Towards IEEE 802.1 Ethernet AVB for Advanced Driver Assistance Systems: a preliminary assessment

Giuliana Alderisi*, Giancarlo Iannizzotto#, Lucia Lo Bello*

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

Department of Sciences for Engineering and Architecture, University of Messina, Italy
ianni@unime.it

Abstract

Advanced Driver Assistance Systems involving cameras and the multimedia domain are both quite challenging for automotive network technologies, due to the bandwidth requirements and latency/jitter constraints they impose. Ethernet as in-car network is expected to breakthrough in both domains and the IEEE Audio Video Bridging (AVB) standard is a promising candidate. This paper reports preliminary performance assessments of AVB for ADAS, multimedia and infotainment, obtained through the simulation of realistic traffic patterns.

1. Introduction and Motivation

Advanced Driver Assistance Systems involving cameras (e.g., lane departure warning systems, signs/traffic lights recognition, collision avoidance systems) and infotainment applications are very bandwidth demanding and impose challenging latency/jitter constraints on automotive networks. Multiple and heterogeneous networks support the different automotive functional domains [1-2]. In the entertainment and infotainment domains as well as in camera-based driver assistance systems, point-to-point Low-Voltage Differential Signaling (LVDS) wires or analogue Color Video Blanking Signal (CVBS) cables can still be found. However, the usage of point-to-point dedicated connections for audio and video content has to be discontinued due to wiring complexity, that affects maintenance, reliability and weight, and costs, in terms of wires, connectors and fuel consumption. For this reason, more recently, the Media Oriented Systems Transport (MOST) protocol [3], that offers data rates of 25, 50 and 150 Mbps, has been introduced in the infotainment domain. MOST 150 provides a very high payload efficiency (defined as the ratio between the payload and the effective sent data [4]), however, there are some limitations. First, spreading of the MOST technology outside the automotive domain is still limited, and the smaller market penetration entails higher production costs. Second, since MOST is a bus system, its total network bandwidth is shared among all connected devices, while switched networks utilize bandwidth more efficiently [4]. For this reason, in this work we make some preliminary assessments on the suitability of a Switched Ethernet network, i.e., the IEEE AVB standard family [5], in the ADAS and infotainment domain. Our interest towards AVB is motivated by the findings in [6] i.e., the enhanced QoS support provided, together with the IEEE standardization, no need for license fees and its cost and quality, that are comparable to those of standard Ethernet. The novelty of this contribution is a quantitative performance evaluation of AVB for ADAS, multimedia and infotainment, obtained through simulation in scenarios that use realistic traffic patterns. The purpose of this work is to demonstrate that, even in conditions of high workload, AVB performance are adequate and promising. This preliminary work is part of an ongoing research, involving both industry and academia, aimed at investigating the case for Ethernet (at large, also involving other real-time extensions) to become the single network for both ADAS and infotainment. The paper is organized as follows. Section 2 gives an overview of related work, while Section 3 deals with the AVB protocol. Section 4 describes the simulation scenario, while Section 5 shows the results obtained. Finally, Section 6 addresses conclusions and future work.

2. Related work

A number of recent studies addressed the performance of Ethernet as in-car network. However, for the sake of brevity, here we will only mention a few of them, i.e. those that we consider the most relevant either because they deal with the same network addressed in this work (i.e. AVB) or because they inspired the network topology here investigated. In [4], MOST 150 and Ethernet AVB for ADAS and infotainment support are addressed and a discussion about the payload efficiency and network utilization of the two networks is provided. The comparison is based on the theoretical calculation of the two performance figures in specific use cases and, differently from our work, no simulation results are given. The outcome of the comparison is that MOST 150 outperforms AVB as far as payload efficiency is concerned, while AVB is the preferred solution in terms of network utilization, thanks to the multiplied bandwidth, that allows it to support ADAS (for instance, multiple live video streams for cameras), while MOST, that shares the available bandwidth, cannot support the same amount of traffic on the same network. These considerations suggest the need for a MOST/Ethernet gateway, that is addressed in [7], with a concept for a migration of MOST to AVB. In [8] the performance of an in-car network with a double star topology under intensive streaming flows is assessed and the outcome of the study is the need for QoS mechanisms. We agree on these results, and here we chose AVB as it provides advanced QoS management. The same authors, in [9], addressed the QoS offered by three different network topologies, i.e. , a star, a daisy-chain and a tree-based one, in a mixed traffic scenario. The outcome of the work is that the star topology outperforms all the others in terms of end-to-end delay.

The studies in [8-9] motivated the choice of the star-based topology for AVB in our simulation scenarios.

3. Overview on the IEEE Audio Video Bridging standard

The AVB standard defined by the IEEE 802.1 Audio/Video Bridging Task Group [2] provides the specifications for time-synchronized low latency streaming services through IEEE 802 networks and includes three specifications, i.e..

- The IEEE 802.1AS Time Synchronization, which provides precise time synchronization of distributed local clocks with a reference that has an accuracy of better than 1 μs.
- The IEEE 802.1Qat Stream Reservation, which allows for the resource reservation within switches (buffers, queues) along the path between sender and receiver.
- The IEEE 802.1Qav Queuing and Forwarding for AV Bridges, which splits time-critical and non-time-critical traffic into different traffic classes extending methods described in the IEEE 802.1Q standard and applies traffic shaping at the output ports of switches and end nodes to avoid traffic bursts.

The AVB standard guarantees a fixed upper bound for latency for seven hops within the network and defines two QoS classes, i.e.

- Class A, that provides a maximum latency of 2ms;
- Class B, that provides a maximum latency of 50ms.

With a careful planning of periodic execution and mapping to the high priority queues within switches, AVB is able to guarantee low jitter.

4. Simulation scenario

In the scenario under study, shown in Fig.1, different traffic types are present:

• Advanced Driver Assistance Services (ADAS)
• High-quality Multimedia Audio.
• HD-Video entertainment.

The ADASs here considered are based on a system composed by 6 IP-cameras, i.e. Front, Night Vision, Left, Right, Rear, and one Lane Departure Warning/Traffic Sign Recognition (LDW/TSR) camera. All cameras generate video streams (one for each camera) and send them to a specialized Driver assistance Electronic Control Unit (ECU), named DA-Cam. The DA-Cam processes the streams and produces both “views” (e.g. Top view, Side view) and navigation warnings. Both the views and navigation warnings are sent to a Head Unit (HU) that is equipped with a monitor, installed on the car’s dashboard, on which the received views and warnings are displayed. In our scenario, the video frame rate generated by an IP-camera is 30 frames per second (fps), while the video resolution selected for displaying the stream on the Head Unit monitor is 640 x 480 pixel. We supposed a maximum video frame size of 27.3 KB, as in [10] and we adopted the small modification in the encoding algorithm described in [11] to reduce the variability in the generated traffic and minimize the delay due to the acquisition and coding of the video frame. The car is also equipped with a multimedia in-car audio system allowing to listen to music at a very high quality and a HD-Video entertainment (e.g., a BluRay) system [9]. The BluRay video stream (encoded with MPEG4 High Definition standard) is directly sent to the rear seats monitor (Rear Seat Entertainment, RSE), while the audio stream is coded with AAC (Advanced Audio Coding) and sent to the multimedia in-car audio system. Alternatively, the audio stream produced by a multimedia audio player can be sent to the in-car audio system instead of the BluRay audio stream. Table 1 shows the requirements for the traffic flows in our scenario. Our performance metrics are:

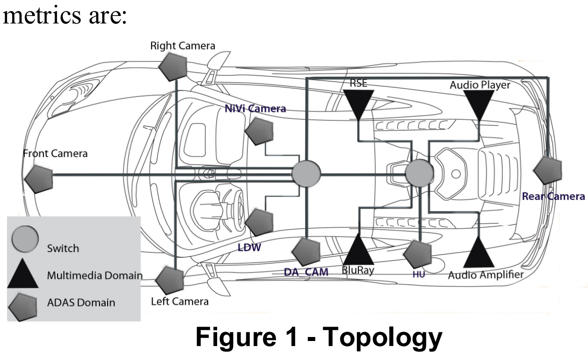

Figure 1 - Topology

Type	Bandwidth [Mbps]	Appl.Payload [Byte]	Service rate [ms]	AVB Priority
Cameras	32.75	910	1.11	AVB SR Class A
LDW/TSR camera	13.10	910	0.555	AVB SR Class A
DA-Cam Video traffic				
Single flow (Case A)	6.55	910	1.11	AVB SR Class A
Aggregated flow (Case B)	32.76	910	0.22	
DA-Cam Warning traffic	0.016	100	50	AVB SR Class A
BluRay	40	1400	0.28	AVB SR Class B
Audio	8	1400	1.4	AVB SR Class B

Table 1 – Characteristics of Traffic Model and configured traffic/priority classes

- *Latency*, defined as the one-way frame end-to-end delay, i.e., the time from the source sending a packet to the destination receiving it.
- *Jitter*, here defined as the absolute value of the difference between two consecutive interarrival times. The Interarrival time is the difference between the arrival times of two consecutive frames of the same stream. This jitter is calculated at the destination as in Equation (1)

$$J_n=|(a_n-a_{n-1})-(a_{n-1}-a_{n-2})| \quad (1)$$

where n>2, The arrival time a_n of the n-th Ethernet frame, and so the latency and the jitter, are measured at the Application level.

4.1 Network topology and traffic flows

Our topology here is a *double-star* where ADAS and entertainment traffic are transmitted on the same physical infrastructure. The network is shown in Fig.1 and consists of two directly connected switches. The first switch connects all the 6 cameras with the DA_Cam, the second one connects all the entertainment units and the Head Unit. The flows that traverse both the switches originate from the DA-Cam and go into the Head Unit. These flows consist of either a single view, i.e. a single flow (corresponding to case A in Section 5), or multiple views aggregated in a single flow (corresponding to case B in Section 5), in addition to navigation warnings present in both cases. The entertainment traffic traversing the second switch is the multimedia video traffic sent by the BlueRay player to be displayed on the RSE and the multimedia audio flow that the relevant audio player streams to the Digital Audio Amplifier. Table 1 summarizes the traffic flows we simulated in our scenario. It can be seen that the network operates under a high workload. Here we simulated the RTP/UDP/IP stack on top of Ethernet, so the Application payload in Table 1 was increased accordingly for all the traffic flows. The network performance was evaluated using the OMNeT++ [12] simulation tool and the INET-framework. The duration of all simulations was 600s.

4.2 AVB Setup

As shown in Table 1, in our simulation all the ADAS-related traffic flows are mapped onto the AVB SR Class A (higher priority, level 3), while the entertainment traffic, i.e., the BluRay stream and multimedia audio, are mapped onto AVB SR Class B (lower priority, level 2). This is because we assumed that all the traffic related to ADAS is more important than the entertainment traffic. According to the IEEE 802.1Qat Standard [5] only 75% of the total bandwidth can be reserved to Class A and Class B, to leave room for best-effort traffic. As the standard [5] specifies that the percentage of bandwidth assigned can be sized according to the traffic needs, and given that there is no best-effort traffic in our scenario, as it is out of the scope of the paper, here we assigned 90% of the Switch 1 bandwidth 1 to Class A and 5% to Class B, and 50% of the Switch 2 bandwidth to Class A and the other 50% to Class B. The medium access scheduling algorithm for both Class A and B is Credit-Based Fair Queuing.

5 Simulation Results

In this Section, we present latency and jitter results, measured for the ADAS video frames (Table 2) and for Ethernet data frames of all the traffic flows (Table 3). As each video frame is split into 30 Ethernet frames for transmission, to measure the latency and jitter of the video frames we are interested in determining the latency of sequences of 30 Ethernet frames. We calculate the time between the instant when the first Ethernet frame of a sequence is sent and the instant when the last Ethernet frame of the sequence arrives at destination. As far as latency measure is concerned, we determine the time elapsed between the arrival time of the first Ethernet frame of a video frame and the arrival time of the first Ethernet frame of the next video frame and obtain an interarrival time. The difference between two interarrival times is the jitter. Here we preliminary assessed that we do not have packet loss. This behaviour was obtained by a suitable configuration of the switch buffers size and thanks to the support provided by the AVB protocol. Moreover, the sum of all incoming traffic on different ports that is directed to the same outgoing port does not exceed the available bandwidth. Table 4, that compares workload and throughput obtained by simulation, confirms our findings. Table 2 shows the results for video frames. The maximum transmission time for a full video frame is 32.68 ms, the maximum experienced jitter is 135μs . The DA-Cam video frames have no jitter even though there is another parallel flow on the same link (i.e., navigation warnings) as both flows originate from the same unit and the navigation flow has a definitely lower bandwidth and service rate.

Traffic type	Latency [ms]				Jitter [μs]			
	Case A		Case B		Case A		Case B	
	Mean	Max	Mean	Max	Mean	Max	Mean	Max
Cameras	32.68	32.68	32.43	32.43	2.80	81	2.86	135
LDW/TSR camera	16.50	16.50	16.51	16.51	6.24	78	7.6	118
DA-Cam	32.41	32.41	32.43	32.43	0	0	0	0

Table 2: Mean and maximum latency and jitter values for the ADAS video frames

The null jitter is due to approximations in the simulator, that does not model the jitter of the hardware implementation. This is not a big issue, as here we want to show the effects of the protocol and not of the hardware. And, such a jitter is well below 1μs. For the Ethernet frames of the cameras (see Table 3) in Case A (i.e. a single flow from the DA-Cam to the HU) the jitter is due to the fact that several video streams are sent from the cameras to the same receiver (DA-Cam), and so the Ethernet frames are delayed in the relevant Switch1 outgoing queue. In Case B (aggregated flow) we see an increase of jitter due to the higher workload (see Table 2), but such an increase is not significant. The simulation results show that the AVB protocol performs well under the scenarios provided.

Traffic type	Latency [ms]				Jitter [μs]			
	Case A		Case B		Case A		Case B	
	Mean	Max	Mean	Max	Mean	Max	Mean	Max
Cameras	0.321	0.487	0.333	0.487	1.45	54	6.02	81
LDW/TSR camera	0.365	0.569	0.385	0.569	555	881	555	960
DA-Cam	0.238	0.238	0.242	0.242	0	0	0	0
Navigation warnings	0.098	0.170	0.105	0.180	0.52	6.2	0.62	7.8
BluRay	0.239	0.239	0.239	0.239	0	0	0	0
Audio	0.239	0.239	0.239	0.239	0	0	0	0

Table 3: Mean and maximum latency and jitter of the Ethernet frames for every traffic class.

	Workload (Mbps)		Throughput (Mbps)	
	Case A	Case B	Case A	Case B
Switch 1	57.64	87.78	57.64	87.78
Switch2	58.27	87.49	58.27	87.49

Table 4: Switch workload and throughput

In particular the latency experienced by the navigation warnings messages is way below the reaction time of a human driver (usually not lower than 0.5s [14]) and therefore fully compliant with the ADAS application constraints. The latency of BluRay traffic meets the requirements stated in [8]. The traffic from the cameras used for direct services (named Cameras in Table 2) experiences a latency which is always less than 33ms, while for this class of traffic the limit can be set to 45ms [8]. Since the LDW/TSR camera has a frame rate twice as the previous ones, its maximum allowed latency must be set to 45/2 = 22.5ms, but its latency is always less than 16.55ms (Table 2), so such a limit is met. The traffic originated by the DA-Cam ECU and addressed to the HU meets the same constraints as the Cameras traffic. Throughput and workload are almost the same in both the topologies and cases, thus confirming that there is no packet loss (see Table 4).

6 Conclusion and Future work

This preliminary performance assessment represents a first step of our research. Future work will deal with other topologies and broader performance assessments under varying workloads, to see how the network reacts to increasing loads. Simulations with cross-domain and best-effort traffic will be also performed, to see the effect of unexpected traffic bursts. Comparative assessments between AVB and other protocols, such as MOST and TTEthernet [15] will be also performed.

7 References

[1] Navet, N., Simonot-Lion, F. 2009. Trends in Automotive Communication Systems. Embedded Systems Handbook, 2nd Edition, Part three. Dr. R. Zurawski, Ed. CRC Press, Taylor and Francis Group, Boca Raton, FL. 13-1, 13-24.

[2]Nolte, T., Hansson, H., Lo Bello, L.. "Automotive Communications - Past, Current and Future", in Proc. of the 10th IEEE Int. Conference on Emerging Technologies and Factory Automation (ETFA'05), 985-992, Catania, Italy, Sept. 2005.

[3] MOST Media Oriented System Transport, Rev. 2.4, OASIS Silicon Systems, 2005.

[4] Noebauer J,. Zinner H. 2011. Payload efficiency and network considerations of MOST and Ethernet. Elektronik automotive. Special issue MOST. Mar 2011, 14-17.

[5] IEEE 802.1 AVB Task Group "IEEE 802.1 Audio/Video Bridging", http://www.ieee802.org/ 1/pages/avbridges.html.

[6] Lo Bello, L., "The case for Ethernet in Automotive Communications", ACMSIGBED Review, 27th Issue, Special Issue on the Tenth International Workshop on Real-Time Networks, Vol. 8, n.4, ISSN: 1551-3688, Dec.2011

[7]Zinner, H., Noebauer, J., Gallner, T., Seitz, J. and Waas, T., "Application and realization of gateways between conventional automotive and IP/Ethernet based networks", in Proceedings of the 48th ACM/EDAC/IEEE Design Automation Conference, DAC'11, 1-6., San Diego, CA, Jun 05-09, 2011.

[8] Lim, H.-T., Weckemann, K., Herrscher, D. Performance Study of an In-Car Switched Ethernet Network without Prioritization. In: Communication Technologies for Vehicles (LCNS 2011), Springer, Heidelberg, 2011.

[9] Lim, H.T., Volker L., Herrscher, D. 2011. Challenges in a future IP/Ethernet-based in-car network for real-time applications. In Proc of the 48th ACM/EDAC/IEEE Design Automation Conference, DAC11, San Diego, CA, June 2011.

[10] Hintermaier, W, Steinbach, E., A Novel Real-Time Video Data Scheduling Approach for Driver Assistance Services, IEEE Intelligent Vehicles Symp, Baden-Baden, Germany, 2011.

[11] Y. G. Lee and B. C. Song, An Intra-Frame Rate Control Algorithm for Ultralow Delay H.264/Advanced Video Coding (AVC), IEEE Trans. on Circuits and Systems for Video Technology, vol. 19, no. 5, may 2009, pp. 747-752.

[12]OMNeT++/OMNEST Network Simulation Framework Version 4.0.[Online].Available: http://www.omnetpp.org/

[13] Green, M. "How Long Does It Take To Stop? Methodological Analysis of Driver Perception-Brake Times" Transportation Human Factors, 2, pp 195-216, 2000.

[14]TTA-Group: TTEthernet-Spezifikation. 2008. http://www.ttagroup.org/ttethernet/specification.htm