Experimental Demonstration of 100 Gb/s Optical Network Transport and Aggregation for Ethernet Fronthaul with Low and Bounded Delay

Raimena Veisllari[1], Steinar Bjornstad[1, 2], Jan P. Braute[1]
[1]*TransPacket, Hoffsveien 21-23,0275 Oslo, Norway*
[2]*Department of Telematics, Norwegian University of Science and Technology, O.S. Bragstads 2b, Trondheim, Norway*
Email:{raimena; steinar}@transpacket.com

Abstract: 3-node integrated packet/circuit network experiment demonstrates 100Gb/s transport and aggregation of five 10Gb/s links with low and bounded delay. 3.4μs maximum end-to-end delay is achieved, even when combining with less delay-sensitive traffic, reaching 98% utilization.
OCIS codes: (060.4253) Networks, circuit-switched; (060.4259) Networks, packet-switched

1. Introduction

An important objective for the next generation 5G mobile networks is extending the range of services and delve into new verticals, e.g. IoT, automation and automotive. While different verticals and applications have different requirements in terms of bandwidth and delay, e.g. from very relaxed web access to strict virtual reality delay requirements, the network must support the full range of old and new services in a cost- effective and profitable way for the service providers. Applying mobile fronthaul (FH) for transport is a promising candidate for delivering cost-efficient high-bandwidth 5G networks. Removing processing, at different degrees, from the radio equipment (RE) to a centralized/virtualized radio equipment controller (REC), lowers the cost of the RE. Thus, operators can cost-efficiently expand the mobile network coverage by deploying lower-cost REs. Mobile FH protocols, as CPRI and OBSAI, have very strict delay requirements: end-to-end one-way delay lower than 100 μs between the RE and the REC [1],[2]. Currently, transport of FH traffic is therefore circuit-switched through dedicated bandwidth resources, e.g. dedicated fibers and wavelengths. Thus, by moving towards packet-based solutions, there is large potential for gain through statistical multiplexing: leveraging their cost efficiency of sharing bandwidth resources. The objective is to aggregate multiple FH streams into the same channels (wavelengths). Furthermore, carriers offering FH transport will benefit from a converged infrastructure supporting it alongside their BH infrastructure, and in future also fixed access for a fully converged network.

These objectives and the wide Ethernet deployment have driven new standardization work as IEEE 1914.3 radio over Ethernet (RoE)[3], IEEE next generation fronthaul interface (NGFI) 1914.1[2], and time-sensitive networking (TSN) for fronthaul IEEE 802.1CM[4], while the CPRI consortium has recently released the first eCPRI over IP/Ethernet specification[1]. Through all these specifications, though, the 100 μs one-way FH delay budget still remains for the high priority class as a maximum absolute end-to-end delay through the network (including processing, queuing and transmission delays) and limits the maximum distance between RE and REC. A main challenge in packet networks, e.g. Ethernet bridges and IP routers, is that they were not designed for strict timing support and delay is dependent on traffic load. Furthermore, to ensure correct processing of the encapsulated data, a smooth PDV-free stream is required. In practice the PDV can be removed by a playout buffer at the receiver side: delay the fastest packet to be equal to the slowest packet. For this purpose, the minimum size of the buffer has to be set for the peak PDV. Therefore, provisioning a FH network requires the delay to be deterministic and as low as possible, i.e. the maximum delay and peak PDV must be known for being able to dimension the FH network.

Integrated hybrid (packet/circuit) optical networks [5] enable a multi-service path supporting: (1) a guaranteed service transport (GST) service class with fixed delay and ultra-low packet delay variation (PDV) independent of load, and (2) high throughput efficiency through a lower-priority statistically multiplexed (SM) service class. In previous work [6] has been demonstrated a two node 100G path suitable for carrying FH traffic through the GST class, while increasing the throughput through added SM traffic. In this work, the IHON architecture and application for fronthaul is further extended, demonstrating aggregation of multiple FH streams with low and bounded delay into a 100G path. In addition, the 100G path's delay is further optimized for FH transport. The aggregated FH traffic is then carried through the 100G path shared with lower-priority SM traffic, with fixed delay and ultra-low PDV in the range of ns. The SM performance is evaluated and delay results show that this is well within the 1ms backhaul traffic bounds. Thus, a full converged fronthaul and backhaul network solution is proposed and evaluated through an experimental IHON demonstration carried in a lab testbed with three nodes.

2. IHON node mechanisms

The bounded delay aggregation mechanism is illustrated in Fig.1(a). The Ethernet FH streams are aggregated in an equal priority round-robin scheduler. The streams are competing only with each other. Thus, the aggregation delay is bounded to- and is a function of- the MTU and number of the aggregated streams S. The maximum aggregation delay is a worst-case delay, i.e. the delay bound D_{agg} , as in Fig.1(a), occurs when a maximum sized packet arrives and has to wait for all other packets in the other (S-1) streams, all of maximum size, to be served before it is transmitted: D_{agg} =MTU/10Gbps+(S-1)*MTU/100Gbps+ MTU/100Gbps, where the latter is its transmission delay (service time) in the 100G network port. The PDV is then defined as the difference between the best case, only one stream, and worst case: PDV=(S-1)*MTU/100Gbps. The GST priority mechanism, referred to as bypass, is illustrated in Fig.1(b). A fixed delay δ of 0.2 µs is applied electronically to the stream of packets bypassing the node. It corresponds to a fixed header processing delay of 80ns plus a maximum SM packet length of 1550Byte (124ns). This enables a gap detector looking into the input bypass stream to identify the length of inter-packet gaps. When a gap is detected, the SM scheduler inspects the SM queues for a packet of suitable size that fits the gap. If such a packet is found, it is inserted in the gap without affecting the timing of the packets in the bypassing stream. Hence, the effect of the bypass mechanism on the high priority class (GST) is deterministic and known through the network. In a path with N IHON nodes, the maximum end-to-end delay of the FH traffic is deterministic and bounded to: D_{agg} + (N-1)δ.

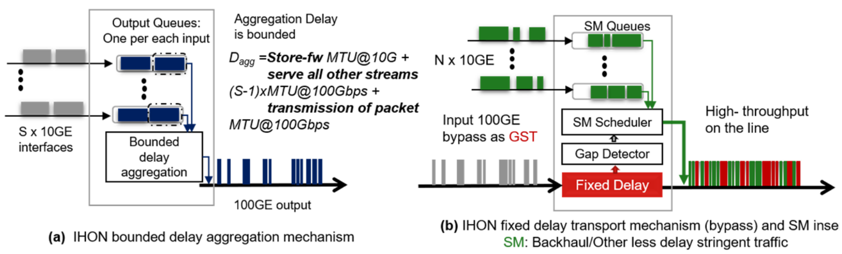

FH Parameter used	Notation and Value
Fixed bypass	δ=0.2 µs
MTU	1522 Byte
Store-forward	St-fw=MTU/10Gbps =1.22 µs
Max Service Time @100G	T=MTU/100Gbps= 0.122 µs
Number of FH streams	S = from 1 up to 5 (in this work)
Number of nodes	N=3
Max Agg. Delay e.g. 5 FH streams	Dag.=st-fw + S*T e.g. 1.83 µs
Max IHON end-to-end (path) Delay	Dp=Dag + (N-1) δ S=5,N=3 Dp=2.24 µs

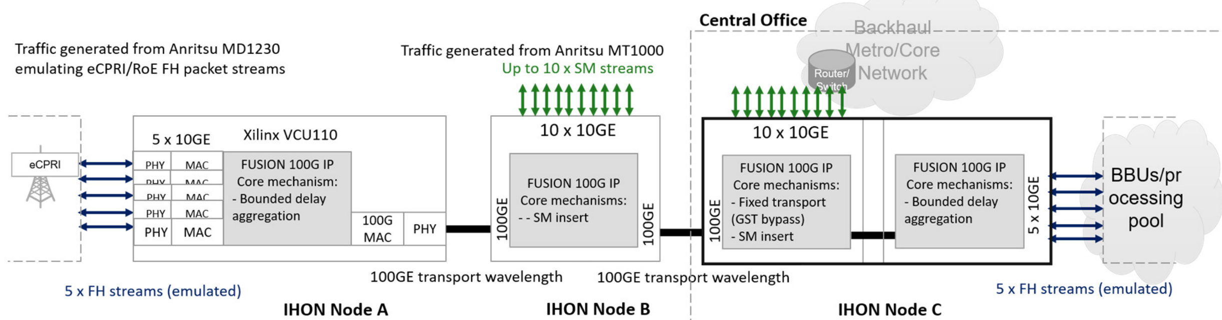

Fig.1: (a) Bounded delay aggregation with a deterministic upper bound; (b) SM insertion in idle time gaps of the 100G input bypass stream; (c) Experimental setup with three network nodes. Note: Node C is a composite of two FPGA boards for the two required IHON mechanisms: bypass and aggregation. Anritsu testers are used for traffic generation (emulating FH sources) and performance characterization

3. Experimental Results and Discussion

The nodes used in the experimental setup, illustrated in more details in node A Fig. 1(c), are Xilinx VCU100 Ultrascale FPGA evaluation boards, which include Ethernet PHY/MACs for the respective 10/100GE ports and the IHON mechanisms implemented in FUSION 100G Ethernet IP Cores from TransPacket. For emulating the FH and BH/Access traffic streams, 10G Anritsu testers are used. In addition, for increasing the offered load through the system, custom traffic generators are implemented on the 10G ports in the FPGA and used when required. The FH traffic MTU used is fixed to 1522B, where 1518Byte is the Ethernet maximum sized packet and 4 Byte the VLAN tag emulating eCPRI over Ethernet [1], i.e. FH streams are already tagged for differentiating the FH traffic classes. In a first test case, the FH streams follow an on-off periodic traffic pattern as in [7]; bursts of 3MTUs for the user data (UD) and 1MTU for the real-time control per period, generate a load approximated to eCPRI Split D downlink of 4Gbps per 10G port [8], Fig.2(a). In the second case, Fig.2(b), a fixed FH load of 0.98 per 10G port is generated, emulating a 10G CPRI/Radio over Ethernet, or eCPRI Split I_D, where the load is approximately 10Gb/s. The SM traffic, for all tests, follows a uniform distribution of 64 to 1522B packets. SM load is then fixed to 0.98 at 10G ports by fixing the interframe gap (IFG). For both cases, first is measured the performance of FH streams, added one by one up to five, then added SM one by one until saturation, i.e. SM experiences losses. A reference end-to-end delay for a single stream added at one 10G port at node A and received at 10G port at C, as in Fig.1(c), with fixed 1522B packets, was measured to a maximum delay of 2.9μs and PDV=0.1μs, with 10G tester resolution of 50ns. This is the best-case (minimum) delay as the FH stream is not competing with other FH traffic. Removing the fixed bypass delay components 2δ, the 1.22 μs store and forward delay and the 0.12μs transmission delay at node A, the Ethernet 10/100G PHY/MACs through the path add a maximum delay of 1.16μs. The maximum delay (upper bound) for up to five streams should be the same for all FH traffic patterns. The aggregation delay at node A is Dagg =1522B/10Gbps + 5*1522B/100Gbps=1.83μs. The end-to-end delay added by IHON mechanisms is thus Dagg+ 2δ≤2.24μs; adding the measured MAC delay component 1.16μs, the maximum FH path delay is then to 3.4 μs. Measured results are consistent with this expected delay, and are the same for both test cases, thus, independent of the FH traffic pattern and load at 10G. I.e. the maximum measured delay through all cases was 3.39 μs. Furthermore, results show that the FH performance remains independent of the SM traffic inserted in the 100G path, even when the path is congested and SM starts experiences losses. Because SM insertion is dependent on the bypass traffic load and pattern (distribution of inter-packet gaps) [5], it experiences losses at different loads, see Fig2(a) and (b). Results show that a high throughput is achievable in the path: 88.8Gbps for FH rate of 50Gbps, and 98.4 Gbps for FH rate of 20Gbps. At these points of high path utilization/SM congestion, without SM losses, the maximum SM delay is comparable to its buffer size of 32KB at 100Gbps~32μs. Thus, SM can be further dimensioned for efficiently utilizing the path while carrying BH traffic within the 1ms delay bounds.

With five aggregated FH streams, the FH peak PDV was measured to 0.6 μs for both FH traffic patterns, independent of the added SM load. This is consistent with the theoretical bound of 0.49μs aggregation PDV, plus the measured reference PDV with one FH stream of 0.1μs. Thus, for a 10G Ethernet port at the RE/REC would be required only a small (lower cost) 8KByte playout buffer for smoothing out the delay variation.

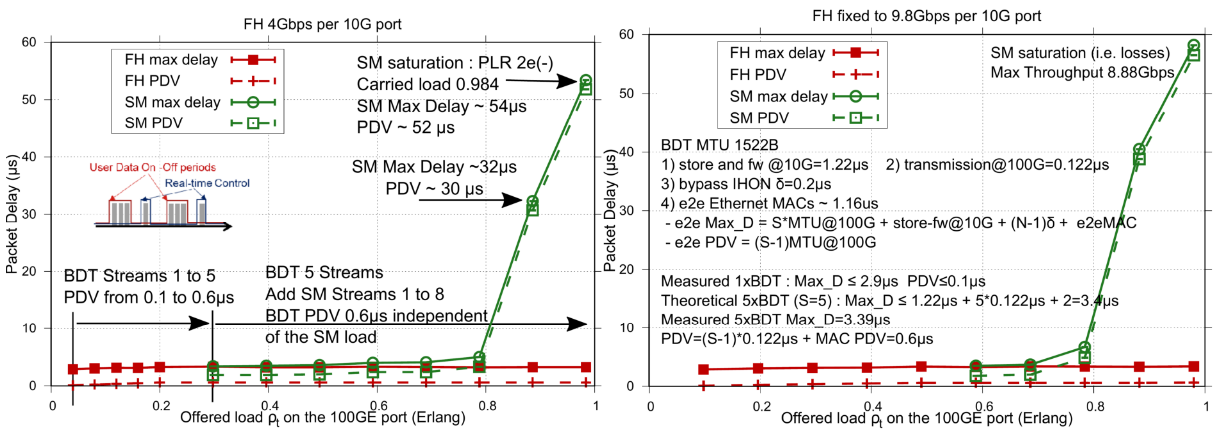

(a) (b)

Fig.2: (a) Maximum delay and PDV as a function of the total offered load in the 100G path, with FH as on/off sources with offered load 0.4 on the 10G ports; (b) FH 0.98 load at 10G ports. In both cases: first add 1 up to 5 FH streams in the path, then add SM one by one at 0.98 load at 10G.

4. Acknowledgement

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

5. Conclusion

It was experimentally demonstrated for the first time IHON fixed delay transport and bounded delay aggregation of 10 Gb/s streams into 100 Gb/s Ethernet, suitable for transport of Ethernet fronthaul streams from multiple radio units. The FH path was bridged through three nodes and shared with less delay sensitive, e.g. BH/access, services. Zero packet loss, a maximum delay of 3.4 μs and peak PDV less than 0.6 μs – independent of all traffic load - was demonstrated for five 10G Ethernet fronthaul streams, leaving 96.6 μs of a 100 μs fronthaul delay budget to transmission delay. 98% and 88% path utilization was demonstrated, respectively with 20% and 50% fronthaul traffic, without BH/access services packet loss.

6. References

[1] CPRI Consortium, "Requirements for the eCPRI Transport Network, D0.1," Aug. 2017.

[2] IEEE Std. P1914.1, "Standard for Packet-based Fronthaul Transport Networks," Sept. 2017.

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

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

[5] S. Bjornstad et al., OFC 2013, post deadline paper PDP5A.8.

[6] R. Veisllari et al., ECOC 2017, paper M2A2.

[7] J. Farkas, "Support of eCPRI in 802.1CM," online: http://www.ieee802.org/1/files/public/docs2017/cm-farkas-eCPRI-support-0917-v01.pdf

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