OPTIONS FOR TIME-SENSITIVE NETWORKING FOR 5G FRONTHAUL

Jim Zou[1], Silviu Adrian Sasu[1], John Messenger[2], Jörg-Peter Elbers[1]

[1]ADVA Optical Networking SE, Martinsried/Munich, Germany
[2]ADVA Optical Networking Ltd., York, United Kingdom
**email: jzou@advaoptical.com*

Keywords: time-sensitive networking, radio access network, 5G, fronthaul

Abstract

Mobile networks towards 5G impose increasingly stringent requirement on high throughput, low packet jitter and low latency due to emerging time-sensitive applications. Ethernet-based mobile fronthaul is low-cost and flexible, but lacking in determinism. We review and discuss the options for time-sensitive networking based on Ethernet.

1 Introduction

5G radio access network (RAN) imposes many unprecedented and challenging requirements on the optical mid- and fronthaul transport. The 3GPP introduces a new NG-RAN architecture in Release 15, consisting of three main functional modules, i.e. the Central Unit (CU), the Distributed Unit (DU), and the Radio Unit (RU) [1]. Since in the 3GPP standards, NG-RAN is more like a logical architecture, these functional modules could be deployed in different combinations. As outlined to enable enhanced Mobile Broad-Band (eMBB), ultra-Reliable Low-Latency Communications (uRLLC), and massive Machine-Type Communications (mMTC) over a single network, the fronthaul segment between DU and RU is more latency-sensitive and bandwidth demanding compared to the CU-DU interface [2]. Therefore, the new NG-RAN and end-user applications drive the need for new transport networks. Although the existing 4G LTE RAN and core can be reused (i.e. non-standalone 5G NR) as the initial step, it will not be efficient to meet the requirements on new physical topologies, functional split options, etc.

Among potential transport technologies, Ethernet is a low-cost and ubiquitous technology used in vast types of networks, through its well-established economies of scale. For 5G RAN, Ethernet as a packet-based system can enable meshed connectivity between RAN functional modules, statistical multiplexing gain, multi-vendor interoperability, and prevent the dependence on wavelength like the CPRI fronthaul.

However, as originally designed, Ethernet frames are treated identically, regardless of the priority of traffic flows, which obstruct a direct adoption of Ethernet especially for the fronthaul. This is because fronthaul is a lower-layer split (LLS) interface that is more sensitive to latency and packet jitter [3]. Moreover, the TDD-based wireless communication, multi-point radio access technologies, and cloud-RAN rely on precise time synchronisation once adopted [4].

Although traffic classes using different priorities and virtual LANs (VLANs) according to IEEE 802.1Q were introduced to separate these traffic flows on the same LAN [5], due to the buffering mechanism, mixing flows and port contention in Ethernet switches cause non-deterministic delay and thus degrades the Quality of Service of particular traffic. For instance, when the switch has already started the transmission of an Ethernet frame, while at that very moment, another frame with even the highest priority also requests transmission on the same egress port, this high-priority frame has to be buffered until the ongoing one is ended.

For the purpose of cell site synchronisation, the most straightforward way would be to get the reference clock (typically Pulse per Second, PPS) and Time of Day (ToD) from the global navigation satellite system (GNSS). Yet this is costly for installation, and access to the GNSS satellite signal cannot be guaranteed at all times. As an alternative, the frequency and time information need to be distributed from one central source through the network. ITU-T specifies Synchronous Ethernet (referred as SyncE) to transfer a synchronisation signal between network elements [6]. It is a bit-level clock recovery solution similar to SDH/SONET and highly robust for long-term frequency distribution, but it only provides frequency synchronisation and every node in the sync path requires physical-layer support for SyncE. To complement SyncE, the IEEE 1588v2 Precision Time Protocol (PTP) provides mechanism for end-to-end phase and time alignment, as well as frequency where SyncE is not available [7]. However, as a packed-based Layer 2/3 protocol, the performance of PTP is extremely prone to the packet delay variation (PDV) across network elements.

Putting the pieces together, *time* plays an important role in the RAN, and the emerging 5G time-sensitive applications draw even more attention to this *time* aspect. Time-sensitive networking (TSN), defined by the task group in IEEE 802.1 as a Layer 2 technology, aims at deterministic data transmission over Ethernet networks. In contrast to the most conventional IEEE 802.3 Ethernet, TSN defines traffic queues based on time through every network element that has to be synchronised, to provide zero loss from congestion and bounded latency for a variety of time-sensitive or time-critical data streams coexisting on a network that also supports best effort traffic.

In this paper, we will review the TSN solution to 5G fronthaul, and discuss existing approaches and ongoing works on scheduling and traffic shaping. We will also present the proof of concept of a 100GbE aggregator that prioritizes the time-sensitive service and ensure an accurate time synchronization for the TSN-enabled fronthaul.

2. Traffic Scheduling and Shaping Techniques

As revealed in the introduction, the key to enabling determinism in TSN is the concept of sharing time. The IEEE 802.1 TSN task group selected a few but most critical options from IEEE 1588-2008, to specify a profile, IEEE 802.1AS (and its revision 802.1AS-Rev), for time synchronisation over a virtual bridged local area network (referred to IEEE 802.1Q). The next step is to make sure that packets will be delivered and received just on time, i.e. *deterministically*.

Since the standard Ethernet switching cannot avoid non-determinism, many amendments or extensions to existing Ethernet standards were incorporated to enhance traffic scheduling and shaping, which will be elaborated in the following subsections.

2.1 IEEE 802.1Qav: Forwarding and Queuing Enhancement

The 802.1Qav uses credit-based shaper to determine the priority of a queued frame and when it can be transmitted. The credit value is constantly increased in case of no frame transmission, whereas it is decreased in case of transmission. A transmission is then only allowed when the credit value is positive. The purpose of 802.1Qav is to distribute burst traffic over the course, which prevents overloading of network bandwidth and hence reduces excessive delay caused by the burst traffic congestion. However, the buffer requirement is topology dependent and may cause the worst-case delay.

2.2 IEEE 802.1Qbv: Time-Aware Shaper

Leveraging the credit-based shaping, the 802.1Qbv reserves a protected time window to only allow the scheduled traffic and block non-scheduled one at each port. On top of that, it uses a transmission gate per traffic class queue to allow only one traffic class queue to transmit at a specific time within this window, which is the same concept as the time-division multiple access (TDMA). In addition, a *guard band* time period needs to be allocated in front of the scheduled time window, to avoid contention between the scheduled traffic and the residual non-scheduled packet tail from the previous periodic cycle. The 802.1Qbv can therefore eliminate the PDV of scheduled traffic (i.e. deterministic latency), but requires a network-wide schedule which is not standardized, and the protected time window and guard band mechanism might cause a certain degree of bandwidth waste.

2.3 IEEE 802.1Qbu/802.3br: Frame Preemption and Interspersing Express Traffic

The initial purpose of frame preemption is to reduce the delay of guard band necessary in the 802.1Qbv, given that the guard band has to be larger than the transmission time of the longest packet. The transmission of non-scheduled or best-effort packet is interrupted in order to yield to the scheduled traffic, and resume the remaining transmission after the end of the scheduled window. To do that, two ends of the link need to activate the preemption support through the Link Layer Discovery Protocol (LLDP), and the Ethernet MAC needs to be instructed to hold back the preemptable traffic. Frame preemption only works hop-by-hop, and can be enabled only in networks that are aware of the 802.1Qbu. The 802.1Qbu specifies its bridge management, while the 802.3br does the MAC part. Since this scheme has a minimum fragment size of 124 octets in length, taking the mandatory interframe gap, preamble, and delimiter into account, an additional delay of 143 octets transmission time might be occurred.

2.4 IEEE 802.1Qat: Stream Reservation Protocol

Even though every switch hop is capable of deterministic latency, if may only work if the switch resources (e.g. throughput, buffer size) are available along the entire path. The idea of 802.1Qat is to provide a standardized protocol to manage the resource allocation or reservation of each switch. It operates on streams that are identified by the source MAC address and a higher level identification like the IP port address to meet bandwidth requirements and latency class.

2.5 IEEE 802.1Qcc: Enhanced Stream Reservation

The 802.1Qcc is still an ongoing project that extends the 802.1Qat, enabling TSN network administration on large scale networks. In addition, the size and the frequency of reservation messages are optimized.

2.6 Reliability and redundancy

For mission-critical use cases, such as autonomous driving and industrial communication and automation systems (Industry 4.0), the fault tolerance is an absolute necessity on top of the timing assurance. The IEEE 802.1CB is currently being specified for this purpose, which allows a selective packet replication on a per-stream basis identified by the IP/MAC address, traffic class, and sequence number. The duplicated stream is sent on a separate path, and at the final destination port where the two identical streams converge, the redundant stream will be recognized and eliminated.

The IEEE 802.1Qci uses the stream identification from 802.1CB to enable per-stream filtering and policing for both bridges and end points. It uses the combination of stream gates and flow meters to allow frame queuing decisions.

2.7 IEEE 802.1Qch: Cyclic Queuing and Forwarding

The 802.1Qch combines the capabilities in the previous standards, i.e. 802.1Qbv and 802.1Qci, to collect and forward the streams in a particular traffic class in one cycle, to guarantee a fixed delay per hop. The benefits of the 802.1Qch cyclic queuing and forwarding approach is that if the cycle time is set appropriately, the latency of a packet through the network can be calculated by computing the sum of the per-hop delays. The per-hop delays are dominated by cycle time and are unaffected by any other topology considerations. This is a clear advantage compared to the 802.1Qbv time-aware shaping where the worst-case delays are topology dependent and not so easily calculable. The 802.1Qch is also expected to be applied together with 802.1Qbu/802.3br to reduce the minimum cycle time.

2.8 IEEE 802.1Qcr: Asynchronous Traffic Shaping

One of the prerequisites for queueing and forwarding mechanisms is network synchronisation. Alternatively, the 802.1Qcr (based on the urgent-based scheduler as an early-stage development) aims to be independent on network-wide cycle synchronisation or scheduling, while still to provide deterministic and bounded low latency. Since the switch implemented with the 802.1Qcr runs on a standalone clock, the shaper calculates and assigns the eligibility time of flows sent out from the queue, and the time determines the transmission scheduling. A frame is eligible for transmission if the assigned eligibility time is less than or equal to the current time. This mechanism is able to control the maximum per-hop delay through the eligibility time assignment, but it also implies additional computation process on each frame.

2.9 FUSION: Gap Preservation

So far, most TSN standards and implementations try to guarantee a deterministic delay through multi-hop network bridging, but not the minimum PDV in case of statistical multiplexing. However, leveraging the statistical multiplexing gain is particularly beneficial to the Ethernet-based fronthaul, because more data streams can be delivered over a single fibre link, saving the precious fibre resources (e.g. in the dense urban area deployment). Therefore, the integrated hybrid (packet/circuit) optical network was proposed to enable a fixed delay traffic class combined with a statistically multiplexed traffic class. A proprietary implementation of such a concept, known as FUSION, was demonstrated experimentally and in the field [8][9]. Its idea is to apply a fixed delay on the high-priority traffic, in order to detect the inter-packet gaps at the ingress and preserve the gaps at the egress, while the statistically multiplexed traffic can be inserted in between as long as it fit each gap spacing. By means of only monitoring the incoming packets, one big advantage is that such a mechanism does not require network-wide synchronization.

3. TSN-based Fronthaul

Although all the aforementioned mechanism tackle the TSN implementation, they provide essentially a powerful toolbox instead of profiling certain combined capabilities to meet requirements of a particular application or use case. In other words, not every mechanism is needed for a given application. Specifically, in light of the fronthaul application, the IEEE 802.1CM selects and specifies TSN profiles to provide bridged Ethernet for the fronthaul network, as illustrated in Fig. 1. The 802.1CM takes requirements from CPRI and eCPRI, and separates them into two Class 1 and Class 2, respectively. The standard also specifies two profiles (A and B) applicable to both classes, where Profile A employs strict priority scheduling for fronthaul traffic, and Profile B extends Profile A with frame pre-emption based on the 802.1Qbu. Although packet timing (e.g. PTP) for distributing time and frequency synchronization to the remote radio equipment is defined in deployment approaches, it is not clear how the timing packet should be treated when being transmitted together with fronthaul traffic at the switch or aggregator.

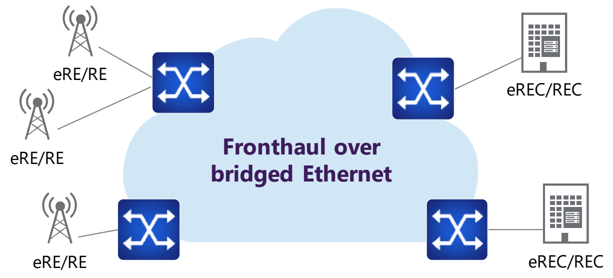

Fig. 1 IEEE 802.1CM defines Ethernet-based fronthaul over a bridged network

Previously, we have demonstrated an Ethernet-based mobile front- and backhaul network in a field trial [9], where the prototyped 10 Gbit/s to 100 Gbit/s Ethernet aggregator employing the FUSION gap preservation scheme was able to minimize the PDV down to 100 ns for the high-priority PTP packets. To extend the work, we later demonstrated a fronthaul link consisting of three aggregation nodes [10], enabling bounded-delay aggregation for the fronthaul traffic combined with less delay sensitive backhaul traffic. PTP packets were guaranteed by means of gap preservation with a fixed latency.

4 Conclusion

It is foreseen that Ethernet will become the convergence layer for future 5G x-haul transport. Although native Ethernet has limitation on the time-sensitive applications, the TSN offers a powerful toolbox to ensure timing accuracy.

5 Acknowledgements

This work has been supported by the European Union’s Horizon 2020 Research and Innovation Program under Grant Agreement No. 762057 (5G-PICTURE).

6 References

[1] 3GPP TS 38.401: ‘NG-RAN; Architecture description’, July 2019

[2] NGMN Alliance: ‘NGMN Overview on 5G RAN Functional Decomposition’, Feb. 2018

[3] P. Sehier, P. Chanclou, N. Benzaoui, et al.: ‘Transport Evolution for the RAN of the Future’, J. Opt. Commun. Netw., 2019, 11, (4), pp B97-B108

[4] N. J. Gomes, P. Sehier, H. Thomas, et al.: ‘Boosting 5G Through Ethernet’, IEEE Veh. Technol. Mag., 2018, 13, (1), pp 74-84

[5] IEEE Std. 802.1Q-1998: ‘IEEE Standards for Local and Metropolitan Area Networks: Virtual Bridged Local Area Networks’, June 1998

[6] ITU-T Std. G.8262/Y.1362: ‘Timing characteristics of synchronous Ethernet equipment slave clock’, Nov. 2018

[7] IEEE Std. 1588-2008: ‘Precision clock synchronization protocol for networked measurement & control systems’, July 2018

[8] R. Veisllari, S. Bjornstad, J. P. Braute, et al.: ‘Experimental Demonstration of 100 Gb/s Optical Network Transport and Aggregation for Ethernet Fronthaul with Low and Bounded Delay’, Proc. OFC, 2018, pp. M4I.2

[9] J. Zou, J.-P. Elbers, R. Veisllari, et al.: ‘Field-Trial Evaluation of Low-Latency and Timing-Accurate 100G Ethernet Aggregator for Converged Mobile X-haul’, Proc. ECOC, 2018, pp. Tu3B.3

[10] J. Zou, R. Veisllari, V. Vassilev, J. P. Braute, et al.: ‘Live Demo of Low-Latency and Timing-Accurate Mobile X-haul Based on SDN-Enabled 100G Ethernet Aggregator’, Proc. ECOC, 2018, pp. TuDS.12