A Reconfigurable Architecture for Packet Based 5G Transport Networks

Raghu M. Rao
Xilinx Inc.
2100, Logic Drive, San Jose, CA, USA
raghu.rao@xilinx.com

Mickaël Fontaine, Raimena Veisllari
TransPacket
Hoffsvein 21-23, Oslo, Norway
{mickael.fontaine; raimena.veisllari}@transpacket.com

Abstract— **The 5G system architecture envisions a converged backhaul and fronthaul network which necessitates the fronthaul links to move to an Ethernet packet based network with preemption capabilities to ensure bounded delays. Furthermore, the introduction of the PDCP aggregation node as part of the Telco Cloud introduces a lower priority midhaul link but still requiring bounded delays. This paper presents a realistic Ethernet packet based 5G transport network with fronthaul express traffic and lower priority, preemptable, midhaul/backhaul traffic in a multi-hop network connecting the remote radio units and the baseband unit to the Telco Cloud and the packet core. The system includes an intelligent scheduler that aggregates and schedules midhaul/backhaul traffic by exploiting the inter-packet gaps without impacting the delay or PDV of the fronthaul streams. Analytical and experimental results are presented to demonstrate the effectiveness of the scheduling algorithm.**

Index Terms—**5G, CRAN, RoE, NGFI, Fronthaul, Midhaul, Backhaul, Xhaul, Time Sensitive Networking, eCPRI, CPRI, Scheduling, Preemption.**

I. INTRODUCTION

As 5G wireless technology gains traction, three use cases have emerged as the key to defining the air interface and the transport network. The enhanced mobile broadband (EMBB), the first of these use cases, drives the throughput at the user equipment (UE) to 10x compared to 4G. The massive machine type communication (MMTC) or massive IoT brings millions of connected devices into the network. While each of these devices could transmit and receive small bursts of data, the high total number of devices communicating small packets of data has the potential to overwhelm the transport network. The third use case is the ultra-reliable low latency communication (URLLC), which is being defined to enable applications such as remote surgery and self-driving cars that require extreme reliability and low latency communications. Given the vastly different traffic patterns that these three use cases generate, a revolutionary change in the radio access network is required, all the way from the air interface to the packet core and the transport network.

The diverse nature of 5G traffic requires a dynamic and flexible network, which the existing network is incapable of providing. In particular, the fronthaul in a 4G network relies on hardwired point-to-point connections in a time division multiplexed system using the CPRI [1] protocol. While CPRI based fronthaul is capable of providing varying transmission rates, it is not dynamically configurable and is not efficient under varying traffic conditions. This has prompted the move to packet-based fronthaul, based on Ethernet technology. This move also offers the possibility to support fronthaul, midhaul and backhaul over the same Ethernet transport infrastructure.

However, Ethernet is a best effort technology and requires Time Sensitive Networking (TSN) [2] to provide a flexible and dynamic transport network with low and bounded delays. In addition, deployments of Ethernet Fronthaul require very accurate timing to serve advanced radio functions and also require encapsulation of the radio signal in an Ethernet frame. There are two competing standards for encapsulating Fronthaul traffic in an Ethernet frame; 1) eCPRI from the CPRI consortium [3] and 2) Radio over Ethernet (RoE) or NGFI from the IEEE 1914 Working Group [4]. While the encapsulations themselves are different, the principle is similar in both cases.

The EMBB use case relies on technologies such as Massive MIMO that greatly improve network capacity by enabling significant frequency reuse and dual/multi connectivity that uses the WiFi spectrum to provide additional links to the user equipment. This comes at the price of increased fronthaul capacity. To contain fronthaul capacity, the baseband signal chain is being partitioned, and a portion of the layer 1 is moving to the remote radio unit (RRU). In addition, other functional splits in layer 2 are also being considered. One important functional split called Option 2 splits the upper L2 with the PDCP layer into an aggregation node to enable dual/multi connectivity. This PDCP aggregation node is also an important part of the multi-access edge computing node and the now emerging Telco Cloud which can provide much quicker hand-offs and also much reduced end-to-end latency. The link from the Telco Cloud to the base station is called midhaul. All this leads to prioritized traffic in the cellular transport network with fronthaul being high priority, a midhaul with lower priority but still requiring bounded delays, and backhaul which connects the Telco Cloud to the packet core with even more relaxed delay bounds. Fronthaul traffic, due to its strict timing constraints, cannot be preempted [2], while midhaul and backhaul are preemptable.

This paper describes a complete Ethernet fronthaul ecosystem providing all the required building blocks to implement the transport network for new 5G use cases. The paper further presents an Ethernet transport network incorporating fronthaul express traffic and lower priority (preemptable) midhaul/backhaul traffic in a multi-hop network

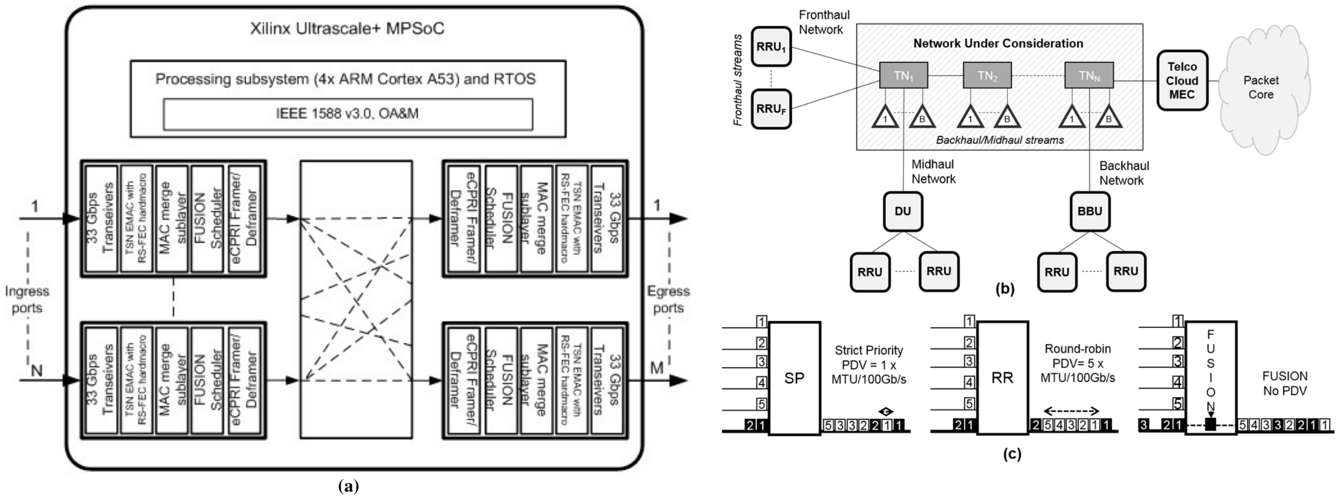

Fig. 1. (a) Block diagram of a multi-port eCPRI switch with TSN capabilities. (b) Example of a seamlessly converged fronthaul, midhaul and backhaul network; TN: transport node. (c) Packet scheduling mechanisms under consideration, from left to right: strict priority, round-robin and FUSION. Fronthaul packets are marked in black, and lower priority midhaul/backhaul packets in white.

connecting the remote radio units (RRU), the baseband unit (BBU), the Telco Cloud and the packet core. While the paper itself presents a system using the Xilinx UltraScale+ family [5] of devices with 32.75Gbps transceivers [6][7], the delay and PDV results presented from the numerical analyses, are verified using results from earlier experiments [8][9], run on a previous generation of Xilinx UltraScale 20nm [5] devices with 30Gbps transceivers.

II. Ethernet-based Fronthaul: Issues and Challenges

Ethernet Fronthaul needs to support strict timing constraints and quality of service required by the encapsulated radio protocol, e.g. eCPRI or 1914.3 RoE, including low delay and low packet delay variation (PDV). Indeed, to ensure correct processing of the encapsulated data, a steady stream without (PDV) is required. In practice, the PDV can be removed by a playout buffer at the receiver side with the minimum buffer size equal to or greater than the peak PDV. However, this requires knowledge of the maximum delay and peak PDV to dimension the network.

The challenge with Ethernet is that packet delay and delay variation is dependent on traffic load. The upcoming TSN standard for Fronthaul (IEEE 802.1CM [2]) provides a mechanism to preempt low priority traffic in favor of high priority Fronthaul traffic [10], but does not include deterministic scheduling. The current TSN scheduling mechanisms (IEEE 802.1Qbv [11], IEEE 802.1Qch [12]) targeting industrial automation applications were not deemed suitable for Fronthaul, and have therefore not been included in the IEEE 802.1CM standard. This paper describes a Fronthaul ecosystem that includes a mapper for CPRI/eCPRI [13],[14] encapsulation, an accurate timing solution, an optional pre-emption mechanism based on the TSN 802.1CM profile, and a deterministic scheduling layer. The scheduling mechanism incorporates a look ahead scheme [9] that efficiently schedules traffic by exploiting the inter packet gaps (IPG) and enables the aggregation and convergence of multiple fronthaul, midhaul and backhaul streams in the same network. The overall cost function to minimize is the packet/frame delay and delay variation even in the presence of multiple traffic sources.

III. System Description

A complete Fronthaul ecosystem has been designed as a set of separate IPs covering radio over Ethernet encapsulation (eCPRI framer/deframer), deterministic scheduling (FUSION Scheduler), preemption (TSN EMAC and PHY), and accurate synchronization (IEEE 1588 v2.0 or 1588HA).

The platform chosen for the system is the Xilinx UltraScale+ MPSoC devices [5] that have 16.3Gbps[6] and 32.75Gbps [7] transceivers to provide the full range of data rates required by CPRI v7.0 [1] and the new packet based eCPRI protocol [3]. These devices also have hard-macro RS-FEC which is optional in these protocols. The block diagram of a multi-port switch based on all the pieces of IP is shown in Fig. 1(a). The Ethernet subsystem (MAC/PCS) is soft IP [15] developed in compliance with the standards defined by the 25G Ethernet Consortium and is also enhanced to support preemption. The MAC merge sublayer is a wrapper around the EMAC that provides the necessary buffering to support preemption. The eCPRI framer and deframer implement the packetization and depacketization of frequency domain or time domain IQ data into the Ethernet frame and support all eCPRI standard protocol requirements. A standard CPRI protocol framer could also be used instead to support legacy traffic as a CPRI-to-eCPRI bridge. The timing synchronization is accomplished with an IEEE1588v3.0 (High Accuracy) system. The look-ahead scheduler called FUSION developed by TransPacket is also optimized for Xilinx devices and achieves minimal or no PDV for fronthaul, and enables multiplexing midhaul and backhaul traffic over the same transport network. These characteristics are achieved by implementing the following scheduling features:

a. Deterministic aggregation: capability to aggregate several fronthaul streams with low delay and controlled PDV, i.e. bounded delay.
b. Deterministic priority: capability to statistically multiplex lower priority traffic, e.g. midhaul and backhaul traffic, with no impact on fronthaul streams.

IV. Mobile Transport Network Description and Analyses

The deterministic aggregation with low and bounded delay of 5 x 10Gb/s Ethernet Fronthaul streams into a 100Gb/s Ethernet transport path, and the transport/switching with deterministic priority through four nodes has recently been experimentally demonstrated in [8]. Furthermore, the deterministic priority on a 100Gb/s path is further studied/demonstrated in [9]. In this work, we present the analysis of a multi-node mobile network with, strict priority (SP) and weighted round-robin (RR) Ethernet schedulers, versus the FUSION scheduler of the current system. The network under consideration is illustrated in Fig. 1(b).

For the numerical analyses of the fronthaul traffic performance in the network of N nodes, the following network nodes, traffic sources and parameters are considered, also described in Table 1:

a. Fronthaul (FH) aggregation node:
   - Input: F * fronthaul streams encapsulated in 10Gb/s Ethernet interfaces.
   - Output: G * 100Gb/s Ethernet aggregated stream.
b. Transport node:
   - Fronthaul traffic to be switched between the input and output 100Gb/s Ethernet interfaces of the transport path. Note there is no traffic contention between FH streams.
   - Backhaul (BH) /midhaul traffic
     - Input: B * 10Gb/s Ethernet interfaces
     - Output: G * 100Gb/s Ethernet interface on the transport path

The end-to-end maximum delay and delay variation of the fronthaul traffic added by the network nodes is calculated as the sum of the delay experienced at the aggregation node with the delay experienced at each transport node in the path.

The output stream of the three scheduling algorithms when aggregating both FH (black packets) and BH (white packets) is illustrated in Fig. 1(c). The strict priority mechanism, increases the FH PDV by one BH MTU transmission time per each transport node. E.g. consider the case when a high priority FH packet, e.g. black packet 2 at SP in Fig. 1(c), arrives just after the SP scheduler has started transmitting a packet of maximum size (MTU) from the low priority BH queue, e.g. white packet 1. The added delay to the FH packets will vary from 117ns, its store-and-forward delay variation for 64 to 1522 Byte packets, up to an additional 122ns of the transmission time of a maximum sized BH packet. Thus, the peak PDV for SP is 239ns per node.

In the case of the round-robin scheduler, the PDV and maximum delay further deteriorate by increasing up to B * MTU midhaul/backhaul packets which can be served before the FH packet. Thus, inducing a PDV up to 718ns per node to the FH packets.

TABLE I. Parameters and Values.

Parameter used	Notation and Value
Minimum transmission unit, i.e. minimum packet length	64 Byte
Maximum transmission unit (MTU)	1522 Byte
Maximum store-forward delay at aggregation node	St-fw=MTU/10Gbps=1.22 μs
Maximum Transmission Time @100G	T=MTU/100Gbps= 0.122 μs
Number of FH streams	F =[1, 5]
Number of Backhaul/Midhaul streams per transport node	B = 5
Number of output transport interfaces	G=1
Number of nodes	N=[2, 10]
Maximum delay at aggregation node	Dag.=st-fw + F*T
Maximum delay at RR transport node	Dt=1*T + B*T=0.721us
Maximum delay at SP transport node	Dt=1*T + 1*T=0.244us
Maximum delay at FUSION transport node	Dt=1*T + 20ns=0.142us
Maximum end-to-end (path) Delay	Dmax=Dag + (N-1) Dt

The FUSION scheduler on the other hand, adds a fixed delay to the FH traffic at each node. This enables the look ahead for identifying gaps within this FH stream and filling/inserting less delay sensitive BH packets only in fitting gaps. Thus, no PDV is induced on the FH stream of packets. This delay is equal to 144ns and corresponds to 122 ns, the transmission time of a lower priority packet MTU at the transport interface plus 20 ns fixed processing time. While all packets experience this delay, i.e. even if there is no BH packet to be transmitted, this enables a fixed PDV per node which simplifies the network and PDV buffer dimensioning challenges for the carrier.

The results on the maximum end-to-end delay and peak PDV of the fronthaul traffic are illustrated respectively in Fig. 2(a) and (b). Note that a First-bit-in First-bit-out (FIFO) delay model is used, i.e. store and forward delay is included only at the receiving (aggregation) node. The FUSION numerical results are equal with the experimental results [8] in a four-node network, where the PDV added by the Ethernet 10G and 100G MAC/PHY is not included. The difference between the analytical and experimental results is less than 10 ns.

Results illustrate the FUSION scheduling advantage regarding fronthaul traffic performance isolation: fronthaul traffic delay and delay variation are independent from the number of midhaul/backhaul traffic streams and the number of transport switching nodes it traverses. This is observed in Fig. 2(b), where the FUSION PDV is constant and independent of number of nodes, while in Fig. 2(a) the maximum delay increases only with a fixed MTU per additional node Conversely, note the proportional increase of the maximum delay and PDV for the two other schedulers with the number of nodes. The worst case is for the round-robin scheduler, where the maximum delay and delay variation increase proportionally with the number of BH streams in the network since each node has a fixed number of ports and in this example, it is fixed to five ports. The strict priority scheduler's delay and delay variation

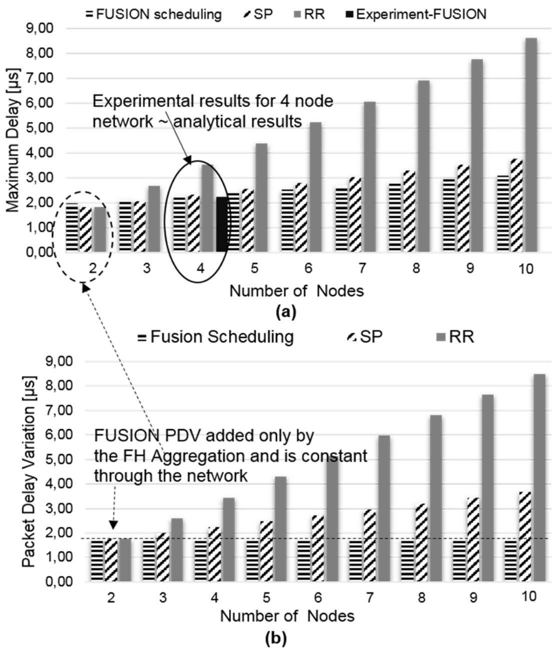

Fig. 2 (a) Maximum end-to-end delay as a function of number of nodes. (b) Peak PDV as a function of number of nodes.

increase proportionally with the number of traversed nodes as each node adds an additional MTU delay variation.

For the case of only two nodes marked with dotted lines in the graph, i.e. no backhaul traffic is being added in the network, the FUSION delay is slightly higher than strict priority, i.e. by 144ns which is the deterministic priority delay that enables the system to add backhaul traffic at any point in the future. This small delay comes with the advantage of a seamless network upgrade/change, ensuring that the FH peak PDV shall not change. The PDV playout buffer size can thus be dimensioned for the maximum number of FH streams that the network shall aggregate. Furthermore, this low PDV of 1.77µs translates into a very low buffer size requirement of 22.5Kbyte, which is important for lowering the costs especially at the access network, i.e. the RRU site.

Furthermore, while the gain after N=3 of FUSION versus SP on the maximum delay is not as high as with RR, the gain not only comes with the fixed PDV through the network but enables FUSION deterministic add/drop of FH in the transport nodes. In future work shall be demonstrated the FH add/drop in these nodes, without impact on the FH traffic on the transport path.

One of the major concerns of mobile operators is to keep the Ethernet fronthaul simple, and to ensure that the support of deterministic behavior does not lead to a significant increase of their operational cost due the complexity of configuration of mechanisms that depend on too many factors, e.g. topology, load, traffic type, and traffic variation [16]. In IEEE 1914.1 Task Force, there is an on-going discussion (initiated by Mobile Network Operators (MNOs) about defining classes of Ethernet transport nodes for fronthaul, associated with a given upper value for processing time [17]. The objective is to have a fixed value, independent of traffic type and load, in order to simplify the network design, and network upgrade/change. The ability to add/drop FH and BH/MH traffic to existing deployment without impacting the deployed FH services is a key added value of the FUSION scheduler, which addresses the MNO requirements for simplicity, and ease of design.

V. CONCLUSION

This paper presented a converged fronthaul, midhaul and backhaul system with TSN enabled EMAC and the FUSION Scheduler on Xilinx Ultrascale+ devices. Furthermore, it was demonstrated that the FUSION switching and scheduling IP Core is able to support deterministic Ethernet as required by a fully converged mobile transport network. Using FUSION IP cores, fronthaul traffic from e.g. CPRI and eCPRI streams can be aggregated with low and bounded delay, and switched with fixed and low delay throughout the network. The scheduler further enables the aggregation of midhaul and backhaul streams into the same network without affecting the timing of the fronthaul traffic streams.

ACKNOWLEDGMENT

This work is part of research project 5G-PICTURE supported by the European Horizon 2020 initiative.

REFERENCES

[1] CPRI Consortium, "CPRI Specification V7.0," Oct. 2015.

[2] IEEE Std. P802.1CM, "Time Sensitive Networking for fronthaul".

[3] CPRI Consortium, "eCPRI Specification V1.0," Aug. 2017.

[4] IEEE Std. P1914.3, "Radio Over Ethernet Encapsulations and Mappings," Sept. 2017.

[5] Xilinx UltraScale+ Architecture and Product Data Sheet: Overview (DS890).

[6] UltraScale Architecture GTH Transceivers User Guide (UG576)

[7] UltraScale Architecture GTY Transceivers User Guide (UG578).

[8] R. Veisllari et al., "Experimental Demonstration of 100 Gb/s Optical Network Transport and Aggregation for Ethernet Fronthaul with Low and Bounded Delay," OFC 2018 (submitted)

[9] R. Veisllari et al., "Experimental Demonstration of 100 Gb/s Optical Packet Network for Mobile Fronthaul with Load-independent Ultra-low Latency," ECOC 2017, paper M2A2.

[10] IEEE Std. P802.1Qbu, "Frame Preemption".

[11] IEEE Std. P802.1Qbv, "Enhancements for Scheduled Traffic".

[12] IEEE Std. P802.1Qch, "Cyclic Queuing and Forwarding".

[13] CPRI v8.8 LogiCore IP Product Guide (PG056)

[14] CPRI Gateway PoC v1.0 (PB047).

[15] 10G/25G High Speed Ethernet Subsystem v2.2 Product Guide .

[16] D. Chen and E. Ryytty, "Qbv Optional for Fronthaul over Ethernet," online: http://www.ieee802.org/1/files/public/docs2016/cm-chen-Qbv-Optional-for-Fronthaul-over-Ethernet-0316-v02.pdf.

[17] A. Checko, "Node classes – initial contribution," online: http://www.ieee802.org/1/files/public/docs2017/cm-farkas-eCPRI-support-0917-v01.pdf.