LEVEL 5 BY LAYER 2: TIME-SENSITIVE NETWORKING FOR AUTONOMOUS VEHICLES

Soheil Samii and Helge Zinner

ABSTRACT

Time-Sensitive Networking standards for Ethernet provide real-time and dependability mechanisms such as traffic shaping and scheduling, time synchronization, and redundancy. This article provides a review of these standards in light of possible future use cases in automotive systems using in-vehicle Ethernet networks.

INTRODUCTION

Today's vehicles are equipped with a multitude of multimedia and infotainment applications, as well as active safety and advanced driver assistance systems (ADASs). These application areas and their corresponding sensors and devices drive bandwidth requirements in the automotive industry at an unprecedented rate. These range from onboard services like navigation systems, amplifiers, and connectivity units to off-board services connecting the car to the Internet and to external devices. Premium vehicles often serve as WiFi hotspots and offer USB interfaces for device charging. These examples all show that IT-based technologies have found their way into the automotive electronic architecture. Ethernet is an established and proven technology that is currently finding its way into automotive systems, with new physical layer standards for 100 Mb/s[1] and 1 Gb/s standardized in the IEEE 802.3 Working Group as 100BASE-T1 and 1000BASE-T1, respectively. As part of this transformation, new layer 2 solutions, standardized by the Time-Sensitive Networking (TSN) Task Group of the IEEE 802.1 Working Group [2], are enabling high-speed Ethernet and IP communication to enter various segments of the electronics and electrical architecture of modern vehicles. Our focus in this article is to highlight the role of these new layer 2 TSN standards in the ongoing industry transformation to support future mobility solutions with autonomous driving technology. As indicated in the title, we argue in this article that the TSN standards provide real-time and dependability enhancements to switched Ethernet technology to support autonomous driving systems up to Level 5, the highest driving automation level defined by the Society of Automotive Engineers (SAE).

ESTABLISHING THE BANDWIDTH FOUNDATION

In addition to increased off-board connectivity, the vehicle's onboard networking requirements demand increased networking capability. Common automotive applications in functional areas like ADASs already require processing of high-resolution sensor data. The amount of this type of real-time transmission of data is increasing tremendously. This increase is driven, for example, by machine-vision video, as well as radar and lidar data. Furthermore, next-generation cars will likely require significantly more sensors than today's ADASs. This may require a combination of more network connections within the vehicle and standardized automotive-grade high-speed data technologies. Today's car infrastructure may need to grow dramatically in order to accommodate the upcoming increase in number and data production of sensors like radars, lidars, and cameras.

On top of an increased number of network elements and bandwidth, existing automotive communication technologies such as controller area networking (CAN), MOST, and FlexRay are running out of steam. The video links used in today's automotive systems are based on low voltage data signal (LVDS) technology, which is a point-to-point link to a hardware module and does not inherently provide packet forwarding or routing capabilities. Packet networking is extremely useful as the video stream is often used by multiple modules in the system, and a network would allow flexible forwarding and/or multicast of video packets. If there was no network, a daisy-chain LVDS connection needs to be used, or application software needs to be added to pass video data to other modules on, for example, an Ethernet network.

CAN is the predominant communication network in use within automotive (electric/electronic) (E/E) architectures [9]. The typical data rate is 2–5 Mb/s for the latest CAN version (CAN with flexible data rate, CAN-FD), and the bandwidth must be shared with other nodes on the bus in a "one-at-a-time" fashion. The automotive FlexRay technology provides a higher bandwidth of 10–20 Mb/s and is utilized wherever there is a demand for hard real-time communication. This time-triggered approach allows for predictable delays independent of the overall bus load. However, it requires a complex configuration process, as well as synchronization of the network communication schedule and software task scheduling across multiple electronic control units (ECUs).

The current and growing set of functionality provided by IEEE Ethernet offers a solution to these issues. The current automotive Ethernet link (IEEE 802.3bw) specifies a data rate of 100 Mb/s full duplex over a single unshielded twisted pair. Its adoption continues to grow in the automo-

---

[1] In IEEE Communications Society magazines, Mb/s indicates megabits per second, and MB/s indicates megabytes per second.

Digital Object Identifier: 10.1109/MCOMSTD.2018.1700079

Soheil Samii is with General Motors/Linköping University; Helge Zinner is with Continental.

2471-2825/18/$25.00 © 2018 IEEE

Project	Purpose	Status	Potential automotive application scenario
P802.1AS-REV	Robust time synchronization	Ongoing	Common notion of time for sensor fusion (already provided in 802.1AS-2011); fault-tolerant time synchronization with backup grand master, necessary for functionality that needs to be available in fallback mode (e.g., in Level 4 and 5 systems).
802.1CB	Redundant communication paths	Finished	Fallback capability, fail-operational applications tolerating connector and wire faults
802.1Qbu	Frame preemption	Finished	Steering and braking actuation
802.1Qbv	Time-triggered, scheduled communication	Finished	Closed loop control (e.g., steering and braking), interprocessor communication, periodic sensor data
802.1Qca	Path reservation	Finished	Adaptive architectures with runtime path reservations
P802.1Qcc	TSN stream and network configuration	Technically stable	Adaptive architectures with runtime stream reservations
802.1Qch	Cyclic scheduling	Finished	Periodic or sporadic traffic from sensors
802.1Qci	Ingress filtering and policing	Finished	Network protection due to random hardware failures; detect and mitigate malicious intrusions or denial of service attacks
P802.1Qcr	Asynchronous traffic shaping	Ongoing	Low-latency communication for sporadic or aperiodic traffic (e.g., radar or lidar sensors)
802.3br	MAC support for frame preemption	Finished	Steering and braking actuation

Table 1. Automotive relevant standards in TSN.

tive industry because of its bandwidth and cabling simplicity. The automotive two-wire Ethernet technology for 1000 Mb/s (IEEE 802.3bp) is standardized, available in silicon, and clearly paves the way for more high-speed data transport. 1000 Mb/s will provide the next level of bandwidth to support sensor stream transport. Furthermore, to address the tremendous bandwidth appetite of raw sensors and core network applications, the IEEE recently began a new project to support 10 Gb/s transmission rates [1]. Another recent IEEE initiative is exploring solutions for 10 Mb/s that aims to support multiple low-speed sensors in a cost-effective multi-drop configuration. The automotive community is beginning to see a convergence of standardization and technology development efforts, ultimately leading to a suite of standard automotive Ethernet products supporting a large range of bandwidth requirements, from low-megabits-per-second to multi-gigabits-per-second application needs.

Ethernet TSN

In addition to the continuously rising demand for bandwidth, future applications like self-driving cars are likely to require a bridged/switched local area network architecture that can provide dependable and real-time data transport guarantees to the applications running in the autonomous vehicle. Ethernet TSN is gaining more and more interest since its scope is to address future automotive communication needs, including precise time synchronization, generally dependable, deterministic communication capability, ultra-low latencies, and fault tolerance.

While the current set of TSN standards address many of the envisioned future automotive requirements, we can see hints of future standardization work needed. For example, as opposed to the previous generation of audio video bridging (AVB) standards, not all of the TSN standards in development are compatible with plug-and-play requirements. To give some specific examples, the time-aware shaper (802.1Qbv) and the preemption standard (802.1Qbu) require static or managed configuration of the network. This means that there is no layer 2 protocol to automatically configure new streams (or delete existing streams) at runtime and automatically decide schedules and preemption parameters. While this is not a limitation considering today's automotive industry practice with engineered, static networks and carefully managed configuration and updates, future vehicle architectures may require dynamic and runtime network management, which could be addressed by future 802.1 standardization efforts. In fact, one of the TSN standards, IEEE 802.1CB for seamless redundancy (frame replication and elimination), already introduced an auto-configuration protocol to establish redundant paths; however, any scheduling-related parameters, such as those for time-aware shaping (802.1Qbv), cyclic queuing and forwarding (802.1Qch), and preemption (802.1Qbu), are managed either statically or by an application above layer 2.

Ethernet TSN is a core subset of the IEEE 802 standards being developed to support reliable communication in 802 local area networks [3]. The 802.1 Ethernet TSN Task Group focused on achieving deterministic, low-latency, and fault-tolerant data transport. TSN grew out of the group that developed the existing IEEE 802.1 AVB standards [11] and enabled the transmission of multimedia content across an automotive E/E architecture. In addition to the standardization work in IEEE 802.1, AVnu has developed AVB profiles for different markets (e.g., for the automotive and industrial automation industries) [15]. At its heart, AVB specified the concept of synchronizing clocks to sub-microsecond and reserving bandwidth together with the credit-based shaper mechanism to ensure sufficient network bandwidth and a guaranteed maximum transmission latency adequate for the transport of streaming audio/video media content in the car, primarily for infotainment purposes. ADAS cameras and sensors also use streamed data, but these streams have different requirements in terms of data transmission reliability. Streams used for the automation of the driving tasks are classified as safety-critical traffic where loss of data may result in a dangerous driving situation.

TSN enables substantial simplification of the challenging implementation of safety-critical automated driving systems by providing common and

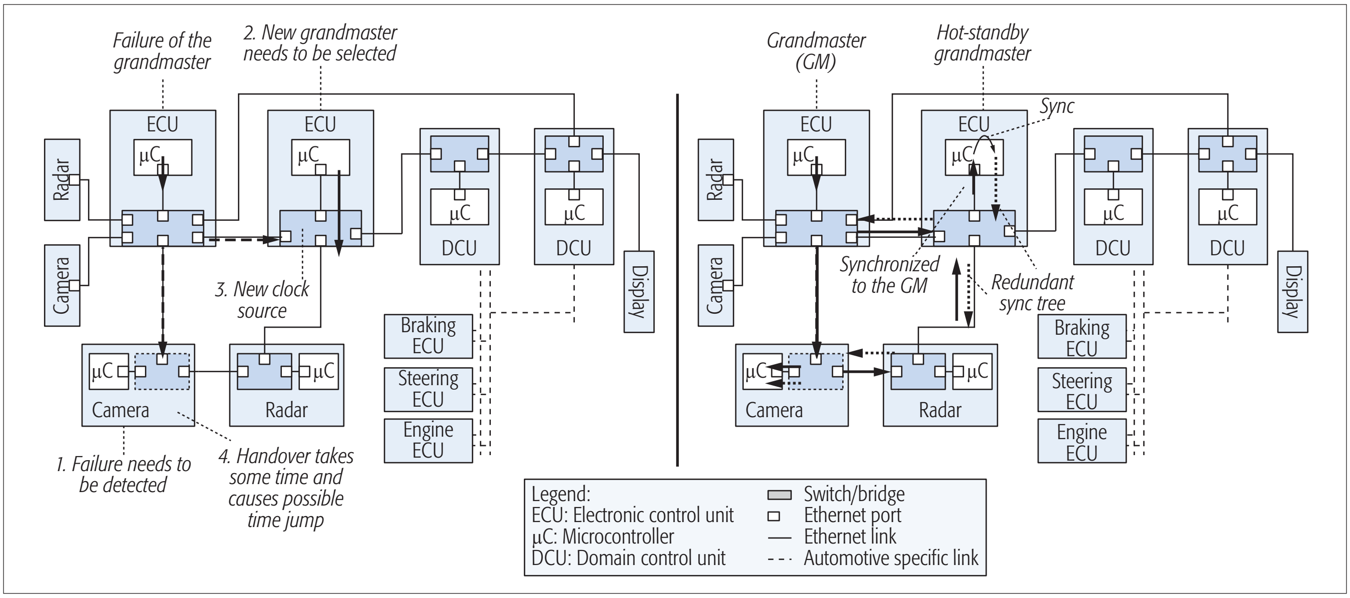

Figure 1. Time-aware network with redundant sync trees and grandmasters (right side).

standard methods for key network functions that all silicon providers, equipment makers, and automotive manufacturers can use to achieve their design goals [4]. Table 1 gives an overview of the current status and of TSN projects and their application to automotive networks.

REDUNDANCY

TSN standards provide mechanisms for redundancy in two key areas for autonomous vehicles: network time masters and data transmission paths. For time source redundancy, the Best Master Clock Selection Algorithm (BMCA) of 802.1AS-2011 is able to handle the loss of a link, port, or the active grandmaster itself. In the case of a grandmaster failure (resulting in loss of the Generalized Precision Time Protocol, gPTP, Announce frame), the remaining grandmaster-capable devices in the network send Announce frames and eventually, based on the distributed agreement protocol defined for BMCA, determine the new grandmaster. This handover to the new grandmaster may take multiple seconds and is slower compared to existing fault-tolerant automotive time synchronization approaches (e.g., FlexRay). Such delays may not be tolerated for some applications. In addition, the handover could result in a time jump at some or all receipients of the new clock reference (Fig. 1). One goal of the new version, 802.1AS-REV, is to provide faster reaction to grandmaster time failures and thus enable high availability for the common global clock, a basic function for the TSN standards. It achieves this by the parallel usage of multiple (and synchronized) grandmaster clocks (hot standbys) in conjunction with multiple clock synchronization paths.

TSN standards also provide support for redundant data transmission. The TSN group recognized that packet retransmission is not always a possible solution for delay-sensitive real-time applications (e.g., sensor data for machine vision systems). Thus, if a transmission error occurs on the link, the original data is lost because retransmission is impractical. IEEE 802.1CB, "Frame Replication and Elimination for Reliability" (FRER), solves this problem by having end stations send two copies of a frame over disjoint paths through the network, thus increasing the probability that this message will be delivered on time. FRER standardizes hardware support for the duplication and elimination of frames in Ethernet switches, enabling silicon vendors to provide hardware support in their products. 802.1CB also provides an optional function for proxy mode of operation (Fig. 2), where the switch performs all duplication and elimination, thus avoiding any hardware impact or intensive software processing on the host microprocessor(s). This is especially useful in many automotive applications where the end station and switch exist on the same board (i.e., all integrated into the same ECU). In other words, 802.1CB can provide redundancy in a transparent manner to the end stations in the network, leaving the microprocessors with only the burden of identifying and marking the streams that shall be transmitted on redundant paths.

NETWORK AND QOS PROTECTION

Future Ethernet-based core architectures will likely support the transfer of a multiplicity of sensor streams. Therefore, multiple streams are likely to be merged into the same switch egress queue, creating the potential for inter-stream interference. If a sensor stream exceeds its configured bandwidth, packets from other streams in that same queue may miss their latency goals — or worse, they might not even be able to enter the queue due to a buffer overrun. 802.1Qci, "Per-Stream Filtering and Policing," constantly checks whether an incoming stream meets its requirements at the ingress to the bridge. It also supports traffic flow metering and monitoring with frame classification based on a configurable set of parameters. Upon detection of excess bandwidth usage, an 802.1Qci filter can apply several actions — such as blocking all future frames of this particular stream — to maintain the QoS requirements of the other compliant streams.

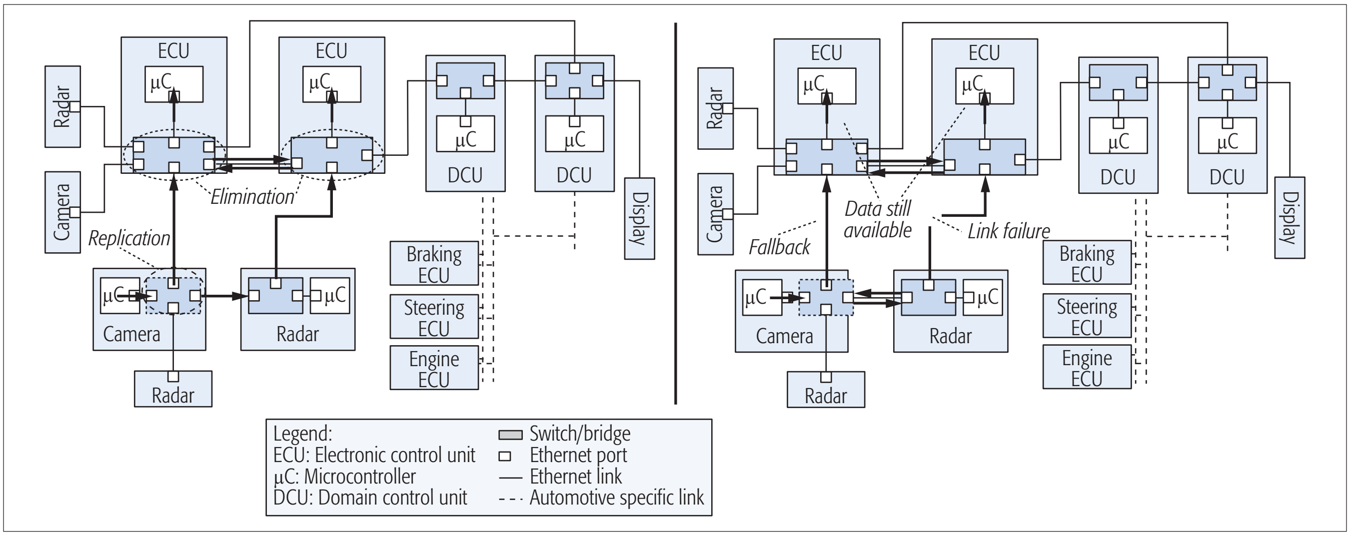

Figure 2. Hardware-based replication and elimination of camera data. Data are still available to the receiver (left side) in case of link failures between the camera and the receiving ECU.

DETERMINISM AND LOW LATENCY

The design goal of the AVB group was a maximum latency of 2 ms over 7 hops in a 100 Mb/s network for audio and video traffic. Recent TSN activities had the goal to set the maximum latency to 100 μs over 5 hops. This objective supports some potential future automotive networking scenarios, particularly for tight control applications such as steering, braking, and propulsion over Ethernet. ECUs providing these functions in a vehicle can be seen as the actuators for autonomous driving control applications. Commands to such ECUs affect the longitudinal and lateral control of the vehicle. These commands benefit from short and deterministic end-to-end latency, especially at high vehicle speeds (Fig. 3). To address this desire, TSN added three more traffic shapers to provide a variety of solutions that allow the network designers to trade off between implementation/configuration complexity and latency [13]. 802.1Qbv and 802.1Qch implement time-aware queue-draining procedures on each queue at the egress ports of every bridge in the path. This enables bridges and end stations to schedule the transmission of frames based on timing derived from the IEEE 802.1AS grandmaster. This so-called time-aware shaper (TAS) implementation can achieve the theoretical lowest possible worst case latency for an engineered automotive network with the help of a common global clock.

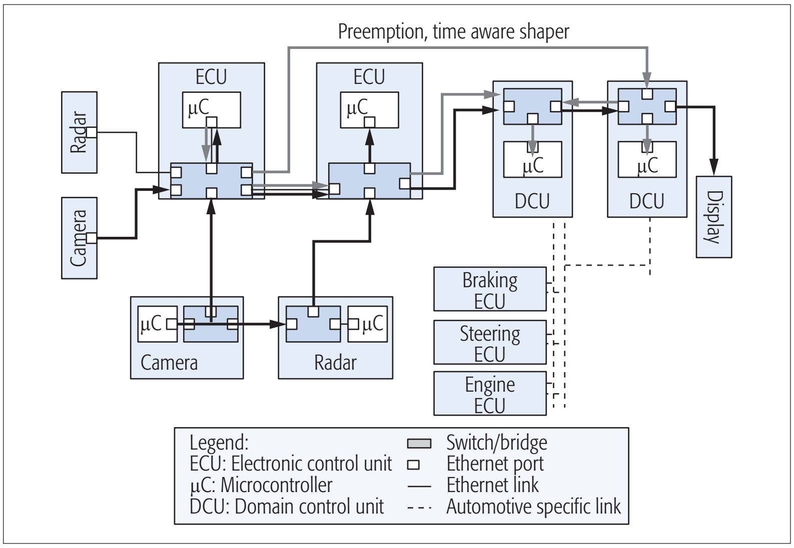

Figure 3. Dependable real-time Ethernet core network supporting sensor fusion control and actuation using the full TSN stack.

The TSN group is currently also defining a less complex but low-latency traffic shaper solution referred to as the asynchronous traffic shaper (ATS) in the P802.1Qcr project, which was started in late 2017. ATS considers a traffic model that supports a large variety of traffic patterns (e.g., periodic, sporadic with a minimum inter-arrival time, and with bursts). Figure 4 demonstrates its operation. An ATS shaper is based on the classic "leaky bucket" principle of network traffic shaping. An ATS shaper accumulates credit at the committed information rate for each flow, up to the capacity of the bucket. When a frame arrives on the shaper's queue, the shaper allows transmission as soon as the accumulated credit is greater than or equal to the size of the frame, and debits its credit balance by the size of the frame. In this way, the ATS shaper spreads out (i.e., shapes) traffic, but also allows bursts of short packets to transmit quickly as long as the burst is less than the size of the shaper's bucket. Figure 4 shows several examples. The blue line shows a flow with relatively large frames occurring at very regular intervals. In contrast, the red line shows a flow with small frames grouped in bursts. The green flow is somewhere in between the blue and red flows. The goal of ATS is to provide deterministic real-time communication without relying on a global notion of time in the network. Therefore, in contrast to TAS, ATS will continue to operate correctly even in the case of time-synchronization errors in the network.

FRAME PREEMPTION

In addition to various real-time scheduling policies at egress, the TSN group developed a frame preemption technology in collaboration with the 802.3 Working Group. The overall goal is to allow

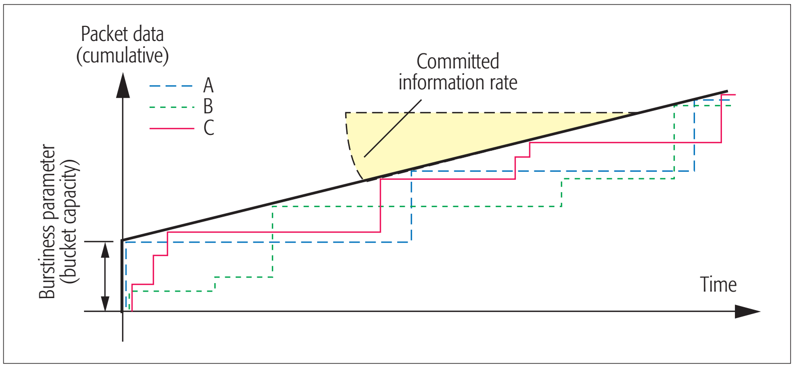

Figure 4. Three traffic patterns (A, B, C) satisfying the same bandwidth model of a flow [6]. The bold black line represents the limit by the burstiness parameter (also known as the bucket capacity) and the committed information rate.

some frames to preempt ongoing frame transmissions in order to accelerate certain time-critical frames through the network to their destinations. A frame that has preempted an ongoing frame transmission cannot in turn be preempted by yet another frame (i.e., multiple levels of preemption, as opposed to real-time operating system task schedulers, are not possible with the TSN preemption standard).

Frame preemption required changes in the medium access control (MAC) layer. The 802.3 Working Group developed the 802.3br-2016 standard that defines two separate interfaces: an express MAC (eMAC) and a preemptable MAC (pMAC). Frames transmitted through the eMAC cannot be preempted, while frames transmitted through the pMAC can be preempted (by frames from the eMAC). On the ingress side, 802.3br defines techniques to merge multiple frame segments received in the pMAC.

The TSN standard 802.1Qbu-2016 defines how frames are allocated to either eMAC or pMAC. For each port of an end station or bridge, 802.1Q defines up to eight egress queues. Each egress queue is allocated statically to either the eMAC or the pMAC of the egress port. The allocation for a given queue can be different from another bridge. The allocation being done at the queue level means that all frames allocated to that particular queue will go through the same MAC interface. In other words, streams of the same traffic class all go through either the eMAC or the pMAC. The preemption standard provides yet another tool to control the timing of certain frames through the network.

SECURITY

While cybersecurity requirements are out of scope for this article, we would like to highlight the already established and standardized IEEE 802.1 solutions for secure Ethernet. The 802.1 Security Task Group has defined standardized solutions and protocols for port authentication, key agreement, integrity and confidentiality, and device identity.

MACsec, which is standardized in IEEE 802.1AE-2006 and its amendments, provides hop-by-hop Ethernet frame integrity and, optionally, confidentiality. The underlying technology relies on symmetric-key cryptography (more precisely, AES with 128- to 256-bit keys in Galois counter mode of operation). MACsec uses a specific EtherType and adds a message integrity code that is verified by the receiving link partner.

Port authentication and key agreement is standardized in IEEE 802.1X-2010 and its amendments. 802.1X defines protocols to include or exclude devices in a network as they connect to ports. The authentication method itself is not prescribed, leaving flexibility to the system and network designer to determine the appropriate level of security. 802.1X also defines the MACsec Key Agreement protocol that establishes the symmetric keys used in 802.1AE. It also defines solutions to re-key to mitigate and avoid replay attacks.

Finally, 802.1AR defines secure device identity based on device identifiers. A device identifier is programmed as a digital certificate by the manufacturer of the network device, all tracing back through the manufacturer public key infrastructure (PKI). The network designer can program a device identifier specific to the particular system of which the network device is part. These device identities can be used as one authentication method in 802.1X.

AUTOMATED DRIVING

Before going through automated driving use cases and their connection to the various TSN standards, let us briefly review the five driving automation levels defined by SAE J3016 [5]. Each step up the sequence of levels has increasing system automation and decreasing human involvement in the driving task [14]. Driving automation levels categorize and classify automotive features based on their capability to perform the dynamic driving task (DDT), which broadly consists of two subtasks:

1. Operational behaviors comprising longitudinal control, lateral control, as well as object and event detection and response
2. Tactical behaviors such as speed and lane selection, and maneuver planning

Note that strategic behaviors such as route and destination planning are excluded in the SAE J3016 definition. The five driving automation levels are defined as follows.

Level 1, "Driver Assistance": This provides sustained longitudinal or lateral control relative to external objects and events (e.g., adaptive cruise control).

Level 2, "Partial Automation": Longitudinal and lateral control for a given operational design domain (ODD). The driver must supervise the system behavior and perform the remainder of the DDT (e.g., detect objects and respond appropriately by taking control). Cadillac's Super Cruise™ is an example of a hands-free driver assistance system [10].

Level 3, "Conditional Automation": Complete automation of the DDT for a given ODD, providing appropriate response to objects and events. The driver is required to take over control if the system is about to exit its ODD or in case of a system failure.

Level 4, "High Automation": Complete DDT within a given ODD. Automatically bring the vehicle to "minimal risk condition" without reliance on the driver if the system is about to exit its ODD, in case of system failure, or in case of vehicle base failure.

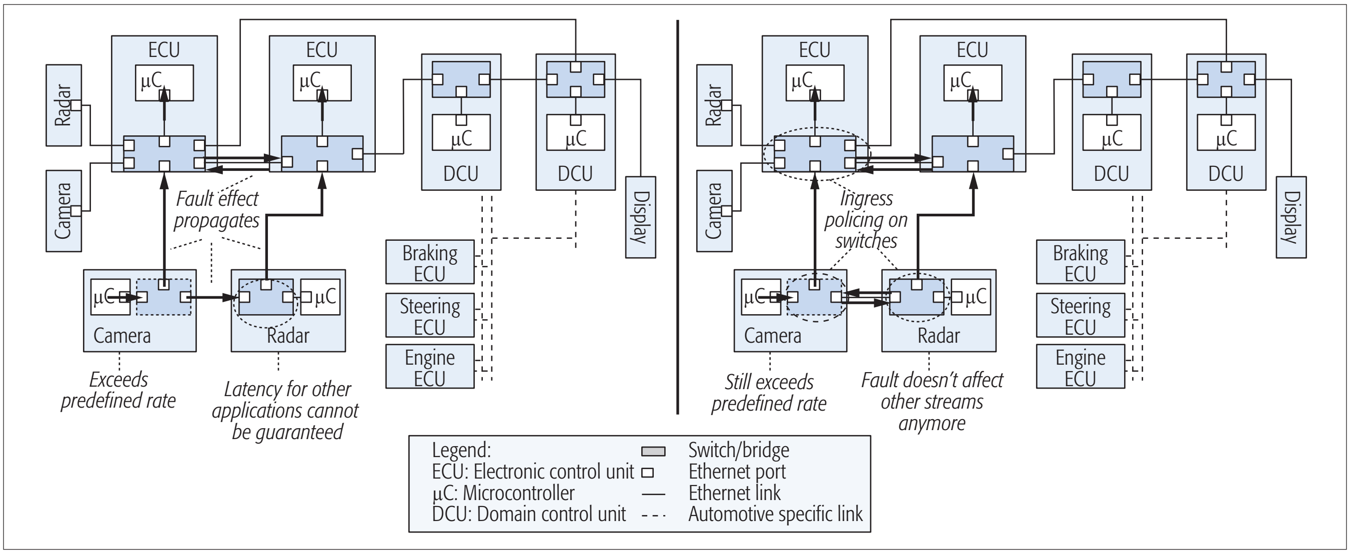

Figure 5. Ingress policing and filtering on the switches will block or limit the traffic, so the faulty sender (camera) is “silenced.”

Level 5, “Full Automation”: This meeans complete DDT under all road conditions in which the operator is legally permitted to operate a vehicle (i.e., no prescribed ODD).

Use Cases and Relations to TSN Standards

Current and future autonomous driving systems consist of many sensors, likely of different modalities. Cameras, radars, lidars [8, 12], and ultrasonic devices are all considered within the industry in the development of autonomous driving. The quantity, modality, and placement of the various sensors depend heavily on the vehicle’s ODD, the desired set of capabilities, and the driving automation level. For example, a system capable of driving on the freeway with automated lane change compared to a system without automated lane change capability will have different requirements on the sensor suite. The former (with lane change capability) may need sensors in the rear of the vehicle to enable automatic determination of safe conditions prior to initiating an automated lane change, whereas the latter (no lane change capability) may or may not strictly require such sensors. Similar arguments can be made for a Level 2 system compared to a Level 4 system, which may require redundancies in the sensor suite to allow the system to automatically reach a minimal risk condition in different failure modes.

Implementation of the variety of sensor suites needed to support the spectrum of autonomous driving levels may use an interesting mix of TSN networking tools. First, a sensor suite for an autonomous driving system, regardless of operational domain and driving automation level, may consist of many sensors of different data rates and traffic patterns. Some traffic is highly regular and periodic, whereas other traffic may include periodic or sporadic bursts. In addition, the temporal requirements of a sensor fusion application are different depending on the sensor and the type of perception function it supports (e.g., objects detected by short-range sensors are closer to the vehicle and therefore may require a quicker response than those detected by long-range sensors). The TSN standards support multiple traffic-shaping classes that can be used to support mixed-criticality real-time communication. For example, 802.1Qbv may be a good choice for some highly regular traffic with consistent frame payloads, whereas AVB credit-based shaping (CBS), or the upcoming asynchronous traffic shaper (P802.1Qcr), may be good choices for shaping the bursty traffic produced by radar and lidar sensors, while still meeting timing requirements [6, 7].

Second, autonomous driving systems may include a sensor fusion component that computes a representation of the observed environment around the host vehicle based on input from a multiplicity of many types of sensors. In addition to the actual sensor data itself (e.g., camera images or radar/lidar scans), a sensor fusion algorithm may need to relate the different sensor data in the temporal domain. It is therefore required that sensors synchronize their time and time stamp their data relative to a common time base. TSN’s 802.1AS gPTP with BMCA provides such a time base with microsecond-level precision. For a Level 4 and 5 system in particular, not only is it important that sensors implement gPTP, but it may also be important that the network implements appropriate redundancies to allow sensors to remain synchronized even if the grandmaster or an Ethernet link fails. 802.1AS-REV will support multiple simultaneous grandmasters to allow for quick recovery of time sync in failure modes (Fig. 1). This supports the Level 4 and 5 definition that the system remains operational and reaches a minimal risk condition even in the presence of failures.

Third, various errors in the network need to be detected and mitigated (whether the errors are due to random hardware faults, systematic software faults, or malicious intrusions or attacks). The TSN standard 802.1Qci supports per-stream ”babbling idiot” error detection and mitigation by blocking a stream or a port to guarantee that errors are contained and not propagated through the Ethernet network. Ingress policing and filtering could also be used toward enforcing the security on a vehicle network. When a non-authorized ECU is detected, or when a significant amount of non-compliant traffic is detected, policing can

While cybersecurity requirements are out of scope for this article, we would like to highlight the already established and standardized IEEE 802.1 solutions for secure Ethernet. The 802.1 Security Task Group has defined standardized solutions and protocols for port authentication, key agreement, integrity and confidentiality, and device identity.

effectively shut down the offender. Figure 5 shows the application scenario to protect the vehicle network against traffic overload. Bandwidth and latency for other streams (e.g., the radar unit in this case) can no longer be guaranteed. Thus, the fault effect propagates through the network, as depicted on the left of the figure.

It may also be important to guarantee that if a link or a switch should stop operating while the autonomous driving system is engaged, at least a subset of sensors and other key communications, such as actuation of steering and braking, remain operational. The TSN standard 802.1CB supports loops/rings in the network and thereby providing redundant communication paths. An important feature in 802.1CB is ″proxy mode″ in which duplication of packets is performed by the first switch in the data path, and subsequent elimination of redundant packets, as well as detection of absence of redundant packets, is typically done by the last switch in the data path. This way, 802.1CB provides full redundancy and error detection at the network layer in a fully transparent manner to the software applications, as shown in Fig. 2.

For illustration purposes, Fig. 3 shows a hypothetical example of a “converged network” where the variety of TSN scheduling techniques are all used together on the same network. The radar and camera sensors use the asynchronous shaping policies (e.g., CBS or ATS), while in the same network, preemption and time-aware shaping are used to precisely control timing and achieve the lowest possible latencies to the steering and braking actuator ECUs. This shows that TSN enables common network resources to host multiple communication patterns of different regularities and mixed criticalities. We can thus observe that the TSN standards provide us with multiple scheduling algorithms to be chosen based on the temporal requirements of the application. These all work together with the TSN standards for synchronization and fault tolerance.

SUMMARY

This article presents an overview of the relevant TSN standards for automotive use cases, with particular focus on the growing needs of autonomous driving systems. It described the various driving automation levels defined by SAE to map and illustrate the importance of each specific TSN standard in realizing autonomous driving systems for different automation levels and different operational domains. TSN enhances switched Ethernet technology with multiple key properties: error detection and mitigation, redundant communication with fast fail-over, precise time synchronization with fail-over and redundancy, as well as a comprehensive suite of real-time scheduling policies including synchronized, time-triggered communication, asynchronous traffic shaping approaches based on credits and token buckets, and selective frame preemption. Ethernet technology and TSN standards thus provide a network foundation with bandwidth, real-time, and dependability properties that support future automotive systems.

REFERENCES

[1] IEEE 802.3 Multi-Gig Automotive Ethernet PHY Task Force; http://www.ieee802.org/3/ch/public/index.html.

[2] IEEE 802.1 Time-Sensitive Networking Task Group; http://www.ieee802.org/1/pages/tsn.html.

[3] J. Farkas, “Introduction to IEEE 802.1 — Focus on the Time-Sensitive Networking Task Group,” Vancouver, BC, Canada, Mar. 2017; http://www.ieee802.org/1/files/public/docs2017/tsn-farkas-intro-0317-v01.pdf.

[4] S. Samii and H. Zinner, “Ethernet TSN as Enabling Technology for ADAS and Automated Driving Systems,” *IEEE-SA 2015 Ethernet & IP @ Automotive Technology Day*, Yokohama, Japan, Oct. 2015.

[5] SAE Int’l., J3016_201609, “Taxonomy and Definitions for Terms Related to On-Road Motor Vehicle Automated Driving Systems”; http://standards.sae.org/j3016_201401.

[6] J. Specht and S. Samii, “Urgency-Based Scheduler for Time-Sensitive Switched Ethernet Networks,” *Euromicro Conf. Real-Time Systems*, Toulouse, France, July 2016.

[7] J. Specht and S. Samii, “Synthesis of Queue and Priority Assignment for Asynchronous Traffic Shaping in Switched Ethernet,” *IEEE Real-Time Systems Symp.*, Paris, France, Dec. 2017.

[8] A. Terzis, *Handbook of Camera Monitor Systems: Video Interface Technology*, Springer, 2016, pp. 253–75. ISBN 978-3-319-29611-1.

[9] CAN in Automation (CiA), “CAN Knowledge,” Nov. 2017; https://www.can-cia.org/can-knowledge.

[10] Cadillac, Super Cruise™; http://www.cadillac.com/world-of-cadillac/innovation/super-cruise.html.

[11] M. D. Johas Teener *et al.*, “Heterogeneous Networks for Audio and Video: Using IEEE 802.1 Audio Video Bridging,” *Proc. IEEE*, 2013.

[12] Voyage, “An Introduction to LIDAR: The Key Self-Driving Car Sensor,” May 2017; https://news.voyage.auto/an-introduction-to-lidar-the-key-self-driving-car-sensor-a7e405590cff.

[13] H. Zinner and S. Samii, “TSN Traffic Shaping for Deterministic Automotive Networks,” Automotive Ethernet Congress, Munich, Germany, 2018.

[14] General Motors, “2018 Self-Driving Safety Report”; https://www.gm.com/content/dam/gm/en_us/english/selfdriving/gmsafetyreport.pdf.

[15] AVnu Alliance; http://avnu.org.

BIOGRAPHIES

SOHEIL SAMII (soheil.samii@gm.com) is with General Motors R&D and works on next-generation automotive systems and software architectures. He earned his Ph.D. degree in 2011 from Linköping University, Sweden, where he currently holds an adjunct senior lecturer position and teaches a graduate course on Ethernet and TSN. He is active in the 802.1 TSN Task Group and has served on the Program Committees of multiple conferences such as DAC, RTSS, and RTAS.

HELGE ZINNER received his Ph.D. in electrical engineering from the Institute for Communication Networks, University of Ilmenau, Germany. He has been with Continental as a development engineer for infotainment systems since 2005. He currently works as a senior system architect for the cross divisional usage of automotive high-speed network technologies like Ethernet and SerDes. He has worked actively in multiple standardization organizations (OPEN Alliance SIG and IEEE) and has served on various automotive committees and working groups. He is one of the first engineers to push for the adoption of Ethernet technology in the automotive industry.