



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

Nicolas Navet, Jörn Migge, Josetxo Villanueva, Marc Boyer

## ► To cite this version:

Nicolas Navet, Jörn Migge, Josetxo Villanueva, Marc Boyer. Pre-Shaping Bursty Transmissions under IEEE802.1Q as a Simple and Efficient QoS Mechanism. SAE International Journal of Passenger Cars - Electronic and Electrical Systems, 2018, 11 (3), 10.4271/2018-01-0756 . hal-02468313

**HAL Id: hal-02468313**

**<https://hal.science/hal-02468313>**

Submitted on 5 Feb 2020

**HAL** is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

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

Nicolas Navet, Jorn Migge, Josetxo Villanueva, Marc Boyer

## 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 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 burst (e.g., a camera frame), all the other characteristics of 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 to meet the timing constraints for the rest of the traffic. We then discuss the limitations of the pre-shaping mechanism and future works needed to facilitate its adoption.

## Introduction

### Context of the paper

There are currently many ongoing initiatives in the automotive industry to design and implement QoS protocols on top of standard Ethernet. This can be explained by the need to support new and diverse in-vehicle communication requirements for audio/video streams, infotainment command and control traffic, ADAS systems, etc. Among the prominent protocols considered for that purpose, IEEE802.1Q which allows the use priority levels for the streams, AVB with the Credit-Based Shaper mechanism (IEEE802.1Qav) and TSN with its Time-Aware Shaper (IEEE802.1Qbv) as well as frame preemption extensions (IEEE803.3br/802.1Qbu).

## 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 8 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 1) that it can lead to starvation for the lower-priority traffic and 2) it does not offer support for bandwidth reservation.

A first solution to overcome these issues is to use *traffic shaping* policies, this is what is done in AVB with the Credit-Based Shaper (CBS) defined in IEEE 801.Qav. The reader can consult [1,5] and [6] for an analysis of the CBS mechanisms.

A different paradigm to manage the interferences between streams is *time-triggered (TT) communication* where time-windows for transmissions are reserved to certain streams. Time-Sensitive Networking is a set of standards under development within the IEEE 802.1 working group 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.

## 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.* large jitters, large maximal delays 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 be 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 [3]). For example, the CMI of AVB class A being 125 $\mu$ s, then it is not possible to specify a flow with a throughput less than 4Mb/s. The use of AVB custom classes helps to get the most out of AVB (see [3]) but, in many cases, it will not be sufficient. In addition, setting the parameters for custom classes requires worst-case schedulability analysis and an optimization algorithm to set CBS IdleSlopes.

- TSN/TAS, especially when used combined 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 which is a complex optimization problem. This problem, to the best of our knowledge, is only partly addressed yet. In addition, 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. 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 on 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 [2], this problem to the best of our knowledge has not been addressed in the performance evaluation of automotive networks.

### Contributions of the paper

This work explores the use of what we refer to as “pre-shaping” strategies for segmented messages under IEEE802.1Q. This strategy, 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. Finally, we discuss the limitations of the technique and its scope of applicability.

### 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 even on Ethernet to message fragmentation. 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 instance for streams used in ADAS, or when decompression must take place on the receiving ends.

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.

If 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 where pre-shaped streams belong to the streams than can be preempted. In the automotive context, pre-shaping can be implemented in software at the middleware 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 basic 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.

Considering these parameters, the last frame is sent  $(I+E)/(N-I)$  time units after the message release. Thus, if the communication latency of the last frame is bounded by  $R_{max}$ , the idle time  $I$  must be chosen between 0 and  $(D-R_{max})/(N-I) - E$ .

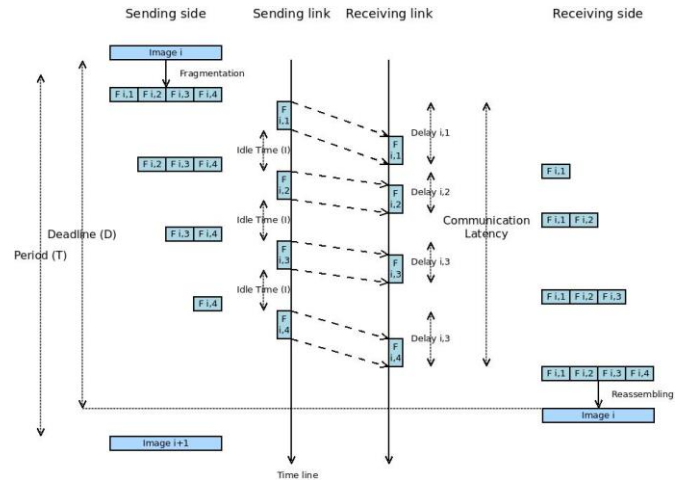


Figure 1. System model for the pre-shaping mechanism. A message, such as a camera frame, is transmitted with a period  $T$ . Each message is transmitted as  $N$  frames which are released for transmission each  $I$  time units. The last frame of the message will be released at time  $(N-1) \cdot I$ 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

### Case-study: Renault prototype Ethernet network

#### Topology and traffic

The case-study is a prototype Ethernet network made 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.

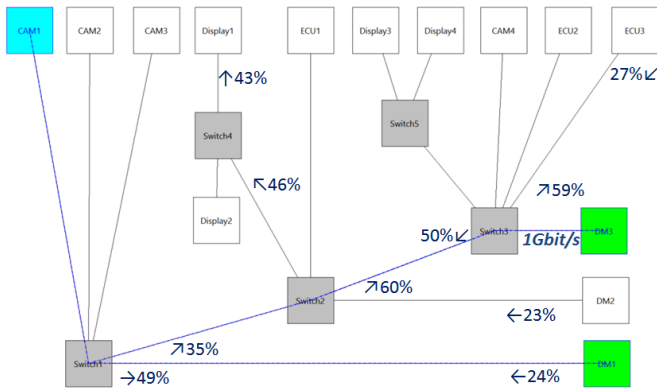


Figure 2. Topology of the prototype network used in the experiments. The multicast stream shown here goes from camera 1 to domain masters 1 and 3 (RTaW-Pegase screenshot). The graphic shows the 10 most loaded links, with a maximum of 60% load, and the single 1Gbit/s link.

The traffic is made up of 4 classes for a total of 14 streams whose characteristics are summarized in Table 1.

Table 1. Characteristics of the 4 types of traffic. The performance constraints is either to meet timing constraints (soft and hard deadline) or throughput constraints.

Audio streams	<ul style="list-style-type: none"> <li>– 8 streams</li> <li>– 128 and 256 byte frames</li> <li>– up to sub-10ms period and deadline</li> <li>– soft deadline constraints</li> </ul>
Video Streams	<ul style="list-style-type: none"> <li>– 2 ADAS + 6 Vision streams</li> <li>– up to 30*1446byte frame each 16ms (60FPS) or each 33ms (30FPS)</li> <li>– 10 ms or 30 ms deadline</li> <li>– hard and soft deadline constraints</li> </ul>
Command & Control	<ul style="list-style-type: none"> <li>– 11 streams, 256 to 1024 byte frames</li> <li>– up to sub-10ms period and deadline</li> <li>– deadline constraints (hard)</li> </ul>
Best-effort: File, data transfer, diagnostics	<ul style="list-style-type: none"> <li>– 14 streams, TFTP traffic pattern</li> <li>– Up to 0.2ms period</li> <li>– Throughput guarantee: up to 20Mbits</li> </ul>

## 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 be achieved for FTP-like streams. The design and timing analysis tool used is RTaW-Pegase v2.4.5 (see [7]), a product of 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  $\pm 200$ ppm.

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 (also referred to as IEEE802.1Q in the following) with priority allocation as follows in decreasing priority order: Command & Control, Audio, Video, and best-effort streams at the lowest priority level.
- Static-priority Ethernet with pre-shaping (also referred to as IEEE802.1Q with pre-shaping) for video-streams. The pre-shaping configuration has been done manually until reaching the configuration shown in Table 2 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 [2]). 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 minimal Idle-Slopes allowing to meet the timing constraints of AVB traffic, minimizing thus 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.

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 8 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.

## Average latencies for best-effort streams

Figure 4 shows the average communication latencies for all best effort streams with the 3 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. In our view, this also improves the robustness of the system since the priority levels reflect the actual criticality of the streams.

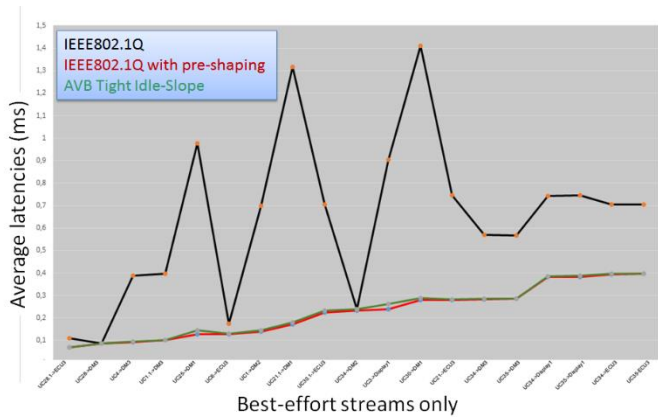


Figure 4. Average communication latencies for best effort under IEEE802.1Q, IEEE802.1Q with pre-shaping for video streams, and AVB configured with the tight idle slope mechanism.

It should be noted that AVB/CBS can be combined with TSN/TAS configured so as to give exclusive bus access to command and control streams. In this case, on the same case-study, the use of CBS combined with TAS allow to outperforms pre-shaping and CBS alone (see [2]). The use of TAS however involves additional complexity in terms of configuration and requires dedicated hardware and software.

### 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 variability of the latencies, and thus the jitters in reception, are also importantly reduced with pre-shaping.

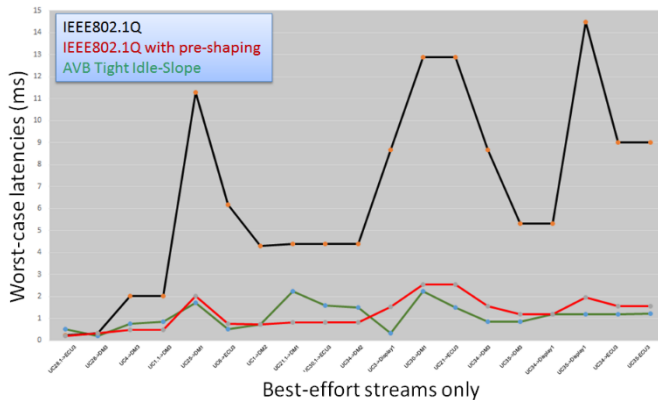


Figure 5. Worst-case communication latencies for best effort under IEEE802.1Q, IEEE802.1Q with pre-shaping for video streams, and AVB configured with the tight idle slope mechanism.

### 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 must, in the current state of the standardization, be at the top two priority levels.

Figure 6 shows the worst-case network traversal times (WCTT) and average network traversal time (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 pre-shaping has no impact on the WCTTs of the C&C traffic with respect to IEEE801.Q without pre-shaping. 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.

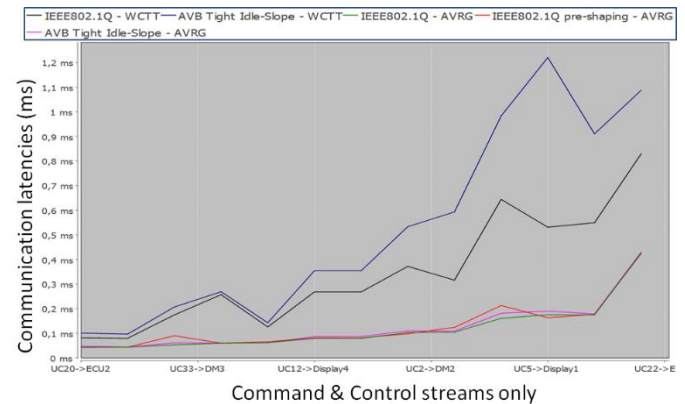


Figure 6. Worst-case and average communication latencies for Command and Control streams under IEEE802.1Q, IEEE802.1Q with pre-shaping for video streams, and AVB/CBS for Audio/Video streams configured with the tight idle slope mechanism. The worst-case latencies for IEEE802.1Q with pre-shaping are all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

### 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 [4]), which has been successfully used for years in the automotive industry.

If simple and effective, the pre-shaping policy with static-priority policy 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 TSN, or a per stream shaping, like in AFDX or in PSFP (802.1Qci).

- 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 but standard AVB.
- The deadline of a message may impose to assign a high priority to a stream to meet the timing constraints while the stream is not important from a safety point of view. This can be typically the case for audio/video streams, while control laws governing the dynamics of the vehicle may tolerate longer communication latencies. This limitation affects also schemes based on static priority.
- Setting the parameters for the flows subject to the pre-shaping mechanism is a time consuming task when done by trial-and-error. To the best of our knowledge, there is no guidelines, such as optimality results, available to guide the designer in this task. The process of setting parameters should be automated which requires further studies and specific 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.

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, while 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. 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.

## References

1. 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.
2. 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.
3. Navet, N., Villanueva, J., Migge, J., and Boyer, M., "Experimental assessment of QoS protocols for in-car Ethernet networks", 2017 IEEE Standards Association (IEEE-SA) Ethernet & IP @ Automotive Technology Day, San-Jose, Ca, October 31-November 2, 2017.
4. Grenier, M., Havet, L., Navet, N., "Scheduling Frames with Offsets in Automotive Systems: a Major Performance Boost", chapter 14 in the Automotive Embedded Systems Handbook, CRC Press / Taylor and Francis, ISBN 978-0849380266, December 2008.
5. R. Queck, R., "Analysis of Ethernet AVB for automotive networks using Network Calculus", 2012 IEEE International Conference on Vehicular Electronics and Safety (ICVES), Istanbul, July, 2012.
6. Li, X., George, L., "Deterministic delay analysis of AVB switched Ethernet networks using an extended Trajectory Approach", Real-Time Systems, Volume 53, Issue 1, pp 121–186, 2016.
7. Navet, H., 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.

## Contact Information

Nicolas Navet  
email: [nicolas.navet@uni.lu](mailto:nicolas.navet@uni.lu)  
web: <http://nicolas.navet.eu>

## Definitions/Abbreviations

<b>AVB</b>	Audio Video Bridging
<b>CBS</b>	Credit-Based Shaping
<b>SFS</b>	Per-Stream Filtering and Policing
<b>TSN</b>	Time Sensitive Networking
<b>TAS</b>	Time Aware Shaper