Novel trends in Automotive Networks: A perspective on Ethernet and the IEEE Audio Video Bridging

Lucia Lo Bello

Department of Electrical, Electronic and Computer Engineering
University of Catania
Catania, Italy
lucia.lobello@dieei.unict.it

Abstract—Ethernet is going to play a major role in automotive communications, thus representing a significant paradigm shift in automotive networking. Ethernet technology will allow for multiple in-vehicle systems (such as, multimedia/infotainment, camera-based advanced driver assistance and on-board diagnostics) to simultaneously access information over a single unshielded twisted pair cable. The leading technology for automotive applications is the IEEE Audio Video Bridging (AVB), which offers several advantages, such as open specification, multiple sources of electronic components, high bandwidth, the compliance with the challenging EMC/EMI automotive requirements, and significant savings on cabling costs, thickness and weight. This paper surveys the state of the art on Ethernet-based automotive communications and especially on the IEEE AVB, with a particular focus on the way to provide support to the so-called scheduled traffic, that is a class of time-sensitive traffic (e.g., control traffic) that is transmitted according to a time schedule.

Keywords—Automotive communications; Ethernet; IEEE Audio Video Bridging (AVB); Scheduled traffic.

1. Introduction

In modern cars, a number of novel applications are steadily being introduced in multiple functional domains, especially for multimedia/infotainment applications (e.g., CD audio, DVD players) and for camera-based Advanced Driver Assistance Systems (ADASs), which provide services such as, Lane Departure Warning, Traffic Sign Recognition, Night vision, bird's-eye view, just to mention a few of them. These new applications call for bounded latencies and increased bandwidth. Existing automotive networks, for instance, CAN or FlexRay, cannot provide enough bandwidth for these applications. In addition, the concept of an in-car backbone interconnecting the gateways of the different functional domains for cross-domain communication is now broadly accepted by the major car makers and automotive suppliers.

In this context, for three years now, it was envisaged that Ethernet, hitherto used only in diagnostic applications, was destined to play a major role in automotive communication [1]. This is because Ethernet allows for design flexibility and scalability, independence and freedom of choice thanks to the large number of available technologies, protocols, manufacturers and system component vendors, and offers the potential for IP-based in-car communications. Nowadays Ethernet has become a recognized network technology for in-car communications, for adoption in specific domains, like

multimedia/infotainment and ADASs, and also as an in-car backbone, thanks to the technology advancements and the solutions already available. One example of existing technology is the Open Alliance BroadR-Reach (OABR), which operates with bidirectional transmission over one-pair Unshielded Twisted Pair (UTP) and provides a robust and cost-effective solution at 100 Mb/s that meets the automotive electromagnetic compatibility (EMC) and electromagnetic interference (EMI) requirements and also the IEEE audio/video time synchronization streaming requirements.

The OPEN Alliance [15] established the OABR specification as an open industry and de facto standard, which is also validated by multiple Original Equipment Manufacturers (OEMs). Multiple sources for OABR are definitely possible: NXP presented its PHY solution and it is likely that others will do the same in the future. In addition, there is also a roadmap for higher data rates, i.e., the IEEE 802.3bp Reduced Twisted Pair Gigabit Ethernet (RTPGE) at 1Gb/s, whose specification is currently in progress within the IEEE.

In-car Ethernet is already a reality today and technical solutions are already available. In the BMW time line, the IEEE Audio Video Bridging was introduced in series cars in 2013, the new X5 being the pilot vehicle. The target vehicle, according to the BMW roadmap, is the new 7 series. In fact, the IEEE Audio Video Bridging (AVB) standard [2][3][4] allows for efficient transport of audio/video traffic with a guaranteed maximum latency. This desirable property, together with other advantages in terms of scalability, cost and performance, is driving the gradual replacement, in the multimedia/infotainment domain, of the Media Oriented Systems Transport (MOST) protocol [5], the current de-facto standard for such applications, with AVB. For similar reasons, i.e., reduction of wiring harness and lower weight, deployment, and maintenance costs, AVB is also going to replace the Low-Voltage Differential Signaling (LVDS) cables that are currently used for camera-based ADASs.

The AVB standard [3] allows for the reservation of resources (buffers, queues) within switches along the path between sender and receiver, and defines two Stream Reservation (SR) classes: Class A, which provides a maximum latency of 2ms, and Class B, which provides a maximum latency of 50ms, for seven hops within the network. SR traffic

undergoes traffic shaping at the output ports of switches and end nodes to prevent traffic bursts.

Driven by the industry interest in assessing the performance of AVB for automotive applications, several recent works in the literature addressed the AVB capabilities through simulations or analysis [6-11][22][23], pinpointing strengths and weaknesses. In [6] the performance of AVB and TTE in a tree-based network under a workload consisting of both control and streaming traffic are assessed. AVB proved to be less deterministic than TTE under heavy load of full-size frames. However, AVB is more flexible and does not require expensive specialized hardware as TTE does.

The study in [8] addresses the performance of AVB and TTE in two different star-based topologies, in a scenario including several traffic types, such as, ADASs, multimedia, and infotainment, under a significantly high overall network load for ADAS traffic. The results obtained by simulation show that both protocols fit the latency requirements of future ADAS systems. The work in [22] shows that the IEEE 802.1AS can provide the same clock quality as FlexRay.

The analysis in [9] describes two effects that have an impact on the latency estimation of AVB. They are the "own-priority and higher-priority blocking", that occurs when several streams share the same port, and the "shaper blocking", that is relevant to the large blocking times that a flow may experience due to traffic shaping. The study in [23] analyses the forwarding policy for AVB automotive networks using the Network Calculus for calculating the worst case delays. The work [25] addresses the additional delay in AVB due to the shaper at the highest priority level.

Despite AVB allows to calculate the worst case latency for all the real-time message classes, further improvements are still needed, as new needs have recently arisen. The Time-Sensitive Networking Task Group of IEEE 802.1 (TSN TG) is therefore working on AVB extensions. Several amendments are in progress within the activities of this group, which deal with enhanced time synchronization, robustness, redundancy, stream reservations, and support for time-sensitive and scheduled traffic.

Scheduled traffic is a traffic class that require to schedule frame transmission based on timing derived from the IEEE 802.1AS standard [2]. This traffic class is typically time-sensitive and requires to be delivered with a bounded latency, so any interference from other traffic classes should be avoided. Scheduled traffic includes, for instance, high-priority delay-sensitive command and control traffic, which requires deterministic and very small delays. To achieve the challenging objective of providing support to this kind of traffic, novel features have to be added to the ones already provided by the AVB standard.

Several approaches are being investigated and discussed both within the TSN TG and in the scientific community to provide support for scheduled traffic over IEEE AVB networks. The aim of this paper is to offer a broad perspective on the adoption of Ethernet as in-car network and on the role of the IEEE AVB in this scenario, with a specific focus on how to support scheduled traffic.

The paper is organized as follows. Sect. II discusses the current state as far as Ethernet in automotive communications is concerned, encompassing current achievements and outlining the next steps. Sect. III addresses the IEEE AVB standard and outlines the ongoing standardization work within the IEEE 802.1 Working Group. Sect. IV reports on cutting-edge results from the IEEE 802.1Qbv project on providing low latency within enhanced AVB bridges. In Sect. V an approach recently proposed in the literature to support scheduled traffic over AVB networks, called AVB_ST [12], is summarized and a performance comparison between such an approach and standard AVB in a realistic automotive scenario is presented. Finally Sect. VI concludes the paper and outlines the ongoing work, within the IEEE P802.1Qbv project [13][14], that addresses enhancements to existing protocols and procedures to enable bridges and end stations to support scheduled traffic.

2. In-car Ethernet

This Section provides a picture of the current state of the art as far as the introduction of Ethernet as in-car network is concerned. The BroadR-Reach Automotive™ PHY, that is the de facto standard for in-car Ethernet is addressed. Current trends and next steps are also outlined.

A. Where we are today

The BroadR-Reach Automotive™ PHY realizes bidirectional transmission over one-pair UTP and is the standard PHY for in-car Ethernet. This technology proved to meet the challenging EMC/EMI requirements of automotive networks. In addition, it complies with the IEEE Audio/Video time synchronization streaming constraints. In addition, the significant savings on cabling costs, thickness and weight provided by BroadR-Reach make it a very cost-effective solution comparing, for instance, with MOST for multimedia/infotainment and LDVS for ADASs.

Multi-sourcing in the automotive industry is mandatory, not only for costs, but also for reliability and availability considerations. For this reason, Broadcom, the owner of the BroadR-Reach, agreed to license the technology under Reasonable And Non-Discriminatory terms (RAND). The factual mean used for licensing was the establishment of a Special Interest Group (SIG), called the OPEN (One-Pair Ether-Net) Alliance [15], a non-profit, open industry alliance including automotive industry and technology providers that aims at fostering the adoption of Ethernet-based networks as the standard in automotive networking applications.

Since the release of the Open Alliance BroadR-Reach Ethernet (OABR) specification, the number of members of the Open Alliance is steadily growing (over 200 members in 2013). As a result, today multiple sources of OABR-compliant automotive electronics products are available on the market (e.g., a number of component vendors already provide the needed cables, connectors, magnetics, etc.) and new players entered the game, like NXP, that is now the second source for the OABR PHY, and Freescale.

B. Next steps

As far as Ethernet in cars is concerned, the next steps deal with achieving higher data rates and energy saving.

The need for high data rates grows with the bandwidth demands of novel automotive applications, such as camera-based ADASs. When operating at high data rates over UTP cabling, several factors have to be considered in the next generation Ethernet PHY, such as internal noise, alien crosstalk from other cables in a bundle, external interference (e.g. radio-frequency interference) [16]. Higher data rates are addressed by the current work on the next-generation automotive Gigabit PHY, i.e., the IEEE 802.3bp Reduced Twisted Pair Gigabit Ethernet (RTPGE), whose specification is currently in progress within the IEEE 802.3bp Task Force.

The next generation of BroadR-Reach technology will include novel features to address energy-efficient operation. In fact, energy saving is an important feature for in-car networks, as it allows for both less fuel consumption for internal-combustion engine vehicles and longer reach for electric vehicles. As indicated in [17], solutions are sought to achieve significant energy savings during periods of inactivity while still complying with the harsh EMC/noise requirements. The work in [24] introduces a Low Frequency Wakeup physical-layer mechanism that is quite effective in reducing the wake-up latency and power consumption of point-to-point automotive networks while offering partial networking support.

One solution provided today is the Power over BroadR-Reach Ethernet, that uses a single 1-pair UTP cable for power and data, thus saving on wiring, weight, and number of connectors. In [16] a use case is described, in which ADAS cameras are powered by Power over BroadR-Reach Ethernet, which enables to monitor the power drained by each network device and to selectively switch on/off cameras and other devices whenever needed, for efficient energy management and overload protection. A further step is represented by the activities relevant to the IEEE P802.3bu 1-Pair Power over Data Lines (PoDL) Task Force (formerly known as the IEEE 802.3 1-Pair Power over Data Lines, 1PPoDL), which aims at extending the benefits of Power over Ethernet to the P802.3bp Reduced Twisted Pair Gigabit Ethernet (RTPGE).

In [17] novel sleep/wakeup concepts allowing for partial networking techniques in automotive networks are described. Among several solutions, the global wakeup via Ethernet line is indicated as the preferred option for partial networking and is presented as possible extension to BroadR-Reach specifications, also extendable to RTPGE.

3. IEEE Audio Video Bridging

The IEEE 802.1 AVB standard consists of a set of technical standards defined by the IEEE 802.1 Working Group that provide the specification for time-synchronized low latency streaming services through IEEE 802.1 networks.

The main specifications are the following:

- IEEE 802.1AS: Timing and Synchronization for Time-Sensitive Applications [2].

- IEEE 802.1Qat: Stream Reservation Protocol [3].

- IEEE802.1Qav: Forwarding and Queuing Enhancements for Time-Sensitive Streams [4].

The IEEE 802.1AS Time Synchronization protocol is a variation of the IEEE 1588 [17] standard. It provides precise time synchronization of the network nodes to a reference time with an accuracy better than 1 us (typical implementations are better than ±300 ns for 7 hops).

The IEEE 802.1Qat Stream Reservation allows for the reservation of resources, such as buffers and queues, within bridges (switches are called bridges in the AVB terminology) along the path between the talker (i.e., the end station that is the source of a stream) and the listener (i.e., the end station that is the destination of a stream).

The IEEE 802.1Qav Queuing and Forwarding for AV Bridges separates time-critical and non-time-critical traffic into different traffic classes extending methods described in the IEEE 802.1Q standard. It also applies a Credit-Based Shaper (CBS) algorithm that performs traffic shaping at the output ports of bridges and end nodes to prevent traffic bursts.

For seven hops within the network, the AVB standard guarantees a fixed upper bound for latency. Two Stream Reservation (SR) classes are defined, i.e.,

- Class A, that provides a maximum latency of 2ms

- Class B, that provides a maximum latency of 50ms.

Each traffic class has an associate *credit* parameter, whose value changes within two limits, called *loCredit* and *hiCredit*, respectively. During the transmission of frames belonging to a given class, *credit* decreases at the *sendSlope* rate defined for that class. Conversely, *credit* increases at the constant rate *idleSlope* defined for the class when the frames of that class are waiting for the transmission or when no more frames of the class are waiting, but *credit* is negative. If *credit* is greater than zero and no more frames of the corresponding traffic class are waiting, *credit* is immediately reset to zero. More details on the AVB CBS are found in [4], Annex L.

AVB can guarantee low jitter if a careful planning of periodic execution and a suitable mapping to the high priority queues within bridges is made. Thanks to the resource reservation protocol, that is able to dynamically handle Quality of Service (QoS), new devices can join the network at any time.

C. On-going standardization activities

The standardization process of the IEEE Audio Video Bridging Generation 2 protocol is in progress. Several AVB improvement projects are ongoing. Among them:

• IEEE 802.1ASbt, which specifies enhancements to IEEE 802.1AS standard, including support for new media types and additional parameter sets for non-Audio/Video applications, e.g., industrial control.

• IEEE 802.1Qca, which extends the IEEE 802.1Qat standard to provide redundancy for data flows and distribution of control parameters for time synchronization and scheduling.

• IEEE 802.1Cb, which specifies protocols for bridges and end stations that provide redundant transmission and identification/elimination of duplicate frames.

• IEEE 802.1Qbv [13], which deals with enhancements to support scheduled traffic. The aim of the project is to define policies for both bridges and end stations to schedule frame transmissions based on timing derived from the IEEE 802.1AS standard.

• IEEE 802.1Qbu, which defines procedures to suspend the transmission of a non-time critical frame and allows for one or multiple time-critical frames to be transmitted. When the time-critical frames have been transmitted, the transmission of the preempted frame is resumed.

• IEEE 802.1Qcc, an amendment to the IEEE 802.1Qat standard that describes protocols to provide support for Configurable Stream Reservation (SR) classes and streams.

In this paper special attention to the IEEE 802.1Qbv will be paid. In fact, further improvements are still needed at both bridges and end stations for achieving the support for scheduled traffic and the latency reduction needed for automotive communications. In particular, the problem of interfering traffic flows has to be solved.

4. The IEEE 802.1Qbv Bridge

To get the lowest possible latency, special bridge design is needed and the whole network has to be properly engineered, managed and controlled, as there is no way to guarantee low latency in non-managed networks. In the following, recent results about the IEEE 802.1Qbv bridges are summarized [19].

D. Latency assessments

Since the IEEE 802.1 Qbv project started, much attention has been paid to quantify the maximum latency for time-sensitive traffic. A detailed analysis (given in [20]) that encompasses all the factors contributing to the delay, clearly showed that, without interfering traffic, it is the maximum length of the time-sensitive frame that affects the maximum latency, especially when store-and-forward bridges are used. As a result, the length of the time-sensitive frame is the parameter that has to be fine-tuned while engineering the network in order to meet the delay bound required by the application. Overall, this finding is a positive result, as time-sensitive flows typically feature small frames.

However, when interfering traffic is present, the analysis in [20] proved that the most determining contribution to the maximum latency for the time-sensitive traffic comes from the maximum size of the *interfering frames*. The motivation behind this finding is the non-preemptive scheduling that is applied at the egress ports of network devices (bridges and end stations). In fact, in the worst case a time-sensitive frame may be delayed in every bridge by the ongoing transmission of a maximum-sized frame not belonging to the time-sensitive class.

To give some quantitative figures, from the analysis in [20], assuming Gigabit Ethernet and a 1522-byte maximum size interfering frame, the bridge latency value that is computed according to formula (1)

$$Interfering\_frame\_size + Bridge\_Delay + Cable\_Delay \quad (1)$$

is equal to 13.898 µs.

Conversely, without interference, assuming a 300-byte long AVB/TSN frame, the bridge latency value that is calculated according to formula (2)

$$AVB\_frame\_size + Bridge\_Delay + Cable\_Delay \quad (2)$$

is about 4.122 µs.

To tackle the problem of interfering frames the concept of Time-Aware Shaper (TAS) was presented.

E. Time-aware shapers

TASs are "smart shapers" that let frames out based on their size and on the knowledge of the next arrival time for scheduled time-sensitive frames. This knowledge is possible, as time-sensitive traffic typically follows a regular pattern (e.g., consisting of small bursts occurring every given number of microseconds). TASs block any lower priority transmission that would interfere with the upcoming transmission of time-sensitive traffic. TABs therefore inhibit, whenever needed, the transmission of SR Class B or best-effort frames, as time-sensitive traffic is usually mapped onto the highest priority, i.e., SR Class A, to provide such a traffic with the best QoS. TASs work as follows. If a given non-time-sensitive frame is ready for transmission, but such a transmission would delay the start of the next transmission of a time-sensitive frame, the transmission of the interfering frame is inhibited. On the other hand, if the non-time-sensitive frame is small enough to complete its transmission before the start time of the upcoming time-sensitive transmission, the frame is allowed to go.

TASs are a mean to temporally isolate time-sensitive flows, as they allow time-sensitive frames to egress from a bridge port without experiencing any interference from other traffic types.

The TAS implementation is feasible thanks to the synchronization provided by the IEEE 802.1AS [2], that makes AVB bridges time-aware nodes, i.e., nodes with network timing information. The modifications needed to realize TASs include the introduction of a suitable mechanism to allow/inhibit transmissions in a given time window and a way to suitably configure the TAS exploiting the information provided by a management information base (MIB).

F. Cut-Through switches

Another delay factor that can be removed according to the ongoing work within the IEEE 802.1Qbv project is the latency introduced by store-and-forward bridges. In fact, while in the presence of interfering frames there is little advantage in using cut-through bridges, if interference is removed by TASs, the target port of the bridge is idle when time-sensitive frames arrive and the benefit of using cut-through bridges becomes significant. In [19], assuming a 64-byte internal buffering (i.e., a 64-byte cut-through point), it is shown that the latency achieved by cut-through bridges is almost half of that obtained using store-and-forward bridges (i.e., 2.074 µs vs. 4.122 µs). Moreover, such a latency can be guaranteed and is also independent of the size of the time-sensitive frame.

G. Additional considerations on priority mapping in AVB traffic classes

TASs help to reduce latency, but TASs alone are not enough to solve the problem of supporting scheduled traffic on AVB networks, as this entails the ability to transmit frames in a time-driven way. As the IEEE AVB provides only two classes of time-sensitive traffic (i.e., SR Class A and B) it is natural to map the time-sensitive flows on the highest priority class, i.e., SR Class A. However, this choice is not enough to support small-size time-sensitive traffic like control traffic, for two reasons. First, according to the IEEE 802.1Qav standard, SR Class frames undergo the CBS algorithm [4] and shaping blocks frame transmission for a given class if the credit level of the class is below zero. Second, if multiple time-sensitive traffic flows are present in the same network and all of them are mapped on the same SR Class they will mutually interfere. The non-negligible effect of mutual interference between traffic flows at the same priority level was shown in [7][8] and in [12]. As a result, if small-size low-latency scheduled traffic, such as control traffic, is handled in the same queue as large video frames mapped on the same class (here we assume Class A), it will suffer from the interference of the other traffic in the same class and its performance will be affected. For this reason, in a recent work [12] a separate traffic class on top of the AVB SR classes, i.e., with a highest priority than Class A, was envisaged as a more effective way of handling scheduled traffic in AVB networks. This approach is summarized in the following.

5. Scheduled Traffic Support through the AVB_ST approach

In [12] a novel approach to support scheduled traffic over IEEE AVB networks, called AVB_ST, was proposed. Here it is briefly resumed for the sake of clarity. Full details are found in the original paper.

H. Overview of the AVB approach

The starting point of the AVB_ST approach is adding a new, separate traffic class on top of the AVB SR Classes A and B. This new class is called the Scheduled Traffic Class (ST Class). ST frames are tagged with the highest priority TAG according to the IEEE 802.1Q standard, while SR Classes A and B take the second and the third highest priority, respectively. This choice comes from the assumption that scheduled traffic includes time-sensitive high-priority flows (e.g., control traffic). ST traffic is handled in a separate queue and does not undergo credit-based shaping, thus avoiding the undesirable effects of traffic shaping on the flow latency. SR Class A and B are handled by Credit-Based Shaping and best-effort traffic by strict priority, as foreseen by the IEEE 802.1Qav standard.

As the model in [12] assumes that ST flows are periodic, with fixed and a priori known period and frame size, it is possible to plan offline the transmission schedule for ST frames. In particular, offset scheduling techniques [21] can be applied in the network engineering phase to guarantee, by design, that transmissions of ST frames will never collide in the whole network (i.e., either in end stations or bridges).
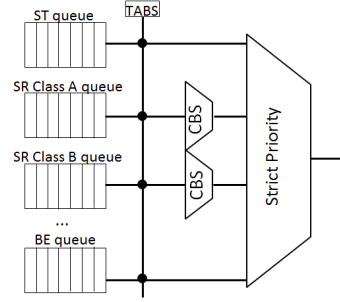

Fig. 1: Traffic Classes in the AVB_ST approach

The support of the ST Class according to the AVB_ST approach requires a reliable synchronization of the network nodes, such that every node knows the right time to transmit its ST traffic. Synchronization is provided by the IEEE 802.1AS standard [2].

In the AVB_ST approach, in order to prevent any interference on ST frames from other traffic classes, TASs enforce a minimal-distance constraint between the transmission of non-ST and ST frames, that inhibits the transmission of non-ST traffic that would delay the upcoming ST one. The messages in the queues associated to the SR Classes undergo both the TAS and credit shaping, while best-effort messages go through the TAS only, as no shaping is foreseen for best-effort traffic in the AVB standard. In the AVB_ST design in [12], traffic shaping for SR Classes is applied after the TAS, so that credits are consumed only when the frame transmission is allowed by the TAS.

As in the AVB_ST approach ST frames are transmitted at known time instants and do not suffer from interference from the same class or from other traffic classes, the reception instant for any ST message can be computed taking into account the synchronization error between the nodes and the drift of each station. The drift represents the clock speed and is expressed in parts per million (ppm). At the receiver, for each ST message, a time window, called ST_Window, within which the ST frame has to be received, is therefore defined. In [12] comparative simulation results of AVB_ST and standard AVB were provided in a factory automation scenario. The AVB_ST approach proved to be very beneficial to support scheduled traffic. In fact, ST traffic obtained low and predictable latency values, without significantly affecting SR traffic. This positive outcome is the result of the combination of three factors: the adoption of offset-based scheduling for ST class; the temporal isolation provided to the ST traffic class by the TAS; no shaping for the ST traffic.

As this paper focuses on automotive networks, in the following new simulations results are presented for both AVB_ST and standard AVB, that were obtained in automotive scenarios utilizing realistic traffic patterns.

I. Comparative assessments between AVB_ST and standard AVB in an automotive scenario

In the realistic scenario under study, the car is equipped with several systems, which provide different services:

Type	Bandwidth [Mbps]	Appl.Payload [Byte]	Service rate [ms]	Act. Time [s]	AVB_ST Priority	AVB Priority
Cameras (3+1)	(4, 12)	uniform(303, 603)	1.6	3 at 0+1 at 400	Class A	Class A
LDW/TSR camera	(2, 6)	uniform(303, 603)	0.8	200	Class A	Class A
DA-Cam Video traffic	(4, 12)	uniform(155, 603)	0.402	3 at 0+1 at 400	Class A	Class A
DA-Cam Warning traffic	0.037	46	10	200	ST Class	Class A
DVD player	11	1050	0.76	500	Class B	Class B
CD Audio player	1.41	1050	5.96	0	Class B	Class B
Cross-domain traffic						
1st flow	0.037	46	10	0	ST Class	Class A
2nd flow	0.074	46	5	300	ST Class	Class A

- Advanced Driver Assistance Systems (ADASs), i.e., a camera-based system composed of 5 IP-cameras that generate video streams (one for each camera).

- High-quality multimedia audio, i.e., a multimedia in-car audio system allowing to listen to music at a very high quality. This system is composed of a CD audio player and an audio amplifier.

- HD-Video entertainment, i.e., a DVD system composed of a DVD player and monitors placed on the back of the front seats, called Rear Seat Entertainment (RSE) monitors.

- Control system, based on information collected in other domains and reported to the driver, e.g., the car engine temperature, the fuel consumption level, etc.

Four cameras, i.e., Front/Night Vision, Left, Right, Rear, are used for direct services, i.e., the services that support the driver with visual information in the form of views displayed on the monitor of a Head Unit (HU) installed on the car dashboard. The video streams captured by the cameras are sent to a central Electronic Control Unit (ECU), named DA-Cam, which analyses and processes them, to produce new output video that are streamed to the Head Unit. These flows are either augmented with additional graphics to assist the driver or resulting from processing multiple camera flows to produce single views (e.g., Top view, Side view, etc.). The video streams generated from the fifth camera, the Traffic Sign Recognition/Lane Departure Warning (TSR/LDW) one, are processed to provide indirect services, which support the driver with Navigation Warnings to improve road safety that are displayed on the Head Unit monitor.

As reported in [11] for the recording, compression, and distribution of high-definition video most recent IP-cameras use H.264/MPEG-4 Part 10 AVC (Advanced Video Coding), a standard for video compression that is widely used and requires a smaller bandwidth (a few Mbps in the case here addressed) than other standards for video compression. However, H.264 generates packets with largely variable size, therefore in many cases, for the sake of performance and analysability, a mechanism is needed to smooth and make more predictable the generated traffic. Here the same approach as in [11] is adopted, which only requires a small modification in the encoding algorithm and is able to reduce the variability in the generated traffic and minimize the delay due to the acquisition and coding of the video frames. In the simulation scenario envisaged here, the video frame rate generated by an IP-camera is assumed equal to 30 frames per second (fps), while the video

resolution selected for displaying the stream on the Head Unit monitor is 640 x 480 pixel. A maximum video frame size of 27.3 KB, as in [11], is assumed.

The high-quality multimedia audio consists of a player that provides an audio stream that is coded with AAC (Advanced Audio Coding) and sent to the multimedia in-car audio amplifier system.

The car is also equipped with a HD-Video DVD entertainment system [11]. The DVD video stream (encoded with MPEG4 High Definition standard) is directly sent to the rear seats entertainment monitors (RSE).

The scenario here addressed also includes a cross-domain traffic service, consisting of periodic data gathered from a gateway device, called a cross-domain Unit (CU), and directed to a Cross-domain Processing Unit (CPU). Such a unit processes data, extracts information useful to the navigation and generates driver warning. When a special condition that requires a driver warning is detected, the Cross-domain Processing Unit turns on the corresponding indicator on the dashboard.

Table I summarizes the traffic flows used in the simulations and shows how the flows were mapped onto the traffic classes provided by the two protocols under study, i.e., AVB_ST and standard AVB. Every station in the network has an activation time that indicates when the relevant traffic starts to be generated. Such a choice has been made to evaluate the protocols under an increasing workload. The performance was assessed in the double-star switched topology shown in Fig. 2.
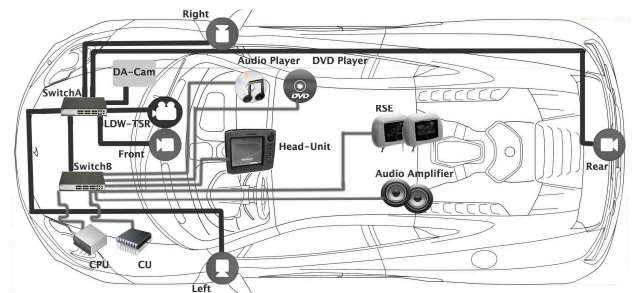

Fig. 2: Network Topology

The two switches, called Switch A and Switch B, are directly connected. Switch A is only traversed by the ADAS traffic that originates from the cameras and is directed to the DA-Cam. Cross-domain and entertainment traffic only cross Switch B, that is also traversed by the ADAS traffic sent from the DA-

Cam to the Head Unit. The latter traffic is the flow resulting from the aggregation of the flows from the cameras that provide direct services. In Fig. 2, the connections related to Switch A are drawn as black lines, while the ones related to Switch B are represented by grey lines.

The network performance was evaluated using the OMNeT++ simulation tool and the INET-framework. For all the simulations, the duration was set to 600s and five simulations for the same scenario were realized to evaluate the results obtained varying the seed of the uniform distribution used to determine the frame length. For the simulation scenario here proposed, as the purpose here is to assess the network behaviour when the cross-domain workload doubles over time, an increasing cross-domain traffic was modelled using two flows. The first one, indicated in Table I as cross-domain 1^{st} flow, starts at t=0s and stops at t=300s. The second one, named cross-domain 2^{nd} flow, is twice the first one. It starts at t=300s and continues until the end of the simulation at t=600s. The performance metrics here adopted are:

- Latency, i.e., the one-way frame end-to-end delay defined as the time interval between the instant at which the frame is sent and the instant at which it is received at the destination. Latency is measured at the MAC level.

- Jitter, defined as the absolute value of the difference between two consecutive inter-arrival times. The inter-arrival time is here defined as the difference between the arrival times of two consecutive frames of the same stream. This jitter is calculated at the destination as in equation (3)

$$J = | (a_n - a_{n-1}) - (a_{n-1} - a_{n-2}) | \qquad (3)$$

where n>2. The arrival time of the Ethernet frames, and so the latency and the jitter, are measured at the MAC level.

Standard AVB Setup

In the simulation of standard AVB, both the traffic classes provided by the standard i.e., the SR Class A (the highest priority) and the SR Class B (the second highest priority) were used. As shown in Table I, all the ADAS-related traffic, i.e., all the video flows from the cameras to the DA-Cam, the video flows from the DA-Cam to the Head Unit and the navigation warnings sent by the DA-Cam to the Head Unit are mapped onto the SR Class A. The entertainment traffic, i.e., the DVD stream and multimedia audio are instead mapped onto SR Class B. This choice is because here the ADAS traffic is assumed to be more critical than the entertainment one. According to the AVB Standard the 75% of the total bandwidth is reserved to Class A and Class B, and both classes undergo the AVB Credit-Based Shaper algorithm [4]. The bridge processing time, similarly to other works in the literature [6][7][8], is here assumed equal to 10 μs.

AVB_ST Setup

In the AVB_ST simulation, both the cross-domain traffic and the navigation warning traffic are mapped on the ST class. This is to reduce the latency for these time-sensitive periodic flows and to protect them from the interference of large frames of ADASs and multimedia flows. All the other traffic flows are mapped on the SR traffic classes in the same way as in the AVB simulation. The bandwidth that is left to the SR Classes is

the difference between the total available bandwidth and the sum of the bandwidth needs of the ST traffic plus the bandwidth required by the synchronization messages to realize a 125 ms synchronization interval, that is about 1.5 Mbit/s according to [10].

Simulation Results

Table II compares the latency and jitter obtained by the different traffic flows with AVB_ST and with standard AVB. The results show that the flows that are mapped on the ST class in AVB_ST, i.e., navigation warnings and cross-domain traffic, benefit from the separation from the other traffic classes comparing with the AVB case, in which all of them are mapped on the SR Class A. With AVB_ST, navigation warnings and cross-domain traffic not only obtain a lower latency than with standard AVB, but also have the mean and maximum latency equal, thus their jitter is zero. The reason for this result is that, in AVB_ST, ST traffic does not experience queuing delays thanks to offset scheduling and TAS, is provided with preferential service while crossing the network and does not undergo shaping, so the latency is low and predictable. Conversely, with standard AVB, these flows may experience higher and non-constant latency, due to the interference from other traffic flows in the same class (Class A) and to the additional delay introduced by the CBS. The positive effect of the AVB_ST approach is more evident for the navigation warning traffic than for the cross-domain one because in the considered scenario navigation warning traffic shares the ports of the bridges along the path between the talker and the listener with the SR flows and therefore suffers from their interference. Conversely, the cross-domain traffic flows do not share the bridge ports with other traffic classes.

Simulation results also show a very small increase of the latency and jitter for the SR classes in AVB_ST comparing with standard AVB. This means that the advantage for the ST classes is obtained without significantly affecting the latency and jitter for SR flows. Table II also shows the confidence intervals obtained, that are small, thus indicating that the number of analysed data is adequate. Table II reports confidence intervals only for the flows that are characterized through probability distributions.

6. Conclusions and On-GOING WORK

This paper outlined the current status about the adoption of Ethernet in automotive communications, with a particular focus on the IEEE AVB and on the current efforts to support scheduled traffic. A comparative performance assessment in a realistic automotive scenario between standard AVB and AVB_ST, a recent approach proposed to deal with scheduled traffic was also presented. Results showed that the introduction in AVB_ST of a separate class for scheduled traffic, combined with offset-based scheduling, the temporal isolation provided to the ST traffic class by the TAS mechanism, and strict priority scheduling, offers low and bounded latencies to scheduled traffic even under a high SR traffic load.

About the ongoing work within the IEEE P802.1Qbv project [13][14], the activities are still in progress and thus they are subject to changes over time. However, as the new ideas are interesting and worth of discussion, here a brief overview is

TABLE II. LATENCY AND JITTER RESULTS (WITH CONFIDENCE INTERVALS)

	Latency (µs)				Jitter (µs)			
	AVB		AVB_ST		AVB		AVB_ST	
	Mean	*Max*	*Mean*	*Max*	*Mean*	*Max*	*Mean*	*Max*
Cameras (3+1)	217 ±4	431±2	217 ±3	431 ±2	396 ±2	446	396 ±4	446
LDW/TSR camera	265 ±3	424 ±2	265 ±2	424 ±1	396 ±3	446	396 ±2	446
DA-Cam Video traffic	87 ±3	118 ±3	96 ±4	165 ±2	112 ±2	147	134 ±3	163
DA-Cam Warning traffic	68 ±2	99 ±3	36.4 -	36.4 -	25 ±2	47	0 -	0
DVD player	167 -	178 -	212 -	220 -	17 -	34	24 -	37 -
CD Audio player	169 -	175 -	220 -	224 -	25 -	27	16 -	35 -
Cross-domain traffic								
1st flow	23 -	25 -	21 -	21 -	0.3 -	0.4	0 -	0
2nd flow	42 -	45 -	31.7 -	31.7 -	0.3 -	0.4	0 -	0

provided. To enable transmissions from each queue of a bridge so as to comply with a known time schedule, specific features have to be added to the bridges. These enhancements consist of a transmission gate associated with each queue and a mechanism to determine whether each queued frame can be selected for transmission or not based on the state of the transmission gate. Two states, with the relevant actions, are specified for a gate. i) Gate Open: In this state, queued frames are selected for transmission, in compliance with the transmission selection algorithm ruling the queue. ii) Gate Closed: In this state, queued frames are not selected for transmission. A frame of a traffic class queue will not be transmitted if the relevant transmission gate is in the "closed" state or if there is not enough time available to transmit the entire frame before the occurrence of the next gate-close event associated with that queue. In fact, each port has an associated list of gate events, which changes the state of the transmission gate relevant to the queue of each traffic class. The gate event list is cyclically repeated from the first event in the list based on the value of a configuration parameter. State changes are triggered by a time-interval associated with a gate event, that is measured relative to the time at which the previous gate event in the list finished its execution. As these enhancements are ongoing work, a deeper discussion at this early stage is not possible and is left for future work.

REFERENCES

[1] L. Lo Bello, "The case for Ethernet in automotive communications," SIGBED Review, vol. 8, n.4, pp. 7–15, Dec. 2011.

[2] IEEE Standard for Local and metropolitan area networks-Timing and Synchronization for Time-Sensitive Applications in Bridged Local Area Networks - IEEE Std 802.1AS-2011.

[3] IEEE Std 802.1Qat, IEEE Standard for Local and metropolitan area networks, Virtual Bridged Local Area Networks, Amendment 14: Stream Reservation Protocol.

[4] IEEE Std. 802.1Qav, IEEE Standard for Local and metropolitan area networks, Virtual Bridged Local AreaNetworks, Amendment 12: Forwarding and Queuing Enhancements for Time-Sensitive Streams.

[5] MOST Media Oriented System Transport, multimedia and Control Networking Technology, Rev. 2.4, OASIS Silicon Systems, 2005.

[6] T. Steinbach, et al. H. T. Lim, F. Korf, T. C. Schmidt, D. Herrscher, A. Wolisz, "Tomorrow's In-Car Interconnect? A Competitive Evaluation of IEEE 802.1 AVB and Time-Triggered Ethernet (AS6802)," in IEEE Vehicular Technology Conf., (VTC Fall), pp.1-5, Sept. 2012.

[7] G. Alderisi, G. Iannizzotto, L. Lo Bello, "Towards IEEE 802.1 Ethernet AVB for Advanced Driver Assistance Systems: A Preliminary Assessment," in Proc. 17th IEEE Conference on Emerging Technologies & Factory Automation (ETFA 2012), Kracow, Poland, Sept. 2012.

[8] G. Alderisi, A. Caltabiano, G. Vasta, G. Iannizzotto, T. Steinbach, L. Lo Bello, "Simulative Assessments of IEEE 802.1 Ethernet AVB and Time-Triggered Ethernet for Advanced Driver Assistance Systems and In-Car Infotainment" ," in Proc. of the IEEE Vehicular Networking Conference (VNC), pp. 187-194, Nov. 2012, Seoul, South Korea.

[9] R. Cummings, K. Richter, R. Ernst, J. Diemer, and A. Ghosal, "Exploring Use of Ethernet for In-Vehicle Control Applications: AFDX, TTEthernet, EtherCAT, and AVB," SAE Int. J. Passeng. Cars - Electron. Electr. Syst. 5(1):72-88, 2012.

[10] H.T. Lim, et al. "IEEE 802.1AS Time Synchronization in a Switched Ethernet based In-car Network," in VNC 2011, IEEE, 2011, 147 –154.

[11] W. Hintermaier, E. Steinbach, "A Novel Real-Time Video Data Scheduling Approach for Driver Assistance Services", 2011 IEEE Intelligent Vehicles Symposium (IV), Germany, June 5-9, 2011.

[12] G.Alderisi, G. Patti, L. Lo Bello, "Introducing Support for Scheduled traffic over IEEE Audio Video Bridging Networks", IEEE 18th International Conference on Emerging Technologies & Factory Automation (ETFA), Cagliari, Italy, 09/2013802.1

[13] 802.1Qbv - Enhancements for Scheduled Traffic. Website: http://www.ieee802.org/1/pages/802.1bv.html.

[14] IEEE P802.1Qbv/D1.2 Draft Standard for Local and Metropolitan Area Networks—Bridges and Bridged Networks—Amendment: Enhancements for Scheduled Traffic, April 21, 2014.

[15] Open Alliance SIG website. [Online].Available: http://www.opensig.org

[16] M. Tazebay, A. Abaye, "From OPEN Alliance BroadR-Reach PHYs to IEEE 802.3bp RTPGE: the development of next generation Automotive Ethernet PHYs", Third Ethernet & IP @ Automotive Technology Day, Leinfelden-Echterdingen Stuttgart, Germany, Sept. 2013.

[17] T. Suermann, "Advanced Sleep and Wakeup concept for Automotive Ethernet", Third Ethernet & IP @ Automotive Technology Day, Leinfelden-Echterdingen Stuttgart, Germany, Sept. 2013.

[18] IEEE Std 1588-2008, Standard for a Precision Clock Synchronization Protocol for Networked Measurement and Control Systems.

[19] D.Pannell,"Overview of IEEE 802.1Qbv Time Aware Shaping", Joint IEEE-SA and ITU Workshop on Ethernet, Geneva, Switzerland, 2013.

[20] D. Pannell, "AVB - Generation 2 Latency Improvement Options", 802.1 AVB group meeting, Singapore, March 2011.

[21] K. Tindell, "Adding Time-Offsets to Schedulability Analysis", Rep. YCS221, Dept. of Computer Science, York University, UK,1994.

[22] H. Zinner, et al. "A Comparison of Time Synchronization in AVB and FlexRay in-vehicle networks", WISES 2011.

[23] R. Queck, "Analysis of Ethernet AVB for Automotive Networks using Network Calculus", ICVES 2012, 61-67, July 2012, Istanbul, Turkey.

[24] J.R. Seyler, et al. "A self-propagating wakeup mechanism for point-to-point networks with partial network support," DATE 2014, March 2014.

[25] S. Stein, R. Racu, "Ethernet Quality of Service @ Volkswagen", 2nd Ethernet Technology Day, Regensburg, Sept.2012.