Report from Academic Salon, 13th–14th October, 2021

Time-Sensitive Networking and Deterministic Applications

Editors

Prof. Dr.-Ing. Georg Carle[1], Prof. Dr. sc.techn. Andreas Herkersdorf[2], Prof. Dr.-Ing. Jörg Ott[3], and Prof. Dr. Sebastian Steinhorst[4]

1 **Technical University of Munich, Department of Informatics, Chair of Network Architectures and Services**
2 **Technical University of Munich, Department of Electrical and Computer Engineering, Chair of Integrated Systems**
3 **Technical University of Munich, Department of Electrical and Computer Engineering, Chair of Connected Mobility**
4 **Technical University of Munich, Department of Electrical and Computer Engineering, Assistant Professorship of Embedded Systems and Internet of Things**

Abstract

This report documents the program and the outcomes of the "Academic Salon on Time-Sensitive Networking and Deterministic Applications", a virtual workshop held on October the 13th and 14th 2021. Participants of this first Academic Salon discussed on topics in the domains of Time-Sensitive Networking, Network Calculus, Deterministic Networking, and others. The two day virtual workshop was partitioned into four sessions. In this report, summaries of impulse talks given during the individual sessions are presented. Remainder of this report closely follows the structure of the Academic Salon itself.

Seminar 13th–14th October, 2021,
URL: https://www.net.in.tum.de/talks/workshops/academic_salon_21.html

Edited in cooperation with Marcin Bosk, Rubi Debnath, Max Helm, Kilian Holzinger, Filip Rezabek, Christoph Schwarzenberg, and Henning Stubbe

Academic Salon
Chair of Network Architectures and Services, Technical University of Munich, Germany

1 Introduction

Georg Carle (Technical University of Munich– Munich, Germany)

This scientific event aimed to bring together scientists from academia and industry with key expertise in Time-Sensitive Networking technologies and applications with requirements met by these network technologies. Of particular interest were network-specific innovations, in particular concerning IEEE Time-Sensitive Networking (TSN) and IETF Deterministic Networking (DetNet), methods for performance assessment, in particular Network Calculus based methods, contributions on implementation, in particular architecture innovations and acceleration techniques including the use of programmable hardware, and insights into industrial and automotive applications.

Time-Sensitive and Deterministic Networking has been receiving significant interest from academia and industry. It addresses various application domains, including industrial networks, in-vehicular networks, mobile and wireless networks, and service provider networks. Different applications require different types of performance guarantees, including guaranteed throughput, maximum latency, ultra-low latency, security and reliability. The goal of this workshop series is to address some of the key challenges in the areas of time-sensitive networking and deterministic applications in relation to:

- Methods for performance analysis, including Network Calculus
- Methods for network planning, capacity dimensioning and resilience
- Protocol advances, such as routing for time-sensitive traffic
- Implementation approaches for low-latency packet processing, including programmable hardware and high-performance software processing
- System architectures for deterministic applications, including in-network computing elements
- Analysis of use cases with deterministic applications
- Machine Learning in support of time-sensitive networking and deterministic applications

2 Contents

3 Time-Sensitive Networked Systems and Network Calculus

The first session of this academic salon consisted of three talks addressing different challenges in context of Time-Sensitive Networking (TSN) and Network Calculus (NC). Following a brief but warm welcome and introduction by the host and session chair, Georg Carle, the invited speakers and their impulse talks sparked discussions on their topics. The session was concluded by an open discussion.

3.1 Regulators, Dampers and Clock Non Idealities in Time Sensitive Networks

Jean-Yves Le Boudec (École Polytechnique Fédérale de Lausanne – Lausanne, Switzerland)

In this talk, the impact of clock non-idealities and sychronization on the computation of NC-based worst-case delays, were addressed. For this, the speaker focused on regulators, also referred to as shapers, and dampers. The talk was based on joint work with Thomas Ludovic and Ehsan Mohammadpour, respectively.

Recalling that in per-class TSN burstiness does cascade, the speaker motivated the need for regulators as mean to counteract burstiness. Moreover, previous work shows that regulators have no negative impact on the worst-case delay. However, the speaker highlighted, these work assumes ideal and synchronized clocks; an assumption that per se does not hold for real systems. With this motivation, the speaker moved to discuss the impact of clocks on regulators. Their findings there are threefold: in networks with regulators and non-synchronized clocks, worst-case delays are potentially unbounded; for synchronized networks, only interleaved regulators may lead to unbounded worst-case delays. In contrast to that, the impact on worst-case delay of per-flow regulators in synchronized networks has an upper bound.

In the talk's final part, the impact of clocks on dampers, as stateless mean to delay packets as needed, was touched on. Here, Le Boudec indicates that dampers lead to reduced jitter. They are less affected by clock non-idealities, since damper do not depend on clock synchronization.

Excerpts from Q&A

1. *Can TSN profit from a clock distribution system such as in Synchronous Optical Networking (SONET)/Synchronous Digital Hierarchy (SDH)?*
   Yes, White Rabbit, which complements Precision Time Protocol (PTP), is doing this. Use of White Rabbit achieves clock synchronization in the range of tens of ns.
2. *Is the discussed damper concept similar to Jitter-EDD (Verma, Zhang and Ferrari, 1991)?*
   Yes, that work presents the first damper!

3.2 TSN Support for Quality of Service in Space

Marc Boyer (Office National d'Etudes et de Recherches Aérospatiales – Paris, France)

Main reference Chaine et al., TSN Support for Quality of Service in Space, 10th European Congress on Embedded Real Time Software and Systems

In the second impulse talk, Marc Boyer, presents information gained while investigating the question how Ethernet-based communication can be used in satellite On-Board Computers (OBCs). Such OBCs typically host applications with, on average, high throughput that require low latency and low jitter, while depending on the guaranteed arrival. Traditionally, satellite OBCs rely, for communication, on MIL-STD-1533 and CAN bus technology. Being in a niche market, products struggle to keep up with the applications growing demands. According to the speaker, this situation motivated the question whether an Ethernet-based technology can be used as replacement. Initially, four Ethernet variants were considered: Ethernet (specifically, IEEE 802.3 and IEEE 802.1Q-2008); ARINC 644, also known as Avionics Full-Duplex Switched Ethernet (AFDX); Time-Triggered Ethernet (TTEthernet), an AFDX extension; and TSN. However, jitter requirements of less than 1 µs disqualified anything but TSN. Thus, next the speaker focused on determining if TSN with a Time-Aware Shaper (TAS) is a suitable replacement. Comparing the architecture of previous used MIL-STD-1553 and TSN, they identify two problems. First, TSN TAS is based on queues, and, thus, in opposition to MIL-STD-1553, assumes an order on messages. Second, in TSN TAS frame loss may occur. Additionally, in TSN TAS, the loss of one frame may impact the behavior of other frames. These two problems make TSN TAS unsuited as replacement for MIL-STD-1553.

The talk was concluded with the remark that, solutions to the raised problems offered by TSN introduce additional unwanted complexity.

Excerpts from Q&A

1. *Why is low jitter such an important requirement?*
   Control traffic and simple devices require cyclic commands for satellite stabilization.
2. *How does TSN compare to TTEthernet for this use case?*
   TTEthernet has a white board semantic, while TSN has a queue semantic. Also, technical and non-technical problems and the size of the community and market are relevant. A smaller community and fewer people lead to higher cost, which is a noticeable issue for TTEthernet.

3.3 Credit-Based Shaper Configuration for Delay-Constrained Flow Allocation in TSN

Lisa Maile (Friedrich-Alexander University Erlangen-Nürnberg – Erlangen, Germany)

The third impulse talk on the topic of NC focused on configuration of Credit-Based Shapers (CBSs) in TSN to guarantee delays. According to the presenter, current TSN standards do not offer attractive means to implement requirements for applications with firm delay demands outside the realm of time-triggered gates. Lisa Maile argues for CBS to address the needs of applications residing between ultra-low-latency and best-effort. Within their research, they developed an online approach for configuration and flow allocation. The steps of this approach can be summarized as follows: finding routes for new flows;

dynamically assigning priorities to each flow; and, finally, setting idle slopes per queue. In contrast to related work, their approach does not rely on well-defined routes, priorities, or sending intervals. Moreover, their approach differs in the ability to manually configure idle slopes, instead of being restricted to leftover service curves. Last, their approach is superior in that it offers CBS shaping on top of the link rate shaping offered by related work. The speaker concluded with an evaluation of their approach in an industrial scenario.

Excerpts from Q&A
In the open discussion following this talk, the speaker clarified the optimization problem tackled in their work and their use of OMNet++ as evaluation environment. Other participants emphasized the relevance of the presented work and were keen on results including concrete results. Lisa Maile indicated that a future publication will likely elaborate on the topic including discussion of the requested results. Other questions touched on the topic of pessimism and overhead introduced by CBS. Here, the speaker referred to ongoing work, e.g. by Le Boudec, that may address this issue. On the topic of reconfigurations, Marc Boyer suggests to try to minimize needed reconfigurations by initially choosing a configuration that tolerates introducing additional flows. Agreeing on the importance of this question, Lisa Maile points out that one of their approaches avoids reconfigurations by trading these for higher delays.

3.4 Modelling Programmable Device Behavior

Max Helm (Technical University of Munich – Munich, Germany)

Main reference Helm et al., Application of Network Calculus Models on Programmable Device Behavior, 33rd International Teletraffic Congress

The sequence of NC impulse talks was completed by a talk on modelling the behavior of programmable devices. Starting with an introduction to the workflow akin to programmable devices, the speaker explained a workflow for modelling. For modelling worst-case latencies, they used NC. Their approach was applied to programmable devices, using P4 and OpenFlow as Software-Defined Networking (SDN) implementations. The speaker highlighted the differences between the latter two implementations and argued that differences counteract a direct comparison. To work around this issue and enable a comparison, they proposed a three step approach. First, they introduce a networking device model whose behavior is understood as combination of the device's basic function plus functionality under investigation. Second, for each modelled function, they evaluate performance of a networking device executing the respective function and, thus, measure each functionality in isolation. Finally, they approximate NC service curves, based on the obtained measurement results.

Subsequently, the speaker presented an evaluation of the proposed method when applied to two different devices: Zodiac FX and NetFPGA SUME. Max Helm noted that the choice of these devices, differing among others in SDN framework and link rate, can be understood as indicator for flexibility of the proposed framework. Their results suggest, according to the speaker, a reasonable predictive power with no high correlation between error and measured functionality. In their conclusion, the speaker hinted that future work could consider more complex service curves or aim for an exact service curve instead of the current approximation.

Excerpts from Q&A

In the discussion after the talk, it was clarified that, during the evaluation measurement, results are expected to remain below the estimated maximum service curve. Further exchanges related to the exact definition of the displayed relative error as well as the importance of other service curve types.

3.5 Panel Discussion on the "Current State and To-Be-Expected Progress of Network Calculus Methods"

The opening question was:

> So many alternatives to achieve real-time requirements — which solutions are likely to be deployed, and for which reasons?

Discussion contributions included:

1. What are break-even points when comparing schedulers, i.e., when considering delay, when does a change of scheduler result in a delay decrease.
2. Participants observed: In academia, the focus is on one shaper and on one mechanism, but real applications have different traffic types and are heterogeneous. How do the different scheduling mechanisms work in such a heterogeneous scenario?
3. The issue of determining which applications have which delay and scheduler requirements was raised as important question.

4 Time-Sensitive Network Technologies

The session on Time-Sensitive Network Technologies started the second day of the academic salon. Chaired by Sebastian Steinhorst, this session touched on topics closely related to TSN. Talks of the four speakers and the discussions following are summarized in the remainder of this section.

4.1 Self-Configuring and Self-Healing Time Sensitive Networking

Paul Pop (Technical University of Denmark – Kongens Lyngby, Denmark)

Main reference Paul Pop, Self-Configuring and Self-Healing Time-Sensitive Networking, 2nd International Workshop on Time-Sensitive and Deterministic Networking

The work presented in this talk is concerned with optimization of TSN sub-standard parameters, which is hard for small problem spaces and becomes intractable for larger problem spaces. An example for this is the synthesis of IEEE 802.1Qbv Gate Control Lists which cannot be solved by a polynomial-time algorithm. Multiple approaches for this synthesis problem have been presented, all of which assume flow isolation and scheduled end systems. Under this assumption, the synthesis problem turns into a combinatorial optimization problem. Adding unscheduled end systems transforms the problem into a schedulability analysis problem. In the talk, NC was briefly mentioned as one instance of this schedulability analysis problem. However, it was also noted that NC is unsuited as cost function in optimisation problems, since involved computations require too much time. Additionally, re-routing of streams is introduced as an option to add unschedulable streams. The challenge of selecting optimal parameters within TSN is further complicated by competing objectives and constraints between substandards, leading to the conclusion that the configuration problems are intractable and interconnected. A vision of a self-configuring TSN ecosystem based on monitoring and machine learning is presented. The goal is to generate good quality solutions in a short time. A possible solution would involve using AI-based optimization with (meta-) heuristics in combination with incremental evaluation and impact analysis.

Excerpts from Q&A

1. *The problem space is too large and execution times for network calculus are too large to be dynamic. With a headroom of capacity, how quick and precise should the assessment be?* Depends on use case and their criticality, e.g. TAS for high priority must offer guarantees at runtime, introduce slack at configuration time.
2. *Comment: From an industry perspective, a promise associated with TSN is self-configuration of the network. A pre-condition for this is absence of any binding between application and network.*

4.2 Performance Comparison of Traffic Shapers in TSN Networks

Luxi Zhao (Beihang University – Beijing, China)

Main reference Luxi Zhao et al., Quantitative Performance Comparison of Various Traffic Shapers in Time-Sensitive Networking

This work utilizes NC to compare the performance of different traffic shapers and their combinations. Additionally, two new combinations (TAS+ATS+CBS and TAS+ATS+SP) are suggested and NC support is provided. The goal of this comparison is to provide TSN users with a basis to decide which combinations to use. Analysis of synthetic test cases with low traffic loads showed that, for traffic shapers in isolation, TAS performed significantly better than all other shapers with respect to latency, backlog, and jitter. Furthermore, Strict Priority (SP) performs better than Asynchronous Traffic Shaper (ATS) when combined with SP for low traffic loads, whereas the reverses can be observed for higher traffic loads. A similar effect can be observed in a comparison between TAS+SP and TAS+ATS+SP. Finally, a real world test case, based on NASA's Orion Crew Exploration Vehicle (CEV), is analyzed. Results of this analysis showed that TAS+ATS+SP gives better bounds on backlog but not end-to-end latency as compared to TAS+SP. Increasing the traffic load showed improvements in both backlog and end-to-end latency bounds.

4.3 Challenges and Solution Directions or Deterministic QoS in Service Provider Networks, Large Scale and Wide Area Networks

Toerless Eckert (Futurewei Technologies – Mountain View, CA, USA)

This talk focused on Deterministic Networking (DetNet) IETF working group progress, namely the architecture, use cases, and IP/MPLS encapsulation. The speaker looked into challenges and solutions for QoS, focusing on bounded latency for DetNet and elaborating on their perspective. In DetNet, achieving bounded latency is hard to achieve due to: per-hop, per-flow state; clock synchronization; and jitter. Also, configuration of TSN shapers on a per-hop, per-state basis is challenging. Therefore, according to the speaker, it would be better to offload the concern form the forwarding plane to the ingress router, which applies a given policy according to per-flow state. With increased network throughput to 400 Gbps in the WAN scale, clock synchronization is challenging. Low jitter might be helpful only for very few use cases. It was noted that, PTP is not commonly deployed unless required, due to involved operational costs. Since standardization is a current area of research, there are no clear solutions. Still, the speaker argues that one should avoid per-hop, per-hop state information with no or "relaxed" clock synchronization with jitter in mind for high-speed, low-cost, large-scale networks with bounded latency.

Excerpts from Q&A

1. *Adoption strategy by service providers to offer critical traffic across WANs?*
   Most likely not, unless TSN is easier to configure, SDN could help with the configuration. Try to bring closer IT/OT.

5 Advances from Ongoing Projects

The third session was opened and chaired by Michael Menth from Eberhard Karl University of Tübingen. In this session the focus was laid on performance evaluation, SDN, resilience management, IP, BIER, service function chaining, WiFi6/7, and how those technologies could cooperate with TSN and IVN networks. As other sessions, discussions are held right after each of the three talks together with Q&A sessions. The session is closed by a brief panel discussion where the importance of SDN in automotive and TSN domains is stressed.

5.1 Early Results from the KITOS Project

Rene Guillaume (Robert Bosch GmbH – Gerlingen, Germany)

This talk introduced the KITOS ("Künstliche Intelligenz für TSN zur Optimierung und Störungserkennung") project as well as showcases the dynamic TSN configuration possibilities and its use cases within TSN. The general motivation, can be summarized as follows: Future industrial networks need to fulfil numerous requirements, examples of which include migration of automation functions towards the edge, production structure adaptation to new demands, and no vendor lock-in. To cope with those requirements, goal of the KITOS project is to enhance the TSN configuration to make it better adaptable to future challenges, as well as to develop a platform that assists with this configuration. Furthermore, the project focuses on newest technologies and generations of Programmable Logic Controllers (PLCs) and investigates the currently well-defined standards. Other aspects are initial network configuration, runtime reconfiguration, monitoring, and, heterogeneous topologies. The use cases consider a control-to-control application with various topologies and migration of the PLC to edge and cloud. Finally, the work identifies that there is a need for new standardization work, extending the existing OPC UA FX and IEC/IEEE 60802 standards, as well as support for multi-vendor networks.

Excerpts from Q&A

1. *From factory operator's perspective what is the benefit of exchanging Profinet to TSN?*
   The migration will be a long-term evolutionary process. It will provide a closed-box solution, but still there are multiple challenges that are being researched right now.
2. *Where do you get your use cases from?*
   No go-to source for that. Requires own research and industry insights help.

5.2 RAP Extensions for the Hybrid Configuration Model

Lukas Osswald (Eberhard Karl University of Tübingen – Tübingen, Germany)

Main reference Osswald et al., RAP Extensions for the Hybrid Configuration Model, 26th IEEE International Conference on Emerging Technologies and Factory Automation

This talk showcased proposed extensions to the Resource Allocation Protocol (RAP) and its applications in the hybrid configuration model. RAP is in early stages of development and aims to overcome the limitations of Stream Reservation Protocol (SRP) and Audio Video Bridging (AVB) protocols. The aim is to include RAP together with the Link-local Registration Protocol (LRP) in the hybrid model. The LRP is used for persistent data distribution throughout the network and allows for stream reservation. The work proposes RAP-Central User Configuration (CUC) using LRP together with RAP signalling to allow the solution to extract information with relevance for admission control from the network. The model proposed in the work includes a protocol connector for user specific signalling, a stream management component for stream life-cycle management, a Central Network Configuration (CNC) connector to convey the stream requirements to the CNC, and a webhook handler for end-stations to receive the synthesized schedule. Partial implementation of the model is open-source and available at https://github.com/uni-tue-kn/rap-cuc.

Excerpts from Q&A

1. *What are the benefits of a hybrid model? Does it simplify or makes it things more complex?*
   The hybrid model does not directly simplify the signalling, but makes it distributed in other scenarios. Also, the hybrid model improves the network without changing the end stations. No needed of OPC UA in such a scenario for stream reservation.
2. *Have you got any information how the maximal latency field is filled by the bridges?*
   Normal use case included it, RAP does not include it. RAP includes the accumulated maximum latency used by the bridges to dynamically compute the max latency. Do not know exactly how it is filled, but something may be in the standard related to it.

5.3 Softwarization of Automotive E/E Architectures: A Software-Defined Networking Approach

Marco Häberle (Eberhard Karl University of Tübingen – Tübingen, Germany)

Main reference Häberle et al., Softwarization of Automotive E/E Architectures: A Software-Defined Networking Approach, 2020 IEEE Vehicular Networking Conference

This talk evaluated how SDN could be applied to automotive networks together with Ethernet and TSN. According to the speaker, current In-Vehicle Networks (IVNs) are based on low bandwidth technologies and usually come with a static configuration. They argue that these properties might result in a future limitation, as IVNs are currently evolving towards zonal architecture deployments. These deployments would then include numerous sensors and run Electronic Control Unit (ECU) functionality as applications on vehicle computers, connected via automotive Ethernet. Automotive SDN promises to solve the challenges of current IVNs by employing the classic SDN control plane for inventory management and calculation of TSN configurations. In the proposed network, the data plane would consist of all switches within a car, interconnecting everything and cooperating with the management system. Such approach allows for employment of redundant backbones within a car and introduction of a discovery mechanism for new devices in the car which can trigger re-configuration of the IVN. Moreover, this approach would also enable use of a hybrid scheduling approach that provides guarantees for critical systems and dynamically adapts to various network loads.

Excerpts from Q&A

1. *We hear that the autonomous cars collect petabytes of data. How does this approach fit with the perspective of so much complexity and data?*
   An approach is to use many different bus networks in the car with very high complexity. Number of ECU in a modern car is in the order of 100. But, if different technologies are

used, such complexity may not be feasible or sustainable in the future. From this stems the need of more softwarization regarding communication.

6 Concepts, Requirements and Applications for Deterministic Services

Jörg Ott introduced the audience to this second to last session on Concepts, Requirements and Application for Deterministic Services. The focus of this session is on various underlying technologies.

6.1 Flexible Research Infrastructure for Reliable and Scalable Intra-Vehicular TSN Networks

Filip Rezabek (Technical University of Munich – Munich, Germany)

Main reference Rezabek et al., EnGINE: Developing a Flexible Research Infrastructure for Reliable and Scalable Intra-Vehicular TSN Networks, 3rd International Workshop on High-Precision, Predictable, and Low-Latency Networking

This talk introduced EnGINE (Environment for Generic In-vehicular Network Experiments), a research infrastructure for in(tra)-vehicular networks. The introduced infrastructure, a testbed used for analysing scenarios, currently consists of 15 nodes built with commercial off-the-shelf hardware. Each node is equipped with Intel NICs with support for different TSN standards and currently one LIDAR sensor. The TSN standards are supported both in hardware and software, in the latter case by utilizing Linux queuing disciplines (qdiscs).

Experiments conducted in this testbed belong to a campaign which is defined by an input dataset. Each experiment is split into four phases:

1. Install – Install a suitable operating system on the nodes,
2. Setup – Install additional packages and copy required script files,
3. Scenario – Setup the test environment and execute the experiment, and
4. Process – Post-processing of scenario artefacts.

A sample scenario is presented which evaluates different traffic shapers along six hops, comparing their latency and jitter. In general, low jitter with few outliers larger than 100 µs can be observed. The end-to-end delay for TAS flows stays mostly within the 2 ms target.

6.2 TSN Plug & Produce, 1000 IA Stations

Günter Steindl (Siemens AG – Nuremberg, Germany)

This talk presented requirements and challenges to be addressed in industry 4.0 environments. From a network perspective, large topologies are frequently encountered with up to multiple thousands nodes and streams. Dynamic topology changes must be considered, as well as large bridges and widely varying link speeds across different media types. For realizing converged networks, applications and middleware need to be decoupled from the network to allow for flexible network utilization. From the perspective of TSN, these requirements limit the application of existing approaches and solutions, as such networks are optimized for flexibility rather than performance.

Excerpts from Q&A

1. *Will there be a fixed data plane model vs. everything being negotiated per deployment?*
   There exist no black and white solutions. When a new packet is scheduled, the network must be able to accept and handle it. As bridges are no longer able to realign traffic from end systems, since this becomes too complex, a portion of the responsibility is moved to the end systems that are required to send packets or packet bursts at the right point in time so the network and CNC can handle the traffic.
2. *This approach gives the network more flexibility, but does it impose too many requirements on the end systems?*
   Current results look promising. The requirements need to be balanced between bridges and end systems. End systems can easily overload bridges in case of faults. Whether this approach is flexible enough is subject of a large discussion.

7 Wrap-Up

The academic salon was ended by a guided discussion and final remarks by Georg Carle. Questions came from general impressions during the workshop and remaining open questions from the various talks.

7.1 Discussion Contributions

1. *TSN is complicated: Are simpler solutions available?*
   - TSN might be complicated, but waiting for other solutions is futile since they might arrive too late.
   - Participants engaged in a discussions about new planning, and new shapers.
2. *Is there guidance to different subgroups picking different solutions?*
   - Some argued that TSN is a perfect topic for research, since it is very complicated and the idea to make it fit for everything, is likely to lead to over-engineering.
   - Others disagreed stating that TSN as topic is not complicated; however, scheduled gates currently include complexity. Hence, scheduled gates are a last resort. It is good that one can make complicated devices, but one should not overdo it.
   - Using aerospace as an example, AFDX or even SpaceWire appear less complicated (and perform good in their respective domains), yet an aerospace profile for TSN has been created. Perhaps industry is moving away from easy but specific technologies to join a more complicated but wider community, with more possible suppliers.
3. *Is there a concept of dampers in TSN yet?*
   The IEEE working group has not decided on it yet, but there is ongoing work.

7.2 Final Remarks

Feedback for the event and its organization was very positive. Several attendees underlined the benefit of the clear focus of the event program. A follow-up event was suggested by several participants.