Report from Academic Salon, 29th–30th September, 2022

Low-Latency Communication, Programmable Network Components, and In-Network Computation

Editors

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

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 Informatics, Chair of Connected Mobility**

Abstract

This report documents the program and the outcomes of the "Academic Salon on Low-Latency Communication, Programmable Network Components and In-Network Computation", a virtual workshop held on September the 29th and 30h 2022. Participants of this academic salon discussed on topics in the domains of Low-Latency Communication, Programmable Network Components, In-Network Computation, and others. The two day virtual workshop was partitioned into four sessions. In this report, summaries of discussions following the various impulse talks given during the individual sessions are presented. The report closely follows the structure of the academic salon itself.

Seminar 29th–30th September, 2022,
URL: https://www.net.in.tum.de/talks/workshops/

Edited in cooperation with Marcin Bosk, Max Helm, Benedikt Jaeger, Filip Rezabek, Henning Stubbe, and Florian Wiedner

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

1 Introduction

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

The academic salon is a scientific event that aims to bring together scientists from academia and industry with key expertise in core technologies on Low-Latency Communication, Programmable Network Components, and use-cases for In-Network Computation applications. In particular, use-cases from the industrial domain, with requirements met by these core technologies, and use-cases that benefit from the capability of In-Network Computation are of interest. Simultaneously, network-related innovations, contributions on implementation, innovative memory and processing technologies, methods and tools for performance assessment, and insights into industrial and automotive applications are in center of this events discussion. That is, the topical scope can be described as follows:

- Network-related innovations, in particular architecture innovations such as acceleration techniques for programmable Network Interface Cards and Switches
- Innovative processing and memory technologies, such as P4-based protocol processing
- Design, specification, verification, implementation, measurement, testing, and analysis of programmable network components
- Methods and tools for performance assessment
- Applications and use-cases, in particular of industrial and the automotive domain.

This academic salon is a sequel to the first academic salon [Car+21] held in the previous year. As before, a major goal of the academic salon is not only the exchange of raw content but also to bring the community closer together. Recordings of this year's academic salon will be made available at: https://www.net.in.tum.de/talks/workshops/, accompanied by the presenters' slides available online at: https://oc.net.in.tum.de/index.php/s/LA4zmzTJAmeiCmt.

2 Contents

Introduction
Georg Carle . . . 2

Network Paradigms that benefit from Programmable Network Components

Scaling Deep Learning and Datacenter Applications with Programmable Networks
Marco Canini . . . 4

P4-Based Implementation of BIER and BIER-FRR for Efficient Multicast
Steffen Lindner . . . 5

Topological Addressing
Luigi Iannone . . . 6

Real-Time Networking approaches and tools for Industrial Application Scenarios

High-capacity and Resilient Large-scale Deterministic IP Networks
Paolo Medagliani . . . 7

Reproducible Layer 3-enabled TSN Experiments
Filip Rezabek, Marcin Bosk . . . 7

Programmable Network Interface Cards and In-Network Computation – In-network computation in industrial scenarios, Processor-based SmartNICs

Packet Ordering for Improving Application Performance using SmartNICs
Dejan Kostic . . . 10

The paradigm shift in automotive zonal gatewaying
Francesc Fons, Angela Gonzalez Mariño, Abdoul Aziz Kane . . . 11

In-network computation and Processor-based SmartNICs
René Glebke, Klaus Wehrle . . . 11

Mechanisms and Methods for Real-Time Networking approaches suitable for Industrial Application Scenarios

Application of Network Calculus for Reliable and Predictable Behavior of IEEE 802.1CB Frame Replication and Elimination in Time-Sensitive Networks
Lisa Maile . . . 12

Tail Latency Estimation and Verification
Max Helm, Florian Wiedner . . . 13

Multi-Level Preemption in TSN
Mubarak Adetunji Ojewale . . . 13

Wrap-Up

3 Network Paradigms that benefit from Programmable Network Components

The first session of this academic salon consisted of three talks addressing different challenges in context of Programmable Network Components. Following a brief 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 panel discussion.

3.1 Scaling Deep Learning and Datacenter Applications with Programmable Networks

Marco Canini (King Abdullah University of Science and Technology – Thuwal, Saudi Arabia)

The talk of Marco Canini entitled "Scaling Deep Learning and Datacenter Applications with Programmable Networks" addresses key questions of the academic salon, including what should be computed in a network. Marco Canini presented projects tackling challenges in this domain, among them SwitchML [Sap+21] and OmniReduce [Fei+21].

The discussion after the talk included the following questions and answers:

QUESTION: What type of failures do are considered for failure recovery? Is there a focus on network malfunction or also malicious actors in the system?

ANSWER: The topic of network malfunction alone is a challenge: TCP is rather complex, if not too complex for some scenarios. Thus, RDMA, in this scenario, cannot rely on its failure recovery mechanism to handle, e.g., network malfunctions. To still provide resistance against such failures, one approach is that the RMDA-capable switch keeps a form of “state” from previous gradients and checks for the updated version of the new gradients. If new entries are missing it will ask corresponding workers for a retransmission.

QUESTION: Are there evaluations how much memory is needed in case of packet reordering or loss? What are the costs of failures in terms of memory? How is a worker failure handled?

ANSWER: Evaluation of worst-case properties was not main concern, so far. Research question and resulting main observation was the task of decoupling storage at the switch and the size of data. That is, the match size of the model. Memory-wise, the switch’s capacity is in the range of 100 kB. To increase the available storage capacity, the bandwidth-delay-product and its size is exploited. Another exploited factor is that the order of function results from workers does not matter. Finally, handling worker failures is considered the application’s duty.

QUESTION: Would an ASIC with support for floating point arithmetic provide a significant advantage?

ANSWER: Yuan et al. [Yua+21] considers this question looking at a number of algorithms. They find a couple of things missing for floating point support in the Intel Tofino, e.g., in terms of bit shifting or byte-ordering. But, with few modifications, support for floating point arithmetic is within reach. In this context, the Mellanox SHARP [Mel19] seems relevant. A question that remains is: would it be worthwhile to add this support. Positive impact of such support would be the possibility to trade-off CPU cycles that could do something else instead of handling floating point operation.

3.2 P4-Based Implementation of BIER and BIER-FRR for Efficient Multicast

Steffen Lindner, Daniel Merling, Michael Menth (Universität Tübingen)

The session's second talk by Steffen Lindner, entitled "P4-Based Implementation of BIER and BIER-FRR for Efficient Multicast", focused on efficient implementation of the Bit Index Explicit Replication (BIER) protocol [MLM20; MLM21].

QUESTION: Is this understanding correct: The current BIER implementation requires `ports` × `external_bandwidth` of internal bandwidth to support BIER multicast?

ANSWER: Indeed, this is currently a requirement. However, the presented implementation is improved on in later work. For these improvements traffic exchanged, as well as its content, is analyzed. This work, which relies to some extent on machine learning, is currently under review.

QUESTION: When dealing with BIER, are there missing function a future switch should provide? Accelerated instructions, such as currently present in CPUs, are expected to come up also in, e.g., programmable switches.

ANSWER: Yes, namely the ability to have dynamic packet copies without the need for multicast groups. At the moment there is no way to have a port bit mask or the like. In other words, one either needs a control plane – to update multicast groups – or some other mean to learn these groups.

QUESTION: The recirculation strategies in current devices, e.g., the Intel Tofino, appear lacking in capacity. Would you agree that future devices should improve on the current state in this regard?

ANSWER: Implementing recirculation on a programmable device without an impact on bandwidth is not feasible. Improvements on the current situation would be well received.

QUESTION: This work relies on network programmability to overcome shortcomings in Intel's Tofino. Is this strategy generalizable for other problems?

ANSWER: This approach is not a general solution. A general problem is that some computation require recirculation when implemented. This, at the end, ends in the main problem that recirculation requires capacity to be realized. Other works approached the implementation of replication without recirculation.

QUESTION: On the issue of bit mask-based replication: if using a bit mask to represent a group, would using a wildcard match (done via TCAMs as generalization of longest prefix match) allow for a fast check which part of the group a packet shall be copied to? This approach seems like a combination of this work with previous work on bit mask-based multicast using existing SDN switches.

ANSWER: Yes, in an advanced version of the presented work something similar to this approach is used. First a specific cluster is retrieved; a cluster representing a subset of required ports. Afterward, a specific multicast group can be retrieved. This is feasible.

3.3 Topological Addressing

Luigi Iannone (Huawei Research Paris – Paris, France)

Final speaker of the first session was Luigi Iannone, elaborating on the topic of topological addressing. Related work by the speaker includes [LLI22].

QUESTION: What would be the highest load imaginable a network element has to handle? Is there a use-case where this processing matters and should be optimized by a device which implements this?

ANSWER: Regarding use-case, the following comes to mind: an industrial IoT scenario with legacy equipment that does not feature an IPv6 stack. The owning company wants to renew its machinery. With the presented approach one can connect the old devices with IPv6. In general, for devices with limited computational power the need to reduce computation in order to save energy arises. The presented approach can reduce communication overhead. Due to the saved computation, energy is saved.

4 Real-Time Networking approaches and tools for Industrial Application Scenarios

During the second session of the second academic salon's first day, focus of the contributions and discussions moved to the area of real-time networking approaches and tools for industrial application scenarios. The session was chaired by Georg Carle.

4.1 High-capacity and Resilient Large-scale Deterministic IP Networks

Paolo Medagliani (Huawei Research Paris – Paris, France)

This session was opened by a talk about high-capacity and resilient large-scale deterministic IP networks by Paolo Medagliani. The talk relates to recently published work by the speaker [Ang+22].

QUESTION: Was an implementation evaluated in an efficient manner? Would an implementation benefit from P4-style network elements or are there hindering limitations of P4?

ANSWER: Implementing the presented approach was not the main objective thus far. P4 used in tests and for first prototypes with real switches. Especially, when looking at industrial control traffic (cf. Programmable Logic Controllers (PLCs)). The current implementation relies on an FPGA-based solution. Still, P4 remains interesting.

QUESTION: Considering a scenario where some network elements have support for this scheme and others do not: How can this be solved?

ANSWER: This is an interesting topic and a research direction that has to be taken. Realistic networks where only some parts of the network have support for Large-scale Deterministic Networks.There, embedded slices could specify QoS requirements. Such an approach could be a solution but seems to be lacking due to its inability to provide reliable jitter.

QUESTION: In general: What is the impact when increasing the number of nodes?

ANSWER: With an increase in number of queues in outgoing ports, a higher delay is observed.

QUESTION: How is routing dealt with? What are relevant constraints in this regard?

ANSWER: Here guarantees can be provided. Using the output of the Integer Linear Program, for each of these paths, multiple paths can be used. Since paths can be assigned weights, for each of the multiple paths a random path can be selected and put in the network. This can be applied to all paths, potentially in parallel. With fast heuristics multiple iterations are feasible. Moreover, pruning of solutions is possible, depending on the concrete circumstances.

4.2 Reproducible Layer 3-enabled TSN Experiments

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

The final talk of this second session was jointly held by Filip Rezabek and Marcin Bosk. They introduced ENGINE [Rez+22], an environment for generic in-vehicular networking experiments.

QUESTION: Is this approach only applicable to vehicular scenarios or does it apply to industrial use-cases as well?

ANSWER: It is definitely applicable in both cases. Both vehicular and industrial settings are very similar on Layer 2, i.e., TSN. There is also a lot of overlap concerning involved infrastructure. But the application stack involved will certainly be changed, e.g., to OPC UA.

Panel Discussion: What challenges of industrial networking remain to be resolved?

The panel discussion was opened by raising a couple of initial questions outlining the following discussion scope: *What (scientific) challenges of industrial networking remain to be resolved? Are there additional network mechanism that can be provided? Or are there enough mechanisms and one just need to select the correct one? Is there a lack of tools that give the right answer? Or are the tools available, but lack something small? Concerning data sets: are sufficient data sets accessible? Is something else is missing?*

Addressing the talk by Paolo Medagliani, it was commented that, when considering the switch to a new solution, the need to roll out new devices arises. This would result in two problems: (1) How can new solutions co-exist with the current network? (2) What are realistic use-cases operators can make money of? To solve these issues, discussions with customers and operations were deemed as needed. The commenter concluded that, networks with low utilization often provide good QoS. Thus, incentive to change to deterministic service is low.

This final comment was rephrased to a question: what is cheaper and more cost-efficient – devices with deterministic IP or cheap devices that can be scaled up to avoid issues? It was remarked that this question matches typical discussions with OEMs. Agreement was reached that cost-effectiveness also has to do with tool support and the need for experts that understanding the use-case before being able to implement solutions. Things get easier if tools and workflows are available where everything is solved automatically. A relevant remark was that when calculating cost effectiveness one should also calculate that when switching to a new technology, a typical problem is that know-how in the company is missing on how to run the hardware. This missing knowledge itself is costly and should not be ignored since it will slow down deployment of new technology. Others remarked that in the past retooling organizations was a big buzzword which fits here. They noted that a problem is speed of the industry and that the current struggle is to move from isolated technologies to IP, while the next change, the move to the cloud, is already around the corner. Their argument is that the industrial networking domain is slower in adoption than the communication domain, and the automotive domain. An argument raised was that this primarily is an organizational problem, not a technical one.

Building on the remark that the move to the cloud is coming, it was pointed out that one of the current problems there is how one can move from hard deadlines to cloud technologies. It was argued that people versed in cloud technologies are unable to grasp the issue of hard deadlines. The suggested solution to the question of solving this discrepancy is to use the competences of domain experts. It was argumented that domain experts for industrial networks are important to help people foreign to this application domains to better understand all the problems to be faced. Other participants agreed, reminding of the example of SmartGrids. There, SmartGrids were misinterpreted as a routing problem by non-domain

experts. An unjust conclusion on transferring an economical problem into a technological one; with SmartGrids having different economics than the Internet. The discussion continued stating that SmartGrids, concerned with distribution of electricity, consider the network only as second aspect. Conclusion of the argument was agreement with the need for domain experts to understand problems, and using their domain language to properly express the problems.

Through the chat, a new theory was voiced: the idea of having a one-fits-all solution for industrial networking and to battle complexity by network automation. Responses tended to agree with the second part but voiced concerns whether a one-fits-all solution exists. The responders' argument included the note that even IP does not address everything, instead it is adapted all the time. The responders countered that instead a specialized solution or ability to do the specialization is required. Rephrased, a solution that fits, after adapting, the problem.

The initially proposed theory was then clarified to a one-fits-all solution that consists of using one network for all kind of traffic, either critical or uncritical traffic. The proposer argued that this would allow also to move Operational Technology (OT) tasks to edge platforms, and make way for the move towards "Software Defined Manufacturing". Responders appreciated the clarification and proposed to use the term *convergence*, reminding that this matches the process the telecommunication world went through many years ago.

Raising the discussion's final point, the question came up whether the work on virtual PLC should be considered a topic that sparks the interest of people in industry. Putting a spotlight on Germany, it was remarked that virtualization is slowly taking up, mainly with startups talking about Software-Defined Automation (SDA). While the run on top of commodity hardware is considered a good factor, participants remarked that the industrial sector is conservative. Thus, there is need to verify multiple times whether new approaches are sound before moving them to production. In other words, many demonstrators are required to convince that it is already possible to fulfill the relevant requirements. This will become subject of further discussions in coming years.

5 Programmable Network Interface Cards and In-Network Computation – In-network computation in industrial scenarios, Processor-based SmartNICs

The third session of the academic salon consists of three talks addressing programmable network interface cards, in-network computation with a particular focus on industrial scenarios, and processor-based SmartNICs. The talks mainly focused on applications and scenarios where SmartNICs can provide significant benefits. The session chair, Michael Menth, welcomes all participants and panelists with a short introduction, then proceeding to the invited speakers and their talks about emerging topics in this area, each followed by an extensive Question and Answer section.

5.1 Packet Ordering for Improving Application Performance using SmartNICs

Dejan Kostic (KTH Royal Institute of Technology – Stockholm, Sweden)

Dejan Kostic presented his view on packet ordering to improve application performance using SmartNICs and their application fields. The talk refers to the following paper presented at the NSDI conference 2022 [Gha+22].

QUESTION: What is your assessment of next-generation devices using, for example, P4 on switches?

ANSWER: Using P4 on switches for next-generation devices or devices such as Intel mount Evans as SmartNIC are used in next-generation much better than an array of ARM cores or other middleware devices.

QUESTION: What is your viewpoint about batching and reordering?

ANSWER: We did not discuss this question, which is the attention between networking protocols and spacing the packets, and the desire to batch packets and send them back to back. We were lucky to get these results by arranging the packets to send to the data center because switches dislike tiny packets arriving back to back; they are not designed for such scenarios. Furthermore, we do not have to worry about stuff on the internet; we can place a reframer on our side of the network.

QUESTION: It looks like this is an excellent addition to an earlier approach with Swedish participation called Hipparch, c.f. http://www0.cs.ucl.ac.uk/staff/jon/hipparch/hipparch.html

ANSWER: Yes, the idea of batching is not new; this is already done a lot in data centers.

QUESTION: The refrainer performed well in your results. When you had more round-trip times, more packets were in flight, but in a data center with shorter round-trip times, it would be interesting to see an effect on combining it with refrainer with a long round-trip time effect.

ANSWER: This is an excellent idea to try as this may have additional influences, and what will happens would be interesting in general. Seeing the results of such an approach is interesting to us.

QUESTION: Is it better when service functions can work efficiently in terms of energy consumption? How significant is the overhead here?

ANSWER: The right place to perform such an experiment is in the same machine setup and check the power consumption there and we are curious about what happens then in for example the L1 cache. Some of the influences come from unnecessary inter-core connections. They had another work called Metron there with big improvements on batching when considering core and cache placements. However, massive efficiency gains are possible on the same hardware when we use the architecture-specific settings in an optimized way.

5.2 The paradigm shift in automotive zonal gatewaying

Francesc Fons, Angela Gonzalez Mariño, Abdoul Aziz Kane (Huawei Research Munich – Munich, Germany)

Francesc Fons, in his talk "The paradigm shift in automotive zonal gatewaying", presented a view on automotive network structures and gateway designs. We are moving from domain-based controllers towards zone gateways using the positions in the automotive vehicle. The speaker's experience with the topic manifests in recent publications [Gon+21; Rez+22]

QUESTION: Is the presented work more of a research idea or already a product?
ANSWER: We share it as an idea and consider it not yet as a product.
QUESTION: Is it a significant trend or more a proposal to the market?
ANSWER: We work on this idea and evaluate if it is competitive enough to be used as a product in the future.

5.3 In-network computation and Processor-based SmartNICs

René Glebke, Klaus Wehrle (RWTH Aachen – Aachen, Germany)

René Glebke, in his talk "In-network computation and Processor-based Smart-NICs", presented insightful views on processing data on-path and related scenarios. The talk builds on and, hence, closely relates to previously published work [Kun+21]. Summarized are the following questions to this talk:

QUESTION: Do you have any applications in mind that could benefit from in-network computing?
ANSWER: So far, we had a hammer and searched for a nail, so this was not our first thing to search for. We are excited at having the chance to collaborate with mechanical and electrical engineering and have problems with overloading the network with additional information needed. We want to find out, for example, the failure rate and additional data post-mortem, which is needed for mechanical engineers to improve their work.

6 Mechanisms and Methods for Real-Time Networking approaches suitable for Industrial Application Scenarios

The fourth and last session of the academic salon consists of three talks and a panel discussion about mechanisms and methods for real-time networking approaches. The focus is on approaches suitable for industrial application scenarios. The session chair was Jörg Ott, who welcomed the participants and panelists to the last session and introduced the talks with emerging ideas and topics and an extensive question and answer section after each talk.

6.1 Application of Network Calculus for Reliable and Predictable Behavior of IEEE 802.1CB Frame Replication and Elimination in Time-Sensitive Networks

Lisa Maile (University of Erlangen-Nuremberg – Erlangen, Germany)

Lisa Maile introduces in her talk applications for Network calculus-based methods on reliable and predictable behavior within frame replication and elimination in Time-Sensitive-Networks. She focuses on the delay calculation, including additional TSN methods to precisely predict worst-case boundaries to be used in network engineering tasks. Related works to this talk include recent work published at the ICC 2022 [Mai+22]. The following questions and answers were given after the talk by Lisa Maile:

QUESTION: Why do we not use Network Coding here?

ANSWER: Yes, you can use network coding to reduce some of the errors here, but this does not help with link failures, as the packets are sent over the same link again. For example, a robot that needs to stop because someone walks over the way is needed to ensure that the other path is used, and network coding adds more redundancy to the sent data.

QUESTION: We could also send the coded packets duplicated over different paths?

ANSWER: What we are doing is that the message is stopped when already arrived, so there is no loss in time, but in the worst case, the network will still have a packet arriving earlier when the redundant packet is arriving.

QUESTION: Do you plan to change the FRER standard to further improve?

ANSWER: The idea of duplicate and redundancy is not new, and things are reinvented from scratch in TSN. Maybe something will be discussed in the future when reaching out to the working group and discussing current developments.

QUESTION: The idea is to improve link layer robustness, but where on the acceptance scale are we with this at the moment?

ANSWER: TSN is generally adapted, and other shapers are used less often. For FRER, it is probably not used at the moment at all since it is still under standardization.

QUESTION: One question always when doing replication on a path, it is obvious that link layer replication gives a faster reaction, but other things can use more lightweight reception mechanisms.

ANSWER: FRER offers safe reception in a fast way; other mechanisms have a higher delay, e.g., detection and retransmission

6.2 Tail Latency Estimation and Verification

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

Max Helm presents an approach using trained Extreme Value Theory to predict flow-level worst-case bounds on end-to-end delay and verify these results. The talk touches on topics the authors elaborate on further in their recent work [HWC22; Gal+21]. After the talk, the results were discussed in a question-and-answer session, summarized here:

QUESTION: Where can this methodology be applied, e.g., TCP timeouts?
ANSWER: The benefit is that it is cost-effective and valuable to extrapolate values gained from a shorter period of measurements.
QUESTION: What do I use this for, so looking at a network topology, what makes sense?
ANSWER: An example is service level agreements with latency requirements, and you can give predictions that it only exceeds once in a particular time, for example, which can be used in such agreements for simpler communication.
QUESTIONER: Yes, for example, we can use it as an operator interesting in SLA requirements to react in advance.

6.3 Multi-Level Preemption in TSN

Mubarak Adetunji Ojewale (University of Porto – Porto, Portugal)

Mubarak Ojewale presents an approach of multi-level preemption in TSN and its applications. The talk relates to the following paper presented at the EFTA conference 2020 [OYN20].

QUESTION: For which link speeds and network loads is this attractive? Are there speeds or loads it is not attractive?
ANSWER: For most of our experiments $100\,\mathrm{Mbit\,s^{-1}}$ are considered. This area is generally attractive because lower priority flows can cause up to several microseconds of additional delay. For multiple $\mathrm{Gbit\,s^{-1}}$, it gets less attractive, but it is not the case for the automotive use case due to higher power consumption and electromagnetic interference. $100\,\mathrm{Mbit\,s^{-1}}$ is more realistic for this use case.
QUESTION: Did you use this multi-level preemption with Time-Aware Shaper (TAS)?
ANSWER: No, we did not apply it to TAS.
QUESTION: Do you know how many frames can fit into the traffic queues in TAS? The problem to be considered here is that if some switch has more than can be fitted in the queues arriving, they need to be dropped. This can result in non-deterministic behavior, which is a problem in TSN. When you have multi-level preemption, how is this done if not all is beautiful?
ANSWER: There are eight priority classes, if we allow more than one priority to share the same queue, we have to consider worst-case backlogs in the network. Such that queues are large enough to hold backlogs and do not have to drop packets at all.
QUESTION: In short, it depends on the switch type to be used; different assumptions must be made. The strict priority still holds, and therefore you have a level violation in this case already.

Panel Discussion

The session chair thanked the previous speakers and led to a more general panel discussion based on major topics in the area of the academic salon.

The panel discussion was opened using an open statement to be discussed:

QoS enabling in networks was done in many different ways in the talks. This is important as technology has advanced. The last time it was important was in complex ATM networks. But, complex switches are reflecting back to those old, complex systems like ATM, which were replaced with a simplistic approach like Ethernet. How much should we add to Ethernet, and when are we done? Are we already striving for too many things? The panel described that there is quite a bit already done, which is satisfying, which leads to the fact that we are at the moment in gigabit networks already done because, in most cases, all is fine with a delay under one millisecond. Furthermore, different approaches are important for different organizations, such as reducing wired or improving security. Security is a major concern, but nobody typically wants to consider it. In general, Ethernet is already suitable, and TSN is an evolution, but it still works in the basics as expected. Furthermore, the discussion evolved into the thesis that we do not need many features and are basically done. Sometimes new features make sense, sometimes not. They are deciding which ones are hard. In the end, we need more features, but the selection should be done carefully.

Lisa Maile raised an additional new question and area of discussion: *Decentral configuration is also something coming up in new standards, and the areas of applications are huge such as train systems. It is not easy to have an overview and information about the system, needed for the configuration. For all decentralized configuration is planned and on the way, but how low can you go with this? Why do we not use over provision instead?*

The panel discussion is that both centralized and decentralized are complex and complicated, but both have cases where they are possible. Currently, we try to combine the worst of two worlds. However, is this leading to an optimal solution? This is a complicated topic we are pushing towards, and maybe an optimal solution will be missed using a decentralized approach.

Hesham Elbakoury raised another topic for discussion: *Is what is defined in TSN good enough? Is it possible to use what is defined on Layer 2 on Layer 3?*

The panel discussed that shaping can be done in both layers and which one is a political decision. Business wise it is a discussion. Huawei has currently both under development for different use cases. Furthermore, some of the shapers, such as Cyclic Queuing and Forwarding on Layer 3, also require changes.

To not come in an infinite loop and because of the time, the panel was stopped here and Jörg Ott thanked all participants of the panel discussion.

7 Wrap-Up

The chairs thanked the participants for the fruitful discussions and interactions. To make sufficient room for these discussions, some agenda items, in particular panel discussion, in future Academic Salons will be reserved, in order to support rich future interactions. At the same time, when considering the goals of community building and fostering exchange for the academic salon series, the extensive discussions of this year's academic salon should be treated as an indicator of success. For the next iteration, the need for longer discussion time slots in the agenda will be taken into account. Represented by Georg Carle, the TUM chairholder appreciated the good feedback and talks shaping this second academic salon. The common understanding was that the academic salon provided a lot of potential for collaborations among the more than 50 particpants, which can be followed up outside the Zoom setting. The next iteration of this event is planned for autumn 2023, and suggestions for this event are welcome.

References

[Ang+22] Vincent Angilella et al. “High Capacity and Resilient Large-Scale Deterministic IP Networks”. en. In: *J Netw Syst Manage* 30.4 (Aug. 2022), p. 71. ISSN: 1573-7705. DOI: 10.1007/s10922-022-09683-3. URL: https://doi.org/10.1007/s10922-022-09683-3.

[Car+21] Georg Carle et al. *Academic Salon on Time-Sensitive Networking and Deterministic Applications*. Oct. 2021. URL: https://net.in.tum.de/talks/workshops/academic_salon_21.html.

[Fei+21] Jiawei Fei et al. “Efficient sparse collective communication and its application to accelerate distributed deep learning”. In: *Proceedings of the 2021 ACM SIGCOMM 2021 Conference*. SIGCOMM ’21. New York, NY, USA: Association for Computing Machinery, Aug. 2021, pp. 676–691. ISBN: 978-1-4503-8383-7. DOI: 10.1145/3452296.3472904. URL: http://doi.org/10.1145/3452296.3472904.

[Gal+21] Sebastian Gallenmüller et al. “Ducked Tails: Trimming the Tail Latency of(f) Packet Processing Systems”. In: *3rd International Workshop on High-Precision, Predictable, and Low-Latency Networking (HiPNet 2021)*. Izmir, Turkey, Oct. 2021.

[Gha+22] Hamid Ghasemirahni et al. “Packet Order Matters! Improving Application Performance by Deliberately Delaying Packets”. In: *19th USENIX Symposium on Networked Systems Design and Implementation (NSDI 22)*. Renton, WA: USENIX Association, Apr. 2022, pp. 807–827. ISBN: 978-1-939133-27-4. URL: https://www.usenix.org/conference/nsdi22/presentation/ghasemirahni.

[Gon+21] Angela Gonzalez Mariño et al. “PDU Normalizer Engine for Heterogeneous In-Vehicle Networks in Automotive Gateways”. en. In: *Applied Reconfigurable Computing. Architectures, Tools, and Applications*. Ed. by Steven Derrien et al. Lecture Notes in Computer Science. Cham: Springer International Publishing, 2021, pp. 140–155. ISBN: 978-3-030-79025-7. DOI: 10.1007/978-3-030-79025-7_10.

[HWC22] Max Helm, Florian Wiedner, and Georg Carle. “Flow-level Tail Latency Estimation and Verification based on Extreme Value Theory”. In: *18th International Conference on Network and Service Management (CNSM 2022)*. Thessaloniki, Greece, Oct. 2022.

[Kun+21] Ike Kunze et al. “Investigating the Applicability of In-Network Computing to Industrial Scenarios”. In: *2021 4th IEEE International Conference on Industrial Cyber-Physical Systems (ICPS)*. May 2021, pp. 334–340. DOI: 10.1109/ICPS49255.2021.9468247.

[LLI22] Guangpeng Li, David Lou, and Luigi Iannone. “Topological addressing enabling energy efficient IoT communication”. In: *Proceedings of the ACM SIGCOMM Workshop on Future of Internet Routing & Addressing*. FIRA ’22. New York, NY, USA: Association for Computing Machinery, Aug. 2022, pp. 12–17. ISBN: 978-1-4503-9328-7. DOI: 10.1145/3527974.3545722. URL: http://doi.org/10.1145/3527974.3545722.

[Mai+22] Lisa Maile et al. “Ensuring Reliable and Predictable Behavior of IEEE 802.1CB Frame Replication and Elimination”. In: *ICC 2022 - IEEE International Conference on Communications*. ISSN: 1938-1883. May 2022, pp. 2706–2712. DOI: 10.1109/ICC45855.2022.9838905.

[Mel19] Mellanox. *Mellanox Scalable Hierarchical Aggregation and Reduction Protocol (SHARP)™ - SHARP v1.8.1 - NVIDIA Networking Docs*. 2019. URL: https://docs.nvidia.com/networking/display/SHARPv181/.

[MLM20] Daniel Merling, Steffen Lindner, and Michael Menth. “P4-based implementation of BIER and BIER-FRR for scalable and resilient multicast”. In: *Journal of Network and Computer Applications* 169 (2020), p. 102764. ISSN: 1084-8045. DOI: https://doi.org/10.1016/j.jnca.2020.102764. URL: https://www.sciencedirect.com/science/article/pii/S1084804520302381.

[MLM21] Daniel Merling, Steffen Lindner, and Michael Menth. “Hardware-Based Evaluation of Scalable and Resilient Multicast With BIER in P4”. In: *IEEE Access* 9 (2021). Conference Name: IEEE Access, pp. 34500–34514. ISSN: 2169-3536. DOI: 10.1109/ACCESS.2021.3061763.

[OYN20] Mubarak Adetunji Ojewale, Patrick Meumeu Yomsi, and Borislav Nikolić. “Multi-Level Preemption in TSN: Feasibility and Requirements Analysis”. In: *2020 IEEE 23rd International Symposium on Real-Time Distributed Computing (ISORC)*. ISSN: 2375-5261. May 2020, pp. 47–55. DOI: 10.1109/ISORC49007.2020.00017.

[Rez+22] Filip Rezabek* et al. “EnGINE: Flexible Research Infrastructure for Reliable and Scalable Time Sensitive Networks”. In: *Journal of Network and Systems Management* 30.4 (Sept. 2022), p. 74. ISSN: 1573-7705. DOI: 10.1007/s10922-022-09686-0. URL: https://doi.org/10.1007/s10922-022-09686-0.

[Sap+21] Amedeo Sapio et al. “Scaling Distributed Machine Learning with In-Network Aggregation”. en. In: 2021, pp. 785–808. ISBN: 978-1-939133-21-2. URL: https://www.usenix.org/conference/nsdi21/presentation/sapio.

[Yua+21] Yifan Yuan et al. *Unlocking the Power of Inline Floating-Point Operations on Programmable Switches*. arXiv:2112.06095 [cs] version: 1. Dec. 2021. URL: http://arxiv.org/abs/2112.06095.