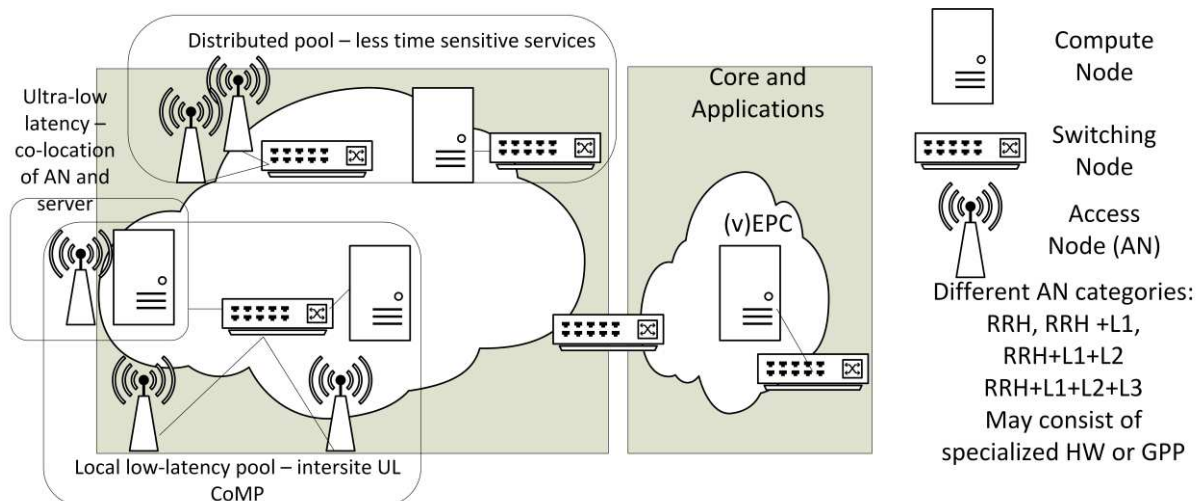


Figure 3: Hybrid 5G RAN architecture

4. Time-Synchronous Networking

DUs will require frequency synchronization to meet radio transmission requirements (center frequency of local oscillators, sampling frequencies of the waveform) and time-alignment between radio bursts transmitted [3]. Frequency synchronization has been inherently available with CPRI, as

the constant bit-rate is locked to a frequency reference. However, for packet-based NGFI solutions, mechanisms for frequency synchronization and time of day (TOD) alignment are now required. IEEE 1588-2008 Precision Time Protocol (PTP) is the likely choice for the latter. PTP assumes symmetrical delay in the forward and reverse directions; hence, one-way-delay measurement and link assignment may be required to minimize asymmetry. Clock frequency offset results in timing error, which may be minimized by exploiting frequency synchronization between master and slave. It is assumed that Synchronous Ethernet (SyncE) will be used for this. Thus, PTP traffic and SyncE are considered alongside any Ethernet-based fronthaul protocol.

As stated in Section 3, while LLS offers increased opportunity for advanced network performance, latency and packet delay variation are major concerns. This section presents a number of methods targeting time sensitive networking over a bridged Ethernet-based fronthaul network, in an ascending order of design complexity.

The IEEE P802.1CM standards group (Time-Sensitive Networks for Fronthaul) is in the process of compiling profiles on time-sensitive networking mechanisms which address the stringent requirements in fronthaul networks [10]. Among other requirements, the standard classifies four categories for time synchronization (A+, A, B and C). For example, category A+ has an absolute time error of 12.5 ns, by far the most stringent requirement, and is applied for MIMO applications. Category C defines a maximum absolute time error in the order of 1.3 μ s, a requirement for time division duplex based operations. Currently, the main focus is on transporting CPRI traffic over Ethernet which is termed as Class 1. Class 1 consists of two profiles (Profile A and Profile B) addressing different requirements.

Profile A comes with the lowest design complexity: it employs no advanced means for time-sensitive networking and employs just strict priority (SP) scheduling for the different transported traffic classes, using an increasing priority from background data, control and management data, I/Q user data to synchronization data. Frame delay variation (FDV) or packet delay variation is caused by queuing (introduced by aggregation and/or by blocking of higher priority (HP) traffic by lower priority (LP) traffic) and the number of hops or switches and frame sizes. Simulation results [11] show that Profile A can meet Category C requirements for very low aggregation levels and small frame sizes. SP is an improvement over no priority-based scheduling or other schedulers such as WRR (Weighted Round Robin). SP can reduce FDV on average, but not the peak delay variation [12].

FUSION, a promising approach presented in [13], combines packet and circuit switching to multiplex HP traffic streams (the circuit switched part) with LP statistically multiplexed streams (the packet switched part) over an Ethernet network. The main idea is to exploit the inter-packet gaps between HP frames to transmit the LP streams (see Fig. 4). As a result, the FDV is significantly reduced and can potentially meet Category B requirements. Furthermore, this approach achieves a significantly improved utilization compared to a fully provisioned circuit switched network and does not require an additional (out-of-band) form of synchronization.




Figure 4: FUSION: Exploiting the inter-packet gaps between HP frames to transmit LP frames

Profile B employs frame pre-emption based on P802.1Qbu [14], where a HP frame can pre-empt a LP one (see Fig. 5). However, pre-emption is not instantaneous and introduces further delay. The worst case of this delay is 124 ns for 10 Gb/s Ethernet (equivalent to the processing time of a 155-octet packet). The advantage of pre-emption is a reduction in the end-to-end latency (compared to profile A and the FUSION approach) for the same number of hops. This leads to an increased reach. But the advantage depends on the traffic mix at each aggregation point. If the traffic is mainly of the same priority, the benefit will be low. With a small frame size (e.g. 300 bytes), frame pre-emption can meet Category C requirements [11].

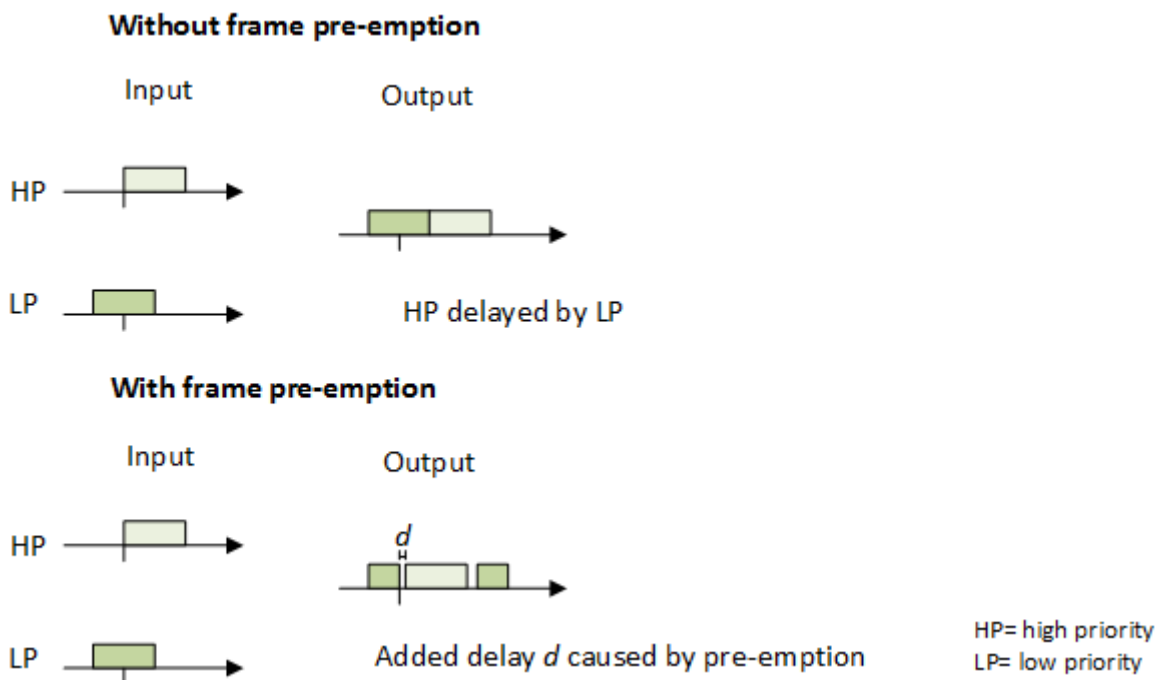


Figure 5: Frame pre-emption: HP frame can pre-empt a LP frame

Time-aware scheduling based on IEEE P802.1Qbv [15] separates traffic into uncontended window sections to reduce FDV further. HP traffic is assigned a HP window section (the protected section) while LP traffic is assigned a best effort section. Transmission through a window is determined by the

scheduler. Simulation results for fronthaul with CPRI [16] and a new functional split [17] show that such a scheduler can completely remove FDV. The complexity of time-aware scheduling increases with the network size. A global scheduler is required to ensure that intra-window contention does not take place within the various network nodes. Further, guard periods are required so that LP traffic does not overrun into the protected section [16], [17]. A combination of time aware-scheduling and pre-emption can be used to reduce the size of required guard periods.

5. Network optimization

Variable split options could facilitate scalable, cost-effective deployments and real-time optimization [3] trading fronthaul resource consumption against radio performance, and serving a fluctuating user demand while meeting Quality of Service (QoS) needs. A Service Level Agreement (SLA) between an operator and the user will define metrics by which a fronthaul service is measured, the methodology by which the metric is verified and penalties if the agreed-upon QoS parameters should not be achieved. SLAs should explicitly consider the new fronthaul configurations, and their reconfiguration and failure modes arising from the introduction of Ethernet-based transport and traffic aggregation.

The life-cycle of a service is usually split into three phases: (1) provisioning and turn-up to verify the SLA, (2) performance management (checking that the service meets the SLA), and (3), fault management (sectionalizing the problem, escalating, and correcting it). Existing, standardized Operations, Administration and Maintenance (OAM) addresses the protocol layering in Ethernet-based networks, with IEEE 802.3ah used for link, IEEE 802.3ag for connection and ITU Y.1731 for service layers [18]. This framework may need to be supplemented to address the challenges presented by the new fronthaul.

Key Performance Indicators (KPIs) for a typical radio service are: availability/downtime, packet delay (latency), packet delay variation (jitter), loss and throughput. Measurements for delay and jitter are of the order of millisecond accuracy based on pings, which is clearly insufficient for the LLSs, as discussed in Section 3. Hence, high-resolution timestamping becomes essential. Historically, vendors have resisted the use of 3rd party probes in favor of “own-brand” devices, but their capability is not always sufficient. Consequently, operators are requesting an open approach that allows the use of 3rd party probes. The new performance metric measurements associated with the fronthaul, together with other data such as RF signal quality and application performance, feed to the SON (self-organizing / self-optimizing network) algorithm that determines the network configuration and associated parameter settings. Fig. 6 illustrates such a SON-controlled system that creates network slices using virtualized network functions (VNFs) and configures the fronthaul based on subscriber service and location. Performance metrics are processed by an analytics service, which feeds to an optimization engine that determines the appropriate network configuration and an orchestration engine that effects network changes through, for example, a software-defined network (SDN) controller and a RAN manager/controller.

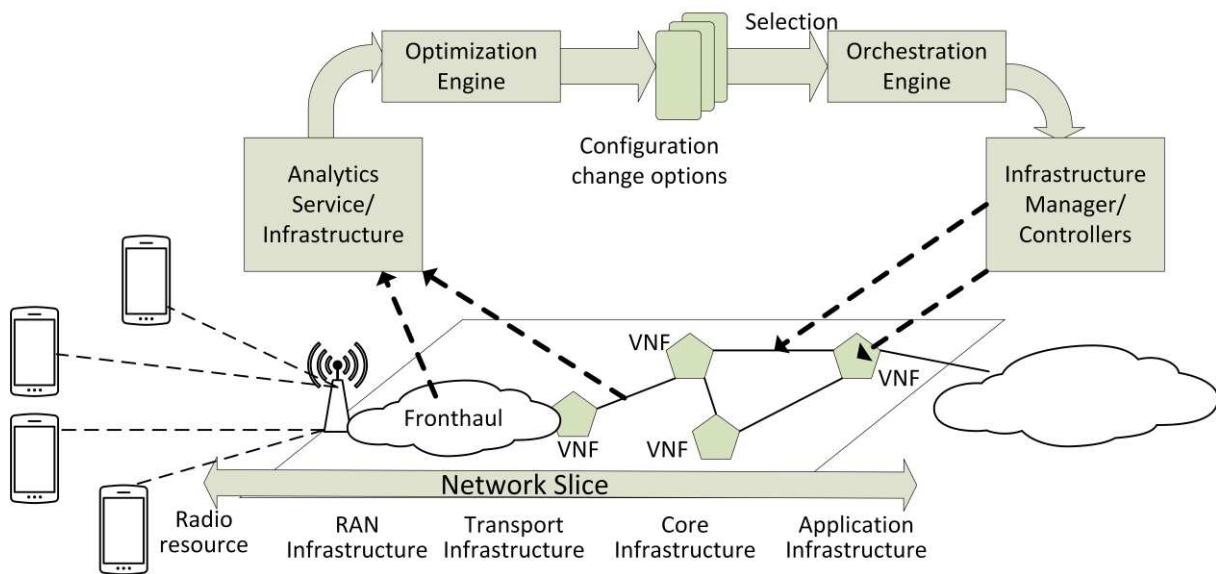


Figure 6 : SON in a 5G network, including orchestration of the RAN

6. Interoperability with PONs

Optical access systems have seen a widespread deployment over a decade or more. Besides active Ethernet, Gigabit-capable Passive Optical Network (G-PON) was introduced in the field by several operators, initially aimed at providing up to 100 Mbit/s rate to multiple end customers connected simultaneously to the same optical fiber. Currently, commercial offers with more than 100 Mbit/s are possible on G-PON and even 1 Gbit/s commercial offers are available. After G-PON, XGS-PON (a PON operating at 10Gbit/s downstream, and 2.5 or 10 Gbit/s upstream) is recognized as the next deployable solution. Following XGS-PON, present standardization is working on 25 Gbit/s line rates as an add-on solution for the deployed Optical Distribution Network (ODN). All of these solutions are based on a wavelength channel pair for the up- and down-streams that coexist on the same ODN. TDM (Time Division Multiplexing) and TDMA (Time Division Multiple Access) are used for sharing the trunk part of the ODN and a single interface at the central office. Since 2015, multi-wavelength PON solutions have been standardized combining time and wavelength (TWDM PON) or only multi-wavelength (WDM PtP PON) approaches.

5G aims at fiber-like experience for mobile users. Undoubtedly, fiber will be the dominant technology solution for backhaul. Optical access systems, particularly for PONs must obviously meet the requirements of both fronthaul and backhaul, and support the different split options. Fig. 7 shows how the existing access solutions can collect Ethernet traffic in two RAN scenarios:

- a) backhaul where CU functions and DUs are co-located at the antenna site
- b) Ethernet-based evolved fronthaul (e-fronthaul) with evolved Ethernet-DUs (DUs with a new RAN function split) localized at the antenna site and virtualized CUs located at a master central office (operator point of presence node).





Figure 7. Optical access solutions (PtP, T(W)DM Pon and WDM PtP PON) for a) backhaul and b) e-fronthaul

The support of low latency and synchronization by T(W)DM or PtP WDM PONs will be the major differentiators from a residential (fixed) optical access system for the e-fronthaul. Here, different

flavors of dynamic bandwidth allocation (DBA) based on time allocation for TDM/TDMA or through combination with wavelength allocation for TWDM could be proposed to accommodate the required timing performance. Coordination between OLT and virtualized CU could be also proposed due to the fact that the virtual CU knows, in advance, the desired time allocation of the radio signal for each DU.

7. Conclusion

The evolution towards the next generation mobile network requires a new, converged radio and fixed access network infrastructure. New functional split options are required between centralized and distributed units in the next generation radio access network to enable more centralized deployment in which radio signals are transported over public, rather than private networks, so that more feasible bit-rates in the fronthaul are required. Transport of the new fronthaul has different requirements for data rates, delay and jitter, compared to the existing fronthaul, and depends on the split option chosen for a particular service. The use of Ethernet as a new transport protocol for the fronthaul is attractive from multiple perspectives: cost, standardized network control and management functions, software-defined networking and extending existing means of network monitoring and tools for optimization. However, new methods and tools for controlling delay and packet jitter will be required to serve the advanced lower-layer split options that are the most promising for improving radio network performance. Finally, compatibility with existing and new optical access technologies needs further research to ensure a future-proof deployment in the face of increasingly demanding radio network requirements.

Acknowledgment

N J Gomes, H Thomas, P Chanclou, D Münch and V Jungnickel would like to acknowledge the iCIRRUS project, funded by the European Union's Horizon 2020 Research and Innovation program under grant agreement no. 644256. P Assimakopoulos acknowledges the NIRVANA project, funded by the UK's Engineering and Physical Sciences Research Council. The authors are grateful to Jörg-Peter Elbers (ADVA Optical Networking) and Rudolf Winkelmann (Nokia Bell Labs) for useful discussions and advice.

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] T. S. Rappaport, R. W. Heath, Jr., R. C. Daniels, and J. N. Murdock, *Millimeter Wave Wireless Communications*. Englewood Cliffs, NJ, USA: Prentice Hall, 2015.
- [3] V. Jungnickel *et al.*, "The role of small cells, coordinated multipoint, and massive MIMO in 5G," in *IEEE Communications Magazine*, vol. 52, no. 5, pp. 44-51, May 2014
- [4] China Mobile Research Institute White Paper, "Cloud-RAN: The road towards green RAN", 2011

- [5] Small Cell Forum, "159.07.02 Small cell virtualization functional splits and use cases", 2016
- [6] 3GPP, "3GPP TR 38.801 v14.0.0 Study on New Radio Access Technologies: Radio Access Architecture and Interfaces", 2017
- [7] IEEE 1914 Next Generation Fronthaul Interface Working Group, <http://sites.ieee.org/sagroups-1914/>
- [8] Common Public Radio Interface: see e.g., "Industry leaders agree to develop new CPRI Specification for 5G", <http://www.cpri.info/press.html>
- [9] C.-L. I, Y. Yuan, J. Huang, S. Ma, C. Cui and R. Duan, "Rethink fronthaul for soft RAN", in *IEEE Communications Magazine*, vol. 53, no. 9, pp. 82-88, September 2015
- [10] IEEE, "IEEE P802.1CM Draft 0.5 Time-Sensitive Networks for Fronthaul", 2016, [Online]. Available: <http://www.ieee802.org/1/pages/802.1cm.html>
- [11] J. Farkas and B. Varga, "P802.1CM Simulation Results for Profiles A & B", 2016, [Online]. Available: <http://ieee802.org/1/files/public/docs2016/cm-farkas-profiles-A-and-B-0316-v01.pdf>
- [12] M. K. Al-Hares, P. Assimakopoulos, S. Hill, A. 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
- [13] R. Veisllari, S. Bjornstad, J. P. Braute, K. Bozorgebrahimi and C. Raffaelli, "Field-Trial Demonstration of Cost Efficient Sub-wavelength Service Through Integrated PacketCircuit Hybrid Network," in *OSA/IEEE Journal of Optical Communications and Networking*, vol.7, no.3, A379-A387, 2015
- [14] IEEE, "IEEE P802.1Qbu Draft 3.1 Frame Preemption", 2015, [Online]. Available: <http://www.ieee802.org/1/pages/802.1bu.html>
- [15] IEEE, "IEEE P802.1Qbv Draft 3.1 Enhancements for Scheduled Traffic", 2015, [Online]. Available: <http://www.ieee802.org/1/pages/802.1bv.html>
- [16] T. Wan and P. Ashwood-Smith, "A Performance Study of CPRI over Ethernet with IEEE 802.1Qbu and 802.1Qbv Enhancements," in *IEEE Global Communications Conference (GLOBECOM)*, 2015
- [17] M.K. Al-Hares, P. Assimakopoulos, D. Münch and N.J. Gomes, "Modeling Time-Aware Shaping in an Ethernet Fronthaul", accepted for presentation in *IEEE Global Communications Conference (GLOBECOM)*, 2017
- [18] R. Vaez-Ghaemi, "Ethernet OAM Test Applications", White Paper, JDSU, 2008, [Online]. Available: http://www.viavisolutions.com/sites/default/files/technical-library-files/ethoam_wp_acc_tm_ae.pdf