Dimensioning Flex Ethernet Groups for the transport of 5G NR fronthaul traffic in C-RAN scenarios

José Alberto Hernández*, Gabriel Otero*, David Larrabeiti* and Óscar González de Dios†
*Universidad Carlos III de Madrid, Spain
Email: {jahgutie, gaoterop, dlarra}@it.uc3m.es
† Telefonica I+D, Madrid, Spain
Email: oscar.gonzalezdedios@telefonica.com

***Abstract*—This article overviews the new Flex Ethernet implementation agreement standardised by the Optical Internetworking Forum and its applicability for the transport of Massive MIMO 5G New Radio fronthaul traffic with deterministic delay. Thanks to the proper dimensioning of FlexE calendar slots to each individual fronthaul traffic flow, bounded delays below the $100\ \mu s$ requirement imposed by the IEEE 802.1CM can be achieved.**

***Index Terms*—Flex-Ethernet; Fronthaul; 5G; IEEE 802.11CM.**

I. INTRODUCTION

In Cloud Radio Access Networks (RAN) scenarios, traditional Radio Base Stations are split into lightweight Remote Radio Units (RRU), spread across the cities and Baseband Units (BBU), centralised and possibly virtualised at the operator's premises. In most cases (depending on the functional split) most of the radio processing is centralised at the BBUs. To support this, a high-speed ultra-low latency fronthaul network must be provided connecting both elements. Recently, the 3GPP has released the 5G New Radio (NR) specification, allowing radio channels of up to 400 MHz and Massive MIMO (up to 64 TX/RX) [1]. These require an even more powerful fronthaul network than that of 4G LTE channels. Indeed, 5G NR radio signals, after sampled and digitised, pose important bandwidth and delay challenges (tens/hundreds of Gb/s with delays below a few hundred microsecs) to their transport across an optical DWDM network.

On the other hand, Flex Ethernet (FlexE) has been recently proposed by the Optical Internetworking Forum (OIF) as a mechanism to decouple MAC and PHY layers of Ethernet clients. It features bonding, sub-rating, and channelisation of 1 to m 100GBase-R PHYs (200G and 400G in the future) [4]. FlexE can be used in *Router to Transport connection* scenarios where the mapping/de-mapping FlexE Shim layer allows to flexibly partition and assign bandwidth groups of 5 Gb/s slots to individual flows [5].

In the context of 5G, the ability to provision dedicated data paths with guaranteed bandwidth and deterministic delay to individual 5G NR fronthaul flows makes FlexE suitable for multiplexing and transporting fronthaul traffic [6]. This article overviews the bandwidth and delay requirements of multiple configurations of 5G NR channels, along with their fronthaul traffic profiles employing Intra-PHY functional split (Split I_U in eCPRI notation [2]); and further shows how to configure FlexE calendars to meet their delay requirements.

II. CHARACTERIZATION OF 5G NR FRONTHAUL TRAFFIC

NR is based on OFDM very much like LTE but allowing a flexible numerology with subcarrier spacings ranging from 15 KHz up to 240 KHz [1]. Table I shows the main parameters for different radio bandwidth channels, and the subsequent bitrates for Functional Split I_U of the eCPRI recommendation. The table indicates the radio bandwidth B_{Radio}, subcarrier bandwidth (Δf), number of subcarriers ($N_{sc} = \frac{0.9 B_{Radio}}{\Delta f}$ since 10% of the subcarriers are used as guard band), frame slot size (T_{slot}) and subsequent OFDM symbol duration ($T_{OFDM} = \frac{T_{slot}}{14}$). As noted, different radio bandwidths have different values of Δf but never exceeding $3,000$ subcarriers in any configuration. The slot duration also varies depending on the radio bandwidth parameters.

With these parameters, the resulting fronthaul bitrate per Antenna-Carrier (i.e. no MIMO assumed) follows [2] for 100% cell load:

$$R_{I_U} = \frac{(2M)N_{sc}}{T_{OFDM}} \quad (b/s) \tag{1}$$

where $2M$ denotes the number of bits per I and Q sample (typically $2M = 30$ bits). The resulting OFDM symbol size in bits arises as:

$$S_{OFDM} = R_{I_U} T_{OFDM} \quad bits. \tag{2}$$

These values of R_{I_U} and S_{OFDM} are included in the table for all NR configurations. For the sake of completeness, CPRI-like split bitrates are also included in the table (Functional Split E, row R_E). As shown, split I_U requires about one half bitrate than that of CPRI split E.

Finally, it is worth remarking that such bitrate values does not include the use of any MIMO (i.e. they are 1x1). Introducing Massive MIMO is well known to introduce important spectral efficiency gains, namely 8x2 MIMO provides 95% extra gain with respect to 2x2, while 16x2 and 64x2 MIMO

978-3-903176-33-1 © 2021 IFIP

B_{Radio} Δf	20 MHz 15 KHz	50 MHz 15 KHz	 30 KHz	100 MHz 30 KHz	 60 KHz	200 MHz 60 KHz	 120 KHz	400 MHz 120 KHz	 240 KHz
N_{sc}	1,200	3,000	1,500	3,000	1,500	3,000	1,500	3,000	1,500
T_{slot} (ms)	1	1	0.5	0.5	0.25	0.25	0.125	0.125	0.0625
T_{OFDM} (μs)	66.67	66.67	33.33	33.33	16.67	16.67	8.33	8.33	4.17
R_{I_U} (Gb/s)	0.54	1.35	1.35	2.7	2.7	5.4	5.4	10.8	10.8
S_{OFDM} (Bytes)	4,500	11,250	4,500	22,500	11,250	45,000	22,500	90,000	45,000
R_E (Gb/s)	0.9216	2.304		4.608		9.216		18.432	

TABLE I
5G NEW RADIO BIT RATES AND SPLIT IU TRAFFIC PROFILES

increases the efficiency to 192% and 199% respectively, as shown in [3]. However, the latter requires to multiply the R_{I_U} bit rates by 16x and 64x respectively.

III. Overview of Flex Ethernet and mapping example of 5G NR clients into FlexE

According to [4], "the Flex Ethernet (i.e. FlexE) implementation agreement provides a mechanism for supporting a variety of Ethernet MAC rates that may or may not correspond to any existing Ethernet PHY rate". Important features of FlexE includes the use of bonding (when the PHYs are smaller than the client rates) and sub-rate and channelisation (for the opposite, when the PHY is larger than the client rates and sub-rating with TDM-based separation of MAC clients is required). The implementation specifies how to partition one or multiple 100GBase-R PHYs into 20 TDM flows with 5 Gb/s resolution using a calendar distribution mechanism with a time-granularity of 66B blocks. Supported Ethernet client rates can operate at 10 Gb/s or 40 Gb/ or m x 25 Gb/s.

As an example, consider an Ethernet link comprising 3 x 100 GBase-R PHYs arranged as a FlexE Group (300 Gb/s total) configured to carry one 64x2 MIMO 100 MHz 5G NR channel (i.e. fronthaul bandwidth 64 x 2.7 Gb/s = 172.8 Gb/s) along with three 8x2 MIMO 20 MHz NR channels (i.e. fronthaul bandwidth 8 x 0.54 Gb/s = 4.32 Gb/s each), as shown in Fig. 1 (a). The first FlexE client requires bonding and channelisation using 175 Gb/s of the total capacity (i.e. the next multiple of 25 after 172.8 Gb/s) while the three 20 MHz channels require sub-rating using 10 Gb/s each flow. In total, the minimum capacity used is 205 Gb/s (as $175+3\times10$).

Fig. 1 (b) shows a possible calendar distribution to carry all 5G NR flows (dark blue for the 175G fronthaul flow and light-blue, pink and orange for the three 10G flows). As shown, the three PHYs are used in parallel (i.e. bonding) to carry both the large and small FH flows. The large flow requires 35 calendar slots (i.e. $\frac{175G}{5G}=35$), each one carries 64 bits (8B) of payload per 12.8 ns (66B/64B). Each small flow requires only 2 calendar slots. According to Table I, the large flow generates OFDM symbols at a constant rate of $11,250$ Bytes every $T_{OFDM}=16.67\ \mu s$, and having 64x2 MIMO into account: $64\times 11,250=720,000\ B$ every T_{OFDM}. Thus, the number of 12.8 ns periods ($N_{cals,FH1}$) required to transmit these data is:

$$N_{cals,FH1}=\frac{64\times 11,250}{35\times 8}=2,571 \text{ periods of } 12.8\ ns \tag{3}$$

which implies a delay per OFDM symbol of:

$$D_{FH1}=2,571\times 12.8\ ns=32.91\ \mu s \tag{4}$$

for the large fronthaul flow FH1.

Clearly, the numbers match to computing this delay as if the whole OFDM symbol of $11,250\ B$ was transported over a dedicated channel operating at $175\ Gb/s$:

$$\frac{64\times 11,250\times 8\ b}{175\cdot 10^9\ b/s}=32,91\ \mu s \tag{5}$$

Similarly, the number of calendars for the small flows, where the OFDM symbol size is $4,500\ B$, are:

$$N_{cals,FH2}=\frac{8\times 4,500}{2\times 8}=2,250 \text{ periods of } 12.8\ ns \tag{6}$$

which implies a fronthaul delay per OFDM symbol of:

$$D_{FH2}=2,250\times 12.8\ ns=28.8\ \mu s \tag{7}$$

which again gives the same number as if we had a dedicated 10 Gb/s channel for each of the three small fronthaul flows:

$$\frac{8\times 4,500\times 8}{10\cdot 10^9}=28.8\mu s \tag{8}$$

Interestingly, delays are comparable in both cases (large and small FH flows) and in the range of several tens of microsecs. Essentially, we observe that some of the calendar slots are unused (therefore wasted) and only 210 Gb/s are being used for the transport of such flows, the remaining 90 Gb/s are empty. Thus, the benefits of statistical multiplexing are not leveraged.

In addition to this, it is worth noticing that, because bandwidth is partitioned and reserved for each individual fronthaul flow, traffic does not coexist on the same link so no queueing delays appear. Each OFDM symbol is allocated in its reserved slot following the calendar structure of Fig. 1, hence delay is deterministic.

The total latency budget for packet switching transport of fronthaul traffic is estimated in 100 microseconds as specified in IEEE 802.1CM, see [8]. In this example FlexE features a deterministic latency of 39.91 μs and 28.8 μs for the first large FH flow and the three small FH flows respectively. This gives a large margin for propagation time and the system designer does not have to budget some additional latency for jitter compensation, as it is the case of packet switching with variable packet delay and jitter.

Using the smallest $m\times 25\ Gb/s$ FlexE configuration for the transport of 5G NR flows above, the resulting transport delays

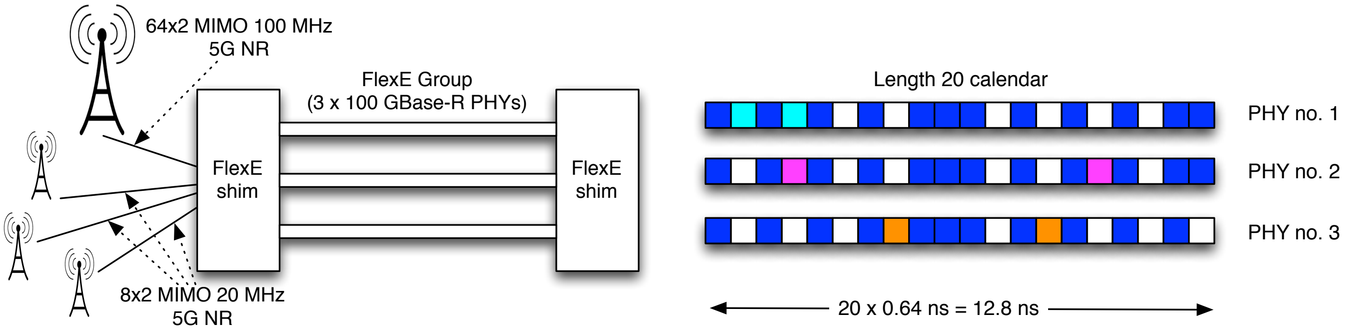

(a) Scenario example

(b) Calendar distribution.

Fig. 1. Use of FlexE for the transport of four 5G NR flows.

	20 MHz	50 MHz	100 MHz	200 MHz	400 MHz
1x	1.44μs @ 25G	3.6μs @ 25G	7.2μs @ 25G	14.4μs @ 25G	28.8μs @ 25G
2x	2.88μs @ 25G	7.2μs @ 25G	14.4μs @ 25G	28.8μs @ 25G	57.6μs @ 25G
4x	5.76μs @ 25G	14.4μs @ 25G	28.8μs @ 25G	57.6μs @ 25G	57.6μs @ 2x25G
8x	11.52μs @ 25G	28.8μs @ 25G	57.6μs @ 25G	57.6μs @ 2x25G	57.6μs @ 4x25G
16x	23.04μs @ 25G	57.6μs @ 25G	57.6μs @ 2x25G	57.6μs @ 4x25G	65.8μs @ 7x25G
32x	46.08μs @ 25G	57.6μs @ 2x25G	57.6μs @ 4x25G	65.8μs @ 7x25G	65.8μs @ 14x25G
64x	46.08μs @ 2x25G	57.6μs @ 4x25G	65.8μs @ 7x25G	65.8μs @ 14x25G	65.8μs @ 28x25G
128x	61.44μs @ 3x25G	57.6μs @ 7x25G	65.8μs @ 14x25G	65.8μs @ 28x25G	65.8μs @ 56x25G

TABLE II

DELAY USING THE MINIMUM $m \times 25\ Gb/s$ FLEXE BITRATE CONFIGURATIONS FOR EACH 5G NR FRONTHAUL FLOW

are those of Table II. The number m of 25G bandwidth slots are computed as:

$$m = \left\lceil \frac{N_{MIMO} \times R_{I_U}}{25\ Gb/s} \right\rceil \tag{9}$$

and the delay is obtained as:

$$D_{FH} = \frac{N_{MIMO} \times S_{OFDM} \times 8\ bit}{\left\lceil \frac{N_{MIMO} \times R_{I_U}}{25\ Gb/s} \right\rceil \times 25 \cdot 10^9\ b/s} \tag{10}$$

As shown, a large majority of cases require 100 Gb/s or less and, in all cases, assuming sufficient capacity is provided, the delay experienced by the OFDM symbols is always below 100 μs as required by IEEE 802.11CM.

IV. CONCLUSIONS

This article has overviewed the new 5G New Radio (NR) and its traffic profile and bandwidth requirements for a C-RAN scenario with Intra-PHY functional split. The standard approach to transport this traffic from RRHs to BBU is the use of regular Ethernet and allocating fronthaul traffic a high priority [9]. This approach suffers from indeterministic packet delays which lead to the set up of conservative jitter buffers to reassemble OFDM symbols. In this paper we propose the use of FlexE as an alternative to regular ethernet to make a simpler network design for fronthaul transport where the topology allows this setting (star). To this end, we have overviewed OIF's Flex Ethernet 2.0 implementation agreement, showing its benefits and suitability for the transport of multiple such fronthaul flows on separate virtual FlexE channels thanks to its bonding, sub-rating and channelization features. We validated the viability of FleXE for the target application through an example where 100G link bonding was employed. The number of calendar slots in each 100G Ethernet PHY required per 5G NR configuration was computed and the total OFDM transmission delay was estimated. The ability to provision dedicated data paths with guaranteed bandwidth and deterministic delay to separated flows makes FlexE very suitable for the transport of 5G NR fronthaul flows.

ACKNOWLEDGEMENTS

The authors would like to acknowledge the support of the EU-funded Metrohaul project (grant no. 761727) and the Spanish projects ACHILLES (PID2019-104207RB-I00) and Go2Edge (RED2018-102585-T).

REFERENCES

[1] S. Parkvall et al: "NR: The new 5G radio access technology", IEEE Comms. Standards Mag., pp. 24–30, 2017.
[2] G. Otero Perez et al: "Fronthaul Network Modeling and Dimensioning Meeting Ultra-low Latency Requirements for 5G", J. Opt. Comm. Netw. 10 (6), 573-581, 2018.
[3] A. Ghosh: "Massive MIMO for the New Radio - Overview and Performance", IEEE 5G Summit, June 2017.
[4] Flex Ethernet 2.0 Implementation Agreement, Optical Internetworking Forum, June 2018.
[5] A. Eira, et al "On the Efficiency of Flexible Ethernet Client Architectures in Optical Transport Networks" J. Opt. Comm. Netw. 10(1) A133–143, 2018
[6] R. Vilalta et al: "Network Slicing Using Dynamic Flex Ethernet over Transport Networks" European Conf. Optical Comms. 2017
[7] Z. He, X. Wang "Using flexible Ethernet to build a flexible network" Int. Conf. Advanced Communication Technology, 2014
[8] G. Otero Pérez, D. Larrabeiti López and J. A. Hernández, "5G New Radio Fronthaul Network Design for eCPRI-IEEE 802.1CM and Extreme Latency Percentiles," in IEEE Access, vol. 7, pp. 82218-82230, 2019
[9] A. Gowda, et al, "Delay analysis of mixed fronthaul and backhaul traffic under strict priority queueing discipline in a 5G packet transport network", Trans. Emerging Telecom. Tech. 28 (6), 1-9, 2017.