A Flexible, Ethernet Fronthaul for 5th Generation Mobile and Beyond

Nathan J. Gomes[1], Volker Jungnickel[2], Philippe Chanclou[3], Joerg-Peter Elbers[4] and Peter Turnbull[4]
1: Communications Research Group, University of Kent, Canterbury, Kent, CT2 7NT, United Kingdom
2: Fraunhofer HHI, Germany, 3: Orange Labs, France, 4: ADVA Optical Networking, Germany/UK
N.J.Gomes@kent.ac.uk

Abstract: Using Ethernet in the fronthaul can deliver the statistical multiplexing gains offered by the new functional splits proposed for the radio access network, but latency and delay variations are challenges that must be overcome.
OCIS codes: (060.2330) Fiber optics communications; (060.4250) Networks

1. Introduction

It is now accepted that the Centralized, Coordinated or Cloud Radio Access Network (C-RAN) will be an essential component of 5th Generation (5G) mobile networks [1, 2] with or without virtualization features. However, it has also become apparent that current fronthaul technology, used in the interconnection of the centralized base stations with the distributed radio units (DRUs), is being challenged to fit the requirements envisaged for 5G, with its wider bandwidths and use of advanced multiple antenna techniques [3]. In order to reduce data rates, alternative splits in the baseband signal processing functionality that resides centrally and that which is placed in the DRUs have been proposed [4-7]. The principal idea is to avoid the transport of (over-)sampled In-phase and Quadrature (IQ) waveforms over the fronthaul and to move all or part of the waveform generation and detection to the DRUs.

Current fronthaul solution are partly proprietary and specific to requirements of the used hardware. From this perspective, the use of Ethernet, as a ubiquitous technology, is an advantage [5]; it can permit convergence of the fronthaul interface with that of existing or new Ethernet fixed access networks. The access Ethernet infrastructure already provides a set of useful features, such as standardized Operations, Administration and Maintenance (OAM) and the provision of physical network resources that are virtualized and can be shared among multiple service providers. The mapping of current fronthaul solutions into Ethernet packets is considered as a first step to permit convergence [8]. More significant benefits are possible with new functional splits, however, as the data is no longer continuous samples, but, depending on the split, may be of variable rate and open to statistical multiplexing gains.

There are, however, challenges in using Ethernet in the fronthaul [5]. The current fronthaul solutions are based on fully synchronous, time-division multiplexed transport of continuous waveform samples which has minimum impact on delay and delay variation of the samples as the radio waveforms are reconstructed at the DRUs. With Ethernet transport, even if we deal with sampled waveforms, these will be placed in (typically) variable length packets which may experience variable packet delay variation over a switched Ethernet infrastructure, especially if these packets share a converged access infrastructure with other types of traffic. Retiming is obviously necessary. The effects and requirements on Ethernet equipment and architecture may be different for different functional splits, but delay variation will still be significant in the construction of radio frames at the DRUs.

In this paper, we provide an overview of different functional split opportunities and a resulting, generalized and converged 5G network architecture using the new fronthaul. Then, we examine the requirements for the transport in terms of bit-rate reductions, latency and delay variations. Finally, we examine the prospects for traffic control and possible orchestration of network functions to optimize the performance of the network.

2. The new, virtualized Radio Access Network architecture

Fig.1 shows the functional split options that have been identified at Layer 1 (the Physical layer) of an LTE-based RAN. Split options within Layers 2 and 3 are also possible. At these higher layers, and certainly at Layer 3, Internet Protocol connections mean that any distributed unit effectively becomes a small base station in itself. Greater centralization of functions and potential for cooperation and virtualization is offered by split points towards the right of Fig. 1. Moving from the right in Fig.1, where the current fronthaul split point, based on sampled waveform transport using protocols such as the Common Public Radio Interface (CPRI), is shown, an immediate bit-rate reduction is offered by Split point 1, which, in the downlink, transports symbols before modulation onto the OFDM subcarriers in the IFFT process. The bit-rate required, however, is constant as all used subcarriers require symbol modulation. Split 2 includes resource mapping at the DRU, which means that unused resource blocks are not transported. From here and increasingly as the split point is moved to the left, statistical multiplexing gains become

significant. Split 3 places the layer mapper at the DRU with precoding of multi-antenna streams carried out there. This reduces the data rate as users are typically served with variable numbers of streams in parallel, depending on the multiple-input multiple-output (MIMO) channel and the opportunities for multiuser MIMO and coordinated multi-point (CoMP). In the uplink, bit-rate reduction may be less significant, and the split point may be different even on a per-user basis, depending on the selected MIMO/CoMP mode for users scheduled on the same radio resource (e.g., soft-bits for better decisions on multi-antenna streams or samples in the frequency-domain for joint detection in CoMP, might be transported). Split 4 places the remaining physical layer functions in the DRU, reducing the bit-rate as FEC bits are not carried but increasing the DRU complexity.
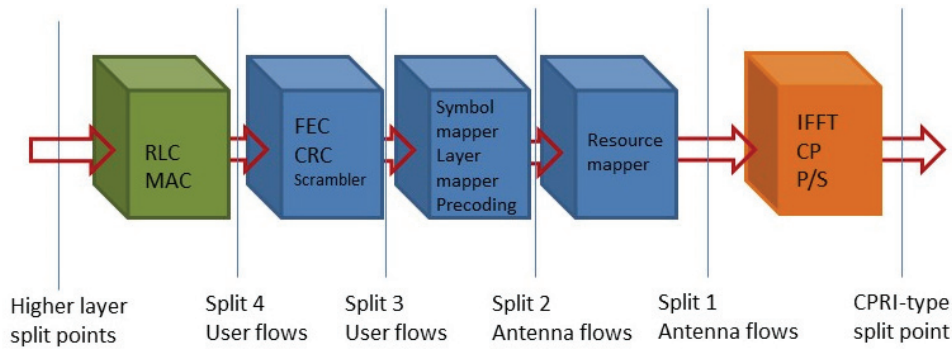

Fig.1 Functional split point options in LTE-like system (downlink example, PHY layer options only, shown)

Fig. 2 shows a generalized switched Ethernet access network that can be used to transport the mobile fronthaul traffic to different types of DRUs from a virtualized base station pool, with different and variable sets of functions implemented there and in the radio units. Some functions may also be distributed into the aggregator units nearer the radio unit sites. The switched Ethernet network can carry in parallel Layer 3 IP network traffic to small base stations and to the mobile network backhaul. Finally, it can also be used to transport fixed access network traffic.


Fig. 2 An X-haul network. Transport can be up to layer 3 (as with micro-base station control) and convergence with fixed access is possible.

3. Fronthaul requirements

As discussed in the previous section, bit-rate requirements reduce as we place more functions in the DRUs, but this opposes the desire for centralization that can lead to benefits in joint processing and sharing of centralized resources. Example calculations have been provided e.g. in [6, 7]. Bit-rate reductions are insignificant when comparing the split points 3 and 4 dealing with user flows; these split points typically require several times lower data rates than split points 1 and 2 dealing with antenna flows [7]. The reduction for these compared to the sampled IQ transport is around one order of magnitude, considering peak traffic rates. It is estimated that signals requiring 10 Gb/s links for CPRI type transport, could use less than 500 Mb/s (and as little as around 100 Mb/s) depending on the split point. Looking forward to requirements for sectors, multiple antenna streams and the envisaged user bit-rates for 5G, it is

clear that 10GbE interfaces will be required, certainly to dimension for peak data rates. Although, some reductions are likely through the significant statistical multiplexing gains, similar to those observed for backhaul traffic, the aggregation to/from multiple DRUs will require the use of 100GbE trunks.

The significant constraint for an Ethernet fronthaul arises from timing and synchronization issues [9], as these are not inherently provided in frame/packet-based protocols. Most carriers will use Synchronous Ethernet in their networks. Carrier-grade equipment means greater expense than consumer products, but Synchronous Ethernet can provide the required frequency synchronization to radio units, while shared infrastructure benefits still accrue. Timing and phase synchronization may be provided by protocols such as IEEE 1588v2 precision timing protocol. There is then a need to conform to maximum delay requirements, in most cases set by the hybrid automatic repeat request (HARQ) link protocols over the air interface; if the split point is within the extent of the HARQ, the fronthaul delay is limited to less than 1 ms for LTE-type networks. If the split point moves the HARQ to the DRU, delays much greater than this are tolerable [9]. Requirements on delay asymmetry impose further constraints [4].

4. Traffic control and orchestration

The converged network of Fig.2 envisages flexibility in provisioning and statistical multiplexing gains through the use of Ethernet switches. These then also become the most significant hindrance to meeting latency and delay variation constraints [9]. While the new functional splits may not demand synchronous reception of each block of samples, all information required to construct a particular radio frame has to arrive within a given time interval. Packet delay variations in switches will be problematic, hence cut-through operation and low-layer switching, based on predefined tables in virtual LANs may be required. Traffic control will be necessary to ensure that there is minimal contention (and therefore queuing) of the mobile fronthaul traffic, and may extend to defining priorities for control traffic, pre-emption to ensure this traffic is received, and orchestration of the software-defined network.

5. Conclusion

The use of Ethernet in a new, flexible fronthaul and of new functional split points to move the processing from the centralized baseband unit to the DRUs in a C-RAN provides significant benefits for convergence, higher efficiency through statistical multiplexing and provision of standard carrier Ethernet OAM. However, transport through a switched network will require software-defined mechanisms for control of the traffic in order to minimize queuing delay and packet delay variations.

6. Acknowledgment

7. References

[1]   China Mobile Research Institute White Paper, "Cloud-RAN: The road towards green RAN", (2011), online [available: http://labs.chinamobile.com/cran/wp-content/uploads/CRAN_white_paper_v2_5_EN.pdf]

[2]   Ericsson White Paper, "Cloud-RAN – the benefits of virtualization, centralization and coordination", (2015). Online [available: http://www.ericsson.com/res/docs/whitepapers/wp-cloud-ran.pdf]

[3]   V. Jungnickel et al. "Software-defined open architecture for front- and backhaul in 5G mobile networks," *16th International Conference on Transparent Optical Networks (ICTON), 2014*, Graz, 2014.

[4]   P. Rost et al, "Cloud technologies for flexible 5G radio access networks", IEEE Comms Mag, vol.52, no.2, pp68-76 (2014).

[5]   N.J. Gomes et al, "Fronthaul evolution: from CPRI to Ethernet", Optical Fiber Technology, vol. 26, part A, pp50-58 (2015).

[6]   U. Doetsch et al., "Quantitative analysis of split base station processing and determination of advantageous architectures for LTE", Bell Labs Tech J, vol.18, no.1, pp.105-128, 2013.

[7]   China Mobile Research Institute White Paper, "Next Generation Fronthaul Interface v1.0", (2015). Online [available: http://labs.chinamobile.com/cran/wp-content/uploads/2015/09/NGFI-Whitepaper_EN_v1.0_201509291.pdf]

[8]   IEEE, "1904.3 Task Force: Standard for Radio over Ethernet encapsulations and mapping", online [available: http://www.ieee1904.org/3/tf3_home.shtml].

[9]   iCIRRUS deliverable D3.1, "Verification of Ethernet as transport protocol for fronthaul/midhaul", (2016), online [available: http://www.icirrus-5gnet.eu] .