Extending IEEE 802.1 AVB with Time-triggered Scheduling: A Simulation Study of the Coexistence of Synchronous and Asynchronous Traffic

Philipp Meyer, Till Steinbach, Franz Korf, and Thomas C. Schmidt
philipp.meyer@haw-hamburg.de, steinbach@ieee.org, korf@informatik.haw-hamburg.de, t.schmidt@ieee.org
Department of Computer Science, Hamburg University of Applied Sciences, Germany

Abstract—**In-car networks based on Ethernet are expected to be the first choice for future applications in the domain of info- and entertainment. However, the full benefit of a technologically integrated in-car network will only become rewarding with an Ethernet-based backbone, unifying several automotive domains in a single infrastructure. Today, there is remarkable interest in the IEEE 802.1 Audio/Video Bridging (AVB) protocol suite, that provides end-to-end performance guarantees in Ethernet networks. But for the strict timing requirements of automotive control-traffic, these guarantees are too weak. An extension of Ethernet AVB with synchronous time-triggered traffic can overcome these limitations. In this paper, we investigate the coexistence of synchronous and asynchronous traffic by experimentally adding time-triggered messages to the credit-based shaper of AVB in a straightforward way. Based on simulations and analytical evaluations, we quantify the impact of such integration concepts for a reasonable design range. Our results demonstrate the feasibility of a shaping strategy with concurrent AVB and time-triggered message, but show a significant impact of the schedule design on the asynchronous AVB streams. Based on our findings, we provide recommendations for configurations that can improve end-to-end network performance for in-car applications by over 100%.**

I. INTRODUCTION

Real-time Ethernet is the most promising solution to increase bandwidth and reduce complexity in next generation in-car networking infrastructures. Today up to 100 electronic control units of different functional domains are interconnected over heterogeneous and sometimes proprietary communication technologies (such as CAN, FlexRay, LIN, MOST or LVDS) resulting in a network structure that is hard to design, manage and maintain. Ethernet is an established, widely deployed and open standard technology with a variable physical layer and a large base of development tools and expertise. 40 years after its invention Ethernet will now enter a new application domain, as BMW starts its series production of the X5 model that contains Ethernet based video cameras.

In the first stage of development, Ethernet will drive in-car multimedia, but also camera and laser-scanner based driver assistance systems. These applications typically require timing guarantees in the order of microseconds. This is a typical domain for the Ethernet Audio/Video Bridging protocol suite [1]. Ethernet AVB is a set of standards, providing quality of service mechanisms for low latency communication.

In the second stage of development, the different networking domains (such as chassis, drive-train, comfort or entertainment) are expected to be interconnected using Ethernet based in-car backbone architectures. In this concept, the backbone replaces the central gateway usually used today, allowing for a more efficient interconnection of domain overlapping functions. A subset of these functions – e.g. for autonomous driving – have rigid real-time requirements, such as an end-to-end latency of less than $100\,\mu s$ [2], or jitter in the order of microseconds. Ethernet AVB was not designed for these regid requirements of control traffic and thus it cannot be used in these scenarios.

There are several proposals to improve AVB for satisfying these requirements of extremely low latency and jitter. Recently the standardisation of scheduled traffic was started in Project Authorization Request (PAR) 802.1Qbv [3]. Scheduled traffic adds synchronous time-triggered messages to the previously only event-based shaping of Ethernet AVB. Time-triggered traffic uses a coordinated time-division multiple access (TDMA) multiplexing strategy to prevent multiple outgoing messages from traversing a line card at the same time. Thereby, time-triggered traffic can meet the latency required with predictable minimal jitter.

In this work, we show simulation results from a conceptual implementation of Ethernet AVB with additional time-triggered traffic. We analyse how both, the implementation of synchronous, time-aware shaping, as well as the designspace of time-triggered schedules in Ethernet AVB, have impact on important real-time related network metrics such as latency, jitter or bandwidth utilisation. Our simulation allows to analyse the design space and extend the analytical worst-case calculations with results from realistic network examples to assess the achievable performance. Due to our modular approach, the simulation model can be adapted to future advances possibly defined in IEEE 802.1Qbv. To the best of our knowledge, we contribute the first open-source simulation framework for the analysis of upcomming shaping algorithms for Ethernet AVB that combines event-based and time-triggered communication.

This paper is organised as follows: In Section II, we introduce the concepts of Ethernet AVB and time-triggered Ethernet and present previous and related work. Section III presents the integration concept and the simulation model. In Section IV, different shaping and scheduling designs are presented, evaluated and discussed. Finally, Section V concludes our work and gives an outlook on our future research.

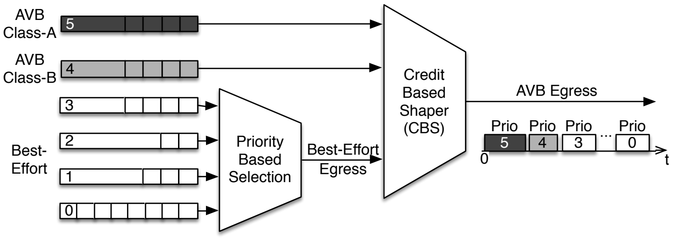

Fig. 1. IEEE 802.1Qav Sender/Forwarding: Transmission selection scheme

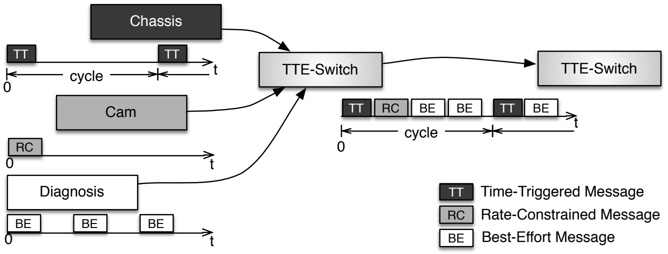

Fig. 2. Prioritising and time-triggered media access in time-triggered Ethernet

II. BACKGROUND & RELATED WORK

This section introduces IEEE 802.1 AVB and time-triggered Ethernet and presents previous and related work.

A. IEEE 802.1 AVB (and Time Sensitive Networking)

The IEEE 802.1 Audio/Video Bridging (AVB) [1] suite, that is specified by the IEEE Time Sensitive Networking Group (formerly Audio Video Bridging (AVB) Task Group), consists of protocols for low latency streaming over 802 networks.

AVB defines the IEEE 802.1AS time synchronisation protocol for the synchronisation of distributed endsystems in Ethernet. It provides a synchronisation error of less than 1 µs over a maximum of seven hops using timestamping [4], [5].

IEEE 802.1Qav defines queuing and forwarding rules for time sensitive applications in Ethernet AVB. For latency requirements up to a maximum of 2 ms over seven hops, stream reservation (SR) class-A was defined. SR class-B guarantees latency requirements of up to 50 ms. Traffic belonging to none of these SR-classes is treated as best-effort traffic. Best-effort messages cover all legacy Ethernet frames.

Transmission selection and traffic shaping in IEEE 802.1Qav is organised by priorities and a credit based shaping (CBS) algorithm (see Figure 1). The transmission of a frame of a stream in a SR class is only allowed when the amount of available credits is greater or equal 0. An upper and lower bound of the credit based shaper limits the streams bandwidth and burstiness. Messages of nodes that are unaware of the IEEE 802.1Qav protocol are mapped to the priorities of best-effort traffic to ensure the real-time capabilities guaranteed by the stream reservation classes.

Ethernet AVB defines a signalling protocol for the dynamic online registration of new real-time streams. IEEE 802.1Qat Stream Reservation Protocol provides a three step signalling process (i.e., stream advertisement, registration and un-registration) to reserve resources along the path between source and sink. At most 75% of the total bandwidth can be reserved. The remaining resources are used for best effort traffic.

B. Time-triggered Ethernet (AS6802)

Several real-time extensions for Standard Ethernet use the concept of synchronous time-triggered messages to provide fully deterministic transmission with low latency and jitter. Popular protocols are PROFINET [6] or TTEthernet [7]. In the following, we discuss the time-triggered concepts based on the TTEthernet protocol. Due to the nature of time-triggered communication the results are applicable to other time-triggered protocols as well and thus should anticipate the expected behaviour in PAR 802.1Qbv.

The TTEthernet (AS6802) specification [7] was standardised in 2011 by the Society of Automotive Engineers (SAE) [8]. It is a compatible extension of IEEE switched Ethernet and uses topologies formed of full-duplex links.

Three different traffic classes can be used in a TTEthernet system: For *time-triggered* (TT) communication, preconfigured schedules assign dedicated transmission slots to each participant. This *coordinated* time-division-multiple-access (TDMA) multiplexing strategy allows for deterministic transmission with predictable delays. It prevents congestion on outgoing line cards and enables isochronous communication with low latency and jitter. To allow for the TDMA concept, a failsafe synchronisation protocol, with an error below 1 µs, implements a global time among all participants.

In addition to synchronised time-triggered messages, two event-triggered message classes are defined: *Rate-constrained* (RC) traffic is intended for the transmission of messages with moderate timing requirements. It limits bandwidth and prioritises according to the strategy of the *ARINC-664 (AFDX)* protocol [9]. RC traffic is comparable with AVB SR Classes A and B. *Best-effort* (BE) traffic conforms to standard Ethernet messages that are transmitted with the lowest priority. The latter allows the integration of hosts that are unaware of the time-triggered protocol and remain unsynchronised. These nodes communicate using best-effort messages. Figure 2 shows the media access policy for messages of different traffic classes.

C. Related & Previous Work

Various work has been dedicated to Ethernet-based communication in cars. Lo Bello [10] provides an overview over different approaches to Ethernet-based automotive communication. Her work argues for deploying IEEE 802.1 AVB and TTEthernet in different application domains. Previous performance assessments of time-triggered Ethernet and Ethernet AVB revealed strengths and weaknesses in different application domains for each of the protocols [11]. As a conclusion an approach with AVB and time-triggered messages on a shared infrastructure – as contributed in this work – was proposed.

There is ongoing research to improve the end-to-end network performance of Ethernet AVB. As AVB uses a priority based scheme, the main challenge is to reduce the time when

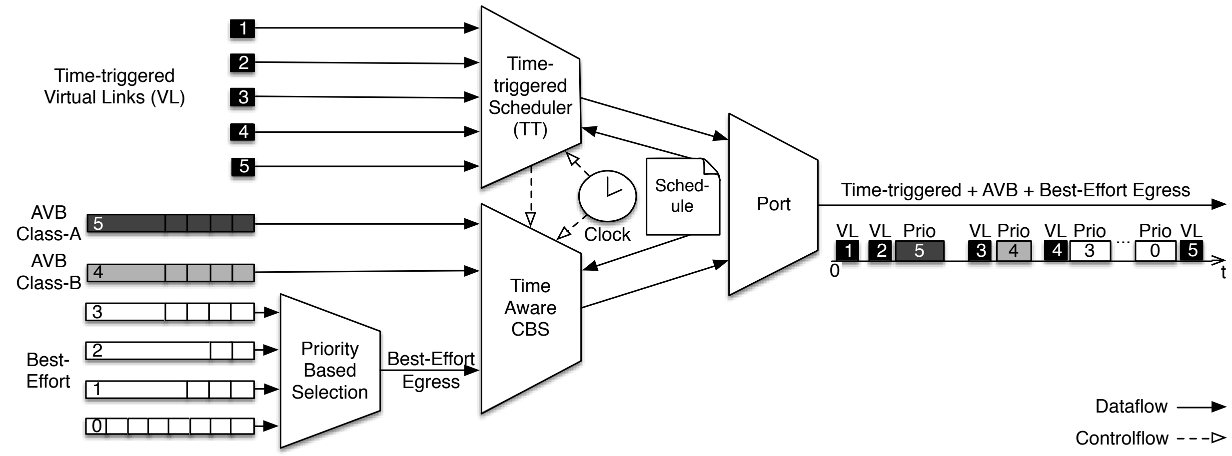

Fig. 3. Transmission Selection Algorithm for AVB and time-triggered traffic. CBS uses clock and schedule, enabling TT communication with no interference

a low priority frame delays a high priority frame. Imtiaz, Jasperneite and Weber [12] analyse the impact of non-real-time cross traffic with varying maximum frame size on the end-to-end latency. Their results can also be applied to the concept of time-triggered traffic in AVB presented in this paper.

The IEEE 802.1 Time Sensitive Networking Group focuses on two approaches to reduce latency of time-sensitive streams. PAR 802.1Qbu [13] introduces frame pre-emption for concurrent streams with different priorities. It defines a service for time-critical frames to suspend the transmission of a non-time-critical frame and resume its transmission afterwards. 802.1Qbu can be only used when both, sender and receiver of a link, are aware of the protocol.

In PAR 802.1Qbv Enhancements for Scheduled Traffic [3], the standardisation of time-triggered traffic was started. The main goal of 802.1Qbv is to provide deterministic communication in so called engineered LANs. An engineered LAN is a network in which schedules can be offline designed. 802.1Qbv is an extension for the scheduled transmission of frames based on timing derived from IEEE 802.1AS (see Section II-A). The proposal uses priority values encoded into the VLAN tag to determine between scheduled and credit-based traffic.

Hillebrand et al. [14] provide a general analysis of the impact of time-triggered and event-triggered traffic in switched networks. The authors focus on the performance evaluation of time-triggered messages, assuming event-triggered frames are less critical. In contrast, this paper analyses the impact of time-triggered communication on competing traffic of other classes on the same infrastructure, as the achievable performance of time-triggered protocols is already well known.

The simulation models for IEEE 802.1 AVB and TTEthernet are both based on the INET-Framework for OMNeT++ (http://inet.omnetpp.org) and have been introduced and validated in previous work [15]. OMNeT++ is an open-source event-based network simulation toolchain. The source code of the real-time Ethernet models is published (http://tte4inet.realmv6.org) and can be used free of charge for simulation-based analyses.

III. EXTENDING AVB WITH TIME-TRIGGERED TRAFFIC

As the latest draft of 802.1Qbv (Draft 0.2) is in a stage too early to implement, we provide a straight forward concept for the integration of time-triggered traffic in Ethernet AVB. Our approach combines IEEE 802.1Qav traffic with time-triggered Ethernet (AS6802). Due to the nature of time-triggered traffic, the results of the simulation-based assessment are easily transferrable to future 802.1Qbv implementations.

Time-triggered (TT) traffic requires to never be delayed by any other frame. Though, TT frames must have the highest priority when time-triggered, AVB, and best-effort frames have to be sent over the same physical network infrastructure. This additional requirement conflicts with the original AVB standard, that demands to assign the highest priority to SR class-A frames. This definition allows for the timing guarantees in Ethernet AVB. Due to the modified priorities in our approach, extensions in queueing and scheduling of AVB frames must be made.

A. *Queueing*

In AS6802, time-triggered frames are sent over statically configured multicast paths, so called Virtual Links (VL). The time-triggered scheduler sends TT frames according to a static schedule, based on a domain-wide synchronised clock. Due to this static schedule, TT frames are never allowed to be delayed by any other frame. The schedule defines a set of TT windows for each line card. During such a window, only TT frames are sent. In the proposed time-aware shaping concept, these TT windows and the domain wide synchronised clock are used as input for the Credit Based Shaping (CBS) algorithm (see Figure 3). The algorithm checks for each frame whether the transmission can be finished before the next TT window starts. If no frame fits, the transmitter remains idle until the transmission of the next scheduled TT frame begins. This idle time – sometimes referred to as guard band – guarantees that AVB and best-effort frames will never interfere with TT frames.

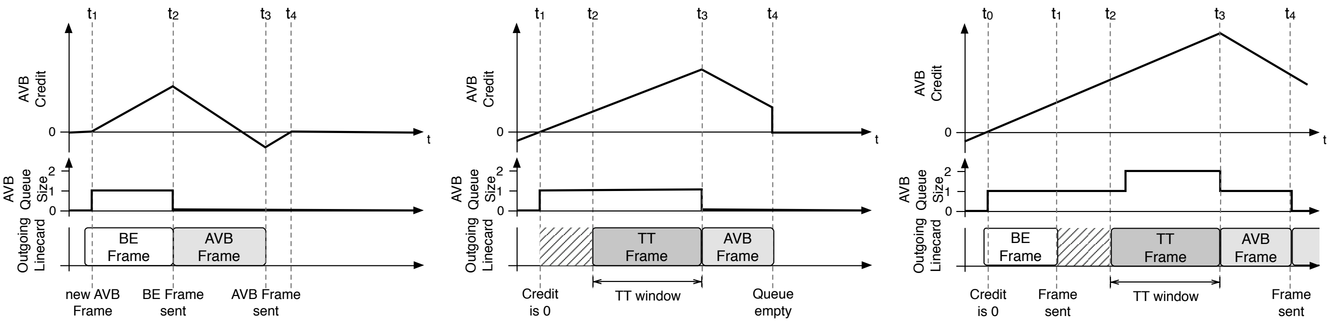

(a) Standard IEEE 802.1Qav (b) Shaping with time-triggered frames (c) Worst case with time-triggered frames

Fig. 4. Credit Based Shaping (CBS) algorithm behaviour in implementation with and without time-triggered traffic

B. Scheduling

The standard CBS algorithm for SR classes A and B is based on a credit value. The credit starts with 0. Whenever the credit is greater or equal to 0, an AVB frame is allowed to be transmitted. When no frame is sent, the credit increases according to the idle slope. If the transmission of an AVB frame is blocked by another frame, the credit increases above 0 (see t_1 in Figure 4a). Since SR class traffic has the highest priority for standard AVB, the AVB frame will be transmitted as soon as the line card gets idle (see t_2 in Figure 4a). Now the credit decreases according to the send slope. If the credit is negative when the transmission of the AVB frame is finished (t_3 in Figure 4a), the credit increases according to the idle slope again until it is 0 (t_4 in Figure 4a).

The proposed time aware shaper may introduce further delays for both AVB Classes: If an AVB frame is ready to be transmitted and too large to be sent before the start of the next TT window, the AVB frame will be queued and the credit will be increased according to the idle slope (see $[t_1, t_2]$ in Figure 4b). When the TT window passed, the queued AVB messages are transferred as long as the credit is greater or equals 0.

C. Analysis of the Worst Case for the New Time Aware Shaper

The proposed time aware shaper causes a new worst case scenario. When an AVB frame is ready to be transmitted, the line card can be already occupied by a BE frame. When the line card becomes idle again, the AVB frame may not be transmittable as it would interfere with a following time-triggered frame. This case is visualized in figure 4c. At t_0, just after the transmission of the BE frame has been started, an AVB frame gets ready to be sent. Since the line card is used, the AVB frame is being queued for the time span $[t_0, t_1]$ and the credit increases. If the transmission time of the AVB frame would be longer than $[t_1, t_2]$, the AVB frame must be further delayed until the TT window finishes (t_3). In the whole time, the credit rises to preserve the bandwidth contingent of the queued AVB stream.

For this reason, latency guarantees for AVB must be recalculated: For the calculation the maximum transmission time of three consecutive Ethernet frames (M) and the intermediate inter frame gaps (T_{ifg}) is used. Hence the maximum interference time rise by factor 3. In case of a 100 Mbit/s (R) network, the highest interference timespan (T_{mi}) now is:

$$\begin{aligned} T_{\text{mi}} &= \frac{3 * M}{R} + 2 * T_{\text{ifg}} \\ &= \frac{3 * 1530Byte}{100Mbit/s} + 2 * 0.96\mu s = 369\mu s \end{aligned} \tag{1}$$

The AVB protocol guarantees that AVB frames will not be delayed more than 125 µs (class measurement interval (T_{ms})) for each switching hop. The outcome of this is the new maximum latency per network device (T_{mdl}):

$$\begin{aligned} T_{\text{mdl}} &= T_{\text{mi}} + T_{\text{ms}} = 369\mu s + 125\mu s \\ &= 494\mu s \approx 500\mu s \end{aligned} \tag{2}$$

Now the maximum latency over 7 hops ($T_{\text{ml}}(7)$) can be calculated for a network with concurrent time-triggered traffic:

$$\begin{aligned} T_{\text{ml}}(N_{\text{hop}}) &= (1 + N_{\text{hop}}) * T_{\text{mdl}} \\ T_{\text{ml}}(7) &= 8 * 500\mu s = 4ms \end{aligned} \tag{3}$$

The result shows a duplication (4 ms) of the AVB maximum latency guarantee of 2 ms when adding time-triggered traffic.

IV. EVALUATION & CASE STUDY

This section analyses the integration of time-triggered and AVB traffic by simulating different network configurations. The simulation results lead to design strategies for networks that compound time-triggered and AVB traffic on the same physical infrastructure.

A. Network Topology & Traffic Configuration

The evaluations are based on a network that generates worst-case scenarios. It consists of ten nodes interconnected via three switches (see Figure 5) to produce a bottleneck with high probability of interference. For all traffic classes, the longest path between sender and receiver is four switching hops.

In the network all links are configured with a bandwidth of 100 Mbit/s. The AVB configuration uses nodes 1 and 2 as talkers, while node 8 is the listener for both streams. Each stream is configured to use 350 B payload in a 125 µs class measurement interval, resulting in a total bandwidth of 50 Mbit/s (2∗25 Mbit/s). The time-triggered messages use a cycle time of 4 ms and full size frames to generate maximum

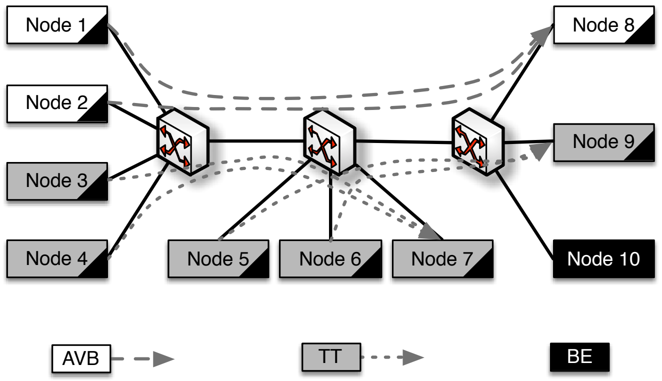

Fig. 5. Network topology and traffic streams for the analyses

interference. For each route, two messages are configured, resulting in a total bandwidth of 25 Mbit/s for time-triggered traffic. Node 10 broadcasts full size best-effort cross-traffic in a 2 ms to 3 ms interval. All nodes also reply with full size best-effort messages to generate traffic bursts.

The case study uses different parameters for time-triggered and best-effort traffic, while the configuration for AVB remains the same. This allows to uphold the comparability of the AVB results through the whole evaluation.

In the following, we analyse end-to-end latency, jitter, AVB credit and queue length behaviour. Over the whole evaluation, the jitter is defined as the maximum difference of end-to-end latency:

$$T_{\text{jitter}} = T_{\max} - T_{\min} \tag{4}$$

This definition of jitter is usually used in real-time systems, as it can be utilised to calculate various application specific parameters, such as buffer sizes or action points of tasks.

B. Validity of Time Aware Shaper

We first compare the end-to-end latency of time-triggered traffic with and without AVB frames. This allows us to assess the validity of our previously presented time aware traffic shaping strategy. The analysis should reveal that AVB traffic never interferes with time-triggered traffic, but still is higher prioritised than best-effort frames.

The simulation setup is configured as explained in section IV-A. The reserved bandwidth for time-triggered messages is 25 Mbit/s and for AVB 50 Mbit/s. Best-effort traffic uses the remaining bandwidth. Figure 6 shows the end-to-end latency distribution as cumulative distribution function (CDF) for a network with only time-triggered and best-effort messages (top) compared with a network with time-triggered, AVB and best-effort messages (bottom).

TABLE I
LATENCY COMPARISON WITH AND WITHOUT AVB TRAFFIC

Network	**Class**	**Min** [µs]	**Max** [µs]	**Mean** [µs]
Without AVB	TT	464.56	465.52	465.04
	BE	244.56	2426.94	750.48
With AVB	TT	464.56	465.52	465.04
	AVB	214.48	1748.83	1403.21
	BE	244.56	9268.56	2485.94

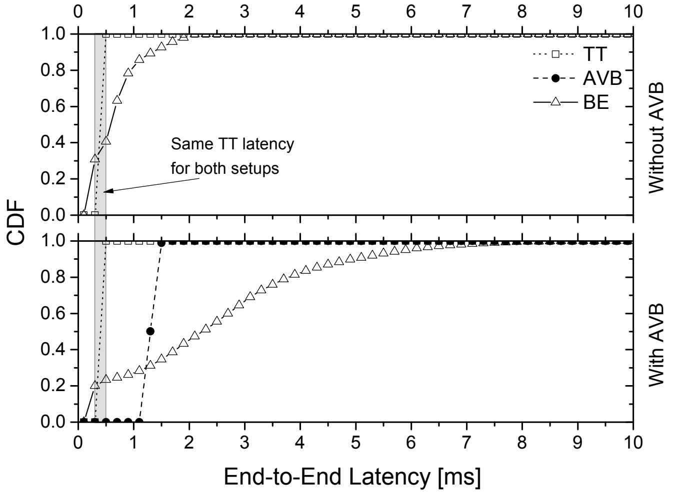

Fig. 6. CDF of latency for network with and without AVB traffic

As expected, the time-triggered behaviour remains independent of concurrent AVB streams (see highlighted area in Figure 6). AVB frames show higher jitter than time-triggered messages when added to the configuration. This is due to their event-triggered scheduling strategy. The latency is bounded ($<$1.8 ms) and though compliant with the specification. The best-effort end-to-end latency significantly suffers when the AVB frames are added to the configuration. The maximum latency increases to nearly 10 ms due to the saturated link between the switches. Table I shows the results in detail.

C. Compact Time-triggered Schedule

One of the most challenging aspects in a network that combines time-triggered and event-triggered (such as AVB) traffic is the design of a schedule. The scheduling of time-triggered messages will not only have side effects on end-to-end latency and jitter, but also the performance of competing messages of other traffic classes. A compact schedule is a schedule without gaps between consecutive time-triggered frames. The more compact a schedule is, the more other messages can be delayed. To show the effect of these compact schedules, we configure a worst-case scenario and evaluate its influence on the AVB communication.

In the worst case scenario, the four time-triggered full size messages are sent on a path that shares the same link with the AVB streams. The time-triggered messages are sent

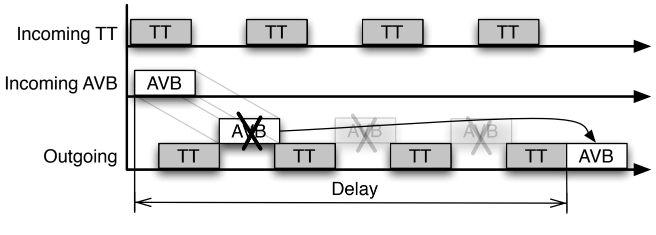

Fig. 7. Compact scheduling of TT frames

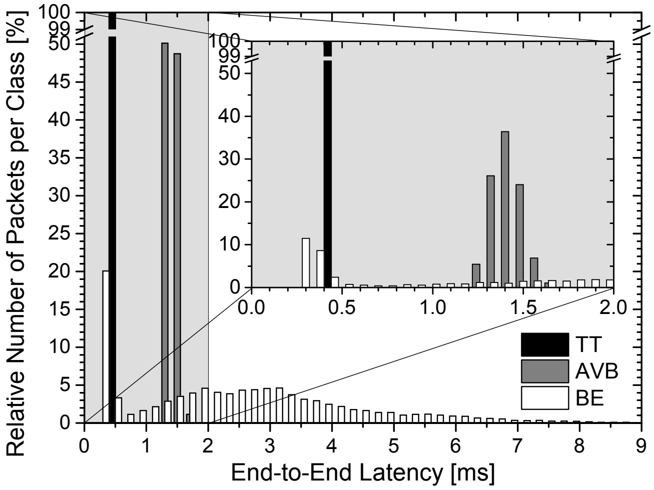

Fig. 8. Latency distribution of the different traffic classes with compact schedule (Please note the different bin sizes in main diagram and inlay)

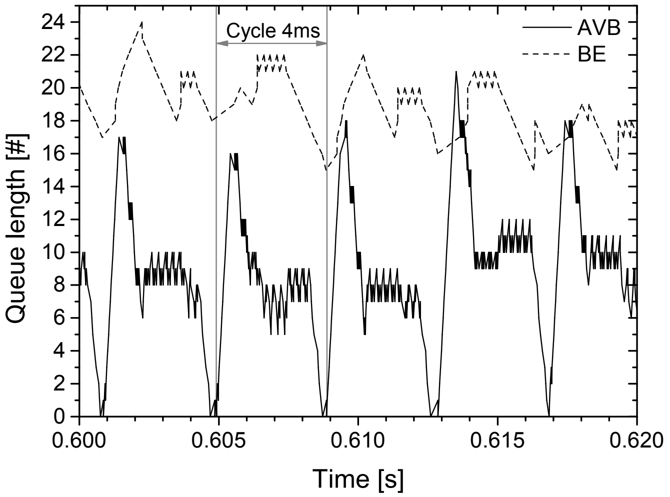

Fig. 9. Queue length of AVB and best-effort frames over simulation time. Detail showing cyclic queuing behaviour due to time-triggered frames.

consecutively with a gap of 23 µs, that is too small for the AVB frames (392 B including header and inter frame gap (IFG)) to fit in. Due to the media reservation for time-triggered messages, the AVB frames are not allowed to be sent until all four time-triggered packets are transmitted. Figure 7 shows the scenario.

The simulation confirms the influence of the time-triggered scheduling that is unfavorable for concurrent AVB streams. Figure 8 shows the latency distribution for all traffic classes. As discussed in the previous section, the time-triggered messages provide both, low latency (<465.52 µs) and jitter (≈940 ns), while best-effort messages suffer from high latency (<9.26 ms) and jitter (<9.02 ms) due to the low priority and heavily saturated link between switch 2 and switch 3. The latency of the AVB streams stays below 1.74 ms, the jitter is 1.40 ms.

For compact schedules the upper bounds for end-to-end latency provided in Section III cannot be achieved. The desired upper bound ($T_{\text{max}}^{\text{AVB}}$) is given in Equation 5:

$$\begin{aligned} T_{\text{max}}^{\text{AVB}} &= T_{\text{max}}^{\text{Node}_{1/2}} + 2 * T_{\text{max}}^{\text{Switch}_{1/2}} + T_{\text{max}}^{\text{Switch}_3} \\ &= 250\mu s + 2 * 500\mu s + 250\mu s = 1.5ms \end{aligned} \tag{5}$$

As there is only interfering time-triggered traffic on the links between the switches, the first and the last switching hop must be calculated with the standard AVB maximum delay of 250 µs. Due to the compact schedule, the simulated end-to-end latency is approximately 240 µs too high compared to the value calculated without regarding scheduling.

To analytically evaluate the maximum end-to-end latency in compact schedules, the calculation for the interference (T_{I}) with time-triggered streams must be adapted to include the gaps (T_{gap_i}) between the time-triggered frames:

$$\begin{aligned} T_{\text{I}} &= \frac{6 * M}{R} + 2 * T_{\text{ifg}} + \sum_{i=0}^{n} T_{\text{gap}_i} \\ &= \frac{6 * 1530Byte}{100Mbit/s} + 2 * 0.96\mu s + 75\mu s = 811.32\mu s \end{aligned} \tag{6}$$

Based on the maximum possible interference the maximum delay for the switches with concurrent time-triggered messages ($T_{\text{max}}^{\text{Switch}_{1/2}}$) can be recalculated:

$$\begin{aligned} T_{\text{max}}^{\text{Switch}_{1/2}} &= T_{\text{I}} + T_{\text{ms}} \\ &= 811.32\mu s + 125\mu s = 936.32\mu s \end{aligned} \tag{7}$$

Equation 8 gives the worst-case delay for the compact schedule:

$$T_{\text{max}}^{\text{AVB}} = 250\mu s + 2 * 936.32\mu s + 250\mu s = 2.4ms \tag{8}$$

The analytical worst-case of approximately 2.4 ms is not reached in this configuration (1.74 ms).

To find the source of the delays, the queue lengths for AVB and best-effort traffic of the saturated outgoing line card of switch 2 to switch 3 are analysed. The length of the queue for time-triggered messages is not of interest, as it cannot exceed 1 by definition. The queue length for best-effort messages rises up to 24. The length for AVB frames is always below 22. Figure 9 shows the queue lengths in detail in the randomly chosen timespan from 600 ms to 620 ms simulation time. The spikes in the queue length of the AVB buffer are cyclic (4 ms cycle time) and caused by the reservation mechanism for time-triggered messages. Up to 22 messages can arrive from the previous switch in the reserved time due to the additional burstiness introduced by the time-triggered shaping. Thanks to the reserved bandwidth, the queue is always processed until the next cycle starts. This behaviour can also be observed in the credit of the port, that rises up to 35 808 credits, allowing the consecutive transmission of up to 22 frames (−1606 credits for one frame).

D. Adjusted Time-triggered Schedule

To show that the network performance of AVB streams in a network with concurrent time-triggered traffic can be significantly improved, the schedule of the previous example is reconfigured. The new schedule contains gaps of 123 µs

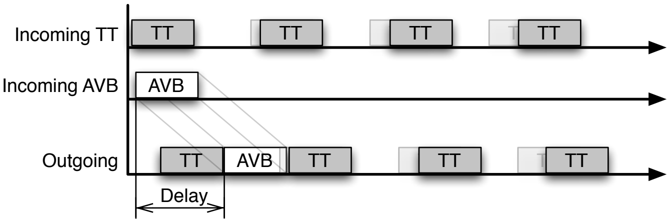

Fig. 10. Scheduling of TT frames adjusted for AVB

between the time-triggered frames that are large enough to fit AVB packets in (see Figure 10).

The results show a significant improvement of the performance for AVB streams with the adjusted schedule. From the latency distribution for all traffic classes (see Figure 11) a better performance for both, AVB and best-effort streams can be read off, while the latency for time-triggered traffic remains the same (464.56 µs to 465.52 µs) compared to the compact schedule (see Section IV-C). The adjusted schedule reduces the maximum latency of AVB and best-effort traffic by 50 %. The mean latency of AVB is even reduced by two thirds (see Table II). The improved schedule is also reflected in the maximum queue lengths, that are reduced for AVB from 22 to 7 messages and for best-effort from 24 to 9 messages.

E. Maximum Transmission Unit (MTU) Analysis

To further improve the AVB performance, we applied the recommendations of Imtiaz et al. [12], and analysed the impact of a reduced Maximum Transmission Unit (MTU).

Exemplarily we changed in the previous configuration the MTU for TT and BE messages from 1500 B to 750 B. Due to the lower probability of interference and shorter maximum transmission time for concurrent BE frames, the latency of the AVB streams is further reduced. The end-to-end latency decreases by another third (max. 555.31 µs; mean 324.34 µs).

Of course, the latency can be further reduced with smaller MTU (see Figure 12). Our analysis reveals that it is important to also analyse the impact of the schedule when determining the best MTU for a network. There are frame sizes that suit better for a given schedule than others, though only reducing the MTU will not inevitably yield overall performance (see linear fit in Figure 12).

TABLE II
LATENCY COMPARISON OF CONFIGURATION WITH COMPACT AND ADJUSTED SCHEDULE. THE RATIO SHOWS FACTOR OF IMPROVEMENT

Network Metric	Traffic Class	Compact Schedule	Adjusted Schedule	Ratio
Min Latency [µs]	TT	464.56	464.56	1:1
	AVB	145.12	145.12	1:1
	BE	244.56	244.56	1:1
Max Latency [µs]	TT	465.52	465.52	1:1
	AVB	1748.83	874.52	1:0.50
	BE	9268.56	4388.67	1:0.47
Mean Latency [µs]	TT	465.04	465.04	1:1
	AVB	1403.21	512.39	1:0.37
	BE	2485.94	1361.65	1:0.55

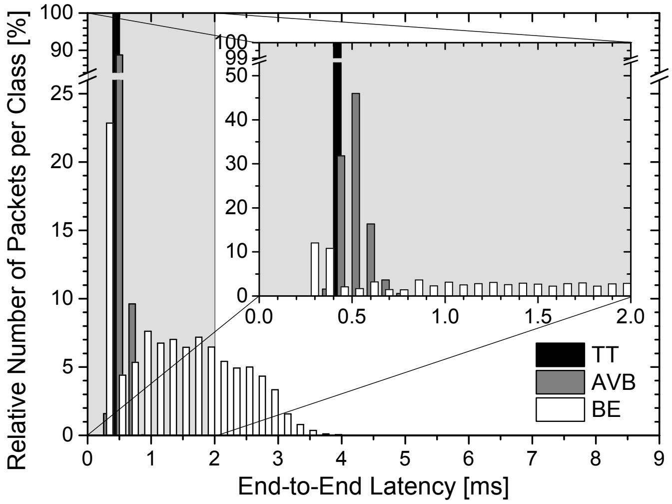

Fig. 11. Latency distribution of the different traffic classes with adjusted schedule (Please note the different bin sizes in main diagram and inlay)

F. Discussion

From the previous evaluation, several recommendations for the configurations of networks with time-triggered and AVB traffic can be made. Scheduling of time-triggered messages has significant impact on the performance of concurrent AVB streams in such a network. When a compact schedule with multiple consecutive TT messages is configured the end-to-end latency of AVB-frames can become high (see Section IV-C). This can even compromise the assured maximum end-to-end latency bounds defined in the specification and thus endanger reliable application behaviour. Schedules that are specifically designed for concurrent AVB traffic allow to reduce this additional delay by more than 50 % (see Table II), without affecting time-triggered traffic.

When reducing the MTU for concurrent traffic (time-triggered and AVB), the delay that is caused by congestion can be further reduced, allowing to achieve almost the same AVB performance in a network with time-triggered frames as without the synchronous traffic (see Section IV-E). In the time-triggered traffic class usually very small control messages are transmitted, and a reduction of the MTU has no influence on the design of these applications. For best-effort messages, segmentation techniques of the transport layer (layer 4) can be utilised to transfer larger information. Table III summarises the relevant network-metrics for the AVB traffic class in all scenarios shown.

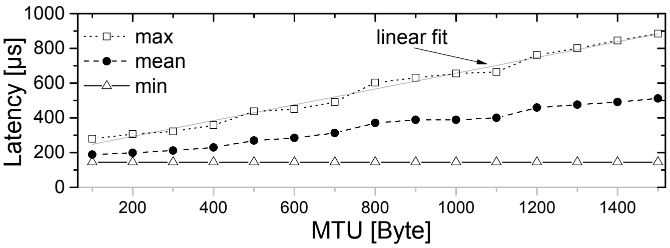

Fig. 12. AVB latency with varying Maximum Transmission Unit (MTU)

TABLE III

COMPARISON OF NETWORK METRICS FOR AVB STREAMS IN THE DIFFERENT CONFIGURATIONS ANALYSED

Network Metric		**Compact Schedule**	**Optimised Schedule**	**Half-size MTU**
AVB Latency	Min	145.12	145.12	145.12
	Max	1748.83	874.52	555.31
	Mean	1403.21	512.39	324.34
AVB Queue Length	Max	22	7	4
	Mean	6.84	1.72	0.99
AVB Credit	Min	−1605	−1606	−1606
	Max	35 808	14 414	7546
	Mean	5067.16	1383.73	−884.85

V. CONCLUSION & OUTLOOK

The Ethernet Audio/Video Bridging protocol suite is one of the favoured technologies for a real-time Ethernet based in-car networking infrastructure. Its dynamic stream reservation mechanism and low configuration effort allows for a fast and flexible network design. But for the rigid timing requirements of control traffic for driver assistance, X-by-wire applications, and autonomous driving, the guaranteed end-to-end latency of 2 ms over 7 hops is insufficient. New standardisation projects propose the integration of time-triggered communication into AVB.

While the achievable performance of time-triggered traffic in Ethernet is already well known, the impact of scheduling on the legacy AVB SR classes is not yet quantified. With this work, we present a simulation study that analyses the effects of concurrent synchronous (time-triggered) and asynchronous (AVB) traffic. Our results show that an implementation of a shaping strategy with time-triggered messages is possible, but imposes an additional delay for AVB messages that can double its end-to-end latency in the worst-case.

By evaluating different configurations, we demonstrate the impact of the schedule design on AVB messages. Compact schedules can significantly delay AVB streams, in contrast to schedules that are specifically designed regarding the concurrent AVB frames. Together with a careful reduction of the MTU for control traffic and best-effort messages, AVB with concurrent time-triggered messages can preserve the same end-to-end latency as without synchronous traffic. We hope that our findings will increase the awareness of the importance of schedule design for the upcoming 802.1Qbu standard. The simulation models used can be downloaded from our site (see http://tte4inet.realmv6.org), to help to assess new network designs.

In our future work, we will analyse realistic in-car network designs based on Ethernet AVB extended with time-triggered traffic. We will examine how well the recommendations made in this paper can be applied on the traffic patterns of real automotive applications and evaluate the metrics in the typical topologies of future in-car networks. For real-world tests in our automotive prototype, we plan on implementing the introduced traffic shaper in hardware. This includes a client implementation based on a SoC design [16], as well as a FPGA-based switch implementation.

For further performance improvements, we will extend the framework with frame preemption (as proposed in PAR 802.1Qbu [13]) to compare the results of both latency reduction strategies and further reduce the necessary guard band.

ACKNOWLEDGEMENTS

This work is funded by the Federal Ministry of Education and Research of Germany (BMBF) within the RECBAR project.

REFERENCES

[1] IEEE 802.1 AVB Task Group, “IEEE 802.1 Audio/Video Bridging (AVB).” [Online]. Available: http://www.ieee802.org/1/pages/avbridges.html

[2] Y. Kim and M. Nakamura, “Automotive Ethernet Network Requirements,” Tech. Rep., Mar. 2011. [Online]. Available: http://www.ieee802.org/1/files/public/docs2011/new-avb-KimNakamura-automotive-network-requirements-0311.pdf

[3] IEEE 802.1 TSN Task Group, “IEEE 802.1Qbv - Enhancements for Scheduled Traffic.” [Online]. Available: http://www.ieee802.org/1/pages/802.1bv.html

[4] M. D. J. Teener and G. M. Garner, “Overview and timing performance of ieee 802.1as,” in *IEEE International Symposium on Precision Clock Synchronization for Measurement, Control and Communication (ISPCS 2008)*. Piscataway, NJ, USA: IEEE Press, 2008, pp. 49–53.

[5] G. M. Garner, A. Gelter, and M. D. J. Teener, “New simulation and test results for IEEE 802.1AS timing performance,” in *International Symposium on Precision Clock Synchronization for Measurement, Control and Communication (ISPCS 2009)*. Piscataway, NJ, USA: IEEE Press, 2009, pp. 1–7.

[6] PROFIBUS & PROFINET International, “Profinet,” Karlsruhe. [Online]. Available: http://www.profibus.com/technology/profinet

[7] W. Steiner, “TTEthernet Specification,” TTTech Computertechnik AG, Nov. 2008. [Online]. Available: http://www.tttech.com

[8] SAE - AS-2D Time Triggered Systems and Architecture Committee, “Time-Triggered Ethernet (AS 6802),” 2009. [Online]. Available: http://www.sae.org

[9] Aeronautical Radio Incorporated, “Aircraft Data Network, Part 7, Avionics Full-Duplex Switched Ethernet Network,” ARINC, Standard ARINC Report 664P7-1, 2009.

[10] L. Lo Bello, “The case for Ethernet in Automotive Communications,” *SIGBED Review, Special Issue on the 10th International Workshop on Real-time Networks (RTN 2011)*, vol. 8, no. 4, Dec. 2011.

[11] T. Steinbach, H.-T. Lim, F. Korf, T. C. Schmidt, D. Herrscher, and A. Wolisz, “Tomorrow’s In-Car Interconnect? A Competitive Evaluation of IEEE 802.1 AVB and Time-Triggered Ethernet (AS6802),” in *2012 IEEE Vehicular Technology Conference (VTC Fall)*. Piscataway, New Jersey: IEEE Press, Sep. 2012.

[12] J. Imtiaz, J. Jasperneite, and K. Weber, “Approaches to reduce the latency for high priority traffic in IEEE 802.1 AVB networks,” in *9th IEEE International Workshop on Factory Communication Systems (WFCS 2012)*, 2012, pp. 161–164.

[13] IEEE 802.1 TSN Task Group, “IEEE 802.1Qbu - Frame Preemption.” [Online]. Available: http://www.ieee802.org/1/pages/802.1bu.html

[14] J. Hillebrand, M. Rahmani, R. Bogenberger, and E. Steinbach, “Coexistence of Time-Triggered and Event-Triggered Traffic in Switched Full-Duplex Ethernet Networks,” in *International Symposium on Industrial Embedded Systems, 2007. SIES ’07.*, Jul. 2007, pp. 217–224.

[15] T. Steinbach, H. Dieumo Kenfack, F. Korf, and T. C. Schmidt, “An Extension of the OMNeT++ INET Framework for Simulating Real-time Ethernet with High Accuracy,” in *Proceedings of the 4th International ICST Conference on Simulation Tools and Techniques*. New York: ACM-DL, Mar. 2011, pp. 375–382.

[16] K. Müller, T. Steinbach, F. Korf, and T. C. Schmidt, “A Real-time Ethernet Prototype Platform for Automotive Applications,” in *2011 IEEE International Conference on Consumer Electronics - Berlin (ICCE-Berlin)*. Piscataway, New Jersey: IEEE Press, Sep. 2011, pp. 221–225.