Towards Safety in Automotive Ethernet-based Networks with Dynamic Workloads

Adam Kostrzewa, Dominik Stöhrmann, Rolf Ernst
Institute of Computer and Network Engineering
TU Braunschweig, Germany
Email: {kostrzewa, stoehrmann, ernst}@ida.ing.tu-bs.de

***Abstract*—The advent of automated and autonomous driving requires decision-making based on high-resolution sensor data. In addition, these systems must react safely and promptly to dynamically changing environmental conditions, e.g. traffic situation, weather etc. In this paper we present the frequently considered isolation mechanisms, standards and concepts for achieving functional safety in automotive systems e.g. 802.1Q, AVB, TSN. However, static configurations, which are introduced by the majority of them, penalizes dynamic communication profiles. Consequently, static network management is no longer sufficient to achieve the required performance level. We show that standardised does not necessarily mean safe "out of the box" nor performant and that the new design vectors must be taken into consideration.**

I. Introduction

The rapid transition from bus-based architectures to switched Ethernet networks in the automotive domain is a result of current design trends that bring new workloads and requirements. These include new infotainment applications with IP traffic via car-to-x communication e.g. transmission between cars or infrastructure, media streaming from cloud. Furthermore, as a result of the transition to autonomous and automated driving, the volume of sensor traffic has increased drastically. We have high resolution sensors (e.g. cameras, lidars and radars) which are transmitting redundantly. For instance, the Apolloscape open dataset from Baidu [1] reports two lidars producing 128 Mbps per device and six video cameras producing 480 Mbps per camera. Safe handling of such workloads exceeds by far bandwidth capacity of the currently used bus-based architectures (e.g. CAN-FD or FlexRay with 10Mbps). Finally, automotive interconnects must still handle increasingly complex low latency traffic from and between different domains, e.g. legacy functions in powertrain or chassis domains, as well as highly interactive driving function and control loops.

Switched Ethernet offers a promising solution for the automotive sector in terms of network capacity, allowing easy scaling of bandwidth, i.e. from 100 Mbps to 1 Gbps and up to 10 Gbps. In addition, it is well standardized in all industries and is used in most products that require networking. This allows easy integration through open network capabilities with other Internet of Things components (e.g. car-to-x communication) and access to many existing products. Another driver for the implementation of Ethernet is the common technology cost achieved by reusing existing research results, mechanisms and solutions, as well as access to the experience of engineering platforms and qualified staff from other areas. However, the transition from bus-based to packet-switched interconnects for automotive purposes is not straightforward as was initially assumed e.g. [2], [3], [4]. In such setups, network must address several challenges which were not initially present in Ethernet performance optimized protocol.

First, in Section II we discuss the problem of isolating data streams with different requirements on network resources and different effects on system safety. To achieve this goal, 802.1 Ethernet has been enhanced with QoS mechanisms such as 802.1Q, AVB or TSN standards, which minimize non-functional dependencies between data streams. However, this is often at the expense of less efficient communication. The frequently used static resource allocation (e.g. traffic shaping, static priorities, predefined time slots) makes communication planning uncomplicated, but leads to a considerable reduction in performance, which is contrary to the requirements of highly dynamic, modern automotive applications. We discuss this challenge in Section III, where we show that achieving safety with dynamic scenarios (e.g. different communication profiles and interference scenarios) is a demanding task. In subsequent subsections we discuss the challenges of dynamic network behavior under permanent and transient fault conditions, security issues, and dynamic energy management.

II. Isolation of Different Traffic Streams

Ethernet-based networks allow multitude of different workloads/transmissions to share links and buffers (following the principle ”route packets not wires”) what has commonly known advantages [5], [6]. However, from a perspective of functional safety, this leads to setups where unlimited interference between packets belonging to different transmissions in switches can endanger the vehicle, its passengers and other transport participants as well as infrastructure. Therefore, the automotive safety standard ISO26262 [7] requires thorough verification of all safety-critical functions (classification through Automotive Safety Integrity Level ASIL [A-D]) before deployment, and "freedom from interference" from all other functions with lower or no ASIL classification (ISO26262: QM). This includes timing interference in the network and protection against transmission faults.

The main challenge arises from the fact that the performance-oriented Ethernet standard, as defined in IEEE

802.1, was not designed with the safety requirement in mind at all. Instead, IEEE 802.1 should promote fairness between traffic flows by offering all equivalent services in an interleaved manner and achieving a high average network performance and utilization.

Consequently, IEEE 802.1 has two major disadvantages from the automotive industry's point of view. Firstly, it is difficult to predict the behavior of the scheduling policy at runtime, e.g. which queue or packet will be selected for transmission on a particular port at a given time. This results in the conservative assumptions about the worst case whenever formal analysis and verification methods must be applied. As a consequence, formal guarantees and analysis results are frequently pessimistic leading to a significant over-allocation/reservation of resources i.e., under-utilization of the network. Additionally, the means to efficiently protect important/critical data streams in IEEE 802.1 are very limited. Therefore it is very difficult, if not impossible in some setups, to achieve low (preferably guaranteed) latencies of highly critical data streams.

These challenges triggered many research and design activities focused on achieving transmission reliability in an Ethernet-based network. Although remarkable progress has been made, there is still a technological gap that must be closed in the future.

A. IEEE 802.1Q and AVB

The safety drawbacks of IEEE 802.1 standard have led to new enhancements to achieve fine-grained QoS. Most of the early efforts were aimed at setups with inherent soft real-time requirements, e.g. audio, video, telecommunications. In such systems, an occasional missed deadline leads to a slight degradation of the user experience without further catastrophic consequences. Because of that, some difficulties in application of these mechanism to safety-critical automotive setups require additional designer attention.

With the IEEE 802.1Q extension of the Ethernet standard, eight priorities were introduced for the allocation of network streams to different traffic classes. Network resource arbitration based on priorities allows a clear separation between traffic classes and is a work conserving policy, i.e. as long as streams/packets are available for transmission, the corresponding switch performs the assignment i.e, individual streams can meet high bandwidth requirements and/or achieve low latencies. However, at the same time designer must account for limited sizes of buffers, and stream characteristics, such as potential bursts, minimum inter-packet distances and maximum arrival rates optimized for evaluation of the worst case behavior. This could drastically limit the performance of the network for all other streams with lower priorities. Additionally, as 802.1Q offers a non-preemptive scheduling transmissions with lower priority may slow down a high-priority packet transmission once per switch and therefore this must be considered at the design stage.

The described drawbacks of IEEE 802.1Q may lead to performance drops and design limitations (especially in the presence of dynamic behaviour of traffic sources), although this scheme is certainly cost-effective and possesses inherent determinism which allows to perform formal worst-case verification analysis and verification.

In order to remedy these shortcomings, the AVB standards were introduced. The AVB is based directly on 802.1Q, but at the same time represents a compromise between the earlier described approaches i.e., IEEE 802.1 and 802.1Q. AVB introduces eight traffic classes: six basic classes (Best Effort) with priority scheduling among them and two stream reservation classes SR-A and SR-B, which are scheduled with Credit Based Shaping (CBS). The CBS approach ensures that the traffic of these classes can only use the transmission resources for a certain period of time. During the remaining time the network resources can be used for transmissions from other, assuring minimum service for all communication participants. low-latency can still be guaranteed for highly critical traffic with tight timing requirements (by putting it into class A with a sufficient bandwidth), while, at the same time, lower-priority streams can be protected from misbehaving of higher-priority streams (by limiting the bandwidth associated with classes A and B). The downside of AVB is complex configuration of switches, especially in setups with higher number of traffic streams. Wrong classification of traffic streams to available traffic classes or errors by configuration of credits (idle and send) slopes for each switch of the network (in between others) can lead to a very poor system performance and effects known as "priority inversion", where less critical traffic receives more resources than a high critical one. Finally, AVB was designed for audio and video streams of infotainment functionalities with maximum delays of 250 us per hop in a carefully designed networks. As shown in e.g. [8], this may not be sufficient for for requirements of ADAS-type communications in some automotive setups.

B. Time-Sensitive Networking

The increasing number of time and safety related requirements for Ethernet-based networks led to the formation of a Time-Sensitive Networking (TSN) task group in IEEE. The main objectives of the group are to improve the QoS of time-critical traffic and to address the issues of fault tolerance (reliability) and security in Ethernet networks. In its work TSN also directly considers the requirements of the ISO26262 automotive safety standard for communication. TSN's efforts have resulted in a series of standards dealing with these real-time requirements, some of which are still in draft form and some of which have already been completed. Some prominent examples are highlighted below, a more detailed description can be found in [9], [10].

Primary efforts of the TSN task group considered the problems of isolation introduced already in previous subsection II-A. IEEE 802.1Qbu adds management and configuration mechanisms for frame preemption. The higher priority frames may preempt the transmission of lower priority frames by splitting the lower priority frame into two or more fragments. This allows further reduction of blocking time by lower priority traffic. IEEE 802.1Qci introduces per-stream filtering and policing for containment of faults and ensure that traffic stays within predefined limits. IEEE 802.1Qcr provides an even stricter separation of traffic classes through either per-stream credit-based shaping. IEEE 802.1Qbv introduces through Time Aware Shaper (TAS) time-driven scheduling for link access between different traffic classes. For each traffic class in the switch, the scheduler at an output port opens or blocks access to the link (a gate) for the predefined amount of time, i.e. the gate is open only during the predefined time intervals and that it is closed otherwise [9],[10]. Additionally, for fault-tolerant communication has been considered e.g., standards for time synchronization (802.1AS-rev), reliable transmission (802.1CB). The latter mechanism introduced by TSN IEEE 802.1CB offers a viable approach for achieving fault tolerance through frame replication and elimination (FRER). According to 802.1CB a frame is copied and transmitted simultaneously via redundant paths (Spatial FRER). Alternatively, frame copies can be send redundantly through the same path (Temporal FRER).

Only this narrow selection of problems covered by the TSN Group shows that there is a wide range of topics which must be considered by these standards. However, with the increasing number of safety mechanisms in switches the configuration of the network becomes to be more and more challenging. Moreover, many of standards are in the draft form e.g. mechanisms for reliability and fault tolerance need additional support as well as the mechanisms for security. Finally, the mechanisms for safety are robust and static, thus efficient TSN configurations for the whole network with plethora of data streams may be very time-consuming and complex. At the same time, high configuration complexity leads to a high potential of design errors and misconfigurations, which increases costs (longer testing time is required, likely service issues after product launch).

Let's consider, for example, the frame replication mechanism mentioned above (FRER, IEEE 802.1CB). Figure 1 provides a schematic diagram of a switch with a time trace of the frames arriving at its ports (frames 1, 2, 3, 1*, 12, 13). In this scenario, copies of the same frame (1 and 1*) are sent through the network on the redundant paths leading to two different ports of this switch, one of the paths being shorter than the other. If the original frame 1 is lost due to the transient error (represented by the lightning), the receiver still receives its copy (frame 1*) on the second port (if both frames would arrive, the switch would eliminate copy 1*.). However, the 802.1CB standard does not prevent the frames (transmissions) from arriving without preserving their order, which could lead to erroneous behavior of the system in case of sensor data. Note that the sequence or frames have been changed (2, 3, 1* instead of 1, 2, 3), as we have to take into account the longer transmission latency on the second path (i.e. higher number of hops and arbitration overhead as well as additional, dynamic interference). This means that in the case of FRER the preservation of the frame arrival order has to be implemented manually, which increases the complexity of the design by an additional vector.

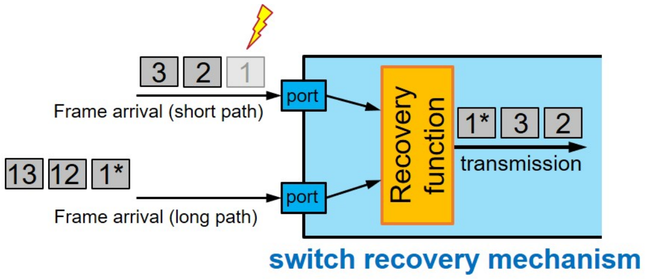
3
2
1
Frame arrival (short path)
port
13
12
1*
Frame arrival (long path)
port
Recovery function
1*
3
2
transmission
switch recovery mechanism

Fig. 1. Out of order transmission of frames in a switch implementing FRER (IEEE 802.1CB) mechanism, where at the presence of error (dropped frame 1) a duplicated frame 1* arrives late what changes the arrival pattern.

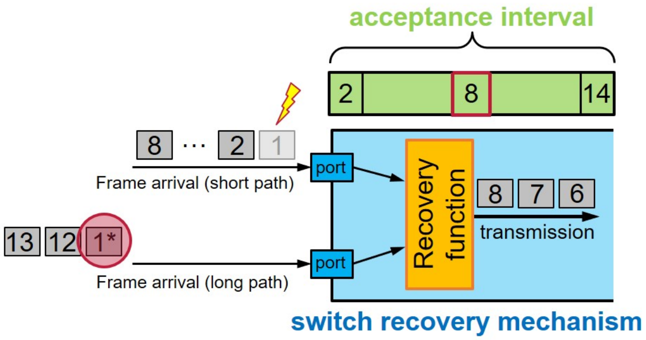
acceptance interval
2
8
14
8
···
2
1
Frame arrival (short path)
port
13
12
1*
Frame arrival (long path)
port
Recovery function
8
7
6
transmission
switch recovery mechanism

Fig. 2. Example of a misconfigured switch implementing FRER (IEEE 802.1CB) mechanism, where at the presence of error (dropped frame 1) a duplicated frame 1* is being drop as it arrives outside of acceptance interval.

Figure 2 shows another example of a possible misconfiguration using the same mechanism (FRER, IEEE 802.1CB), which could have potentially fatal consequences for the safety of the network and the entire system. We are looking at the same switch as in the previous Figure 1. Due to limited hardware resources, any switch implementing FRER must have a certain acceptance interval during which it keeps track of which packets have been seen and which have not, and during which it could eliminate the duplicates. Therefore, if the copy of frame 1* were to arrive outside of this interval (without this interval being set correctly), it could be dropped or treated as a new independent transmission. Consequently, the network would get rid of the valid frame. Such design flaws are difficult to debug.

In summary, TSN presents a promising set of features that could significantly improve the safety of automotive designs. At the same time, it should not be forgotten that standardization concerns compatibility and does not limit diversity. Designers have to compete with increased protocol and circuit complexity as well as switch costs. And finally, they must confront a multitude of configuration and misconfiguration options, i.e. standardized does not necessarily mean ”out of the box” in a safe way (see examples in the Figures 1 and 2).

III. Integration of Dynamic Workloads

The majority of engineering and research efforts in the existing automotive systems concentrate on the critical communication requiring hard real-time guarantees. Therefore, critical network workloads are well defined and tested e.g. known transmission lengths and jitters, predefined access patterns (usually periodic). However, contemporary automotive network design must consider also efficient integration of dynamics in system behavior. For instance, workloads which originate from the new and old applications in automotive systems e.g. cloud communication, over-the-air updates, car-to-x communication, selected modules from the ADAS introduce sporadic transmissions and the access patterns are often unknown at design time – consequently they introduce dynamics to the interconnect traffic patterns. Additionally the amount of data transmitted by senders may vary significantly at runtime. Frequently also these transmissions would require some temporal guarantees. These guarantees however are not that strict as in case of the critical communication e.g. looser deadlines, service guarantees per burst of packets, allowed retransmissions etc..

In many present and future automotive setups, e.g. cross domain centralized architecture from Bosch [11], these two traffic classes (critical and dynamic) must co-exist in the same network infrastructure. Therefore, they compete for the network resources. Unfortunately, the frequently used static allocation of network resources (e.g. traffic shaping, static priorities, predefined time slots) makes communication scheduling uncomplicated, but leads to a considerable reduction in performance whenever dynamics must be considered. For instance, application of the TSN standard allows solving many problems of the critical communication. However, at the same time guarantees offered by the TSN standards come at the high performance penalty for the traffic from dynamic communication domain i.e., transmissions which cannot be optimized w.r.t to time-aware shaper settings. If stream characteristics are modified at runtime, or new streams are added to the network, additional configuration of the system and its TSN components is necessary. Therefore, TSN provides a very limited support for an adaptive behaviour of the system, which is a crucial requirement implementing fail-operational behaviour and high-end driver assistance functionalities. When it comes to applying TSN to even a static workload use-case, system-wide network optimizations might be necessary, see [12]. Hence, TSN standards present powerful mechanisms, which unfortunately, do not offer out-of-the box solutions to all challenges in the automotive industry, but rather present only the tools that should be carefully and thoughtfully applied to the system, preferably in conjunction with other mechanisms, such as service-oriented protocols (e.g. SOME/IP and DDS).

Similarly, the standard Ethernet (IEEE 802.1Q standard) addresses the problem of integrating different workloads by introducing eight prioritized traffic classes. Since there are typically more traffic flows than traffic classes in a network, the flows must share the traffic classes. This leads to security risks because arbitration between frames of the same traffic class is usually done in FIFO order. Therefore, such arbitration could lead to denial of service attacks due to network saturation or missed deadlines due to incorrect node behavior. This also applies to Ethernet AVB and the upcoming Ethernet TSN (IEEE 802.1Qbv), which introduce further traffic shapers in addition to IEEE 802.1Q, e.g. to prevent the starvation of lower priorities or to implement a time-controlled transmission of latency-critical traffic. The additional major disadvantage of this solution is that the rates in the shapers are statically adjusted according to the worst-case scenario, i.e. assuming that all traffic sources are running simultaneously at the maximum possible transmission arrival rates. Therefore, this approach is not work preserving and sacrifices network utilization if the senders expose dynamics in execution time, release jitter or communication volume and the system is not heavily loaded. Even if time-dependent shaping is applied, the large number of different messages almost inevitably leads to slot sharing and thus to the already mentioned head-of-line blocking threats.

For instance, an over the air update (e.g. as applied by Tesla Inc.) would be sporadic with respect to other communication ongoing in the network. Assigning statically resources for communication with the cloud using one of the mechanisms from the previous section e.g. AVB or TSN TAS, would reassure that there are no lost frames (due to bufferoverflows or timeouts). At the same time, however, the assigned resources would not be used by the majority of network uptime (updates should be sporadic in nature). Similar scenario could happen with traffic originating from car-to-car or car-to-infrastructure communication where transmissions are scheduled adhoc and time of communication establishment is hard to predict. Also ADAS workloads in the network could have different profiles, e.g. depending on the weather conditions, or situation of the road - city or highway modes.

A. Dynamics due to faults in the network

Furthermore, not all sources of dynamics can be easily taken into account in the design phase of the automotive system. A good example are transient and permanent errors in the network. As discussed in many related papers [2] and standards, e.g. BroadR-Reach [13], it is necessary to consider transient and permanent errors in the Ethernet infrastructure. Figure 3 presents error rates typically considered in automotive domain. For example, for the two-wire Ethernet a Physical Layer Transceiver (PHY) chip is compliant with the BroadR-Reach standard [13] if it has a bit error rate (BER) of less than or equal to 10^{-10}. Although this may seem relatively low for a single connection, for a 100 Mbps Ethernet, this corresponds to a bit error approximately every two minutes. Such a high transient error rate is only acceptable if transient error handling is part of normal operation, including time constraints.

The solution frequently chosen here is end-to-end error control with ARQs. It allows to limit the overhead to critical messages/streams and covers all error types (link, tail drop etc.). Various forms of ARQ protocols has been proposed in the literature and research (see [5] or [14]), the most popular

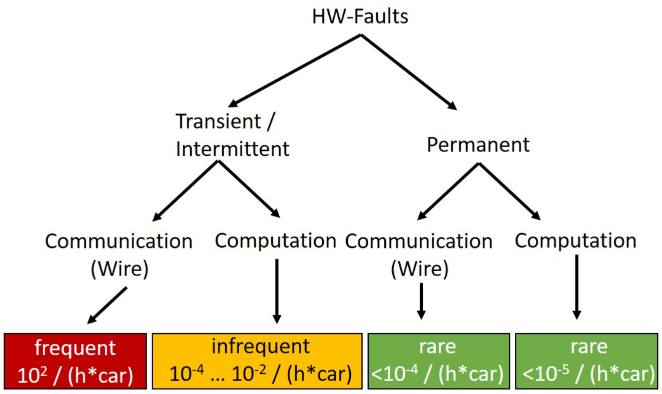

Fig. 3. Estimated error rates for Ethernet-based communication based on the BroadR-Reach [13] standard. Note: The resulting calculation errors depend heavily on the state protection (memory).

of which are Stop and Wait, Go-back-N and Selective Repeat. However, end-to-end re-transmissions and/or software roll-backs could introduce dynamic and burst-like traffic flows (e.g. Automatic Repeat Request - ACK N) into the network. Similar to the scenarios described above, permanent consideration of these dynamic streams in the shaper settings would introduce a static and constant overhead for other traffic classes i.e. decrease performance of the network.

While the high frequency of transient faults is one of the main design concerns in automotive networks, achieving the required functional safety also requires a treatment of permanent faults effects. Although permanent faults are many orders of magnitude less likely then transient failures, dealing with them may require a change of the network operation to a fault operation mode. In the Ethernet context, permanent faults are ones that leave a certain path (selected switches and links) in an error state, so that even if messages are continuously re-transmitted, e.g. due to aforementioned an ARQ protocols, the message cannot reach the destination. The affected element of such a path can be either the switch or the link between switches, i.e. the cabling including the connectors.

The most straightforward solution is to use a fully redundant network hardware infrastructure (e.g. duplication of all cables and switches). Although it is effective, the application of this method in the automotive domain is usually not possible for cost or overhead reasons (power consumption, infrastructure weight, cable placement, production and service costs etc.). Therefore protocol-based alternatives has been considered. One possible countermeasure could be a mechanism similar to the TSN FRER protocol (IEEE 802.1CB), in which frames can be copied and sent simultaneously over redundant paths. This allows fast recovery, but at the same time results in a constant overhead - the same traffic appears twice in the network. On the other hand, the automatic path detection and routing known from the performance-optimized Ethernet networks [5] leads to high latencies and unpredictable (possibly unbounded) interference. Newly redirected traffic can endanger the safety of other transmissions over available paths.

Consequently, the major challenge of run-time adoption to dynamic workloads requires new mechanisms introducing robust network operation with real-time error handling.

B. Security Challenges

Security challenges, including malicious attacks on interconnects, are a well-known and important design aspect of vehicular safety, cf. ISO26262 standard [7]. In legacy vehicle systems, extensive testing of software functions and preventing of attacker's physical access to the components would significantly reduce the associated security risks connected with communication due to the relatively low complexity of bus-based connections. With the emergence of new functions, e.g. automated and autonomous driving, car-to-x communication, the complexity of the switched network and the entire system has increased significantly. The developers have to deal with a large number of new communication channels and a variety of mechanisms, as we have shown in the previous sections.

This trend, which transforms the vehicle electronics known from bus-based control ECUs into complex, interconnected distributed systems based on switched networks, brings new set of security challenges that affect the safety of vehicle [15]. Internal threats such as functions that have not been fully tested, misconfigurations or malicious software can not only lead to unexpected behaviour or system failures, but also be exploited by attackers. Although some of these threats are known from legacy architectures, their handling is becoming increasingly difficult. For example, the data stream initiated by sensor fusion for body or chassis domains often needs to be routed over a switched network as well as a gateway component and finally over a conventional bus network. Each of these system elements offers new attack vectors and opportunities for security breaches. In addition, there is a new set of external threats associated with the communication of electronic components in the car with the outside world. For example, car-to-car, car-to-infrastructure or car-to-cloud connections open up new opportunities for intruders to attack, such as manipulating over-the-air updates of automotive software. The collection of data from the outside world and the traffic infrastructure by various sensors (radar, lidar, video cameras) also makes the car vulnerable to manipulation of the received data, e.g. sensor spoofing.

For this reason, the intrusion detection as well as isolation through various policies (usually stream/class dependent) are currently an important topic of research and development work. TSN standards, cf. Section II-B, also consider security as a primary topic, which can have a considerable influence on the functional safety of the vehicle. The IEEE 802.1Qci standard, for example, introduces mechanisms for detecting and mitigation of interfering transmissions (which could be initiated by attacker) using filtering applied per data stream. However, as in the case of other mechanisms described above, security countermeasures typically introduce a set of static rules and policies at the expense of network performance and utilization. As a result, pessimistic assumptions are made during system design, expecting traffic from attackers or

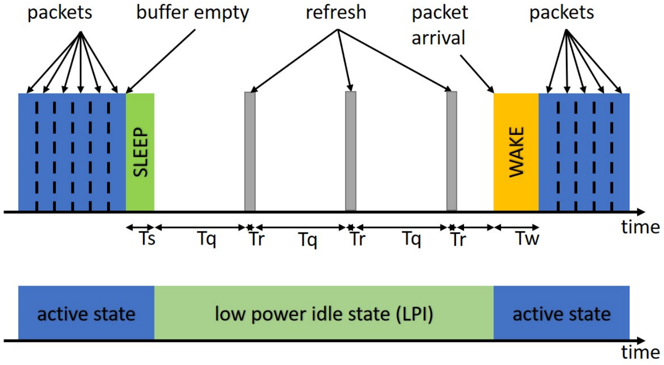

Fig. 4. Energy efficiency protocol sequence while packet transmissions

misbehaving software in the worst case. The dynamic approach, based on monitoring and adjusting policies at runtime, introduces additional delays during which malicious activity (e.g., denial of service traffic) could disrupt security-critical communications. Also countermeasures, such as rollbacks and re-configurations may introduce dynamically new traffic into the network causing packet loss or congested connections for other streams.

C. Dynamic Energy Management

Low energy consumption is an essential factor for most electrical devices and is becoming to be a concern also in automotive domain, in between others, due to the constantly increasing number of hybrid and electric vehicles. Therefore, energy efficiency for the automotive network is a promising research topic, which has to take into account functional requirements such as the provision of significant processing and networking power and availability in all modes of system work.

So far, adaptive power management has been considered for data centers and home networks. The IEEE 802.3az standard used in non-automotive networks allows inactive network elements (e.g. switches) to be put into ”sleep” mode with low power consumption. Figure 4 shows a schematic diagram of this protocol sequence with a representation of the device state and a time trace of the packet arrival. Once the buffer in the switch is empty for a period of Ts, the device can enter the low-power idle (LPI) state. While it is in LPI state, it periodically sends an update signal to indicate that this network component is still active and not broken. As soon as a new packet arrives at the device, the device wakes up (additional delay) and sends data after the time Tw.

From the perspective of safety this LPI state rises new challenges which must be handled. For instance, this behavior (i.e. longer processing of the first packet) can cause dynamic bursts of accumulated packets and therefore overloaded buffers forming an important factor introducing dynamic workloads. Moreover, in case of high network load the packet inter-arrival latency is irregular and can cause unpredictable packet delay. Therefore it could be very difficult, if not impossible in some setups, to achieve low (preferably guaranteed) latencies of highly critical data streams.

IV. Summary

Switched Ethernet can be the key enabling networking technology for the automotive domain. However, we have shown that the safety mechanisms introduced by AVB or TSN, while effective, are robust and static, so efficient network configurations can be very time consuming and complex, especially when dynamic workloads must be considered. At the same time, high configuration complexity leads to a high potential for design errors and misconfigurations, which increases costs, so a well thought-out application is required. A plethora of standards can lead to an ever-increasing protocol and circuit complexity as well as increasing production costs. Therefore the application of Ethernet in the automotive domain requires a systematic approach and thoughtful consideration in order not to lose its advantages.

References

[1] P. Wang, X. Huang, X. Cheng, D. Zhou, Q. Geng, and R. Yang, “The apolloscape open dataset for autonomous driving and its application,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, pp. 1–1, 2019.

[2] M. Möstl, D. Thiele, and R. Ernst, “Invited: Towards fail-operational ethernet based in-vehicle networks,” in *2016 53nd ACM/EDAC/IEEE Design Automation Conference (DAC)*, pp. 1–6, June 2016.

[3] F. L. Soares, D. R. Campelo, Y. Yan, S. Ruepp, L. Dittmann, and L. Ellegard, “Reliability in automotive ethernet networks,” in *2015 11th International Conference on the Design of Reliable Communication Networks (DRCN)*, pp. 85–86, March 2015.

[4] J. Huang, M. Zhao, Y. Zhou, and C. Xing, “In-vehicle networking: Protocols, challenges, and solutions,” *IEEE Network*, vol. 33, pp. 92–98, January 2019.

[5] A. S. Tanenbaum and D. J. Wetherall, *Computer Networks*. USA: Prentice Hall Press, 5th ed., 2010.

[6] W. Dally and B. Towles, *Principles and Practices of Interconnection Networks*. San Francisco, CA, USA: Morgan Kaufmann Publishers Inc., 2003.

[7] ISO26262, “Road vehicles – Functional safety,” 2011.

[8] J. Migge, J. Villanueva, N. Navet, and M. Boyer, “Insights on the performance and configuration of avb and tsn in automotive ethernet networks,” *Proc. Embedded Real-Time Software and Systems (ERTS 2018)*, 2018.

[9] “P802.1Qbv (Draft3.0)-Enhancements for Scheduled Traffic http://www.ieee802.org/1/pages/802.1bv.html,” standard, IEEE Time-Sensitive Networking Task Group, USA, Mar. 2016.

[10] D. Thiele, R. Ernst, and J. Diemer, “Formal worst-case timing analysis of ethernet tsn’s time-aware and peristaltic shapers,” in *2015 IEEE Vehicular Networking Conference (VNC)*, pp. 251–258, Dec 2015.

[11] W. Haas and P. Langjahr, *Cross-domain vehicle control units in modern E/E architectures*, pp. 1619–1627. 01 2016.

[12] D. Ziegenbein and A. Hamann, “Timing-aware control software design for automotive systems,” in *2015 52nd ACM/EDAC/IEEE Design Automation Conference (DAC)*, pp. 1–6, June 2015.

[13] “BroadR-Reach Physical Layer Transceiver Specification For Automotive Applications - Version 3.0,” standard, Broadcom Corporation, San Jose, USA, May 2014.

[14] Shu Lin, D. J. Costello, and M. J. Miller, “Automatic-repeat-request error-control schemes,” *IEEE Communications Magazine*, vol. 22, pp. 5–17, December 1984.

[15] I. Studnia, V. Nicomette, E. Alata, Y. Deswarte, M. Kaâniche, and Y. Laarouchi, “Survey on security threats and protection mechanisms in embedded automotive networks,” in *2013 43rd Annual IEEE/IFIP Conference on Dependable Systems and Networks Workshop (DSN-W)*, pp. 1–12, June 2013.