Pre-shaping Bursty Transmissions under IEEE802.1Q as a Simple and Efficient QoS Mechanism

Nicolas NAVET, University of Luxembourg
Jörn MIGGE, RealTime-at-Work (RTaW)
Josetxo VILLANUEVA, Groupe Renault
Marc BOYER, Onera

Abstract: The automotive industry is swiftly moving towards Ethernet as the high-speed communication network for in-vehicle communication. There is nonetheless a need for protocols that go beyond what standard Ethernet has to offer in order to provide additional QoS to demanding applications such as ADAS systems (Advanced Driver-Assistance Systems) or audio/video streaming. The main protocols currently considered for that purpose are IEEE802.1Q, AVB with the Credit Based Shaper mechanism (IEEE802.1Qav) and TSN with its Time-Aware Shaper (IEEE802.1Qbv). AVB/CBS and TSN/TAS both provide efficient QoS mechanisms and they can be used in a combined manner, which offers many possibilities to the designer. Their use however requires dedicated hardware and software components, and clock synchronization in the case of TAS. Previous studies have also shown that the efficiency of these protocols depends much on the application at hand and the value of the configuration parameters. In this work, we explore the use of "pre-shaping" strategies under IEEE802.1Q for bursty traffic such as audio/video streams as a simple and efficient alternative to AVB/CBS and TSN/TAS. Pre-shaping means inserting on the sender side "well-chosen" pauses between successive frames of a transmission burst (e.g., as it happens when sending a camera frame), all the other characteristics of the traffic remaining unchanged. We show on an automotive case-study how the use of pre-shaping for audio/video streams leads to a drastic reduction of the communication latencies for the best-effort streams while enabling meeting the timing constraints for the rest of the traffic. We then discuss the limitations of the pre-shaping mechanism and what is needed to facilitate its adoption.

Keywords: Time Sensitive Networking ; TSN ; IEEE802.1Q ; in-vehicle networks ; Credit-based Shaper ; IEEE802.1Qav ; traffic shaping ; Quality of Service.

Contents

1 Introduction

1.1 Context of the paper

There are currently several ongoing initiatives to design and implement QoS protocols on top of standard Ethernet that are of interest to the automotive industry. This can be explained by the need to support new and diverse in-vehicle communication requirements for audio/video and infotainment streams, command and control traffic, ADAS systems, etc. Among the prominent protocols considered for that purpose, IEEE802.1Q [1] which allows priority-based frame scheduling, AVB with the Credit-Based Shaper mechanism (IEEE802.1Qav) and TSN with its Time-Aware Shaper (IEEE802.1Qbv) as well as the frame preemption mechanisms (IEEE803.3br/802.1Qbu). The reader interested in a survey of the TSN standards related to low-latency communication, and the ongoing works within the TSN working groups can refer to [2].

1.2 Quality of Service protocols for Ethernet

Temporal Quality-of-Service (QoS) in full-duplex Ethernet implies managing the interfering traffic both in the nodes and in the switches. *Priorities*, as implemented in IEEE802.1Q with eight distinct priority levels, is a conceptually simple and widely used solution. Static priorities have been used for instance in AFDX networks deployed in planes for over a decade. Two inherent limitations of static-priority scheduling are that 1) it can lead to starvation for the lower-priority traffic and 2) it does not offer support for bandwidth reservation. A worst-case timing analysis of a set of streams scheduled with priorities is proposed in [3].

A first solution to overcome these issues is *time-triggered (TT) communication* where transmission time-windows are reserved to certain streams. Time-Sensitive Networking is a set of standards under development within the IEEE 802.1 working groups that includes the definition of QoS mechanisms. An important such mechanism is the Time-Aware Shaper (TAS, IEEE801.Qbv) enabling TT communication for a chosen subset of the traffic. The reader can consult [4] and [5] for a description and an analysis of TAS.

A different paradigm to manage the interferences between streams is the use of *traffic shaping* policies, that is delaying some packets, typically bursty video packets, to give bandwidth to lower priority streams. This is what is done in AVB with the Credit-Based Shaper (CBS) defined in IEEE801.Qav. The analysis of CBS mechanisms, and addressing its limitations, has been an active line of work during the last ten years. The reader can for instance refer to [6,7,8,9,10,11,12,13] and [14]. Another shaping policy related to CBS is the Asynchronous Traffic Shaping (IEEE802.1Qcr), which offers per-stream shaping instead of per-class shaping like CBS. ATS, at the time of writing, is still under development [15].

1.3 Limits of existing solutions

If the QoS protocols listed above are effective in certain contexts, they each possess drawbacks and limitations:

- The use of priorities alone leads to poor performance, *i.e.* important jitters and maximum latencies, and possibly starvation for the low-priority traffic (also referred to as best-effort traffic in the following). In addition, when the traffic is bursty, such as video streams, the memory needed in the switches to avoid packet losses can become important.
- As it is now well documented, AVB/CBS ensures much better performances for best-effort traffic but standard AVB classes are not sufficiently flexible to be an answer for all communication needs (see [16]). The use of AVB custom classes helps to get the most out

of AVB (see [16]) but it will not be always sufficient. In addition, defining the parameters for custom classes requires worst-case schedulability analysis and an optimization algorithm to set CBS IdleSlopes.
- TSN/TAS, especially when used in combination with CBS, provides a lot of possibilities but, to be efficient, the configuration of TAS gate scheduling tables must be done jointly for all senders and switches leads to a complex optimization problem. This problem, to the best of our knowledge, has only been partly addressed yet (see [17] for a starting point). Also, TSN/TAS requires a synchronization protocol to build and maintain a global clock, which induces some overhead and complexity, and reduce the overall robustness of the system. In addition, like in all TT protocols, for maximal freshness of the data in reception, there should be some form of synchronization between the production of the data by the tasks and the transmission of the frames on the network.

The transmission of segmented messages, such as ADAS video streams, changes the shape of the real-time streams and their associated timing constraints. Indeed, since a single message (e.g., a camera frame) is fragmented into several Ethernet frames, the evaluation of the latency of a single Ethernet frame is not suited to assess whether timing constraints are met. Except in a few works such as [18] in the context of CAN networks, this problem to the best of our knowledge has not been addressed in the performance evaluation of automotive networks.

1.4 Contributions of the paper

This work explores the use of what we refer to as the "pre-shaping" strategy for segmented messages under IEEE802.1Q. This mechanism, applied on the sending nodes on a per flow basis, is conceptually simple and easy to implement in software. Insights in the performance that can be expected from it are obtained through a case-study. The main result of this work is to show evidence that, in domains like automotive where the number of switches is small, simple low-overhead software-implemented shaping mechanisms can provide the same level of performances as AVB/CBS. Finally, we discuss the limitations of pre-shaping and its scope of applicability.

2 The Pre-shaping Mechanism

A noteworthy evolution in the traffic exchanged between automotive ECUs is that not only the number of messages but also their size steadily grow, leading to message fragmentation even on Ethernet. This is in particular due to increasing communication needs for audio, video and infotainment streams. For instance, in the case-study considered in the paper, there are several 30FPS cameras each generating a burst of 30 Ethernet frames with 1446 bytes of data every 33ms. These 30 Ethernet frames are making up a single camera frame. The timing constraints expressed as a deadline is on the last packet only, and not on each of the packet. The deadline is typically equal to the period of the message but it can be more stringent for streams used in ADAS for instance, or if decompression must take place on the receiving end.

The pre-shaping mechanism combines standard static priority scheduling with traffic shaping introduced by inserting idle times, pauses, between the times at which the successive frames of a segmented message are enqueued for transmission. All the other characteristics of the traffic remain unchanged. Pre-shaping allows lower or same priority frames that cross the path of pre-shaped stream to be transmitted sooner, taking advantage of the inserted idle times.

Pre-shaping is not targeted at improving the communication latency for the higher-priority traffic but it can be used in conjunction with frame preemption, configured so that the "pre-shaped" streams belong to the set of streams than can be preempted by high-priority frames. In the automotive context, pre-shaping can be implemented in software at the middleware level or communication driver level.

Following notations are needed to describe the system model:

- T is the period of the segmented message,
- N is the number of frames making up the message,
- D is the relative deadline of the message, that is the time after the release of the message by which the last frame of the message must have been received by all receiving stations,
- I is the idle time that is inserted between each frame of the message,
- E is the longest transmission time for a frame of a message ($E=L/C$, when C is the link speed and L the frame length, including the inter-frame gap and preamble).

The number of frames N forming the message depends on the data payload contained in each of the frames. This parameter can also be decided by the designer in the interval permitted by the protocol (i.e., 46 to 1500 bytes). Smaller data payloads induce higher overhead but in many cases will lead to less interferences to the rest of the traffic. The most simple and most practical approach, which is the one experimented in this study, is to not change the size of the frames and only use an idle time between successive frames of the message to implement traffic shaping.

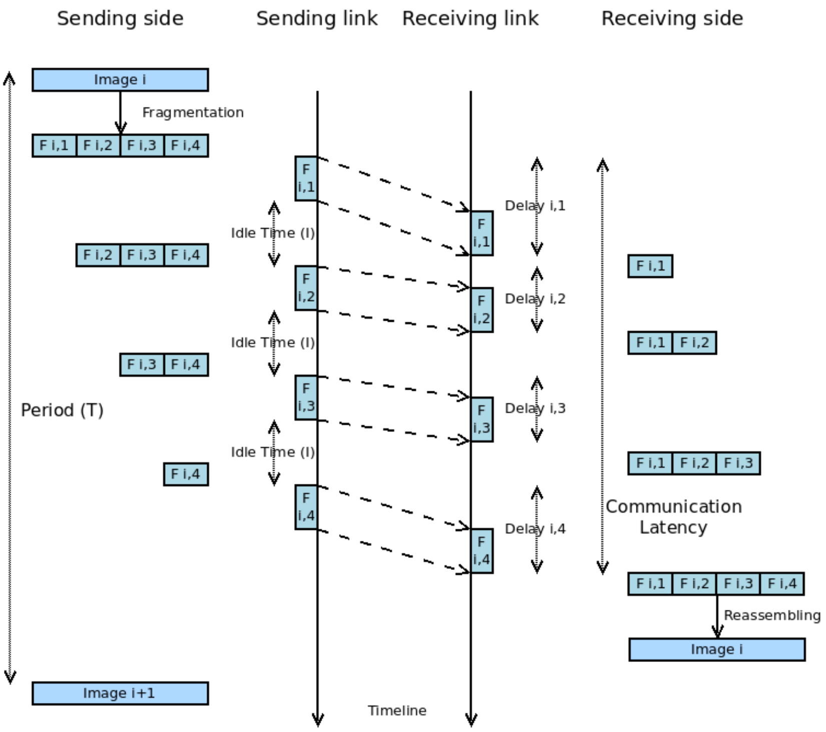

Figure 1. System model for the pre-shaping mechanism. A message, such as a camera frame, is transmitted with a period T. Each message is sent as N frames which are released for transmission each I time units. The last frame of the message will be released at time $(N-1)\cdot I$ and must be received by the deadline.

Considering these parameters, the last frame is enqueued $(I+E)\cdot(N-1)$ time units after the message release. Thus, if the communication latency of the last frame is bounded by $Rmax$, the idle time I must be chosen between 0 and $(D-Rmax)/(N-1) - E$. This latter upper bound spreads the successive transmissions over the longest time interval ensuring deadline respect, giving thus the maximum possible bandwidth to the frames located in lower priority traffic classes. This is the strategy underlying the *PRESH* algorithm, available in the tool *RTaW-Pegase*, that has been used in this study.

3 Case-study: Renault prototype Ethernet network

3.1 Topology and traffic

The case-study is a prototype Ethernet network comprising 5 switches and 14 nodes: 4 cameras, 4 displays, 3 control units and 3 (functional) domain masters, as shown in Figure 2. The data transmission rate is 100Mbit/s on all links except 1Gbit/s on link between domain master 3 (DM3) and switch 3.

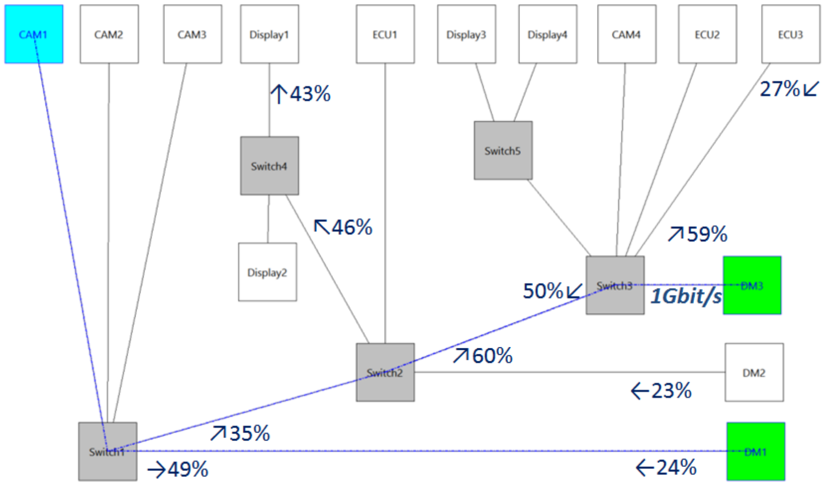

Figure 2. Topology of the prototype network used in the experiments. The multicast stream shown here goes from camera 1 to domain masters 1 and 3 (RTaW-Pegase screenshot). The graphic shows the 10 most loaded links, with a maximum of 60% load, and the single 1Gbit/s link.

The traffic is made up of four classes for a total of 41 streams whose characteristics are summarized in Table 1.

Table 1. Characteristics of the four types of traffic. The performance constraints is either to meet timing constraints (soft and hard deadline) or throughput constraints.

Audio streams	– 8 streams – 128 and 256 byte frames – up to sub-10ms periods and deadlines – soft deadline constraints
Video Streams	– 2 ADAS + 6 Vision streams – up to 30*1446byte frame each 16ms (60FPS) or each 33ms (30FPS) – 10ms (ADAS) or 30ms deadline (Vision) – hard and soft deadline constraints
Command & Control	–11 streams, 256 to 1024 byte frames – up to sub-10ms periods and deadlines – deadline constraints (hard)
Best-effort: File, data transfer, diagnostics	– 14 streams including TFTP traffic pattern – up to 0.2ms periods – both throughput guarantees (up to 20Mbits per stream) and deadline constraints (soft)

3.2 Verification techniques and protocols configuration

This study has been conducted using both timing-accurate simulation and worst-case traversal time (WCTT) analysis using a state-of-the-art network calculus implementation. Both techniques are complementary. Indeed, if WCTT is the safest approach, it is inherently pessimistic. In addition, it does not provide statistics such as the distribution of the latencies or, for instance, an accurate evaluation of the throughput that can achieved for FTP-like streams. The design and timing analysis tool used is *RTaW-Pegase v2.4.5* (see [19]), a product of the company RealTime-at-Work developed in partnership with ONERA research institute. The simulation samples were collected over long simulations (2 days of uninterrupted functioning, about 350 000 transmissions for the lowest frequency frames at 500ms) with the clock drift of each station set to a random value in ±200ppm.

In the rest of the study, we compare the performances of the following QoS protocols on the case-study:

- Static-Priority Ethernet without pre-shaping (referred to as IEEE802.1Q in the following) with priority allocation in decreasing priority order as follows: Command & Control (highest priority), then Audio, then Video, and finally Best-Effort streams at the lowest priority level.
- Static-priority Ethernet with pre-shaping (referred to as IEEE802.1Q with pre-shaping) for video-streams. The pre-shaping configuration has been done using the strategy described in paragraph “Pre-shaping mechanism” leading to the configuration shown in Figure 3 that meets all performance constraints. The priority allocation remains unchanged with respect to the solution without pre-shaping.
- AVB/CBS with custom classes, that is not using the standard 125/250us CMI and standard Idle Slopes which do not lead to a feasible solution (see [16]). CBS is used both in the switches and in the sending nodes. The CBS Idle Slopes on each output port along the path have been set with the *Tight Idle-Slope* algorithm implemented in RTaW-Pegase. This algorithm computes the smallest possible Idle-Slopes allowing to meet the timing constraints of AVB traffic, minimizing the interferences induced to lower-priority streams. In terms of priority, the audio streams are at the highest priority level (AVB top priority) followed by video streams (AVB second priority), then Command & Control, and finally best-effort streams. The AVB classes are placed at priority levels above the rest of the traffic, as, the IEEE802.1Q standard in force at the time of writing, imposes it (see [1]).

Name	Priority	MinDistance	MaxSize	Sender	Receiver
UC9	2	3 ms / 32 ms	10 x 1246 byte	DM3	Display2
UC8	2	1 ms / 32 ms	30 x 1446 byte	DM3	Display1
UC10	2	1 ms / 32 ms	30 x 1046 byte	DM3	Display3
UC11	2	1 ms / 32 ms	30 x 1046 byte	DM3	Display4
UC26	2	1 ms / 32 ms	30 x 1446 byte	CAM1	DM3
UC32	2	0,5 ms / 16 ms	30 x 1446 byte	CAM4	DM3
UC36	2	0,324 ms / 32 ms	30 x 1446 byte	CAM3	DM1
UC37	2	0,324 ms / 32 ms	30 x 1446 byte	CAM2	DM1

Figure 3. Pre-shaping configuration for the eight video streams. The first duration in the *MinDistance* column indicates the idle time between two packet transmissions, while the second duration is the time between two successive camera frames.

3.3 Average latencies for best-effort streams

Figure 4 shows the average communication latencies for all best effort streams with the three protocols under study. Compared to standard IEEE802.1Q (black curve on Figure 4), pre-shaping (red curve) improves the average latencies for best-effort streams by 54% on average, and up to 86%. Without pre-shaping, IEEE802.1Q is not a feasible solution since the throughput constraints for best-effort streams are not met. Both pre-shaping and AVB custom classes are feasible solutions here, and they perform almost identically for the average latencies of best effort streams. However, besides not requiring dedicated hardware, pre-shaping has the advantage over AVB that the Command and Control streams are sent at the highest priority level, which reduces their latencies. It is also beneficial with respect to the robustness of the system. Indeed, since the priority levels reflect the actual criticality of the streams, the critical streams at the highest priority levels will be better protected against low-priority streams not respecting their traffic contract (e.g., “babbling idiot” behavior).

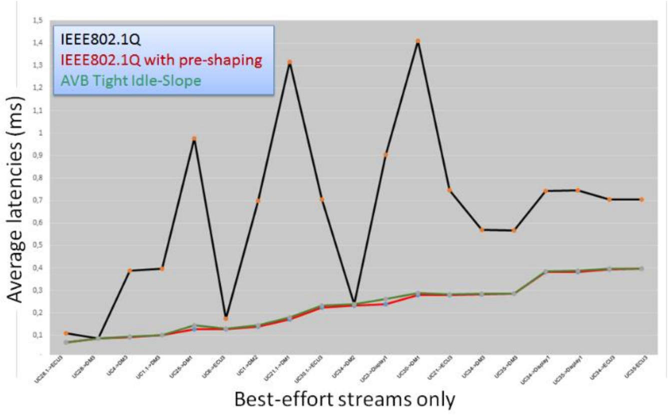

Figure 4. Average communication latencies for best effort under IEEE802.1Q, IEEE802.1Q with pre-shaping for video streams, and AVB/CBS configured with the tight idle-slope algorithm.

It should be noted that AVB/CBS and pre-shaping can be combined with TSN/TAS configured in such a way as to give exclusive bus access to command and control streams. The use of TAS however involves additional complexity in terms of configuration and requires dedicated hardware and software.

3.4 Worst-case latencies for best-effort streams

Figure 5 shows the worst-case communication latencies for all best effort streams. Pre-shaping under IEEE802.1Q improves worst-case latencies for best-effort streams by 66% on average, and up to 90%. Again, we observe similar performances between pre-shaping and AVB custom classes. This experiment shows that the variabilities of the latencies, and thus the jitters in reception, are also importantly reduced with pre-shaping.

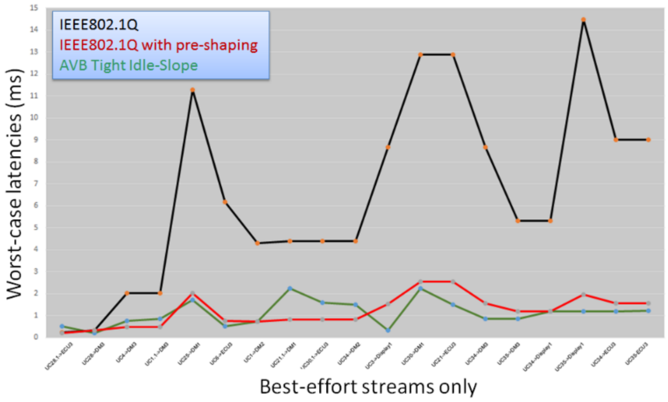

Figure 5. Worst-case communication latencies for best effort under IEEE802.1Q, IEEE802.1Q with pre-shaping for video streams, and AVB configured with the tight idle slope mechanism.

3.5 Impact on Command & Control traffic

We now study the impact of pre-shaping on the Command and Control traffic, which is of higher priority than the video streams under IEEE802.1Q (w/o pre-shaping) and at the immediate lower priority under AVB as AVB classes, in the current state of the standardization [1], must be at the top two priority levels.

Figure 6 shows the worst-case network traversal times (WCTT) and average network traversal times (AVRG) of the C&C streams under:

- IEEE802.1Q with and without pre-shaping,
- AVB/CBS for Audio/Video streams configured with the tight idle slope mechanism.

The relative priorities of the traffic classes are as defined in the "protocols configuration" paragraph. What we observe first is that pre-shaping has no impact on the WCTTs of the C&C traffic. This can be explained since the interference of lower-priority frames in the WCTT calculation is only through the blocking factor, that is the size of the largest lower priority frame whose value remains unchanged with pre-shaping. The WCTTs of C&C when AVB tight IdleSlope is used for audio/video streams are significantly larger than under IEEE802.1Q (+42% on average, and up to 129%). This can be explained by the interference brought by the AVB traffic classes, which are of higher priority than C&C traffic. In terms of the average communication latencies, keeping in mind that this is often not the most important metric for C&C frames, the three solutions performs very well and are almost equivalent.

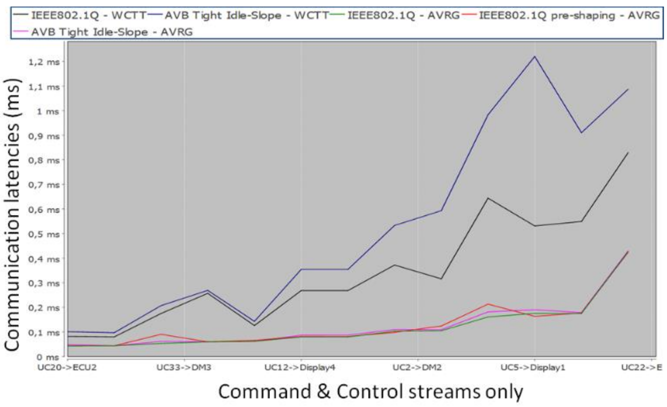

Figure 6. Worst-case and average communication latencies for Command and Control streams under IEEE802.1Q, IEEE802.1Q with pre-shaping for video streams, and AVB/CBS for Audio/Video configured with the tight idle-slope algorithm. The worst-case latencies for IEEE802.1Q with pre-shaping are all strictly equal to the ones obtained with pre-shaping.

3.6 Memory usage in the switches

It has been assumed so far that no packet loss occurs due to insufficient memory to store packets awaiting transmission, be it in end-systems or switches. In practice, dimensioning the amount of memory so that there is no packet loss is especially critical for switches. Figure 7 shows upper bounds on the memory usage in the output ports of the switches obtained by network-calculus analysis. AVB/CBS and IEEE802.1Q with pre-shaping, which both shape the traffic in an efficient manner lead to the lowest memory usage. On the other end of the spectrum, IEEE802.1Q without pre-shaping creates bursts of frames, which accumulate in the switches. IEEE802.1Q with pre-shaping in transmission improves the memory usage by a factor two on average over IEEE802.1Q without pre-shaping. AVB Tight Idle-Slope may insert delays between transmissions on egress ports and thus requires more memory than IEEE802.1Q with pre-shaping (+28% on average).

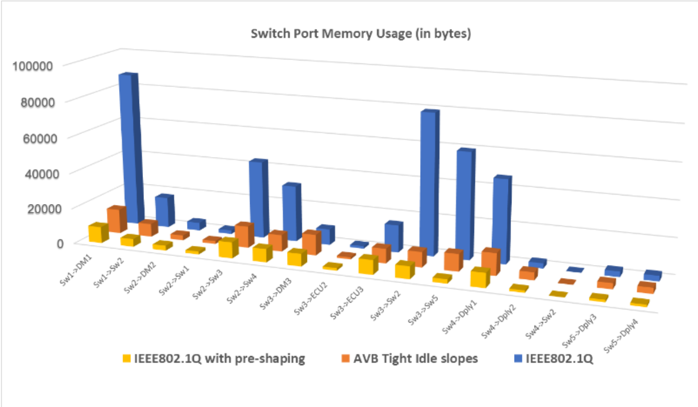

Figure 7. Upper bounds on the memory usage in the output ports of the switches with IEEE802.1Q (w/wo pre-shaping) and AVB Tight Idle slopes. The figures shown represent the sums for all the switches in the network.

4 Discussion & Conclusions

The experiments conducted on a realistic case-study shows that pre-shaping applied to streams generating burst of frames is an effective mechanisms to reduce the communication latencies of the lower-priority streams. In addition, pre-shaping does not require dedicated hardware and can be implemented in software with minimal overhead. In that regard, it shares similarities with the offsets mechanism in CAN (see [20]), which has been successfully used for years in the automotive industry.

If simple and effective, the pre-shaping policy with static-priority scheduling possesses some limitations:

- It does not offer protection against a "babbling idiot", that is a node that would send outside its specification. For instance, a node which, due to a hardware or software fault, would keep on sending frames and flood the network. Two solutions may be used: either a per class shaping, like with CBS in AVB, or a per stream shaping, like in AFDX or in PSFP (IEEE802.1Qci).
- Adding a new function or a new ECU, which results in adding frames to the system, may require a reconfiguration of the pre-shaping parameters for all the flows since the maximal communication latencies will change. This limitation is not specific to pre-shaping and affects most of the QoS protocols except standard AVB with AVB classes at the highest priority levels.
- When done manually by trial-and-error, setting the parameters for the flows subject to the pre-shaping mechanism is a time consuming task, and may not lead to optimal results. The process of setting parameters requires dedicated tool support.
- As there is no re-shaping along the path of a message, unlike for instance in AVB/CBS or TSN/TAS, the efficiency of the pre-shaping will decrease with the number of hops and thus with the size of the network.
- From the OEM perspective, pre-shaping imposes additional requests to ECU suppliers, which has a cost. However, just like transmission offsets in CAN, pre-shaping can be implemented only on a reduced subset of the nodes. For instance, in our case-study only 5 nodes out of 14 were using pre-shaping in transmission.

Proposing algorithms to choose the parameters of the pre-shaping mechanism discussed in this paper is to the best of our knowledge still an open problem. When there is a single stream per class on which pre-shaping is to be applied, a policy that is optimal in terms of meeting the deadlines is to start from the highest priority and set the idle times between transmissions to the longest possible value that still allows meeting the deadline. The idle times values derived with this strategy are however not robust to modifications of the stream sets: if higher or equal priority streams are added, or if a lower priority stream with larger frames is added then some deadlines will be missed. Further work includes thus proposing trade-offs between schedulability optimality and robustness to evolutions of the communication requirements that fit the OEM design process. More generally, there has been over the last 5 years many studies about the individual QoS protocols on top of Ethernet but the literature is still scarce on how to best configure them and use them in a combined manner. If the use of priorities without pre-shaping is now well understood, this is to a much lesser extent the case for the configuration of AVB's CBS parameters when outside the strict scope of SR-A and SR-B. Similarly, the strategies to use TSN/TAS (w/o CBS) and preemption mechanisms remains largely unexplored. Future work includes developing algorithms to automate the choice of configuration parameters considering all the communication constraints, as the "Zero-Config TSN" algorithm proposed in [21]. To ease an incremental design process and variants management, these configuration algorithms, should be able to integrate margins so as to allow the addition of new ECUs, switches and streams without requiring an entire reconfiguration of the communication architecture.

5 References

1. "IEEE Standard for Local and metropolitan area networks–Bridges and Bridged Networks," IEEE Std 802.1Q-2014 (Revision of IEEE Std802.1Q-2011), pp. 1–1832, Dec. 2014.
2. Nasrallah, A., Thyagaturu, A., Alharbi, Z., Wang, C., Shao X., Reisslein, M., ElBakoury, H., "Ultra-Low Latency (ULL) Networks: The IEEE TSN and IETF DetNet Standards and Related 5G ULL Research", IEEE Communications Surveys & Tutorials, in print, 2019. Available on arXiv.org as arXiv:1803.07673v2, 2018.
3. Diemer, J., Thiele, D. and Ernst, R., "Formal Worst-Case Timing Analysis of Ethernet Topologies with Strict-Priority and AVB Switching", 7th IEEE International Symposium on Industrial Embedded Systems (SIES'12), Karlsruhe, pp. 1-10, 2012.
4. Thiele, D., Ernst, R., and Diemer, J., "Formal worst-case timing analysis of Ethernet TSN's time-aware and peristaltic shapers", IEEE Vehicular Networking Conference (VNC), Kyoto, December 16-18, 2015.
5. Ashjaei, M., Patti, G. Behnam, M., Nolte, T., Alderisi, G. and Lo Bello, L., "Schedulability analysis of Ethernet Audio Video Bridging networks with scheduled traffic support", Real-Time Systems, vol. 53, pp526-577, 2017.
6. Cao, J., Cuijpers, P., Bril, R., and Lukkien, J., "Independent yet Tight WCRT Analysis for Individual Priority Classes in Ethernet AVB", Proc. 24th Intl. Conference on Real-Time Networks and Systems (RTNS '16), Brest, Oct. 2016.
7. Queck, R., "Analysis of Ethernet AVB for automotive networks using Network Calculus", IEEE International Conference on Vehicular Electronics and Safety (ICVES), Istanbul, July, 2012.
8. Li, X., George, L., "Deterministic delay analysis of AVB switched Ethernet networks using an extended Trajectory Approach", Real-Time Systems, 53(1), pp 121–186, 2016.
9. Maxim, D., Song, Y.-Q., "Delay Analysis of AVB traffic in Time-Sensitive Networks (TSN)", International Conference on Real-Time Networks and Systems (RTNS2017), Grenoble, 2017.
10. He, F., Zhao, L. and Li, E., "Impact Analysis of Flow Shaping in Ethernet-AVB/TSN and AFDX from Network Calculus and Simulation Perspective", Sensors 2017, MDPI, 17(5), 1181, https://doi.org/10.3390/s17051181, 2017.
11. Meyer, P., Steinbach, T., Korf, F. and Schmidt, T. C., "Extending IEEE 802.1 AVB with Time-triggered Scheduling: A Simulation Study of the Coexistence of Synchronous and Asynchronous Traffic", 2013 IEEE Vehicular Networking Conference, doi: 10.1109/VNC.2013.6737589, Boston, MA, pp. 47-54, 2013.
12. Imtiaz, J., Jasperneite J. and Han, H., "A performance study of Ethernet Audio Video Bridging (AVB) for Industrial real-time communication," 2009 IEEE International Conference on Emerging Technologies & Factory Automation (ETFA2009), doi: 10.1109/ETFA.2009.5347126, Mallorca, 2009.
13. Lim, H., Herrscher D. and Chaari, F., "Performance comparison of IEEE 802.1Q and IEEE 802.1 AVB in an Ethernet-based in-vehicle network," 8th International Conference on Computing Technology and Information Management (NCM and ICNIT), Seoul, 2012.
14. Alderisi, G., Iannizzotto G., and Bello, L. L., "Towards IEEE 802.1 Ethernet AVB for Advanced Driver Assistance Systems: A preliminary assessment," 2012 IEEE International Conference on Emerging Technologies & Factory Automation (ETFA2012), doi: 10.1109/ETFA.2012.6489775, Krakow, 2012.
15. IEEE P802.1Qcr™ Draft Standard for Local and Metropolitan Area Networks –Bridges and Bridged Networks Amendment: Asynchronous Traffic Shaping, D0.5, June 18, 2018.

16. Migge, J., Villanueva, J., Navet, N. and Boyer, M., “Insights on the performance and configuration of AVB and TSN in automotive networks”, Proc. Embedded Real-Time Software and Systems (ERTSS 2018), Toulouse, France, January, 2018.
17. Serna Oliver, R., Craciunas, S. and Steiner, W., “IEEE 802.1Qbv Gate Control List Synthesis using Array Theory Encoding”, IEEE Real-Time and Embedded Technology and Applications Symposium (RTAS), doi: 10.1109/RTAS.2018.00008, Porto, April 2018.
18. Azketa, E., Javier Gutiérrez, J., Carlos Palencia, J., González Harbour, M., Almeida, L., and Marcos M. ,“Schedulability Analysis of Multi-Packet Messages in Segmented CAN”, in IEEE 17th Conference on Emerging Technologies & Factory Automation (ETFA), 2012, doi: 10.1109/ETFA.2012.6489578.
19. Navet, N., Seyler, J. and Migge, J., "Timing verification of real-time automotive Ethernet networks: what can we expect from simulation?", Embedded Real-Time Software and Systems (ERTS 2016), Toulouse, France, January 27-29, 2016.
20. Grenier, M., Havet, L., Navet, N., "Scheduling Frames with Offsets in Automotive Systems: a Major Performance Boost", Chap. 14 in the Automotive Embedded Systems Handbook, Taylor and Francis, ISBN 978-0849380266, December 2008.
21. Migge, J., Villanueva, J., Navet, N. and Boyer, M., “Performance assessment of configuration strategies for automotive Ethernet quality-of-service protocols”, Automotive Ethernet Congress, Munich, January 30-31, 2018.