

# Kent Academic Repository

## Full text document (pdf)

### Citation for published version

Gomes, Nathan J. and Chanclou, Philippe and Turnbull, Peter and Magee, Anthony and Jungnickel, Volker (2015) Fronthaul evolution: From CPRI to Ethernet. *Optical Fiber Technology* . ISSN 1068-5200.

### DOI

<https://doi.org/10.1016/j.yofte.2015.07.009>

### Link to record in KAR

<https://kar.kent.ac.uk/50278/>

### Document Version

UNSPECIFIED

#### 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](mailto: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>

Invited Paper

## **Fronthaul Evolution: from CPRI to Ethernet**

Nathan J. Gomes<sup>a</sup>, Philippe Chanclou<sup>b</sup>, Peter Turnbull<sup>c</sup>, Anthony Magee<sup>c</sup>, and Volker Jungnickel<sup>d</sup>

Corresponding author email: [N.J.Gomes@kent.ac.uk](mailto:N.J.Gomes@kent.ac.uk); Telephone: +44 1227 823719

### Highlights

- Extensively referenced requirements (using key standards) for future fronthaul
- Covers advantages and challenges of using Ethernet in the fronthaul
- Ethernet-based “fronthaul lite” proposed to meet 5G demands
- Proposals for new functional split between baseband unit and radio head
- Caters for both co-ordinated multipoint and massive MIMO

### Abstract

It is proposed that using Ethernet in the fronthaul, between base station baseband unit (BBU) pools and remote radio heads (RRHs), can bring a number of advantages, from use of lower-cost equipment, shared use of infrastructure with fixed access networks, to obtaining statistical multiplexing and optimized performance through probe-based monitoring and software-defined networking. However, a number of challenges exist: ultra-high-bit-rate requirements from the transport of increased bandwidth radio streams for multiple antennas in future mobile networks, and low latency and jitter to meet delay requirements and the demands of joint processing. A new fronthaul functional division is proposed which can alleviate the most demanding bit-rate requirements by transport of baseband signals instead of sampled radio waveforms, and enable statistical multiplexing gains. Delay and synchronization issues remain to be solved.

### Keywords

Cloud-Radio Access Network (C-RAN); Fronthaul; Common Public Radio Interface (CPRI); Synchronous Ethernet; Network Synchronization; Optical transmission

## 1. Introduction

### Background

Next-generation (5<sup>th</sup> generation, 5G) mobile networks are faced with providing a step-change in capability due to the explosion of mobile device usage and the ever-increasing capabilities of the end-user devices. The requirements for 5G are also manifold, as it is envisaged that it will cater for high-bandwidth high-definition streaming and conferencing, to machine interconnectivity and data collection for the Internet of Things, including ultra-low latency applications such as driverless cars.

The higher wireless user data-rates envisaged in next-generation mobile networks (up to 10 Gb/s in proposals for 5G systems [1]) generally demand the use of shorter radio

transmission distances. These shorter distances in cellular mobile communications can be achieved through smaller cells (each with its own small base station) [2] or by distributing the antenna heads, usually termed remote radio heads (RRHs) into each cell [3], see Figure 1. The latter approach is seen to provide certain advantages [4]: (1) it can lead to a greater degree of centralization, with pools of base station baseband units (BBUs) connected to large numbers of RRHs, this enabling flexibility in the connections to meet varying traffic demands in the coverage area, a reduced demand in the need for base station sites, and improved energy efficiency due to the sharing of the power requirements of the co-located base stations; (2) improved coverage in the cell; (3) virtualisation and “cloudification” of base station functions in cloud-radio access networks (C-RANs); (4) enhanced possibilities for joint processing of signals transmitted from and received by different RRHs, such as the enablement of co-ordinated multipoint (CoMP), through low-latency interconnection of the co-located BBUs.

There are challenges, however, in providing the *fronthaul* distribution links between BBUs and RRHs, and these will increase for next-generation mobile networks [5]. The principal challenge lies in the increased bandwidth/bit-rate. The information transported between BBUs and RRHs is generally in the form of sampled radio signals. Already, for long-term evolution-advanced (LTE-A) signals which may have bandwidths up to 100 MHz, a single uncompressed sampled radio waveform requires a link bit-rate of over 5 Gb/s (assuming 16-bit samples). Up to now, compression factors of 2 have been shown to be lossless with some tolerable loss for a compression factor of 3 [6]. The bit-rate requirements for the much higher aggregate bandwidths expected in 5G can be expected to increase by more than an order of magnitude – of the order of tens of Gb/s. Further, this represents the bit-rate requirements of one radio stream, whereas RRHs will typically have a number of physical antenna elements requiring different radio signals.

The second significant challenge for the BBU-RRH design is the more exacting demands being placed on latency and jitter in current proposals for 5G [1]. Requirements for minimum latency may come from user applications or from the needs of maintaining the relevance of channel state information; such latency definitions will not be restricted simply to the fronthaul, so the contribution of the fronthaul and the various components of delay within it may not be clearly defined. Compression, advantageous for reducing bit-rate requirements, may add to latency. Jitter may manifest itself in variations in the transmission times of the radio signals, effectively increasing the phase noise of the transmitted signals.

In order to make best use of the potential flexibility of the BBU-RRH fronthaul, an open and accommodating topology is required. This also provides a challenge. The most widely used current transport standard, the Common Public Radio Interface (CPRI) [7], indicated in Fig. 1, was originally defined as an internal base station interface to allow antenna functions to be moved to rooftops and mastheads, away from the baseband processing in the cabinet – the demand now is to use it over links of several km. If existing fibre network infrastructure is to be used, compatibility with Ethernet based technology is required. Although wavelength division multiplexing (WDM) overlays are possible, much of the mobile radio access network hardware is being built specifically for purpose.

Thus, current fronthaul standards such as CPRI, OBSAI (Open Base Station Architecture Initiative) [8] and ORI (Open Radio Interface) [9] provide dedicated transport protocols, specifically designed for and suited to the requirements of sampled radio waveform transport. The framing carried out, for example, is done at regular intervals with frame sizes matched to specific slices of the wireless system frames (the fronthaul transport frame length in bits expands with increasing bandwidth). This form of framing enables precise

synchronisation of the RRH. However, as waveforms are transported, the bit rates are proportional to radio channel bandwidths and numbers of antenna elements (not, necessarily, user data rates). Further, the *continuous* transport of time-domain radio waveform samples leads to the absence of any possible statistical multiplexing gains as these signals are distributed over a fronthaul network. For operators, there are also questions about how to manage and provide service level agreements around the fronthaul service.

### Architectural Evolution towards Fronthaul

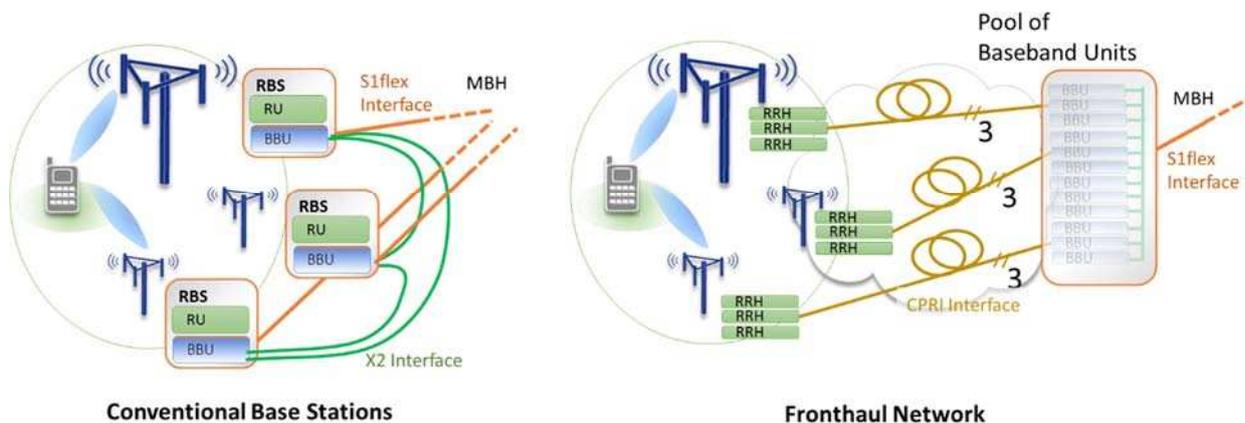


Figure 1 Conventional, separate base stations and C-RAN with BBUs hosted, X2 channels connected locally and a fronthaul network connecting BBUs with RRHs (over CPRI links).

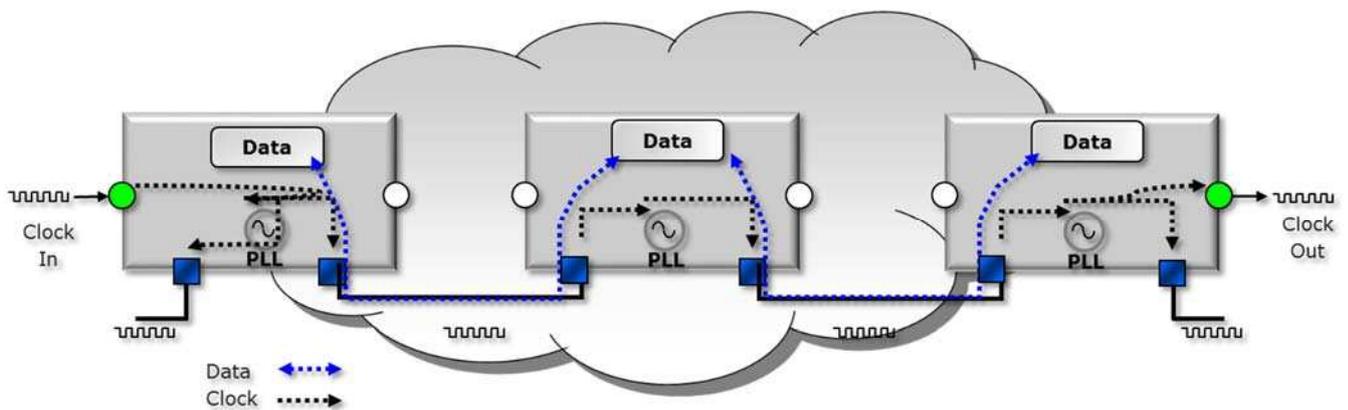
## 2. Ethernet in fronthaul

The use of Ethernet for transport in the fronthaul is appealing as a result of the maturity of the technology and its adoption within access networks. The operations, administration and maintenance (OAM) capabilities of Ethernet [10] offer a standardised means of managing, fault finding and supporting performance monitoring of the network.

It is possible to transport Ethernet in the fronthaul over CPRI, with Ethernet used to provide OAM functions missing in CPRI. The Ethernet “control channel” could be used to support network management protocols. The Ethernet signals are extracted from the CPRI/ORI equipment, which can retain their synchronisation mechanisms. However, apart from the use of OAM mechanisms, most of the advantages possible with the move to an open fronthaul transport protocol are not realised when using CPRI as the underlying transport layer. The use of Ethernet as the underlying transport layer may bring the following:

- Use of commodity equipment, or at least lower-cost, industry-standard equipment.
- Sharing of equipment with fixed access networks, enabling greater convergence and cost reductions.
- Use of switches/routers to enable statistical multiplexing gains and lower the aggregate bit-rate requirements of some links.
- Use of standard IP/Ethernet network switching/routing functionality, including moves to functional virtualisation and overall network orchestration.
- Monitoring through compatible hardware probes.

Thus, Ethernet could be used to transport CPRI/ORI frames. This would have the advantage of some backwards compatibility with existing equipment, with CPRI-Ethernet “mappers” and “de-mappers” at the edges of a fronthaul distribution system still allowing legacy equipment to transfer CPRI signals. However, it does lead to an additional framing overhead. An alternative, then, would be to place the radio waveform samples directly in Ethernet frames. The concern, in both cases, is the possible loss of the frame synchronism that is inherent in CPRI-type transport once placement in Ethernet frames is used. In Synchronous Ethernet, see Fig. 2, nodes extract the reference clock from received data rather than each node using its own internal oscillator. This mode of operation can certainly be helpful in telecoms networks, and standardisation has been led by ITU-T [11], [12], [13]. However, while frequency synchronism between devices in the network can be gained, additional mechanisms are required for precise time and phase synchronism. IEEE 1588, or Precision Time Protocol (PTP), is significant in this context, achieving synchronism through the exchange of time-stamped packets [14]. ITU-T, IEEE and others are continuing to develop standards for the use of both Synchronous Ethernet and PTP, individually or in combination, in various applications [15].



*Figure 2 Synchronous Ethernet operation: clock extraction and resynchronization is carried out at each node*

Thus, bearing these possible solutions in mind, we can briefly review the previously listed possibilities.

First, it seems not possible to use low-cost, commodity Ethernet equipment due to the timing requirements in the transport of radio signal waveforms. Synchronous Ethernet, with timing protocols such as PTP, leads to a requirement for carrier-grade equipment. However, this still may be better than using CPRI alone, because it allows for convergence with fixed access network deployments.

Then, while statistical multiplexing gains and lower aggregate bit-rate advantages are targeted, it is clear that low-latency Ethernet switching is required. Ethernet switches can operate in two modes: store-and-forward, where the entire Ethernet frame is stored and checked at the switch before forwarding on the appropriate output, and cut-through in which the frame is forwarded on the appropriate output as soon as the destination address field (near the start of the frame) has been read. In some cases, a port-to-port mapping, using for example virtual local area networking setups, may be used to reduce latency even further. The disadvantage of erroneous frames being forwarded is not seen to be significant as the

error-rate is usually very low. Examining Fig. 3, some statistical multiplexing gain may be possible with knowledge of the overall highest provision required – option (a) for the links to the RRHs. But, if there is any queuing on output links, delay variation would cause significant performance problems. A port-to-port mapping could be carried out with the regular use of frames by different radio signals (transmitted at regular intervals) and the network knowing that no link is loaded to an extent that there would be queuing. The 10 Gb/s links would be sufficient for two uncompressed 100-MHz bandwidth radio signals, more if compressed radio signals were used. The problem in this case is that cut-through operation cannot be used, as incoming 100 Gb/s frames would need to be buffered, as they cannot be transmitted fast enough on the 10 Gb/s outputs. Similarly, in the uplink, the 10 Gb/s frames would need to be stored first and then transmitted at 100 Gb/s. Cut-through switches require the option shown as (b) in Fig.3, where there is no bit-rate conversion, but there is a need, thereby, for more ultra-high bit-rate links. Cut-through operation could be important if large Ethernet frames are used (e.g. jumbo frames), which may be more efficient in the transport of radio waveform slices. For store-and-forward, large frames will cause larger delays. What is essential in either case is the need for uncontended outputs to avoid variable queuing delays, and this will require management of paths setup through the network.

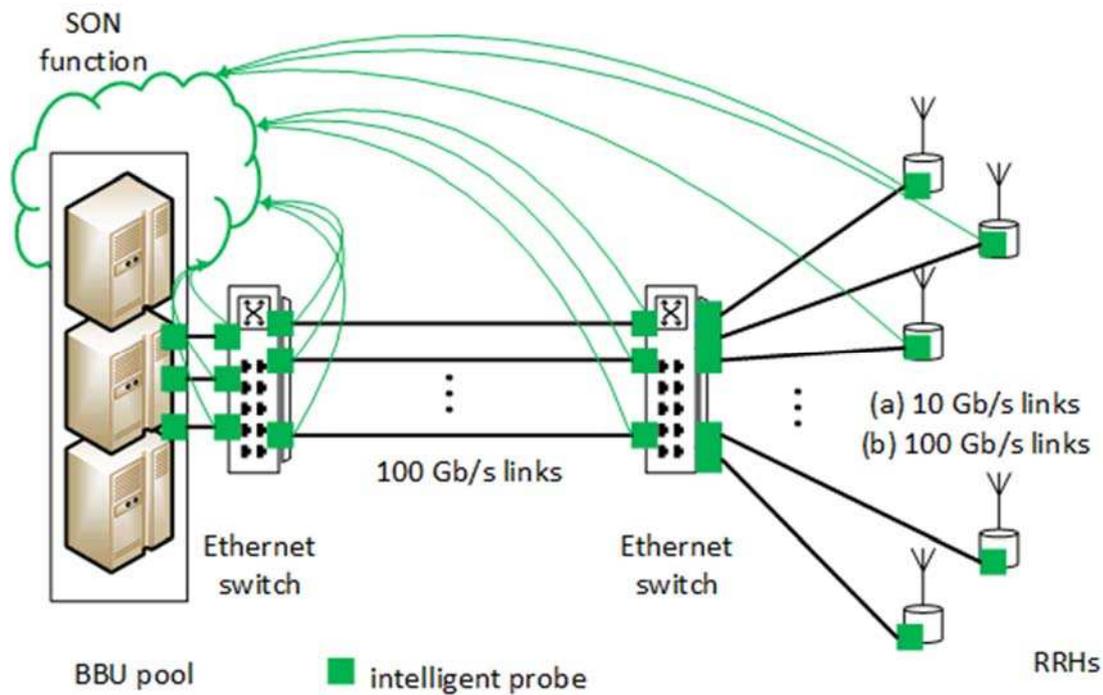


Figure 3 – BBU pool connected to RRHs through Ethernet switches. In option (a), lower-speed links are used beyond an aggregation point that may make use of statistical multiplexing gains. In option (b), the links have the same line rate, permitting the use of cut-through operation.

Path management through the fronthaul network brings us to the possibility of virtualisation and software-defined (and even software-optimised) networking. A network control function can set up the paths between the BBU pool and RRHs with knowledge that there will be no contention. As software-defined networking (SDN) is already being developed for IP/Ethernet networks, the use of Ethernet in the fronthaul lends itself to this far more easily

than using CPRI-based transport. Further, the use of hardware probes within an Ethernet fronthaul could aid the optimisation of performance (the use of the self-optimizing network (SON) function shown in Fig.3).

Having introduced the possibilities enabled by Ethernet networking in the fronthaul, the remainder of this paper is organised as follows. In Section 3, we discuss the network topologies and architectures possible for implementing very high-speed Ethernet access networks that can be suitable for the fronthaul and fixed-mobile convergence, and particularly examine bit-rate requirements as radio channel bandwidth and numbers of antennas increase.. In Section 4, we then examine the challenges faced by fronthaul distribution networks for 5G in terms of the latency, jitter and synchronisation requirements and how these might be helped through current standards activities. In Section 5, owing to the challenges faced by fronthaul radio waveform transport, we discuss the advantages brought by a new functional split in the fronthaul, through the movement of certain baseband functions into the RRH (now sometimes referred to more generically as a remote radio unit – RRU). Section 6 extends the discussion of the functional split to comparisons and compatibility with backhaul and mid-haul, and Section 7 presents the conclusions.

### 3. Fronthaul network architecture and topology

Based on Metro-Ethernet Forum specifications [16], Ethernet is increasingly being used in mobile backhaul, the links/network that connects base stations with the core, with Ethernet service layer functions [17] allowing support of Carrier Ethernet services [18]. The Next Generation Mobile Network (NGMN) Alliance makes an underlying assumption that the backhaul will utilize an all-packet (Ethernet/IP) architecture [19]. The NGMN Alliance further requires that future networks will enable end-to-end packet transport, using a shared, and hence reduced-cost, access- and service-agnostic transport network. This reasoning for the backhaul has initiated interest in including the fronthaul over such a transport network, as has been discussed.

Let us consider the fronthaul network shown in Fig.4. The number of fronthaul interfaces required depends on radio site configuration. Radio sites can be classified into macro cells and micro or small cells. Macro cells generally have three to six sectors. Additionally, for each sector, several RAT on different bands can be present e.g. 2G, 3G at 1800MHz and/or 2100MHz, LTE at 800 MHz and/or 2600MHz. Typical configurations in urban areas with 3 sectors for each RAT can then yield up to 18 radio heads per cell site. The total aggregated digitised radio traffic may then be 18 x 2.5 Gb/s, or around 45 Gb/s. This leads to the need for multiplexing (in time or wavelength) to reduce the number of required fibers required to the CO. For micro/small cells the antennas are usually omnidirectional, thus only one radio head for each RAT and carrier frequency may be required.

Wavelength division multiplexing (WDM), as shown in Fig. 4, is the most obvious way to save fiber resource: one wavelength channel pair can be associated to each BBU and RRH link. WDM can be implemented through pluggable/changeable transceiver modules in the RRH and BBU ports and/or by including WDM equipment which manages the wavelength assignment. The signals transported would already be time-division multiplexed radio samples, and compression could be used to save wavelength resource, further, as previously discussed, but at the expense of latency/jitter. Coarse WDM (CWDM) can be used to reduce costs, and a single fibre with bidirectional transmission could be employed to save fiber/wavelength resource further, but new standards are needed for this. Eventually,

for larger aggregate transport requirements, dense WDM (DWDM) may be required, but with colorless operation (wavelength defined at CO only) to make it compatible with outdoor operation.

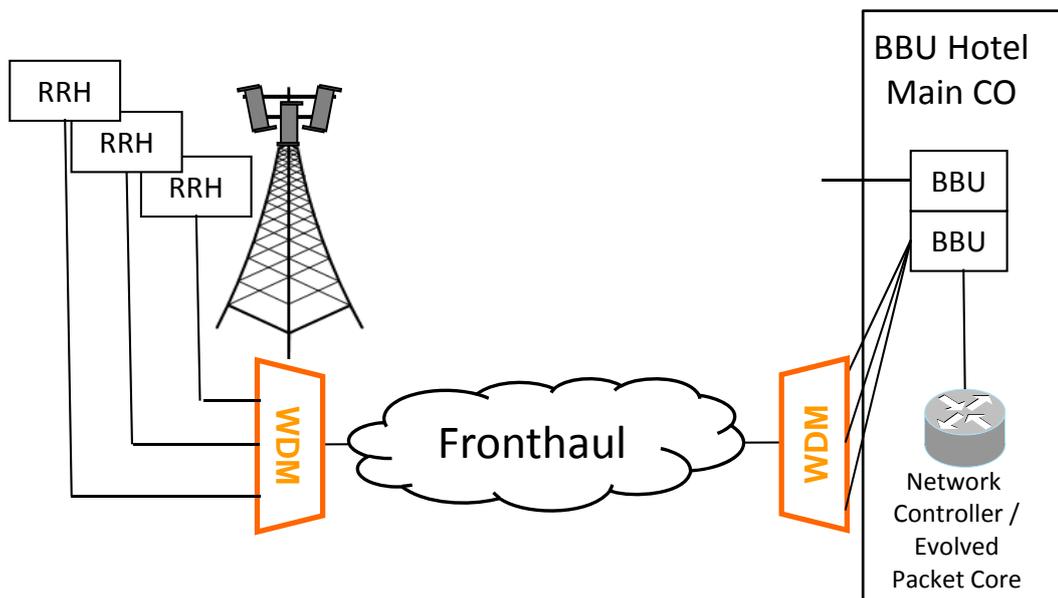


Figure 4 Fronthaul network based on WDM (Wavelength Division Multiplexing)

Returning to the bit-rate requirements, Table 1 summarises the current CPRI/ORI line rates (3 example rates are given) and projects the required bit-rates for 5G. Currently, CPRI/ORI oversample significantly. The projected rates assume Nyquist rates, and only a 10% overhead for framing and line coding. No compression is assumed in either set of rates (so, the requirements could be reduced to 50% of those indicated with currently specified techniques). It is evident that while current line rates of up to 10 Gb/s may support 20 MHz bandwidth LTE, or even up to 100 MHz bandwidth LTE-A systems, wider bandwidths and increased use of multi-antenna techniques, which require multiple radio streams per RRH, will necessitate moves to 100 Gb/s links, and beyond.

Table 1. Current CPRI/ORI interface specifications and projected bit-rates for RRH links

Current CPRI/ORI interfaces		Projected requirements	
Line rate	Example Use	Possible uses	Approx. line rate
614.4 Mb/s	10 MHz LTE channel, with 8B10B encoding	100 MHz, 8 antennas (sectors/MIMO/CoMP)	28 Gb/s
4.9152 Gb/s	8 x 10MHz (multiple antennas, 8B10B)	500 MHz, 8 antennas (sectors/MIMO/CoMP)	141 Gb/s
10.1376 Gb/s	10 x 20 MHz (multiple antennas, 64B66B)	500 MHz, 16x8 massive MIMO	2.25 Tb/s

#### 4. Latency, jitter and synchronisation requirements

Existing fronthaul interfaces have evolved from interfaces internal to the mobile base station equipment. Delay requirements were correspondingly relaxed. For example, CPRI requirement R-26 states that the maximum absolute round trip delay per link, excluding cable length, shall be 5  $\mu$ s [7]. However, as distances and network complexity increase such that this 5  $\mu$ s no longer dominates overall link latency, such specifications may need to be

revisited. Additionally, as the fronthaul standards were initially written with the assumption of symmetry, there are as yet no standardised limits on the allowed asymmetry in the delay. Various values have been suggested, up to the order of 125 ns, but this is an active area of investigation.

These round trip times include not just the round trip delay of the fronthaul optical network, but also any required processing time by the baseband unit, any intermediate nodes and the radio head. Also, while overall delay requirements may be specified for a given air interface, the split between the various elements is not usually defined. This leads to a range of values for the permissible fronthaul optical network latency, based on different radio vendors' equipment and application scenarios. Values for permissible Round Trip Time of the fronthaul range from 100  $\mu$ s [20] to up to 400  $\mu$ s, though values of around 250  $\mu$ s, equivalent to a passive optical network reach of 25 km, are often quoted [21]. Given that reduction in overall latency is a key objective for 5G, there are *suggestions* that a further tightening of requirements is likely.

Figure 5 shows results from analysis, simulation verified by measurements of prototype devices, and published data for a commercially available large port-count store-and-forward switch [22]. With all ports operating at 10 Gb/s, in the simple aggregator case, with no packet inspection, latencies of 100 ns can be achieved provided no contention occurs. Layer 2 cut through switching adds approximately 10 ns to the latency (total of 110 ns). These latencies are independent of packet size.

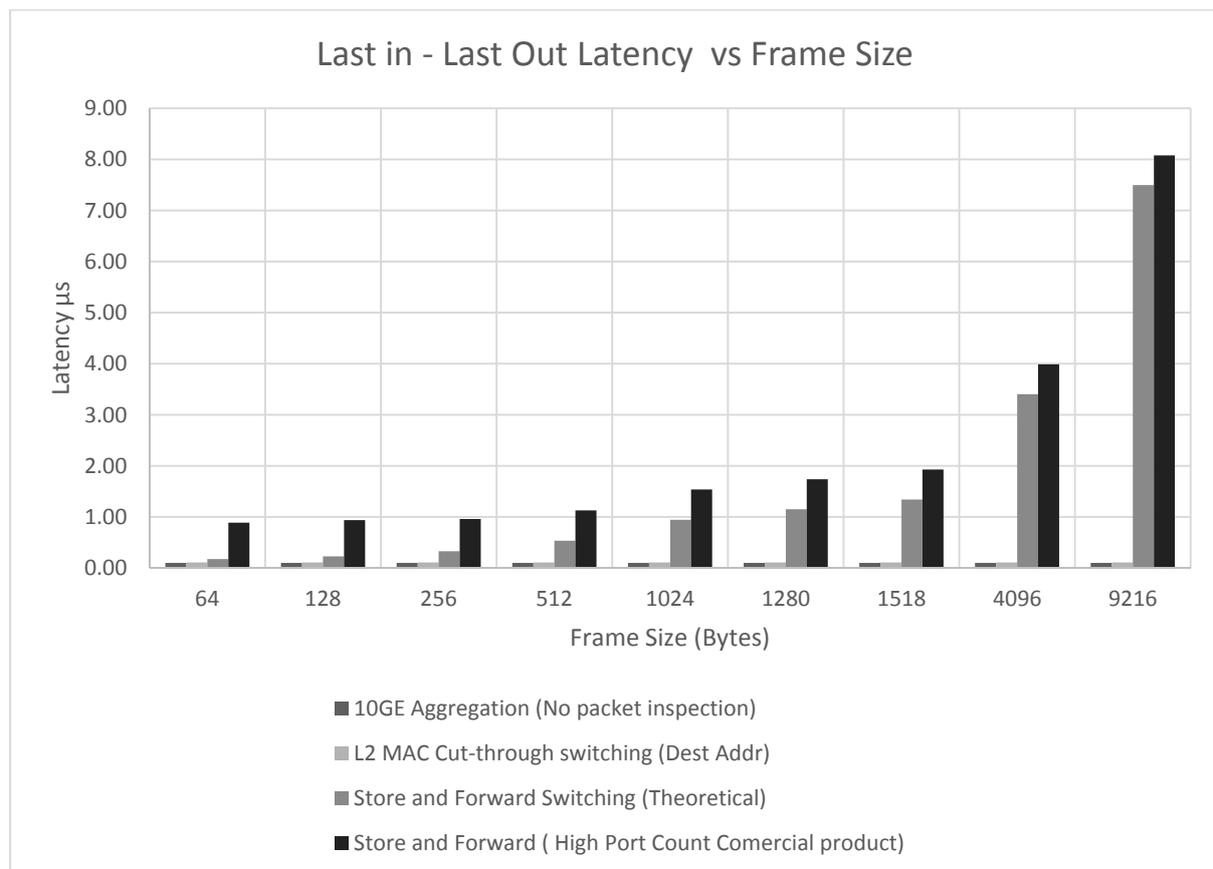


Figure 5 Last bit in to last bit out latency for cut-through and store-and-forward Ethernet switching (one-way, uncontended packets)

In contrast, store-and-forward switches impose an additional delay due to the need to store the entire frame, perform error checking and retrieve the frame for transmission. This leads to typical latencies which depend on frame size, ranging from  $<1 \mu\text{s}$  (64 byte frame) to  $>8 \mu\text{s}$  (9216 byte frame). Larger frames may be desirable in order to minimise the transmission bandwidth overhead, but as can be seen in the store-and-forward case, there may be a significant penalty in latency terms. The two-way latency would thus represent at least 16% of a  $100 \mu\text{s}$  round-trip time limit (assuming only a single switch in the fronthaul network). It should be noted that the datasheet latency of store-and-forward devices omits the delay due to the storage of the entire packet, thus quoting only any excess delay introduced due to internal processing, pipelining etc

As well as latencies (be they due to active equipment or passive network elements) being below a permitted maximum, they must also be known to ns level accuracy. For example CPRI R-19 specifies that the fronthaul downlink latency be known to an accuracy of  $\pm 8.138 \text{ ns}$ , and R-20 further specifies that the round trip be known to an accuracy of  $\pm 16.276 \text{ ns}$  [7]. There are various drivers for these requirements relating to the need for phase or time synchronisation on the air interfaces. CPRI / ORI do include procedures for making measurements of round trip delays, but then use the assumption of symmetry in the calculation of one way latencies.

It must be emphasised that the best practically possible latency performance is shown in Figure 5. Even in the cut-through case, queueing due to multiple packets being eligible for transmission, as may occur if attempting statistical multiplexing, will increase the latency of queued packets. This can lead to significant Packet Delay Variation (PDV) which can seriously degrade the performance of IEEE1588 as well as directly impacting operation of the radio data interface. This PDV is generally increased with larger frame/packet sizes.

Achieving similar levels of performance using Ethernet, supported by PTP, will be difficult without specific network design measures, both in terms of physical topology, devices and traffic management to limit the generation and mitigate the effects of PDV, and is an active area of study.

Latency is of course not the only time related performance metric. A complete fronthaul solution must provide synchronisation of frequency, phase and (increasingly) time between the baseband unit (or controller) and one or more remote radio heads.

Synchronisation performance can be quantified by a number of parameters including fractional frequency offset, jitter/wander and phase/time error. Each of these characteristics has a direct impact on the air-interface performance, but as the fronthaul is only one link in a chain of elements which contribute to the overall synchronisation path from Network Master Clock to Antenna, a path budget must be developed with each link allocated a maximum contribution for each parameter. However, the relationship between parameters can be complex. Jitter (short term variations in the timing of the significant instances of a repetitive signal from the ideal) of the fronthaul synchronisation signal may contribute to phase noise on the air interface, which in turn may lead to degradation of the transmitted modulated signal [23]. As a further example, jitter must also be taken into account when considering fractional frequency offset budgets. The relation depends on not only the magnitude of the jitter, but also the frequency of the jitter and the bandwidth of the clock recovery/synthesis system [7]. For example, E-UTRA wide area base stations must transmit each modulated RF carrier with a minimum frequency accuracy of  $\pm 50 \text{ ppb}$  [24], yet CPRI mandates that the contribution of a CPRI fronthaul link's jitter shall be no more than  $2 \text{ ppb}$  [7].

Time synchronisation has similar issues relating to the need to budget for the entire signal chain in order to derive fronthaul requirements from the overall system requirements. Thus, [24] specifies that in certain modes of operation the Time Alignment Error at the BS antenna ports shall not exceed 65ns. This requirement drives the CPRI requirements on latency accuracy mentioned previously as well as placing upper bounds on uncorrected asymmetry, though sources of this outside of the fronthaul optical network need to be better understood in order to derive a robust specification.

Synchronisation and latency performance to these levels in a general purpose packet switched network are at best challenging, and with 5G expected to have even tighter overall performance requirements, that challenge will become more difficult. However some projects have already demonstrated what can be achieved in a network specifically designed to optimise synchronisation performance.

The current state of the art for wide area synchronisation across an Ethernet network is arguably demonstrated by CERN's White Rabbit project [25]. White Rabbit (WR) networks use a combination of Synchronous Ethernet, IEEE1588 and "Digital Dual-Mixer Time Difference" (DDMTM) phase detection to provide sub-nanosecond accuracies to thousands of nodes, with fibre distances of 10's of km. Currently used predominantly in the scientific arena, optimum performance is only achieved after calibration of equipment and fibre, but many of the techniques employed could equally be deployed in fronthaul (and WR can interoperate with standard 1 Gb/s Ethernet and PTP implementations). Features from WR are likely to appear in new and revised synchronisation standards over the coming years.

## 5. New functional split ("fronthaul lite")

Previous sections have shown that, while there are important advantages for deployment, using Ethernet in the fronthaul faces significant implementation challenges. In Section 3, it was found that the bit-rate requirements and thus the transport costs become intolerably high, especially if we reserve the use of the wavelength domain for the distribution between different physical network elements only. Section 4 highlighted significant delay and delay variation challenges when introducing switching equipment into the fronthaul path, due to the contention with other traffic passing the same node. Moreover, when transporting digitized radio signals, there is no statistical multiplexing gain, which is one of the reasons for adopting switching equipment in the fronthaul: burst-like transmission should be exploited to reduce the overall data rates. Moreover, one cannot switch off or re-route a stream of frames carrying a continuous radio signal to/from a RRH.

### *1. SISO link*

Recently a new functional split between the BBU pool/CO and the RRHs, which no longer transports sampled waveforms as end or start results of the baseband processing at the transmitter and receiver, respectively, has been discussed [26], [27]. Rather, intermediate signals are serialized and then transmitted over Ethernet. This implies that more processing is moved from the BBU pool/CO to the RRH with the aim of reducing the bit-rate requirements and enabling statistical multiplexing gains with other traffic over the same Ethernet link.

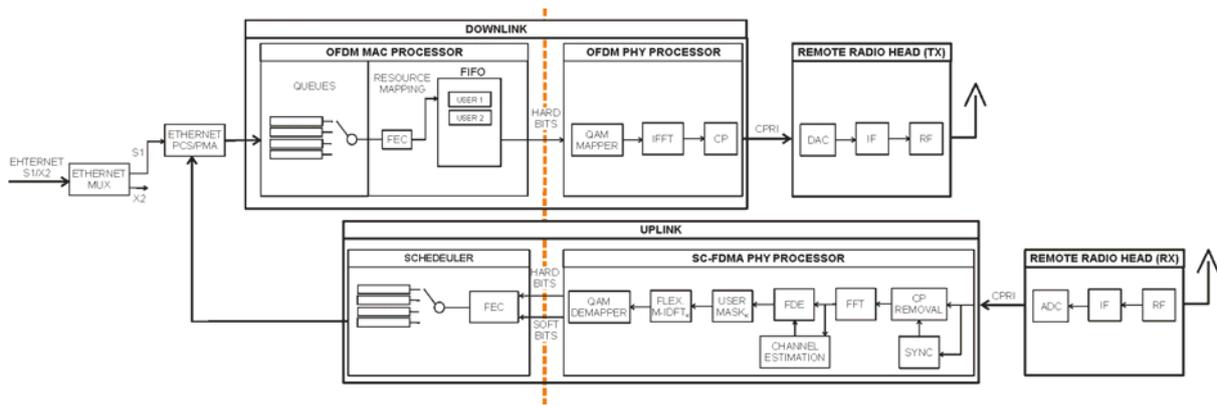


Fig. 6 Signal processing in an LTE base station.

The signal processing in a LTE base station is shown in Fig 6. The base station is connected to the network over two logical links, S1 and X2. While S1 is the feeder link to the advanced gateway (aGW), which is the demarcation point between the network operator and the Internet, X2 is relevant for information exchange between base stations, e.g. during a handover. S1 and X2 may share the same physical network connection.

In the downlink, S1 data are first classified and the header information is used to store packets into user queues. Data transmitted in one 1ms slot to one user is then packed into a so-called transport block. After FEC, they are mapped onto the assigned radio resources in the frequency domain. Even if the assigned resources are non-contiguously mapped in the frequency domain, all data for one user are jointly encoded yielding a single transport block. A so-called tuple is the set of bits yielding a single complex-valued constellation point transmitted on a single radio resource element (i.e. one sub-carrier in one OFDM symbol). Such constellations are fed block-wise into an adaptive modulator, which supports variable constellation alphabets, and the resulting block of quasi-analogue IQ signals is then fed into an inverse fast Fourier transform (IFFT). Finally, the cyclic prefix is added.

In the uplink chain, the cyclic prefix is removed and the channel is equalized in the frequency domain. The signal of a single user is then extracted from the overall received signal and passed through an inverse discrete Fourier transform, the block length of which is a multiple of 12 in LTE. The use of DFT spreading at the terminal reduces the peak-to-average power ratio of the waveform and it exploits the multipath diversity in the channel. When demapping the tuples from the resulting signal constellations, both hard and soft bits are derived in order to inform the decoder about the reliability of hard decisions. Soft-decision decoding yields better results, in particular in combination with bit-interleaved coded modulation (BICM). Afterwards, the data are fed into the user queues, from where they are packed into Ethernet packets and send to the aGW.

With respect to the functional split, today the sampled waveform is transmitted over the optical link in CPRI/ORI. However, if we shift for instance the whole physical layer processing to the RRH, just the transport blocks need to be transmitted between BBU/CO and RRH, which leads to significantly reduced data rates. Moreover, the traffic is proportional to the data load; if there is no data for some users, less transport capacity will be needed in the fronthaul. In this way, statistical multiplexing gains become useful.

## II Support for massive MIMO

One of the reasons to rethink the fronthaul design is the use of more and more antennas at the base stations, as is proposed for massive MIMO. By using CPRI, this would lead to a dramatic increase of the data rate, proportional to the number of antennas. However, these antennas are used to form beams partly carrying the same data towards an intended user. Therefore, it may be possible to also shift the beamforming operation to the RRH, as indicated in Fig. 7. There is no fundamental change of the interfaces and thus the same functional split as in the case of Fig.6 can be used. The scheduling information has to include additional information about the assigned beam, which is included in the precoding matrix indicator (PMI) in LTE. If multiple beams are used in parallel to transmit data to multiple users (multiuser MIMO), then the data rate between MAC and PHY processor will be increased, however this is very untypical for most network utilization that the maximum number of beams will be used.

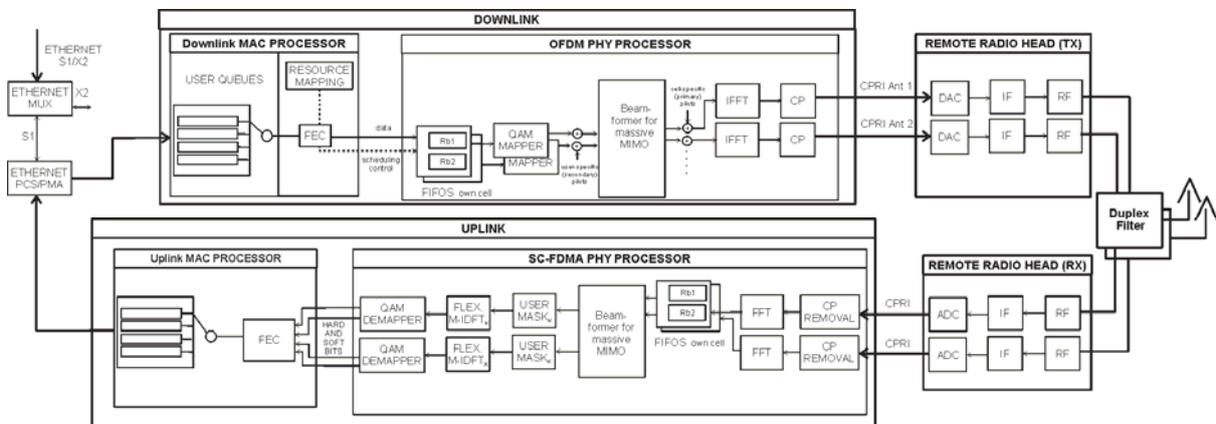


Fig. 7 The beamforming operation for massive MIMO can be shifted from the BBU/CO to the RRH in order to reduce the data rate in the fronthaul.

## III. Support for massive MIMO and CoMP

Massive MIMO and Co-ordinated Multipoint (CoMP) is the most complex case envisaged for future mobile radio systems. CoMP has been investigated in great detail in [28] to reduce the inter-cell interference and to realize more consistently high data rates at the cell edge. The cellular system is considered as a distributed MIMO link where, in the downlink, the base stations are the inputs and the terminals the outputs. Base stations transmit cell-specific pilot signals and terminals feedback their estimated channel state information (CSI). This information is exchanged among the base stations and the data likewise so that each base station becomes able to suppress the unwanted interference.

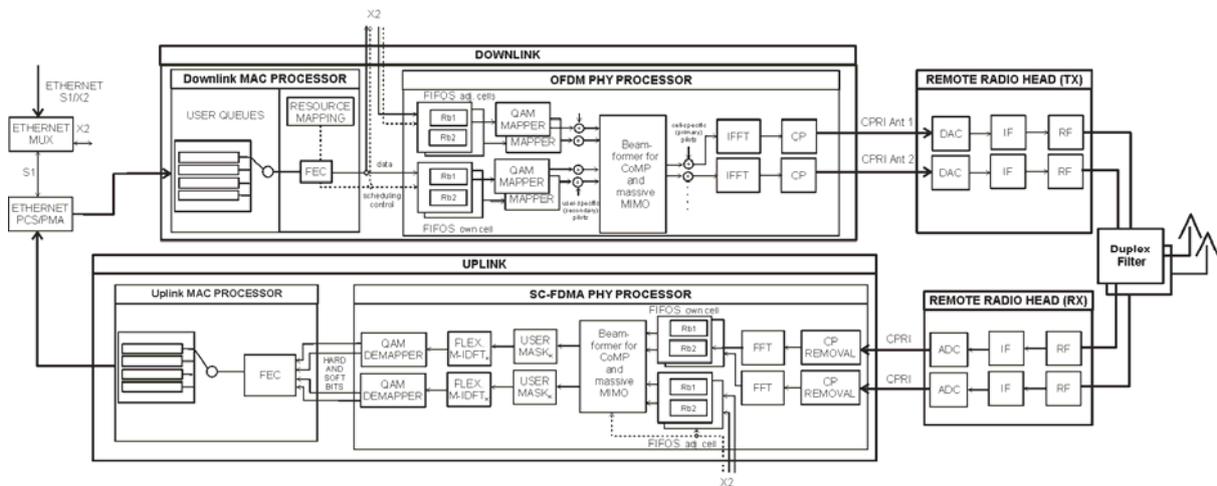


Fig. 8 For CoMP, too, the required beamforming operation can be shifted from the BBU/CO to the RRH in order to reduce the data rate in the fronthaul.

CoMP has been added to the system in Fig. 8. It can be observed that locally, the data from other base stations are made available to be processed jointly with the desired data in the downlink. In this way, the desired signal quality can be improved, essentially by subtracting the interference from other cells. In the uplink, received signals from other cells are also made available so that the interference can be reduced in the reciprocal manner.

However, we need to take into account that now, besides the light fronthaul signals, the exchange of data and CSI between cells needs to be supported, which is commonly a task of the X2 interface over which the base stations are interconnected. In the case of a fronthaul, the exchange of information is within a base station or BBU pool, so the X2 interface is internalised (if it exists). Functionally, however, the CSI information exchange is still required. Fortunately, it has been shown that this information is reduced if massive MIMO and CoMP are combined [2].

## 6. Discussion

We return to the convergence of the use of Ethernet in the backhaul and fronthaul in Fig. 9. For the transport of radio samples, we may need a mapper/demapper function between CPRI/ORI transport (or directly from radio samples) into Ethernet frames, as we have discussed. This can be compared to the backhaul transport using baseband signals in the lower part of the figure. However, the baseband signals which would be transported from the base station to a new digital baseband and RRH unit, could be transported in a similar manner to backhaul using Ethernet switches. Such a “fronthaul” is now termed X-haul as it can be used to transport any type of signal. These are the Ethernet links shown in orange in Fig. 9.

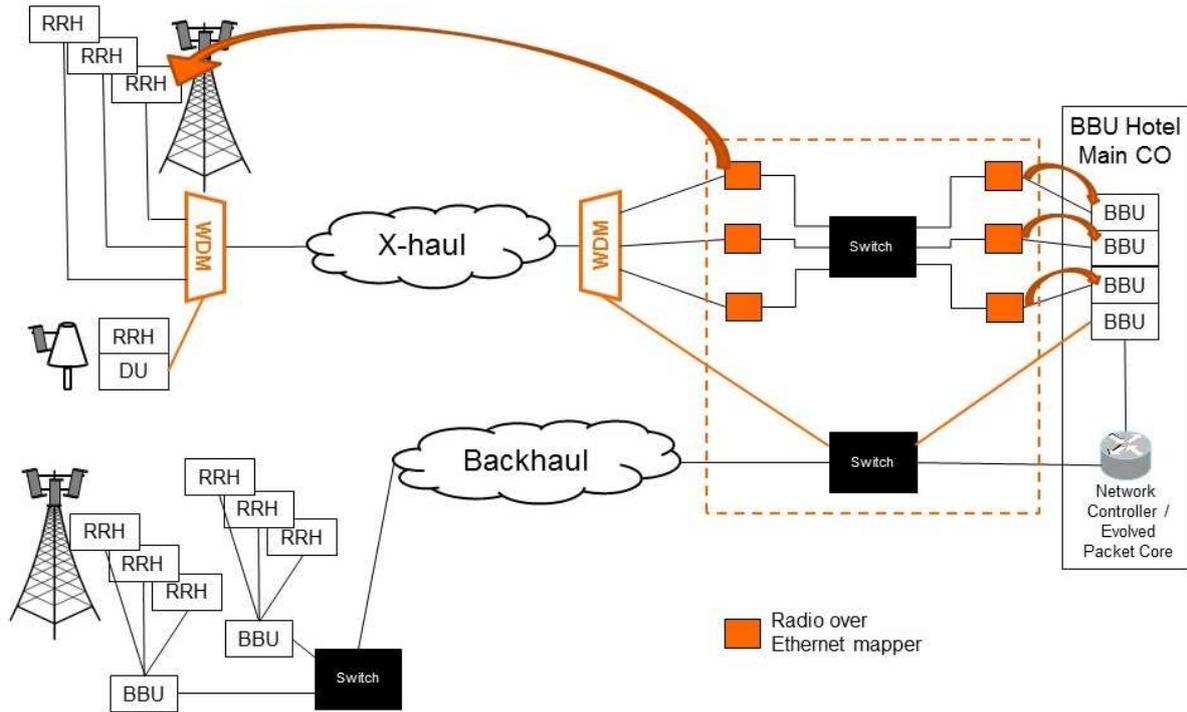


Figure 9: Global picture of Ethernet based network including back- and front-haul

Finally, we note that different terminology has been used to describe a functional split such as that proposed. The Metro-Ethernet Forum (MEF) define midhaul as the interconnection of small cells into the macrocell via Ethernet links, on the assumption that the small cell is covered by a packet-based BS. Although, there are functional similarities to the new proposed functional split, we concentrate on a fronthaul solution with lower bit-rate requirements (our “fronthaul lite”) and multiplexing possibilities, which may not require layer 3 networking for the data transport.

## 7. Conclusion

A new functional split between the BBU and RRH to permit baseband signal transport instead of the transport of sampled radio streams has been proposed to enable a realisable fronthaul. Such a fronthaul can also make use of Ethernet switches and networking statistical multiplexing gains, as it transports relatively bursty data instead of continuous radio waveforms. It can make use of software-defined (and optimized) networking and intelligent probe monitoring developed for Ethernet networks, and can better allow convergence with Ethernet-based fixed access networks. Delay and loss of synchronization remain challenges for a frame-/packet-based fronthaul, and these will be a focus of future work in the iCIRRUS project.

## Acknowledgments

This work has received funding from the European Union’s Horizon 2020 Research and Innovation programme under grant agreement no. 644526 (iCIRRUS).

## References

- [1] NGMN 5G White Paper, NGMN Alliance, February 2015
- [2] V. Jungnickel, K. Manolakis, W. Zirwas, B. Panzner, V. Braun, M. Lossow, M. Sternad, R. Apelfröjd, T. Svensson, "The Role of Small Cells, Coordinated Multipoint, and Massive MIMO in 5G", IEEE Comm. Magazine, p. 44-51, May 2014
- [3] N.J. Gomes, P. Monteiro, A. Gameiro, Next-generation wireless communications using radio over fiber, Wiley, 2012
- [4] C.-L. I, J. Huang, R. Duan, C. Cui, X. Jiang, L. Li, "Recent progress on C-RAN centralization and cloudification", IEEE Access, vol. 2, pp. 1030-1039, September 2014
- [5] T. Pfeiffer, "Next Generation Mobile Fronthaul Architectures", Optical Fiber Communications Conference, OFC 2015, M2J.7, Los Angeles, March 2015
- [6] S.-H. Park, O. Simeone, O. Sahin, S. Shamai, "Fronthaul compression for cloud radio access networks", IEEE Signal Processing Magazine, pp. 69- 79, November 2014
- [7] Ericsson AB, Huawei Technologies Co.Ltd, NEC Corporation, Nokia Networks, Alcatel-Lucent, "CPRI Specification V6.1(2014-07-01): Common Public Radio Interface (CPRI); Interface Specification," <http://www.cpri.info>, 2014.
- [8] Open Base Station Architecture Initiative, "Reference Point 3 Specification V4.2" Available: <http://www.obsai.com/specifications.htm>
- [9] ETSI, "ETSI ORI (Open Radio Interface)," [Online]. Available: <https://portal.etsi.org/tb.aspx?tbid=738&SubTb=738>
- [10] M Prins, R Malhotra, "Ethernet operation administration and maintenance - Opportunities for the NREN community," Terena Networking Conference, May 2011
- [11] ITU-T , "Recommendation ITU-T G.8261/Y.1361: Timing and synchronization aspects in packet networks". Available: <https://www.itu.int/rec/T-REC-G.8261/en>
- [12] ITU-T, "G.8262/Y.1362: Timing characteristics of a synchronous Ethernet equipment slave clock" Available: <https://www.itu.int/rec/T-REC-G.8262/en>
- [13] ITU-T, "G.8264/Y.1364: Distribution of timing information through packet networks" Available: <https://www.itu.int/rec/T-REC-G.8264/en>
- [14] IEEE, "IEEE Std1588-2008: IEEE Standard for a Precision Clock Synchronization Protocol for Networked Measurement and Control Systems," IEEE, 2008
- [15] ITU-T, "G.8271.1/Y.1366.1 Network limits for time synchronization in packet networks," Available: <https://www.itu.int/rec/T-REC-G.8271/en>
- [16] Implementation agreement MEF 22.1, Mobile Backhaul Phase 2, MEF, January 2012
- [17] Metro-Ethernet Network Architecture Framework Part 2: Ethernet Services Layer, MEF 12.1
- [18] Ethernet Services Definition – Phase 2, MEF 6.1
- [19] NGMN Alliance, "NGMN Optimized Backhaul Requirements", August 2008
- [20] Irvine (Ericson AB) "Mobile Networks: Transport Impacts" September 2014. Available: <http://www.ieee802.org/1/files/public/docs2014/new-irvine-mobile-networks-fronthaul-0914.pdf>
- [21] H J Son, S M Shin "Fronthaul Size: Calculation of maximum distance between RRH and BBU" 1 April 2014. Available: [http://www.netmanias.com/en/?m=view&id=blog\\_mbl&no=6276](http://www.netmanias.com/en/?m=view&id=blog_mbl&no=6276)

- [22] Y Yang, (Cisco Systems Inc) "Understanding Switch Latency" June 2012 Available: [http://www.cisco.com/c/en/us/products/collateral/switches/nexus-3000-series-switches/white\\_paper\\_c11-661939.pdf](http://www.cisco.com/c/en/us/products/collateral/switches/nexus-3000-series-switches/white_paper_c11-661939.pdf)
- [23] M.R. Khanzadi, D. Kuylenstierna, A. Panahi, T. Eriksson, and H. Zirath, "Calculation of the performance of communication systems from measured oscillator phase noise," *IEEE Trans. Circuits Syst. I, Reg. Papers*, vol. 61, no. 5, pp. 1553–1565, May 2014.
- [24] ETSI 3GPP, "3GPP TS36.104 Evolved Universal Terrestrial Radio Access (E-UTRA); Base Station (BS) radio transmission and reception" Available: <http://www.3gpp.org/DynaReport/36104.htm>
- [25] CERN Open Hardware Repository: White Rabbit. [www.ohw.org/projects/white-rabbit](http://www.ohw.org/projects/white-rabbit)
- [26] U. Doetsch, 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 Technical J.*, vol. 18 (1), pp. 105-128, June 2013
- [27] 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 Communicatinos Magazine*, Vol. 52, issue: 5, pp. 68-76, May 2014
- [28] R. Irmer, H. Droste, P. Marsch, M. Grieger, G. Fettweis, S. Brueck, H.-P. Mayer, L. Thiele, V. Jungnickel, "Coordinated multipoint: Concepts, performance, and field trial results", *IEEE Communications Magazine*, Vol. 49 (2), pp. 102-111, Feb. 2011

Affiliation a: Communications Research Group, University of Kent, Canterbury, Kent, UK

Affiliation b: Orange Labs Network Team, Orange S.A., Lannion, France

Affiliation c: Advanced Technology Team, ADVA Optical Networking Ltd., York, UK

Affiliation d: Photonic Network and Systems Department, Heinrich-Hertz-Institut, Berlin, Germany