








TSN-FlexTest: Flexible TSN Measurement Testbed (Extended Version)

Marian Ulbricht , Stefan Senk , Hosein K. Nazari , How-Hang Liu ,
Martin Reisslein , Giang T. Nguyen , and Frank H. P. Fitzek 

Abstract—Robust, reliable, and deterministic networks are essential for a variety of applications. In order to provide guaranteed communication network services, Time-Sensitive Networking (TSN) unites a set of standards for time-synchronization, flow control, enhanced reliability, and management. We design the TSN-FlexTest testbed with generic commodity hardware and open-source software components to enable flexible TSN measurements. We have conducted extensive measurements to validate the TSN-FlexTest testbed and to examine TSN characteristics. The measurements provide insights into the effects of TSN configurations, such as increasing the number of synchronization messages for the Precision Time Protocol, indicating that a measurement precision of 15 ns can be achieved. The TSN measurements included extensive evaluations of the Time-aware Shaper (TAS) for sets of Tactile Internet (TI) packet traffic streams. The measurements elucidate the effects of different scheduling and shaping approaches, while revealing the need for pervasive network control that synchronizes the sending nodes with the network switches. We present the first measurements of distributed TAS with synchronized senders on a commodity hardware testbed, demonstrating the same Quality-of-Service as with dedicated wires for high-priority TI streams despite a 200% over-saturation cross traffic load. The testbed is provided as an open-source project to facilitate future TSN research.

Index Terms—Ethernet, Industrial Communication, Quality-of-Service, Testbed, Time-Sensitive Networking.

I. INTRODUCTION

There are various reasons of unexpected behaviors in communication networks, such as changing numbers of connected devices, sudden spikes in network utilization, or even loss of entire transmission paths due to unpredictable events. For a wide range of applications and domains it is necessary to avoid these unexpected behaviors. Among others, the industrial, automotive, medical, and avionic domains have strict requirements on the underlying communication infrastructure [2]–[7].

An abridged preliminary version appeared in [1].

Marian Ulbricht and Stefan Senk contributed equally to this work.

M. Ulbricht, S. Senk, H. K. Nazari, H.-H. Liu, and F. H. P. Fitzek are with Technische Universität Dresden (TUD), Deutsche Telekom Chair of Communication Networks

M. Reisslein is with the School of Electrical, Computer and Energy Eng., Arizona State Univ., Tempe, AZ 85287-5706, USA, Email: reisslein@asu.edu

G. T. Nguyen is with Technische Universität Dresden, Junior Professorship of Haptic Communication Systems

G. T. Nguyen and F. H. P. Fitzek are with the Cluster of Excellence “Centre for Tactile Internet with Human-in-the-Loop” (CeTI) of TUD

To mitigate problems in the communication infrastructure, deterministic data transmission is preferred. Determinism in networks can be described as maintaining full-knowledge and control of packet-based transmissions. However, achieving determinism in data communication is challenging. Depending on the use case, extensive efforts must be made to ensure the network quality for critical services. Time-Sensitive Networking (TSN) unites different mechanisms to achieve deterministic communication over Ethernet networks. TSN is managed by the IEEE TSN Task Group (TG) which defines the underlying standards. TSN emerged from the former Audio/Video Bridging (AVB) Task Group which had a focus on streaming services. With the new TSN TG, the attention moved towards more generic applications, such as industrial communication.

TSN mainly encompasses four categories: time-synchronization, latency and packet delay variation reduction through flow control, ultra-reliability, and resource management. Each category improves specific aspects of data communication. The standards can be flexibly combined to adapt to particular use cases. Although the standards can be used individually, the greatest benefit is obtained by a judicious combination of selected features. For instance, using time-aware traffic shaping algorithms usually requires tight time-synchronization. More specifically, there is a TSN standard that provides mechanisms to shape traffic based on a cyclic behavior, similarly to Time-Division Multiple Access (TDMA). Packet transmissions can then be optimized by synchronizing end-stations with time-aware bridges. Hence, employing a second TSN standard for precise time-synchronization can enable new opportunities or enhance existing mechanisms for deterministic data transmissions.

Recent studies have mainly focused on integrating TSN with wireless technologies, as well as frameworks for managing and optimizing TSN core functionalities, such as the Time-aware Shaper (TAS) [8]. Also, enhancing end-to-end deterministic transmission [9], and the integration of TSN with the software-defined networking (SDN) paradigm has been pursued [10]–[12].

A. Limitations of Existing TSN Evaluation Frameworks

TSN focuses not only on the reliability of packet transmissions, but also on the timing aspects of the packet trans-

missions. It is crucial to precisely measure timing-related metrics, such as one-way delay and Packet Delay Variation (PDV). Three primary methods are used for evaluating TSN protocols and systems: theoretical mathematical analysis, simulation and emulation, and hardware testbeds. Mathematical analysis frameworks, such as [13], [14], have been developed to evaluate TSN. For instance, He et al. [15] analyzed the worst-case travel time of Ethernet frames with the Credit-based Shaper (CBS) in the AVB context, and the TAS from the TSN domain. Guo et al. [16] studied strict minimum criteria in use scenarios. However, the mathematical analysis frameworks abstract the behaviors of TSN systems compared to real-world systems. Simulation frameworks, such as OMNET++ [17] and NS-3 [18], have been widely used in networking research, including TSN research. The advantages of simulations include flexibility, reduced cost, and scalability. A main drawback of simulations in TSN research is that they do not involve real experimental hardware networking components, making it impossible to showcase the applicability and demonstrate the developed TSN systems and protocols with real operating networks. Emulation frameworks and tools were developed to address the applicability issue of simulation software. Ulbricht et al. [19] implemented TSN in the Mininet emulation software [20], operating the entire emulation and TSN framework on Linux. Nodes communicate with one another via virtual networking interfaces. The scale of the emulated system is limited to the computing resources of the PC running the emulation. Therefore, timing-related emulation measurements of large and complex systems are unreliable.

Using dedicated hardware is, in most cases, more precise while having the same complexity as the targeted actual network deployment. Even though scalability is a limiting factor for designing hardware-based TSN testbeds caused by limited resources, dedicated TSN devices offer the full range of settings and parameters to be considered. However, testbeds have complexities: the measuring techniques, setup, and other elements can strongly influence the outcomes. Therefore, building a testbed with high precision and flexibility is challenging, yet desirable and facilitates reproducible experiments.

As TSN is a family of standards, TSN-related testbeds have been built to study various isolated TSN aspects, namely i) packet processing, ii) evaluating the Precision Time Protocol, iii) communication over-the-air; iv) TSN performance; and v) TSN management. One of the primary goals of packet processing frameworks, such as MoonGen [21] and P4STA [22], is to generate packets precisely and to precisely timestamp both transmitted and received packets. Even though MoonGen requires Data Plane Development Kit (DPDK)-supported hardware, it offers a high degree of flexibility and high performance.

Nevertheless, a testbed for studying TSN needs more than packet generation and processing. Next, several existing testbeds focus on time synchronization, primarily with Precision Time Protocol (PTP), such as [23]–[25]. However, they solely evaluate the synchronization precision without assessing TSN mechanisms. More comprehensive testbeds have been built to study the performance of TSN in combination with wireless communication, such as WiFi and 5G [26]–[28]. A

few testbeds target specific application contexts, such as for benchmarking TSN performance of the controller on multi-domain networks [29] and for intra-vehicle networking [30]. However, they either neglect the per-packet delay or have limited precision due to inconsistent hardware and software timestamping. Overall, the existing testbed studies focus on specific narrow sets of Key Performance Indicators (KPIs). Also, the existing studies commonly consider only limited (coarse) precision levels and often neglect detailed investigations of the effects experienced by a TAS scheduled data stream.

B. Contributions: Open-Source TSN-FlexTest Testbed

We address a root impediment in TSN research: the need for high-precision reproducible measurement studies in comprehensive TSN testbeds. We design and develop TSN-FlexTest, an open-source, flexible, high-precision, yet affordable measurement testbed framework to comprehensively study TSN. Our key innovation is a well-engineered combination of affordable Commercial-of-the-Shelf (COTS) hardware with an automated measurement workflow from packet generation to data collection and visualizing measurement results. Our open-source measurement framework facilitates the sharing of experiment scenarios and setups. Furthermore, the testbed workflow is organized into a centralized management and control entity, making the testbed workflow very easy to change. Last, introducing key abstraction entities, such as the Device-under-Test (DuT) and the traffic generator, make the entire testbed flexible to change. Striving towards high precision and affordability at the same time, we advocate the use of commodity hardware, which also facilitates reproducibility.

This paper makes the following two main contributions:

- 1) section IV first reviews the general design principles for a TSN testbed, including the available options for hardware and software components and their trade-offs. Subsequently, we detail the TSN-FlexTest testbed architecture, design, and implementation, including the workflow of our measurement framework. We also share the lessons we learned while developing our measurement framework. In section V, we validate our TSN-FlexTest testbed, which produced results with nanosecond precision. A preliminary version of this contribution is presented in [1], focusing on the testbed design. We make TSN-FlexTest openly available at <https://github.com/5GCampus/tsn-testbed>.
- 2) The second main contribution of this paper is the investigation of critical aspects of TSN leveraging the developed and validated TSN-FlexTest testbed in section VI. These are the significant differences of this extended article compared to the preliminary conference version [1]. There are four main findings: 1) We evaluate a typical configuration in a net-neutral Internet, without specialized packet treatment and with limited resources in the backbone networks. We demonstrate that it is nearly impossible to ensure prescribed service quality in such a net-neutral Internet. 2) We investigate the impacts of soft Quality-of-Service (QoS) with Strict

Priority Queuing (Strict Priority Queuing (SPQ)) and conclude that SPQ cannot enforce strict guarantees. 3) We measure the one-way packet latency for TAS. The results reveal inconsistent advantages of traffic shaping and hard scheduling. The TAS sometimes performs better, especially for the average latency. 4) We demonstrate that it is possible to synchronize a generic application running on COTS hardware with the TAS by scheduling on the transmitter side itself.

II. BACKGROUND

Time-Sensitive Networking, is a collection of standards that enable guaranteed latency and deterministic communication over wired Ethernet. These standards can be categorized according to their functions: time-synchronization, flow control, ultra-reliability, and network management. This background section highlights the main approaches of TSN evaluation. Next, background on high-precision time-synchronization, TSN flow control, and TSN KPIs is provided.

A. TSN Evaluation Methodologies

There is a set of tools in order to evaluate these recent developments in the area of TSN. Each tool provides different advantages and disadvantages, such as low costs, high reproducibility, or significance for real-world deployments.

1) *Analysis, Simulation, and Emulation*: Studies often employ a formal evaluation based on mathematical analysis that is typically coupled with synthetic test cases [31]. The coupled evaluation approach ensures reproducibility: Equations do not rely on hardware nor software, and evaluation results should be the same irrespective of the underlying computing hardware. Depending on the degree of modeling details in the implementation of algorithms and abstracting “real-world” behaviors, simulation can offer a solid foundation for research. Furthermore, simulation provides the ground to investigate various architectures, influences, and interdependencies, that would be difficult to examine in real testbeds. This can be, e.g., due to missing hardware: setting up large networks is expensive and incurs high costs for setup and maintenance; also, interoperability issues may arise. Furthermore, for recent developments as well as new designs and algorithms, hardware implementations may not be available yet.

An alternative to simulation and formal evaluation is emulation. Emulation uses real implementations, but without high investments in hardware infrastructure. The network emulator MiniNet, for instance, can be configured to include TSN features.

2) *Hardware Testbed Measurement*: TSN network behaviors can be studied through measurements on actual hardware testbeds [32]–[36]. Generally, testbed measurements offer authentic results in terms of closeness to real systems [37]–[44]. The main drawback of testbed measurements is in most cases, besides the actual hardware costs, the complexity of the testbed system [45]–[48]. Not only software can have bugs, but the actual hardware implementation may have design flaws as well. Investigating real hardware also introduces dependencies

on vendors. If there is no open hardware platform available, which is—seen with eyes of independent researchers—unfortunately common, then the hardware must be treated as is. Consequently, it is usually impossible to tweak, modify, or gain insight into the implementation specifics of the hardware from commercial vendors. Only self-developments, e.g., with Field-programmable Gate Array (FPGA)-based boards, offer the freedom to modify the hardware, but at a much higher cost of initial implementation.

Studying real-world systems with testbed measurements poses additional challenges. First, researchers must rely on the capabilities of the Device-under-Test (DuT). Hence, a DuT is often treated as a *black box*. Therefore, if the DuT itself does not offer any evaluation metrics, such as a packet counter, then measurements must be taken from the outside, necessitating additional measurement equipment. Further complexity is introduced by the fact that measurement probes in networks are often distributed. For instance, measuring time-dependent parameters, e.g., one-way delay, requires precise synchronization. Also, measurement tools can be expensive. If a certain use case should be investigated, not all tools may offer the degree of freedom of adaptation, resulting in similar dependencies on vendor implementation as for the DuT.

B. High-precision Time-synchronization

1) *Overview*: In a TSN network, network components need to be time-aware and have a shared understanding of time with high precision. This knowledge of time might be either absolute or relative. While absolute timing is typically used to associate data points with specific real-life occurrences, relative timing is primarily concerned with synchronizing various components and maintaining them synchronized to coordinate several processes that are running simultaneously. The needed level of synchronization varies depending on the use case and ranges from tens of milliseconds to tens of nanoseconds. The Precision Time Protocol (PTP) is recommended in the TSN standards for achieving this level of synchronization.

2) *PTP Protocol*: The PTP protocol can distribute precise timing, phase, and frequency over packet networks with sub-microseconds precision. The time source in the PTP protocol is known as the Grand Master (GM) clock. The GM clock exchanges time information with other nodes, which can themselves be a boundary, ordinary, or transparent clock. The GM in the PTP protocol can be chosen either automatically or manually within a PTP time domain. The automated selection process leverages the Best Master Clock (BMC) algorithm. Generally, the clock with the highest precision in the clock domain is the BMC, which is detected by using ID, class, priority, or MAC address. GM clocks are frequently synchronized using Global Navigation Satellite Systems (GNSS). Satellite constellations offering global or regional positioning, navigation, and timing services are referred to as GNSS (such as GPS or GLONASS). Another option is employing highly precise oscillators (e.g., cesium or rubidium clock). Ordinary clocks only have one port for receiving the time information, boundary clocks have two or more PTP ports and can operate as a time source to other connected nodes.

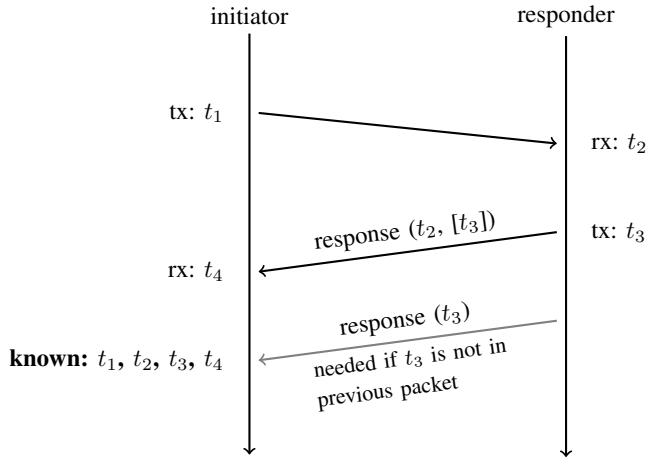


Figure 1: Illustration of Precision Time Protocol synchronization according to IEEE 802.1AS [49]: Via a message exchange in a master-slave system, two clocks can be synchronized with sub-microseconds precision using path-delay measurements.

To distribute time correctly across all clocks in the network, the path-delay between two clocks has to be measured precisely. The synchronization thereby leverages timestamping, i.e., the exact time needs to be taken when a packet leaves at the transmitter (TX) and arrives at the receiver (RX).

Distributing time from the GM is insufficient to synchronize clocks, and the influence of path delay on the synchronization should be minimized. Path delay techniques leverage the timestamps supplied in frames. The path delay can be evaluated in a one-way mode, where the master node sends sync and follow-up messages, or in a two-way mode, where downstream clocks also send delay request packets to master clocks. The path delay can be evaluated in the unicast mode, where the master clock generates sync and follow-up unicast messages, or in a mixed mode, where the master clock sends multicast messages but responses are unicast. The path delay must be evaluated periodically since it is dynamic and fluctuates over time. The length of this period depends on the required level of precision. This period length and other PTP parameters can have an influence on the performance of the PTP protocol.

3) *Timestamps—Step by Step*: To ensure sub-microsecond precision, some degree of hardware support is needed. The RX hardware timestamp is latched by the Network Interface Card (NIC) when the packet arrives and can be used by the software layer later on. The TX timestamp can be implemented in three ways: one-step mode, two-step mode, and a combination of both one- and two-step modes. In two-step mode, the NIC reports the TX timestamp of the currently to-be-sent packet with the next frame, in a so-called *follow-up* message. This can be done solely in software. In one-step mode, the NIC has the capability of correcting the time information of outgoing frames right before their transmission. An offset register specifies the location within the Ethernet frame. Finally, there is a 1.5-step mode, where the Central Processing Unit (CPU) generate a dummy *follow-up* message, and the NIC only updates the few fields to form a real *follow-up* message.

Each method has its advantages and disadvantages: While the one-step mode is less complex for the receiving clock implementation, the complexity of the sender NIC increases because of calculations to provide the timestamp at the point of time the frame will leave the Physical Layer Chip (PHY). On the other hand, the two-step mode does not require the computation on the TX clock. Furthermore, the one-step mode can only be employed if the link speed does not exceed a certain threshold. Because of the timestamp calculations that need to be done in real-time, at some point the link speed is higher than the hardware capabilities.

4) *Path Delay Evaluation*: Figure 1 illustrates the transmission steps for the PTP time-synchronization and the corresponding path-delay calculation. For calculating the path delay, four time values are needed. The initiator and the responder of the path delay measurement exchange two to three messages, to obtain the packet runtime. The initiator generates a message and stores its transmission time t_1 . The responder replies to the message and includes the time t_2 when the first message was received. Depending on the one- or two-step procedure, an additional message transfers the transmission time of the reply t_3 back to the initiator. Using $d = [(t_2 - t_1) + (t_4 - t_3)]/2$, the four timestamps give the path delay between initiator and responder.

5) *The Art of Timestamp Provisioning*: PTP can be deployed on different hardware, such as commercial devices, e.g., Ethernet switches, as well as COTS computer hardware and general-purpose operating systems, e.g., GNU Linux. LinuxPTP [50] provides a commonly used PTP implementation for the Linux platform. LinuxPTP offers software and hardware support for timestamping. Depending on the implementation of the hardware timestamping mechanism, different packet throughputs can be achieved. The hardware-based timestamping in the RX and TX directions can be enabled by packet filters applied in the network driver layer (corresponding to the TX direction) or in the NIC (corresponding to the RX direction). Typical filters timestamp *only* PTP packets or timestamp *all* packets. However, filters are necessary because timestamping *all* packets can at some point exceed the hardware capabilities.

To describe the bottlenecks, the TX and RX directions have to be considered separately. TX timestamping can be enabled operating system-wide by setting a flag in the NIC's configuration space. This enables the timestamping for all packets, or more precisely: packet by packet. A bit in the packets descriptor data structure, which is used to pass the memory location of the data to the network hardware, can be set for this purpose. If two-step TX timestamping is enabled for all or just one packet, then the NIC will latch the clock time instant when the first bit is on wire (practically it is done before, with a time offset), and stores the value in internal memory. This memory structure needs to be read by high-level software. This timestamp storage can be a bottleneck for the TX direction: The software needs to read the timestamps before the limited memory in the NIC is overwritten [21], [51]. Moreover, the high-level software is responsible to match the read-out timestamps with the Socket Buffer (SKB) structure of the transmitted packet, which contains the pointers to the

packet data, and reporting back the TX timestamp to the higher layers of the network protocol stack. In some implementations, the matching is supported by an ID that corresponds to the transmitted packet [52]. An improvement to this procedure is one-step timestamping: the NIC writes the latched TX timestamp directly into the transmitted packet. After this operation, the Ethernet checksum is repaired. These operations need to be conducted on the fly, during the transmission process of the packet through the NIC.

In the receiving direction, there are also different implementations: The most scalable solution is to deliver the RX timestamp within the packets descriptor structure, so that the memory pointer and the timestamp information are stored in the same structure and do not need to be aggregated later on [51], [52]. Other solutions may be limited in performance: Some implementations provide the RX timestamp in an extra register, which has to be read out by software fast enough to align the timestamp with the corresponding packet [51].

6) *Profiling*: PTP is being developed since 2002 [53] and with an updated version in 2008 [54], the IEEE Std. 1588 supports *profiles*. IEEE 1588-2019, i.e., PTPv2.1 [55], was proposed 11 years later. In contrast to the IEEE 1588-2002 and IEEE 1588-2008 which are not compatible, IEEE 1588-2019 is backward compatible with the previous versions, and new features are completely optional. The primary goal of IEEE 1588-2019 is to improve the PTP protocol's precision, flexibility, and resilience. For instance, PTPv2.1 introduces modular transparent clocks, hybrid multicast/unicast operation to improve flexibility, and a new range of security guidelines for improving robustness. The IEEE 1588-2019 standard provides a variety of options, not all of them are necessary for each use case. As a result, each profile chooses a subset of parameters for achieving different degrees of precision, for various use cases and network topologies. Each profile needs to set a default value and range for PTP attributes and mention prohibited, permitted, and required options as well as allowed node types. The profile settings can directly impact the performance of PTP.

IEEE 802.1AS [49] is a standard of the TSN standards set that enables the network to meet the rigorous requirements of time-sensitive applications, such as tactile feedback, or audio and video data streaming. The IEEE 802.1AS standard includes a PTP profile, called generalized Precision Time Protocol (gPTP). A gPTP domain consists of end station and bridge nodes that can communicate PTP messages directly with each other.

C. Time-sensitive Flow Control

According to [56], the TSN flow control mechanisms can be categorized into three different categories: (i) traffic shaping, (ii) traffic scheduling, and (iii) Frame Preemption (FP) [57]. A TSN network uses a combination of them to fulfill the requirements of different traffic profiles. These methods are described in detail below, with an emphasis on comparing different traffic shapers.

Traffic shapers strive to give appropriate packet transmission opportunities to streams of the different priority classes [31],




Figure 2: General illustration of queuing operation. In modern Ethernet switches, NICs, or similar devices, there are usually multiple queues in which incoming packets are stored. This can be done based on different metrics, e.g., the Priority Code Point (PCP), or with a matching approach, e.g., by using ingress/egress ports to queue mappings. Within each queue, a Transmission Selection Algorithm (TSA) can be applied to change the order of packets. In most cases, First-In-First-Out (FIFO) queuing is employed. Between queues, a second TSA can be used, e.g., Strict Priority Queuing (SPQ). Further it is possible to shape the traffic with a traffic shaper, e.g., with a Time-aware Shaper (TAS), which operates by deploying *gates* in front of queues.

[58]–[60]. The primary distinction between traffic shaping and traffic policy is that shapers do not discard packets but rather delay them. TSN traffic shapers aid in assigning resources to TSN traffic streams, based on priorities, to as to isolate the streams from interference caused by cross traffic. Shapers also attempt to limit the residence time of enqueued traffic. Their benefit is often intended for traffic with higher priority, such as Tactile Internet (TI) data streams. If these traffic classes, particularly those with hard deadlines, are mixed with non-scheduled traffic, there are no assurances for best-effort streams. Different types of shapers have been introduced for TSN. They differ in terms of implementation complexity for software and hardware (e.g., buffer size or required chip logic), introduced PDV (often also referred to as *jitter*), and delay, as well as *burstiness* reduction.

Traffic scheduling determines which frame to transmit based on a scheduling algorithm or policy. The appropriate scheduling policy is usually determined by the use case requirements. A scheduling algorithm may merely examine arrival time, as in FIFO and Last-Come-First-Served (LCFS). Alternatively, to spread bandwidth between traffic classes, a variable quantity of transmission opportunities, such as weighted round robin, might be considered. In some applications, strict priority is used to protect the highest-priority traffic, and the packet with the highest priority can be sent first.

1) *Credit Based Shaper*: TSN inherits the CBS (IEEE 802.1Qav [61]) from the AVB TG. This concept was initially intended to give a guarantee of limited delay and PDV. CBS distributes packets in a queue over time so that the limit of the allowed rate is not exceeded. The

standard defines two classes, A and B, each of which gets a credit. When a frame of a class waits for transmission, its credit grows at a rate known as *idle slope*, and it falls at a rate known as *send slope*. CBS was shown to give unsatisfactory delay assurances at high link utilization [14], [59]. Furthermore, the maximum latency of CBS varies according to topology and hop count. Overall, up to seven network hops, IEEE 802.1Qav ensures worst-case latencies of less than 2ms for class A and less than 50ms for class B [61].

2) *Time-Aware Shaper*: IEEE 802.1Qbv [62] proposes a Time-aware Shaper (TAS) based on the notion of TDMA, to solve the limitations of the CBS. TAS is based on a cyclic transmission with time-triggered window slots. Thus, the nodes must have a precise knowledge of time. IEEE 802.1Qbv establishes a set of queues for each TSN bridge output port. Queues are managed via *gates* that may be opened or closed. Throughout a cycle, the Gate Control List (GCL) specifies which queue at what precise moment can transmit its data. Each of these queues usually represents a level of priority. In addition, the time length of each slot can be identified. The cycle time is the accumulated time of all defined slots in a GCL. After the final entry, a new cycle of the GCL begins. The TAS enables network devices to precisely plan the transmission of enqueued frames. In general, for TAS, the output port is a shared resource that can only be used from one frame at once. The scheduler decides if this resource is full and which frame should be selected next for transmission. If a frame is still in transmission, then it blocks the line and no other packet can be scheduled. This issue results in the case that long low prior frames can block high prior frames because they are already in transmission. The effect could be mitigated by inserting blocking slots into the TAS GCL to block every traffic to be transmitted before the time slot of a high-priority stream. Unfortunately, these guard bands are wasted time, because in most cases no frame is transmitted within. Frame Preemption (FP) is a technique that solves the issue without inserting a guard band. With preemption, the transmission on the Ethernet layer could be interrupted, another frame could be inserted, and the old transmission is continued afterward. Therefore, special features in the PHYs are necessary.

The main limitation of TAS is that it requires network-wide nearly perfect time-synchronization. An unsynchronized stream endpoint may cause a delay, where the streams will have to wait for the next time-triggered window. Moreover, TAS cannot control individual streams if they are with the same priority. Therefore, to achieve the same level of per-flow QoS, additional mechanisms are required, such as Per-Stream Filtering and Policing (PSFP) and Frame Replication and Elimination for Reliability (FRER) [60]. A PSFP enabled switch can discard non-essential streams but schedule high-priority streams on a per-stream, per-priority, and per-frame basis, and FRER duplicates frames on multiple different paths and eliminates them at the receiver to achieve reliability. In addition, it is challenging to design coordinated TAS schedules when multiple TSN switches are in the same TSN domain [63].

3) *Cyclic Queuing and Forwarding*: The standard IEEE 802.1Qch [64] proposes Cyclic Queuing and Forwarding (CQF). CQF makes TSN switch design easier by inserting static configurations into the TAS GCL [65]. In this method, multiple queues are used for each traffic class. All arriving frames are held in the closed queue of a traffic class, whereas previously received frames might be forwarded from the open queue of that class. The technique employs synchronized operations, allowing Local Area Network (LAN) bridges to organize cyclic frame transmissions to achieve zero congestion loss and constrained latency regardless of network structure.

The disadvantage of CQF is the extra delay of at least one cycle for all incoming frames. Furthermore, delay in transmission and processing can result in a frame being received at an incorrect cycle. Although CQF and TAS offer ultra-low latency for TSN traffic, they rely on network-wide time-synchronization since packet transmission is enforced at periodic intervals. Both are best suited for deterministic communication and isolation between traffic streams, but they make inefficient use of network resources.

4) *Asynchronous Traffic Shaper*: The Asynchronous Traffic Shaping (ATS) [66] does not require gates or a global clock. Instead, it makes use of the idea of *leaky* and token buckets, which is an algorithm that regulates discrete event rates. Urgency-Based Scheduler (UBS) is a feature that ATS adds to the bridge architecture in an effort to improve traffic flow. It controls the *burstiness* of the streams, redesigns the traffic at each hop with a per-stream token bucket scheduler, and gives urgent traffic precedence over other traffic classes. A set of periods known as eligibility periods are assigned to frames of certain streams by ATS shapers. This information is used in the scheduling process by the Transmission Selection Algorithm (TSA). Even when running with high link usage and a combination of periodic and sporadic traffic, ATS can efficiently use the bandwidth with comparatively lower complexity than TAS. Using UBS in conjunction with ATS cannot replace the need for TAS, but ATS can be seen as an improvement over the CBS [67]. In [68] it is shown that UBS can enhance link usage by reducing contention, but it does not provide the same level of TAS protection for various priority streams from one another [67].

Shapers should be selected based on the application of the system and the required performance. Generally, real-time TSN systems can be categorized into three groups: (i) event-triggered (e.g., ATS), (ii) time-triggered (e.g., TAS), and (iii) mixed systems. Furthermore, traffic shaping can also be used in a combination with other standards to attain specific outcomes. Table I compares several shaper technologies based on various metrics.

D. Key Performance Indicators for Time-Sensitive Networking

In general, traffic in TSN can be categorized either as time-triggered/cyclic or event-driven/acyclic. 80 percent of data communication in industrial applications can be grouped into the latter group [69]. The requirements and expectations from a TSN system vary depending on the use case and traffic profile of the network, however, they can be broadly characterized.

TABLE I: Comparison of TSN shapers: Precise time-synchronization between entities relates to the necessity for precisely synchronized devices, otherwise performing poorly. Awareness for low priority indicates if the method can protect lower priority traffic against burstiness of highest priority traffic. Deterministic data transmission implies that the shaper can precisely guarantee the requirements of the highest-priority traffic.

Standard Name	Precise Time-Sync. Between Entities	Awareness for Low Priority	Deterministic Data Transmission	Stream-level QoS Control	PDV Control	Higher CAPEX
Credit-based Shaper (CBS)	✗	✓	✗	✓	✗	✗
Time-aware Shaper (TAS)	✓	✗	✓	✗	✗	✓
Cyclic Queuing and Forwarding (CQF)	✓	✓	✗	✓	✓	✓
Asynchronous Traffic Shaping (ATS)	✗	✓	✗	✓	✗	✗

The IEC 61784-1:2019 standard [70] defines KPIs for industrial communication networks:

- Cycle time of transmitted frame
- End-to-end frame latency between devices
- Time-synchronization precision
- PDV within data streams
- Packet-loss ratio

Specific performance has to be attained with regard to these metrics when developing a TSN testbed. Throughout, we use the terms "frame" and "packet" interchangeably. We proceed to explain the importance and relevance of each indicator.

1) *Traffic Profile and Cycle Time*: In a converged network, the cycle time and frame size of high-priority traffic impact channel utilization and can delay low-priority packets [71]. Real-time applications in industrial field bus communication require strict control of timing between controllers and controlled devices. Control packets are typically short frames that are sent periodically at a fixed rate. Motion control and robotics can have control loops in the sub-hundred microseconds range, and in factory automation, this range can be as short as a few microseconds [72].

2) *End-to-end Frame Latency Between Devices*: The main target of TSN is not about minimizing latency for all traffic classes, but rather focusing on providing a deterministic communication network for certain priority streams. In some cases, the TSN system therefore needs to delay lower priority traffic for the benefit of higher priority traffic classes.

3) *Time-synchronization*: The performance of a TSN network is correlated with the degree of precision of its time-synchronization. In the presence of clock inconsistencies, the entire process of packet time-stamping and resource allocation utilizing techniques that depend on precise time will be significantly impacted. Therefore, PTP is a continuous process in a TSN network that aims to regulate clock drifts within a time domain. The configuration of the PTP profile based on available resources, requirements, and network topology has a considerable impact on the TSN network as well.

4) *Packet Delay Variation within Data Streams*: PDV should be maintained under control since it may have a negative effect on user experience and applications that depend on low delay variation. High PDV in industrial applications correspond to unstable tool movements, and most machines will stop working in such circumstances [73]. PDV or late delivery of a packet may result in losing a packet or a packet being discarded by the application leading to an increase in the packet-loss ratio [69].

5) *Packet-loss Ratio*: The packet-loss ratio is another crucial element for Time-Sensitive Networking and almost zero packet loss reliability is required for TSN Networks [8]. For instance, two consecutive packet losses in industrial use cases may harm equipment. Packet loss may be caused by high PDV or equipment failure, but generally, the primary cause of packet loss is network congestion [74].

III. RELATED WORK

This section reviews the related work on TSN network performance evaluation.

A. Formal Mathematical Analysis

Some of the TSN core functionalities, such as traffic shaping, are tractable in formal mathematical analysis. Mathematical analysis can explore a wide range of configurations for an equation-based model of the TSN system. As a result, performance can be analyzed more thoroughly than with simulations or empirical measurements for specific configurations.

Several studies have used formal mathematical analysis to evaluate TSN, such as [13], [14]. He et al. [15] concentrate on the CBS in AVB context, and the TAS from the TSN domain. They evaluated the worst-case travel time of Ethernet frames to ensure that deadline requirements are satisfied. Guo et al. [16] describe formal methods for traffic shaping to validate strict minimum criteria in use scenarios. They employed the *UPPAAL verifier* which is a real-time system modeling, verification, and validation tool. They also compared shapers in terms of utilization and delay.

Network calculus is a specific formal mathematical framework comprising a collection of mathematical principles for analyzing network performance. Network calculus can be used for evaluating TSN standards due to its capability to analyze the maximum buffer size and worst-case delay [75]. Zhao et al. [31] analyze the performance of TSN shapers, such as ATS, CBS, and Strict Priority Queueing (SPQ), as well as their combination, with network calculus. Thomas et al. [76] use network calculus to evaluate the worst-case latency for Frame Replication and Elimination for Reliability (FRER). Li et al. [59] used network calculus to evaluate the performance of the Stream Reservation Protocol (SRP) with CBS and demonstrate that increasing reserved bandwidth does not always improve the tightness of the delay bound. Generally, formal mathematical analysis avoids the issue of configuring specific simulation or measurement scenarios, but

the formal modeling may tend to reflect pessimistic timing behaviors [77].

B. Simulation and Emulation

Simulation tools have played a pivotal role in TSN evaluations since new protocols can be designed and validated quickly and with relatively low costs in simulation models. The OMNeT++ [17] and INET [78] frameworks have been used in several TSN studies, e.g., [58], [79]–[83]. OMNeT++ is a modular framework, implemented in C++, and is extensible by combining multiple libraries to create TSN simulators. A freely available open-source library known as INET extends OMNeT++ by adding protocols and mechanisms common to current Ethernet networks. Moreover, INET encompasses wired and wireless connection layer models, e.g., Ethernet and internet stack protocols (such as IPv4, IPv6, TCP, and UDP), facilitating TSN simulations.

Heise et al. [79] present an open-source framework that focuses on a TSN industry profile, including FP and PSFP. Importantly, [79] does not aim to model time-based algorithms or mechanisms. In contrast, local clock models are implemented with the TSN network components in the simulator introduced by Steinbach et al. [80]. The clocks are synchronized using a fail-safe two-step synchronization protocol.

Specht et al. [81] introduced the UBS to deal with asynchronous traffic, whereby UBS fulfills the delay requirement even if time-synchronization fails. Falk et al. [82] implemented a standard-compatible simulation model based on OMNeT++ and INET, which includes VLAN tagging, forwarding and queuing enhancements for time-sensitive streams (Credit-based Shaper), Frame Preemption, and scheduled traffic (TAS). Moreover, Huang et al. [58] introduced an alternative to CBS called time-aware cyclic-queueing (TACQ), and employed OMNeT++ and INET to demonstrate the benefits of TACQ for scheduling mix-flows and reducing the PDV for isochronous traffic compared to CBS.

Apart from OMNeT++, NS-3 [18] is a well-known network simulator into which TSN has recently been integrated [84], [85]. The powerful modeling features for wireless channels make NS-3 a promising candidate for combining 5G and TSN in future simulation models [84], [85]. However, [84], [85] mainly focus on TAS and neglect other TSN features, such as time synchronization.

In summary, none of the existing TSN simulators includes all TSN features. Although the latest INET library supports most of these features, the correctness of combining them has not yet been validated. Moreover, the existing TSN simulators lack a comprehensive framework that accommodates large-scale heterogeneous network architectures [56].

Emulators provide a cost-efficient alternative to dedicated hardware testbeds by mimicking the underlying hardware and software environments to test code in a variety of settings. Thus, emulations are commonly an important intermediate evaluation and developmental phase towards real-world experiments with actual hardware.

One example of software for emulation is *Mininet* [20]. Mininet uses Linux *namespaces* to emulate a set of nodes

on a host system. The connections between the nodes are made via virtual interfaces, so that the behavior is similar to a Virtual Machine (VM)-based virtual environment but with less overhead. Mininet provides an `ssh` connection into each node, thus each node can be handled as a standalone Linux system. Ulbricht et al. [19] implemented TSN in Mininet. The scalability of network emulation is limited by the resources of the host system; therefore, the precision of emulated TSN networks typically decreases as the number of used nodes increases.

C. Measurements in Hardware Testbeds

Simulation, mathematical analysis, and emulation are important stepping stones on the way towards designing real-world testbeds. However, using dedicated hardware is in most cases more precise while having the same complexity as the targeted real network deployment. Whereas simulation models reduce the system complexity, dedicated TSN devices offer the full range of settings and parameters to be considered. Caused by limited resources, scalability is a limiting factor for designing hardware-based TSN testbeds. Employing hardware testbeds comes with its own set of complexities: the measuring techniques, setup, and other elements can strongly influence the outcome. For instance, simulation models usually have built-in tools for collecting evaluation data. However, aside from some exceptions, such as development kits or explicit measurement tools, commercial hardware is typically not designed with built-in evaluation tools.

Measurement evaluations of proposed TSN methods have been conducted on dedicated hardware testbeds for several purposes, as reviewed next.

1) *Packet Processing*: A set of software and hardware should be employed for conducting the intended evaluation in testbeds. For example, when utilizing time-aware gates, the precision of producing scheduled traffic can be crucial. We briefly review options for generating and time-stamping packets. MoonGen [21] is a high-speed packet generator that takes advantage of hardware capabilities to reliably regulate the pace of arbitrary traffic patterns on commodity hardware and to timestamp data packets with sub-microsecond precision. For efficient operation, MoonGen requires the Data Plane Development Kit (DPDK), which comprises libraries to speed up workloads operating on a broad range of CPU architectures.

P4STA [22] operates on programmable hardware, such as smartNICs and FPGA [86]–[88], as well as P4-based technology [89] to achieve measurement precisions of a few nanoseconds. Programmable network hardware is important for flexible and precise measurements [90]. Runge et al. [91] and Beifuß et al. [92] focused on QoS aspects of packet processing in COTS hardware and investigated the effects of the Linux network stacks and related queuing issues.

2) *Evaluating the Precision Time Protocol (PTP)*: Some hardware testbed studies concentrate on improving the hardware and software aspects of time-synchronization, mostly with PTP. Clock synchronization precision in smart-grid systems is investigated in [23]. For these systems to operate reliably, timestamping is necessary, and it is crucial to maintain

performance standards even when communication bandwidth is limited. To achieve an optimal PTP system for a Wide Area Network (WAN), Kassouf et al. investigate the effects of various network factors and clock settings on synchronization precision. They report a 6.38 ns back-to-back master-slave deviation, compared to our 3.4 ns standard deviation and 11 ns maximum deviation, see Figure 6. Ferencz et al. [24] built a testbed using inexpensive COTS components, such as generally available CPUs with x86 architecture and typical NICs. Using IEEE 1588v2 and hardware-assisted COTS devices, they aim to attain sub-1 μ s precision. In a PTP-limited performance evaluation, they observed a synchronization precision of 482.71 ns. Similarly, Kyriakakis et al. [25] employ hardware-assisted clocks for synchronization. To increase the timestamp precision, they integrate a PTP hardware-assist unit into a multicore FPGA platform. Through studies on a testbed consisting of two FPGAs that implement the suggested design and are interconnected through a COTS switch, they examined the worst-case execution time, achieving a worst-case offset of 138 ns.

A completely centralized IEEE 802.1Qcc [93] architecture for configuring the gPTP profile is used by Thi et al. [94] in an Software-defined Networking (SDN)-based automotive and industrial application setting. Thi et al. examined specialized hardware from the manufacturer NXP and limited their evaluations to temporal synchronization, ignoring packet scheduling.

The existing HW testbed studies on evaluating PTP typically evaluate only the synchronization precision. In contrast, this TSN-FlexTest study evaluates the PTP synchronization precision, and broadly evaluates TSN mechanisms.

3) *Communication Over-the-Air*: The performance and design of TSN in complex structures as well as their integration with other communication technologies, such as WiFi and 5G, have attracted the interest of both academia and industry. Sudhakaran et al. [26] designed an industrial testbed that combines traditional TSN in the wired domain with IEEE 802.11. Sudhakaran et al. [26] are primarily concerned with the feasibility and testing of an industrial collaborative robotics use case. Kehl et al. [28] built an over-the-air testbed on an industrial shop floor with a prototype of 5G-TSN integration to investigate a typical industrial use case with cloud-controlled mobile robots.

Agarwal et al. [27] investigated the benefit of employing TSN in microgrid monitoring. They used a commercial Cisco switch to conduct measurements on a laboratory-scale microgrid with four nodes. Agarwal et al. [27] found that TSN can enhance QoS through increased data rates and traffic shaping while retaining connection reliability.

4) *TSN Measurement*: Bohm et al. [29] investigated the controller performance of TSN networks on a three-domain TSN testbed. They found that a negotiation mechanism for provisioning real-time end-to-end connections across various TSN domains can be beneficial. In contrast to the TSN-FlexTest measurement technique, Bohm et al. [29] did not measure the latencies of individual packets and did not consider oversaturated network conditions.

Bosk et al. [30] devised a methodology to assess TSN networks, specifically in the context of Intra-Vehicle Networking

(IVN). Through systematic analysis of IVN traffic patterns for TSN and Best-Effort (BE) traffic, Bosk et al. [30] evaluate different TSN standards in a configuration using COTS hardware and open-source solutions. They try to optimize the system to meet the IVN application requirements. To achieve this, Bosk et al. employ the EnGene [95] framework, which provides a reproducible and scalable TSN experimentation environment. However, the measurement environment in [30], [95] mixes hardware- and software-based time stamping in combination with low-precision `iperf` packet generation which results in relatively long cycle times for the TAS configurations and underscore the need of a well-evaluated TSN measurement methodology.

5) *Time-sensitive Networking Management*: Several testbed measurement studies investigated the impact of dynamic re-configuration of the TSN network and optimizing TSN main functionalities, e.g., shapers, frame preemption, and stream reliability, as reviewed in the following. Jiang et al. [83] provided not only the TSN simulation but also built a hardware testbed. The simulation results verify that the real-world testbed matches the TSN specification, even when they have just μ s resolution. A self-configuring testbed for QoS management has been developed by Garbugli et al. [96] and employed for testing numerous scenarios with varying packet sizes. However, the details of the testbed, such as the specifications of the switches in the testbed, are not disclosed. Groß et al. [97] proposed a specialized hardware extension for standard Ethernet controllers. The flexible and scalable architecture proposed in [97] can run solely in hardware, solely in software, or in both hardware and software. The proposed FPGA-based component can reduce CPU load and software complexity and is particularly beneficial for applications with mixed cyclic and acyclic traffic types.

Coleman et al. [98] addressed the performance impact of the degree of precision of network and CPU clock synchronization on a real-time TSN network. Coleman et al. presented an enhancement utilizing the PCIe bus interface to improve the network-to-CPU synchronization and assessed the improvement using COTS hardware. Vlk et al. [99] developed a method for increasing the schedulability and throughput of time-triggered traffic in IEEE 802.1Qbv. This technique relaxes the scheduling constraints and throughput of time-triggered traffic in a TSN network while maintaining the deterministic nature and timing guarantees. Related routing techniques for IEEE 802.1Qbv have been examined by Nayak et al. [100].

Miranda et al. [101] demonstrate the setup and operation of a cloud-based Linux testbed for TSN experimentation. Using a Central Network Controller (CNC) controller prototype, TSN bridges and nodes can be initialized and managed in the Linux testbed.

6) *Improved Hardware Design for Deterministic Communication*: A few research studies have focused on hardware designs to improve the TSN performance. A very early discussion on TSN testbeds in [102] provides some guidelines for building a measurement testbed and notes some pitfalls while building a testbed. Pruski et al. [103] elaborate hardware constraints in switch input/output queues and recommend a

design architecture for high-capacity switches. The proposed switch simulation can decrease residency time and resource costs by supporting FP, CBS, and TSN, as well as other scheduling methods. A time-triggered Switch-Memory-Switch (SMS) design for memory-efficient TSN switches is presented by Li et al. [104]. SMS can achieve efficiency by facilitating complicated TSN tasks with scalable schedulability and fault-tolerance support. In addition, some measurements using NetFPGA hardware have been reported in [102].

The studies on designing and evaluating TSN switches based on FPGA, e.g., [102], [104], are interesting because FPGA can achieve high performance, programmability, and customization with a lower entry cost. However, FPGA is not cost-efficient for mass production. Also, the FPGA design procedure is time-consuming and requires a relatively high effort for achieving high precision levels. In contrast, the proposed TSN-FlexTest testbed is based on COTS hardware, which can be purchased and set up with relatively low development effort, while achieving a high degree of precision.

Overall, the existing testbed studies focus on specific narrow sets of KPIs. Also, the existing studies commonly consider low precision levels and often lack detailed investigations of the effects that are experienced by a TAS scheduled data stream. With our TSN-FlexTest testbed, we evaluate the measurement methods against the PTP synchronization precision and we extend the measurement method with exact time stamping to a resolution that is in the same range as the PTP synchronization to enable comprehensive detailed measurements of the TSN mechanisms, such as TAS.

IV. TSN-FLEXTEST TESTBED DESIGN

This section provides insights into the hardware and software selected for the testbed. Also, the procedure flow and the used test data are described.

A. Testbed Overview

The purpose of the described TSN testbed is the examination of the behavior of different network stream patterns in one single TSN switch. Figure 3 shows the testbed topology. Five dedicated nodes are connected to one TSN switch, which is the DuT, with 1 Gbit s^{-1} links of negligible length. A node is an end-station in a TSN domain, e.g., talker or listener, and is comprised of COTS hardware, including an x86 CPU and multiple NICs. Node1, node2, and node3 replay and transmit different, high-priority packet streams through the DuT to one sink node (node0) which captures all received traffic. Node4 generates two streams of interfering cross traffic. All five generated packet streams from the transmitting nodes share the same bottleneck by traversing the output port of the DuT to the sink node (node0). On this bottleneck, different scheduling strategies can be applied.

The topology is designed for the unidirectional measurement of the packet latency from the time instant when a transmitting node (node1, node2, or node3) transmits the first bit of a packet onto the physical link to the time instant when the first byte of the packet is received by the receive-NIC on the receiving node (node0). The testbed thus enables the




Figure 3: The TSN testbed architecture is comprised of five nodes in a star topology: one sink (node0) and four source nodes are used to evaluate the Device-under-Test (DuT). The solid line arrows indicate the direction of the (payload) data communication. A control PC manages the measurements and is connected (see dashed line arrows) via a management switch with the nodes. "T" represent the logical timestamping units. Node4 generates cross traffic with MoonGen on two links at once. All link bitrates are 1 Gbit s^{-1} .

measurement of the sojourn (residence) time of a packet inside the switch DuT, which always operates in cut-through mode in our evaluations, with a precision of a few tens of nanoseconds (see Section V-A). The packet sojourn time in the switch DuT accounts for the switch packet processing delay and packet queuing delay, and the packet transmission delay for the cut-through portion of the packet (see subsection V-C); in our setup, the link propagation delays are negligible due to the short physical cables. For the measurement, all devices are synchronized using PTP, whereby one-step timestamping is used for one-way delay measurements. The synchronization traffic is transferred in-band to be closely aligned with real TSN scenarios where the main benefit of TSN is the convergence of link technologies.

B. Analysis of Available and Utilized Hardware

In preparation for the testbed design, we conducted an extensive market analysis of the available TSN hardware components, which is summarized in Table II. We consider TSN Intellectual Property Cores (IP-Cores), System on Chips (SoCs), Switched Endpoints (SEPs), switches, and NICs. Some vendors provide combined setups as TSN kits. We categorize the devices by available TSN features. All devices support 802.1Qbv and PTP time synchronization. The Intel NICs does not have a dedicated hardware TAS, but supports the hardware-accelerated TX-transmission time, which enhances all software schedulers with additional hardware support and improves their precision.

In comparison to industry-grade products, basic COTS (generic commodity) hardware components in combination with Linux facilitate the flexible extraction of measurement data. On the other hand, industry-grade hardware products provide typically state-of-the-art reliable performance. For example, the TrustNode has an industry-grade version and

a research version. The research version, which we purchased, has a highly precise and reliable packet time stamping method which makes additional time synchronization for source and sink unnecessary [19]. Unfortunately, with enabled timestamping of every packet, the device was not stable enough to handle two 1 Gbit s^{-1} streams of MoonGen which makes the TrustNode unsuitable for our test setup.

The integration and debugging of IP-Cores is a very complex and resource-consuming task and requires Very High Speed Integrated Circuit Hardware Description Language (VHDL) programming skills and a VHDL based device. Therefore, the IP-Core solution should be considered as a last resort if simpler methods are not applicable. Also, the availability is a blocking point for hardware selection, Intel has proposed a special TSN CPU with Time Coordinated Computing (TCC) [105] which will probably reduce the effects of CPU load on timing and measurements, as described in section III. Unfortunately, the device was not yet available for order. Thus, the testbed hardware selection is a trade-off of feature completeness, measurement capability, stability, and availability.

As summarized in Table III, we selected the Intel NIC as flexible COTS device, and the stable and available switch from FibroLan. Building on the insights from section III, we designed the testing nodes with enough resources to minimize CPU scheduling effects on the measurements.

C. Analysis of Available and Utilized Software

1) *Precise Time-Synchronization*: A stable timebase is an important requirement for a viable TSN testbed. There are two main approaches for measuring the timing of network packets: The measurement system can be centralized in one location in the network; thus, the transmitted and received packets can be easily correlated in the same device without any requirements to synchronize clocks [107], [108]. The centralized setup is limited in that it does not allow measurements on different devices which are typical for Industry 4.0 (I4.0) or other TSN environments. Conducting all evaluations on one device also creates new issues regarding the shared resources for RX and TX.

The distributed method allows realistic measurement at different network locations. The probe locations require a common timebase which can be provided by clock synchronization. In comparison to the centralized method, the synchronization precision is the most important issue in the distributed configuration. As described in subsection II-B, the PTP protocol can provide a precise and stable synchronization through two clocks via Ethernet.

For validating the reliability of our testbed we conducted several investigations regarding the clock stability. Emmerich et al. [21] investigated the synchronization capabilities of the Intel 82599 NIC. The internal clock of the Intel 82599 is 64 bit wide, whereby the Least Significant Bit (LSB) represents 1 ns. A timer event periodically increments the clock value. Unfortunately, this period depends on the link speed and for 1 Gbit s^{-1} is fixed to 64 ns [109]. Emmerich et al. [21] measured that the reported clock values from the NIC are consistently multiples of 64 ns.

Newer NICs feature improved clock resolution. The I210 NIC, which evolved from the Intel 82599, has an internal 96 bit wide timer which is incremented every 8 ns, independent of the link speed. [51] Due to the larger clock register, the I210 has 32 bit in the sub nanosecond range, allowing the clock to be adjusted in fine-granular steps. Every 8 ns the NIC adds an increment value to the clock register. This increment value can be slightly greater or smaller than 8 ns. The value itself is provided during the clock synchronization process from ptp41 via the *ptp_adjfreq()* kernel interface callback function. By manipulating the increment value, the clock can be forced to run slightly faster or slower. The clock control is thus controlled by the ptp41 clock servo. Due to this implementation, the investigated clock values reported from the I210 NIC are not multiples of the clock granularity, but can differ in multiples of 1 ns.

There are several effects that influence the PTP precision. As described in subsection II-B, ptp41 periodically measures the path delay and adjusts the local hardware clock based on the received master clock values corrected by the path delay. Thus, the jitter of the path delay and the corresponding servo-loop settings are important factors for the PTP precision.

2) *Packet Generation*: There are basically two types of traffic generators: a) (raw) socket based, and b) DPDK based. Both approaches have pros and cons: using sockets and generating packets in user space is highly flexible and allows packet generation and other network usage (e.g., PTP) at once. The disadvantage of sockets are the limited performance for generating enough traffic to fully load an Ethernet link. Emmerich et al. [110] and Wong et al. [111] found that not all packet generators can saturate a link, especially if small packets are used. If precision is not important, the usage of Transmission Control Protocol (TCP) or User Datagram Protocol (UDP) sockets within *iperf* is a cheap solution for generating traffic [30]. Even if *iperf* has not an optimal performance as packet generator [111].

A more powerful approach is provided by DPDK, by making the configuration memory space of the NIC accessible from user space. This enables the user to write Ethernet data into memory and to pass the memory pointer directly to the NIC's TX queue without utilizing the Linux network stack. The problem with DPDK is that it replaces the NIC's driver with a user space pass-through driver and disconnects the NIC from the Linux network stack. Thus, common software, such as ptp41, cannot be used on this interface anymore.

Next to the achievable data rate, the generator's precision is an important KPI. A fixed cyclic stream should transmit the packets periodically exactly according to the predefined cycle time. However, due to CPU load or other software effects, the packets transmission time can have a random or constant offset. DPDK based traffic generators achieve higher periodicity precision than socket-based generators [110]. MoonGen [21] (which was compared with *tcpreplay* in [1]) provides a lua framework for generating high data rate packet load with the help of DPDK. We use MoonGen for generating high-load cross traffic on a dedicated node (node4) that does not require PTP synchronization. For the replay of our test traffic we employ *tcpreplay* [1], which uses raw sockets to replay

TABLE II: Summary of available TSN IP-Core, SoC, SEPs, TSN kits, switches, and NICs. Features marked as (Y) are only available on devices with hardware-assisted TX-injection on the NICs. This table extends [8] and is up-to-date as of 10/2022.

ID	Type	Manufacturer	Product Name	Supported features							
				TAS	ATS	FP	PSFP	PTP	CBS	SRP	CQF
1	IP-Core	Fraunhofer IPMS	IP-Core for TSN	Y	N	Y	Y	Y	Y	Y	Y
2	SoC	Broadcom	BCM53570	Y	N	Y	Y	Y	N	Y	Y
3	SEP	Hilscher	netX 90 Starter-kit	Y	N	Y	N	Y	N	N	N
4	SEP	TTTech Industrial	Edge IP Solution EVM	Y	N	Y	Y	Y	Y	N	N
4	TSN Kit	Analog Device	RAPID-TSNEK-v0001	Y	N	Y	Y	Y	N	Y	N
6	TSN Kit	Kontron	KBox C-102-2 TSN starter kit	Y	N	Y	N	Y	N	Y	N
7	TSN Kit	XPX	LS1021A-TSN-R	Y	N	N	Y	Y	Y	N	N
8	Switch	Advantech	EKI-8510G-2FI	Y	N	Y	Y	Y	N	N	N
9	Switch	Cisco	IE-4000	Y	N	N	N	Y	N	N	N
10	Switch	Comcores	Manticore	Y	N	Y	Y	Y	Y	N	N
11	Switch	FibroLAN	Falcon-RX/G	Y	N	Y	Y	Y	N	N	N
12	Switch	Hirschmann	RSPE	Y	N	N	N	Y	N	N	N
13	Switch	Hirschmann	BOBCAT	Y	N	N	N	Y	N	N	N
14	Switch	InnoRoute	TrustNode industrial	Y	N	N	Y	Y	N	N	N
15	Switch	Kontron	KSwitch D10 MMT Series	Y	N	Y	N	Y	Y	N	N
16	Switch	Marvell	88Q6113	Y	N	N	N	Y	Y	N	N
17	Switch	SoC Engineering	MTSN Kit 1G Multiport	Y	N	Y	Y	Y	Y	Y	N
18	NIC	ASIX Electronics	AXM57104	Y	N	N	Y	Y	Y	Y	N
19	NIC	Intel	I210	(Y)	(Y) [106]	N	N	Y	(Y)	Y	N
20	NIC	Intel	I225	(Y)	(Y) [106]	N	N	Y	(Y)	Y	N
21	NIC	Kontron	PCIe-0400	Y	(Y) [106]	Y	N	Y	(Y)	Y	N
22	NIC/SEP	InnoRoute	Raspberry RealTime HAT	Y	(Y) [106]	N	N	Y	(Y)	Y	N

TABLE III: Utilized Hardware in the Testbed.

Component	Description
CPU	Intel Core i7 6700
Mainboard	Gigabyte H270N-WIFI
RAM	2×16 GB
NIC Source (Prio)	Intel I210
NIC Source (cross traffic)	Intel X520
NIC Sink (onboard)	Intel I211
Operating System	Ubuntu 20.04.3 (GNU/Linux 5.11.4-rt11)
DuT	FibroLAN Falcon-RX/G (SW:8.0.2.4)

arbitrary traffic from prerecorded files.

3) *Scheduler – GCL Configuration*: Since version 5.4, the Linux kernel provides several features for hardware timestamping and hardware-assisted packet transmission. One feature is an (additional) transmission timestamp. In contrast to the TX timestamp, this is the time when the frame should be transmitted, i.e., a value in the future. In theory, the packet with an additional transmission timestamp is handed over from the network driver to the NIC hardware, which stores the packet in the transmission queue and will release the packet when the transmission time is reached.

Dealing with limited resources in the hardware layer of the NIC this procedure needs to be supported by the higher layers of the network protocol stack. The NIC hardware cannot store a large amount of data for transmission, the internal memory is limited to a few packets. The NIC queue is a FIFO buffer; therefore, the packets cannot be reordered. To respect the limited hardware resources, the Linux network stack has to be aware of these characteristics. The qdisc environment provides an interface in the Linux network protocol stack to insert and manage queuing and packet scheduling. One loadable module is the Earliest Time First (ETF) scheduler which reorders incoming packets according to their transmission timestamp. To respect the limited hardware queue size, the packets need

to be temporarily stored in a software queue and handed to the hardware just shortly before the transmission time. This time offset needs to be determined experimentally for each NIC and system configuration. To set the transmission time of each packet, one possible module is *ta-prio* which implements a software TAS for scheduling packets.

D. Testbed General Procedure Flow

The TSN testbed supports the automatic test of the packet transmission capabilities of a dedicated DuT. Several nodes are controlled to replay measurement traffic and send it through the DuT. Each transmission node inserts a precise timestamp into the payload section of TX packets, following Section II-B. Additionally, nodes can be configured to replay a cross traffic to simulate a busy link. One node (node0 in Fig. 3) captures the measurement packets after they have passed the DuT. By comparing the TX and RX timestamps, the packets transit time through the DuT can be evaluated. A detailed description of the testbed setup is provided in [1].

1) *Packet Header/Our Streams*: To determine the packet handling behaviors of the DuT, the selection of characteristic test data is important. In the current scenario, we use three types of test data: a) generic generated cyclic streams, b) generated data streams which reflect typical I4.0 characteristics, and c) real captured data of a robot control environment, as summarized in Table IV.

The generic generated streams (stream set Θ) have maximum sized frames, but distinct packet generation cycle times to represent three cyclic TSN streams. The cycle times are prime numbers to enforce statistical collisions of all streams, even without cross traffic.

The stream set Ψ corresponds to a typical I4.0 scenario where a machine delivers tactile data in combination with video and audio streams for human control and mechanical

TABLE IV: Application layer (payload) traffic stream cycle duration and audio or video frame size mean (and standard deviation), as well as average bitrate (at link layer, incl. headers), of high-, medium-, and low-priority TSN streams of stream sets Θ , Ψ , and Ω [with corresponding Priority Code Points (PCPs) 6, 5, and 4]; generic streams and best-effort cross traffic send maximum-sized Ethernet frame payload.

Stream Set	Traffic Type	PCP	Periodicity	t_{Cycle} [ms]	Appl. Data Size [B]	Bitr. [b/s]
Θ_{high}	Generic 1	6	cyclic	0.2	1472 (0.0)	61M
Θ_{med}	Generic 2	5	cyclic	0.3	1472 (0.0)	41M
Θ_{low}	Generic 3	4	cyclic	0.5	1472 (0.0)	24M
Ψ_{high}	Tactile	6	cyclic	1	82 (0.0)	1.1M
Ψ_{med}	Audio VBR	5	cyclic	24	480 (0.0)	177k
Ψ_{low}	Video VBR	4	cyclic	16.67	8336.4 (24283.8)	4.1M
Ω_{high}	Spot CTRL	6	acyclic	-	501.6 (984.4)	175k
Ω_{med}	Audio CBR	5	cyclic	20	153.9 (46.6)	82k
Ω_{low}	Video CBR	4	cyclic	25	62412.5 (60355.9)	20M
	Cross Traffic	0	acyclic	-	1472 (0.0)	2G

error monitoring [112]. The audio and video data are lossy compressed therefore the data is cyclic and bursty with different burst sizes.

For stream set Ω , we captured the data streams of the well-known robot *spot* [113]. The streams contain the control data stream of the spot robot in combination with the spot camera stream. For audio, we included a lossy compressed voice audio file.

To be replayed in the TSN testbed, the packet header needs to be modified. Figure 4 shows the packet structure of stream type Θ , Ψ , and Ω . The testbed uses several fields in the packet header for matching. The source Medium Access Control Layer (MAC) address is used for matching the source node name of a received packet. The DuT itself is configured to apply different queuing and scheduling strategies depending on the packet’s virtual LAN (VLAN) PCP ID. According to subsection II-B5, each transmitted packet is timestamped automatically by the TX network hardware, which adds a 64 bit timestamp into the packet payload.

2) *Measurement Procedure*: The flowchart in Figure 5 illustrates the test-bed measurement procedure. The testbed software, including data stream generators, are openly available from <https://github.com/5GCampus/tsn-testbed>. The measurement procedure is structured into reconfiguration, traffic replay, and record results for plotting. In the reconfiguration phase, the control PC resets all network connections and sets the nodes into a dedicated state. Depending on the test configuration, the software TAS is configured at the transmitting nodes. To finish the preparation, the control Personal Computer (PC) transfers the TX pcaps to the nodes.

In the measurement phase, the control PC starts the capturing process at the receiving node and initiates the pcap replay at the transmitting nodes. If the replay or capture has finished, the received data is recorded and transferred to the control PC for further processing. We spotted that some DuT drop the PTP connection sporadically. Because the PTP synchronization is the backbone of the testbed we check the log for PTP errors or to high clock derivations during the measurement and repeat the measurement if an error was detected.

Listing 1: Highlighted part of PTP configuration file with relation to the gPTP profile.

```
[PTP config file]
#Default Data Set#
gmCapable          1
twoStepFlag        1
slaveOnly           1
domainNumber       0
[...]
#Port Data Set#
logAnnounceInterval 0
logSyncInterval    -6
operLogSyncInterval 0
announceReceiptTimeout 3
syncReceiptTimeout 3 #from gPTP
neighborPropDelayThresh 800 #from gPTP
min_neighbor_prop_delay -20000000 #from gPTP
BMCA ptp
[...]
#Run Time Options#
logging_level      7
verbose           0
[...]
#Servo Options#
pi_proportional_const 0.0
pi_integral_const    0.0
pi_proportional_scale 0.0
pi_proportional_exponent -0.3
pi_proportional_norm_max 0.7
pi_integral_scale    0.0
[...]
max_frequency      900000000
[...]
#Transport Options#
[...]
#Default Interface Options#
clock_type         OC
network_transport  L2 #from gPTP
delay_mechanism    P2P #from gPTP
time_stamping      hardware
[...]
```

V. TSN-FLEXTEST TESTBED VALIDATION

Generally, a measurement is only as good as the metering setup. This section evaluates the measurement precision of our TSN-FlexTest testbed.

A. Evaluating the Precision Time Protocol

As noted in subsection IV-C1, we designed the TSN-FlexTest testbed as a distributed architecture. In the distributed setting, clock synchronization via PTP is a key factor for precise delay measurement.

For validating the testbed precision we conducted measurements for several servo settings and examined the influence of the cross traffic on the `ptp41` precision. To extract the clock synchronization metrics, we patched `ptp41` to provide the actual clock deviation values and not only an root mean square (RMS) value. Previous measurements demonstrated that the synchronization metrics of `ptp41` are as precise as physical measurements with an oscilloscope [114]. Figure 6 shows the results of the clock precision measurement. For simplicity




Figure 4: Frame structure of high-priority test data streams. The data is sent using VLAN-tagged Ethernet frames. The priority is set by modifying the PCP according to Table IV. Inside the IP header, the ID field is leveraged to provide a sequence number. The payload holds the Timestamp (TS), which is split into integer and decimal part for seconds and nanoseconds, respectively.

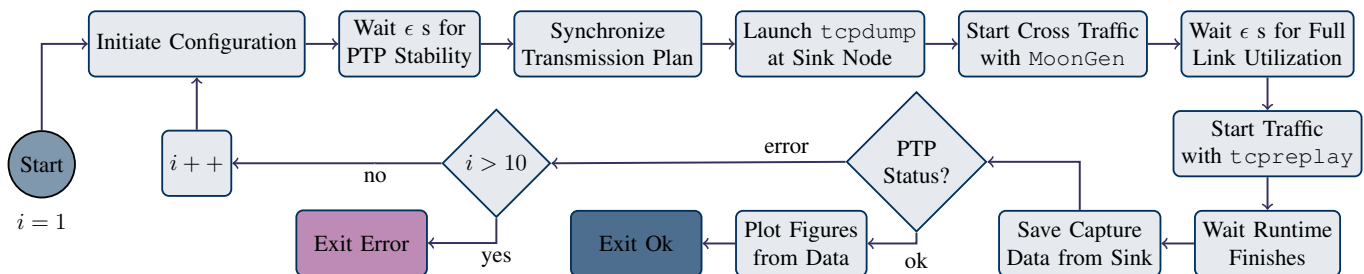


Figure 5: Flowchart of the automated measurement process: The process flow should ensure that all measurements run successfully even if there are a few errors, e.g., due to timeouts during the PTP-synchronization. Therefore, each test is performed at most ten times until it stops with an error. The process includes steps for initialization, synchronization, the measurement itself, as well as storing and analyzing the data.




Figure 6: Boxplot of slave clock deviation to the GM in nanoseconds as a function of number of synch messages per second over a measurement period of 30 min.

and to avoid clutter in this evaluation, we manipulate only one `ptp4l` servo parameter, namely the synchronization frequency with the number of sync messages per second in the range from [8, ..., 128].

We measured the clock deviation between one slave node and the DuT switch for several configurations with and without cross traffic in order to investigate the independence of the clock deviation from the cross traffic. Due to the activated hardware time stamping, the transmission of additional (cross

traffic) packets increases the queuing delay of the PTP packets. However, this additional queuing delay is compensated for by the timestamps (which are applied after the TX queue and before the RX queue). Thus, cross traffic does *not* affect the testbed precision. Generally, the transmission channel (link) between a node and switch has a jitter value, which needs to be compensated by the servo. The results in Figure 6 indicate that the synchronization frequencies of 8 to 128 sync messages per second result in a reasonably small maximum clock deviation of less than ± 14 ns, indicating a precision of approximately 30 ns. As the synchronization frequency to 16 sync messages/s does not exhibit any clock deviation outliers, we use 16 synch messages/s for the remainder of this article.

B. Sending Cyclic Data Precisely

As described in subsection IV-C2 different packet generation approaches have different advantages and precision levels. In Table V we extend the measurements from [1]. The results in Table V clearly indicate that all transmission techniques that use the Linux network stack (socket based and `TA_PRIORITY`) are less precise than sending packets via DPDK based generators e.g., MoonGen. The substantial difference between the two `tcpreplay` versions is caused by a bug in the lower version, which results in a constantly too short cycle, due to the implementation this error increases with the runtime.

C. Cut-through Switching: Effects of Ethernet Frame Size

Figure 7 shows the DuT one-way packet delay for cut-through switching. We generated a random homogeneous dataset with packets of sizes from 64 B to 1522 B and sent them through the DuT. With the testbed, we measured the

TABLE V: Measurement results of achievable cycle time precision: deviations from expected cycle time (of 1 ms) in nanoseconds; with 16 synch messages/second. Type characterizes the software that generates (or subsequently shapes) the packets on sending nodes 1, 2, and 3.

Type	Tool	Version	Min	Mean	Median	Stdev	P99%	Max
Linux NAPI	tcpreplay	4.3.2	-997176	149881	160763	29630	166339	447074
Linux NAPI	tcpreplay	4.3.4	-998312	1	48	11001	15960	2254136
DPDK	MoonGen	25c61ee	-6848	19	0	138.7	392	6895
Sender TAS (15 μ s)	TA_PRIORIO	Kernel 5.15	-997904	18	-126	19205	24822	2958199
Sender TAS (500 μ s)	TA_PRIORIO	Kernel 5.15	-998784	336	2	1023658	2467535	6460670




Figure 7: Packet latency of cut-through switch DuT as a function of packets size. The switch forwards packets smaller than 337 B when they have been completely received; for larger packets, only 337 B are accumulated before forwarding.

packet residence time in the DuT and plot the measured delay as a function of the packet size. Packets of size 64 B to 337 B experience a linearly increasing delay. This is because the FibroLAN switch completely receives and stores these small packets before forwarding, incurring the equivalent of the transmission delay (which the packet experienced as the sending node transmitted the packet bits into the physical wire) as the packet bits are accumulated at the switch. For larger packets, the considered FibroLAN switch accumulates 337 B of data before starting to forward the packet; thus, incurring a constant equivalent transmission delay component for accumulating 337 B of data. The cut-through experiment should be conducted for every new DuT which is considered in the testbed, because the forwarding characteristics for different packet sizes depend on the DuT’s internal architecture and will affect all delay measurements [115].

D. Comparisons with Time-Sensitive Network Simulations

The validation of a physical testbed with simulated data seems to be unconventionally, because hardware measurements are true by definition; however, this is only valid if the hardware is configured in the right way. Simulations have often less complex settings and the setup is less complicated than configuring hardware measurements and retrieving hardware measurement results. So, achieving similar results with hardware and software is a good indicator for data reliability.

To validate our testbed, we implement the same network architecture as in Figure 3 based on OMNet++ 6.0 with the INET framework version 4.4. Four different GCL configurations are measured to compare with the behavior of our real testbed.

1) *Simulation Setup*: The GCL configurations are shown in Table VI, where the 15 μ s slot size is used, including the Guard Band. The first configuration is a simple round-robin. The second configuration always opens the gate to the highest priority traffic, and other traffic holds the transmission time equally. In the third configuration, the GCL gate duration time corresponds to the priority, whereby the higher the priority, the longer the gate remains open. The fourth configuration modifies the second configuration by compensating for high-priority traffic and avoiding the opening of more than two queues simultaneously.

The simulation settings are aligned with the generic generated streams (stream set Θ) to verify our testbed. Since our main focus is on the real testbed measurements, only the generic stream set is simulated. To simulate the cross traffic, packets with uniformly distributed sizes [64 B to 1522 B] are generated according to a Poisson process. The simulated PCP value of stream Θ_{low} is different from the value of the real testbed (4) since, in the INET framework version 4.4, PCP value with five and four are mapped to the same queue. Therefore, we set the simulated PCP value three to fulfill the one-to-one mapping between streams and queues so that the TAS can control each stream. The cut-through switching is not enabled in the simulation due to a bug in the current INET version.

We summarize the clock and synchronization settings as follows. A GM clock is configured to align with the simulation time, and other clocks have a constant drifting rate uniformly selected from -100 Parts Per Million (ppm) to 100 ppm in the initial phase of the simulation. A clock with one ppm constant drift means it would be one second faster if one million seconds passed by. To synchronize the devices with the GM, the INET framework version 4.4 provides the gPTP tools according to IEEE 802.1AS. The switch is configured to transmit the synchronization and peer-delay measurement messages every 0.5 ms.

The Cumulative Distribution Function (CDF) characterizes the distribution of discrete random variables, allowing one to derive the probability of a random variable smaller than a specific value. However, we are mainly interested in how frequently the latency exceeds a prescribed level; therefore, the simulation results are presented as Complementary CDF




Figure 8: Simulation results for generic stream set Θ . The three colors correspond to three streams, where *blue* is stream Θ_{high} , *purple* is Θ_{med} , and *magenta* is Θ_{low} , respectively. Moreover, the configuration of GCL 1 to 4 corresponds to the GCLs in Table VI from left to right. In addition, w/o means without, and w/ represents with. These conventions apply analogously to the subsequent figures.

TABLE VI: The four figures show the four considered Gate Control List (GCL) configurations (1, 2, 3, and 4, from left to right) with optional guard bands, i.e., the measurements are conducted with and without guard bands. Each GCL configuration consists of time slots (**Time**-axis) and different streams with associated priorities (**Priority**-axis). Each time slot is equi-sized since all frames in the stream set Θ have the same frame size of 1522 B. However, the slot size itself can vary in length; i.e., in measurement A, a slot is $15\mu s$ long; in measurement B, it is $3 \times 15\mu s = 45\mu s$. This allows to detect impacts of too short or too long slot sizes and to draw conclusions for possible timing issues. The time where each stream is permitted to send is indicated by a dark colored box. Configurations 2 and 3 do not need a dedicated time slot for a guard band, since the to-be-protected *high* priority stream is always allowed to send data.

Time \ Priority	t_1	t_2	t_3	t_4	t_5
high	X				G
medium		X			G
low			X		G
cross traffic				X	G

Time \ Priority	t_1	t_2	t_3	t_4
high	X	X	X	X
medium		X		
low			X	
cross traffic				X

Time \ Priority	t_1	t_2	t_3	t_4
high	X	X	X	X
medium		X	X	X
low			X	X
cross traffic				X

Time \ Priority	t_1	t_2	t_3	t_4	t_5	t_6	t_7	t_8
high	X		X		X		X	G
medium		X		X		X		G
low			X		X			G
cross traffic				X				G

(CDF) curves.

Each generic stream measurement scenario was run for 30 minutes (covering $3.6 \cdot 10^7$ cycles of duration $50\mu s$). For selected simulated data and spot robot scenarios, we conducted pilot measurements over one minute (covering $1.2 \cdot 10^6$ cycles of duration $50\mu s$) and over ten minutes (covering $1.2 \cdot 10^7$ cycles of duration $50\mu s$). After confirming that the one minute and 10 minute measurement runs gave equivalent measurement results, we ran each simulated data and spot robot measurement scenario for one minute. The corresponding validating simulations were run for one minute ($1.2 \cdot 10^6$ cycles of duration $50\mu s$) and 5 seconds ($1 \cdot 10^5$ cycles of duration $50\mu s$),

which gave equivalent result; then the simulations were run for 5 seconds.

2) *Results*: In Figure 8, the y-axis represents the complementary cumulative probability, and the x-axis represents the one-way packet latency of the streams in the DuT. An ideal switch without latency variations corresponds to a vertical line, see Figure 9a. If the switch applies an increasing constant delay to all packets, then the vertical line will move to the right. However, in most cases, the delay of each packet changes, and small delays have a higher probability (i.e., are more likely to occur). These packet latency variations bend the vertical line into a curved shape, see e.g., Figure 8a. The curve

may not be smooth when packet delays are not continuously distributed. For example, a step shape arises in Figure 8d, because packet bursts are transmitted within consecutive time slots. Generally, the bottom part of the CCDF curve signifies the maximum packet latency, whereas the top denotes the minimum packet latency.

In Figure 8a, the packet delay curves completely overlap because each stream has the same transmission time slot. However, in Figure 12a, a tail of Θ_{high} appears. The cut-through switching may cause this tail phenomenon in our testbed since Θ_{high} streams should wait until the lower priority stream finishes transmission even when the gate is open. This tail issue can be addressed by increasing the time slot size. Therefore, the simulation result of Figure 8a is similar to the testbed measurement result presented in Figure 12e. Moreover, the difference between Figure 8a and Figure 8b is the guard band implementation, where Figure 8b shifts slightly to the right compared to Figure 8a. Once the guard band is introduced, all the streams must wait a certain time to obtain an open gate again, which increases the latency. The same behavior occurred in the testbed measurements depicted in Figure 12e and Figure 12j.

We turn now to Figure 8c and Figure 8d, which were obtained for the same configuration as the testbed measurement results in Figure 12b and Figure 12c. The testbed results generally follow the simulation results, whereby Θ_{med} and Θ_{low} overlap for GCL 2, but separate for the GCL 3 configuration. The behavior of the step shape in both Figure 12b and Figure 12c will be explained in the following section and can be smoothed by enlarging the slot size. Therefore, the simulation results of Figure 8c and Figure 8d are more similar to Figure 12f and Figure 12g.

Next, we consider Figure 8e and Figure 8f. Compared to the testbed results, the ranges between streams are close to each other. This is because the cross traffic frame size varies in the simulation, whereas the full Ethernet frame size is utilized in the testbed all the time. A large frame size leads to a blocking effect for other streams and further increases their latency. However, the sequence or the priority of streams still holds, which means that the order of the latency curves from left to right follows the order of Θ_{high} , Θ_{med} , and Θ_{low} . The order is presented at the top of Figure 8e and is readily observed in Figure 12d, which also validates the correct operation of our testbed.

Applying a guard band has a significant influence to GCL 4, as observed by comparing Figure 8e and Figure 8f. This is because the guard band protects the highest priority stream so that stream set Θ can transmit at the beginning of GCL 4 without interference from other streams. The same result can be observed from the testbed measurements in the comparison of Figure 12d and Figure 12k, where the tail effect is mitigated at the bottom side of Figure 12k.

We conclude this section by noting that our testbed works properly compared to the simulation results. However, we only emphasized the similarity and differences between the simulation and the testbed in this section. The comparison of the four GCL configurations and their shortage, how to improve the latency, and the impact of different types of input

streams would be covered in the following sections.

VI. TSN EVALUATION WITH TSN-FLEXTEST TESTBED

In this section, the TSN-FlexTest testbed is used to evaluate Time-Sensitive Networking (TSN). We examine how the Packet Delay Variation (PDV) can be reduced by leveraging multiple TSN mechanisms. We investigate the stream characteristics with respect to frame size, packet frequency, burstiness and the impact of different scheduler settings, e.g., TAS slot size and guard band. Also, we compare TSN performance for the generic data streams with recorded real-world data sets.

In order to provide QoS *guarantees* for network streams and to reduce the PDV with Time-Sensitive Networking, there are multiple possible options: First, the architecture can be modified, e.g., a dedicated connection can be established. With dedicated connections, the delay increases caused by other network participants (streams) can almost completely be avoided, which provides a *baseline* benchmark for the following improvements. Second, if additional dedicated connections are either not applicable, e.g., due to physical circumstances, or economically not viable, then existing technologies for data prioritization can be used. For instance, the Ethernet VLAN tag provides the PCP field with 3 bit to distinguish 8 network packet priorities. In conjunction with SPQ, it is possible to prioritize packets (see subsection II-C). We subsequently refer to the PCP approach as *soft* QoS. Finally, there is a third choice for enhanced network control: With a TAS, it is possible to achieve deterministic transmission behavior, which we refer to as *hard* QoS in the following.

We have conducted extensive measurements in the TSN-FlexTest testbed to elucidate the advantages and disadvantages of the different technologies. Furthermore, we examine for the first time the Time-aware Shaper with synchronized senders using COTS hardware. While TAS with synchronized senders has been employed with industry-grade equipment, to the best of our knowledge, we are the first to conduct detailed measurements for TAS with synchronized senders in a testbed built from basic COTS hardware. Our measurements demonstrate that TAS with synchronized senders improves the performance of conventional "switch-only" TAS to provide near-optimal performance on generic commodity (basic COTS) hardware. Our measurements indicate that TAS with synchronized senders on generic commodity hardware is comparable to the baseline while still allowing for the sharing of the network links by multiple streams.

The following measurements primarily use stream set Θ , allowing better insights into the transmission and forwarding characteristics, since effects within the stream set itself are negligible. In the discussion, we only involve stream set Ψ and stream set Ω for specific peculiarities and refer the interested reader for the corresponding figures to the Appendix.

A. Baseline

The baseline measurement considers only *one* sender (actually, three senders, but only one is sending in a given measurement run) and one receiver with the DuT connecting both. The baseline scenario replays all streams, defined in Table IV, one

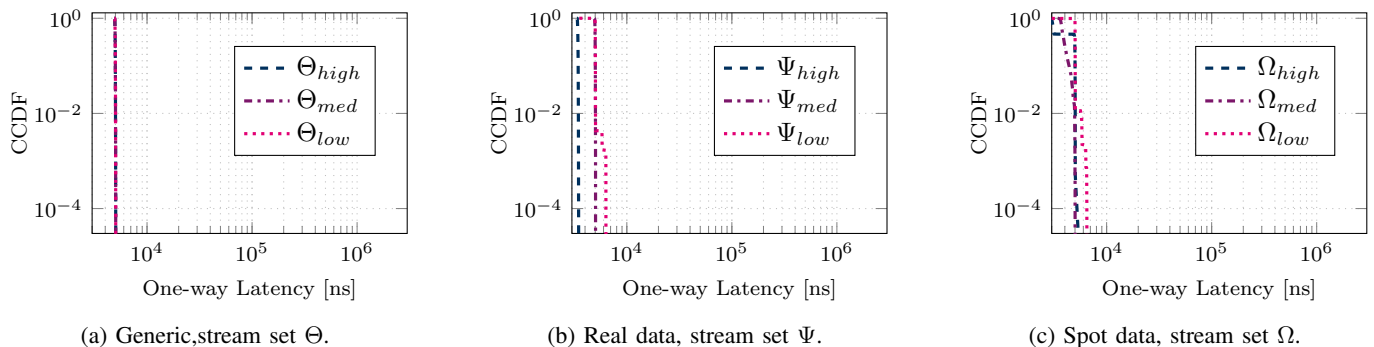


Figure 9: Baseline measurement of three stream sets for testbed validation. The figures show from left to right the independent replay of stream set Θ , stream set Ψ , and stream set Ω , each in a combined plot. Specifically, for each of the three stream sets, each of the three streams in the set was replayed individually by a single transmission node (without any transmissions from the other nodes). The figures thus show the absolute minimum delay for traversing the DuT, whereby this minimum DuT traversal delay is mainly governed by the queuing delay, while switch processing delay and transmission delay are small and nearly constant for the constant frame size. Collisions (resource contention) can arise due to the in-band PTP synchronization, whereby PTP frames are sent with the highest priority (7, which is higher than the highest test stream PCP of 6), and thus the PTP frames can delay the data frame transmissions.

by one, without any cross traffic. The packet latencies for the three streams within each stream set are combined together in one plot, although in fact three separate measurements have been performed. According to the stream priorities, it is expected that the colored lines in the CCDF always follow the order: *blue* < *purple* < *magenta*. An additional objective is that the *blue* curve should show a minimal deterministic latency, since the GCL configurations are optimized towards the highest priority. This objective is valid for all following CCDFs.

Figure 9 shows the measured packet latency, i.e., packet residence time inside the DuT. An identical residence time for all packets in stream set Θ can be observed in Figure 9a, indicated by a vertical line. With a standard deviation of 43.5 ns, 41.5 ns, and 38.9 ns for stream Θ_{high} , Θ_{med} , and Θ_{low} respectively, the one-way delay distribution is very narrow (the values are in the range of the PHY tolerances), an almost deterministic behavior. It is also expected that all three streams exhibit the same distribution since only the packet rate varies. In contrast, Figure 9b and Figure 9c exhibit varying gradients of the packet latencies. This can be explained by the frame size: Whereas the streams in stream set Θ have a fixed frame size of 1522 B, and the streams Ψ_{high} and Ψ_{med} have a constant frame sizes of 128 B and 526 B, respectively (see Table IV), all other streams have variable frame sizes. Therefore, the CCDFs of stream set Ψ and stream set Ω exhibit the actual distribution of the frame sizes and the corresponding transmission delays. For instance, a few frames of stream Ψ_{low} are smaller than 337 B (see subsection V-C) and therefore are forwarded faster. Additionally, the video stream has packet bursts, i.e., a lot of frames are transmitted in a short period of time, which causes a filled queue. Though the DuT could handle the stream, the in-band time-synchronization causes a slightly longer delay for a fraction of the frames. A shift between stream Ψ_{med} and Ψ_{low} by 1286 ns can be seen above the 99 percentile – this equals to an Ethernet frame of about

140 B. The same applies to the other streams with a variable frame size.

B. No Prioritization (Internet Scenario)

After establishing a baseline, we investigated the one-way delay in a typical *internet* scenario, i.e., the data streams are transmitted simultaneously with the *same* priority. In this case, the streams compete equally for resources in the DuT. Figure 10 shows two measurements: Figure 10a illustrates the case where only the three streams of stream set Θ are transmitted simultaneously, and Figure 10b shows the more typical case with added interference caused by cross traffic. In all following measurements, the cross traffic consists of full Ethernet frames, with a throughput of 2 Gbit s^{-1} (200% loaded link), and a variable inter-packet delay based on a Poisson distribution. In this internet scenario, the packets of the three simultaneously transmitted streams will have collisions and are affected by queuing latency at the egress port towards node0 at the DuT. This scenario reveals the scheduling behavior of the DuT if packets from several streams need to be transferred through the same output port (our synthetically created bottleneck). Thereby, we can observe the queuing mechanisms of the DuT in an oversaturation scenario that forces the switch to drop packets. The DuT is configured to ignore the VLAN PCP. We observe from Figure 10 that the one-way delay is not aligned with the priority set in the packet headers. In Figure 10a, we can identify that for 97.96% of all sent packets, the one-way latency is identical to the baseline measurement (see Figure 9a). Above, we can notice that there is a latency increase, independent of the stream priority. We observe a clear step caused by statistical collisions of frames, i.e., arriving frames compete for resources. These *collisions* (*resource contentions*) need to be considered for applications with hard real-time requirements. Statistical collision of frames results in steps in the CCDF, because sometimes a frames is forwarded immediately, whereas sometimes it has to wait.




Figure 10: Measurement of stream set Θ with disabled prioritization settings at the DuT (all streams are treated equally). The streams are sent at the same time, so that packet collisions at the end of the DuT’s queue cause additional delay which is increased in case of added cross traffic. Because of the higher frequency of Generic 1, this stream is affected more often by collisions than others and has an higher delay than Generic 3.

This *blocking* issue can be solved by frame preemption, see subsection II-C2. Moreover, we observe this staircase formation in all following figures where collisions cannot be prevented. Additionally, cross traffic affects all streams equally and increases the packet delay in general (see Figure 10b). Now, only a very small part of the sent packets are immediately forwarded by the DuT and almost all other packets are delayed by up to 135.4 μ s, 109.6 μ s, and 85.3 μ s for stream Θ_{high} , Θ_{med} , and Θ_{low} respectively.

The frame rate, i.e., the frequency of frame transmissions [frames/second], of a given stream influences the statistical incidence of collisions. If no special scheduling is applied, then frames of streams with a high frame rate have to wait more often. This is the reason why the stream latencies in Figure 10b are ordered in the opposite direction as their priorities.

This is the typical configuration in a *net-neutral* internet, without specialized packet treatment and with limited resources in the backbone networks. This is especially true due to *rush hours* in the internet, e.g., in the evening or at large events (black Friday, popular broadcasting events). For Tactile Internet (TI) scenarios [116]–[118], it is nearly impossible to ensure a prescribed service quality in such a net-neutral internet.

C. Soft Quality-of-Service with Strict Priority Queuing

In a first step towards providing QoS, we examine the improvement of the service quality achieved by applying queuing disciplines that take the stream priorities into consideration. One common approach is Strict Priority Queuing (SPQ), defined by IEEE 802.1Q, where always the queue that holds the highest priority packets is emptied first. Figure 11 shows this SPQ approach with stream set Θ for measurements without and with added cross traffic. We observe from Figure 11a, that the distribution of the one-way latency has the lowest latency values for the highest stream priority, the second lowest latency for the second highest stream priority, and the




Figure 11: Measurement of stream set Θ with Strict Priority Queuing (SPQ) at DuT. Similarly to Figure 10, packet collisions cause additional queuing delay. With SPQ, the DuT scheduler preserves the stream priorities defined in Table IV, i.e., the queue with the highest priority is emptied first, then the second highest priority queue, and so on. Added cross traffic increases the queuing delay, but generally lowers the overall one-way latency for high-priority streams compared to Figure 10b.

highest latency values for the lowest stream priority. Thus, we conclude that SPQ achieves a certain priority isolation between the three streams. However, Figure 11b shows that the overall one-way latency still increases from 12.0 μ s, 12.4 μ s, and 14.2 μ s to 16.7 μ s, 20.8 μ s, and 17.8 μ s for stream Θ_{high} , Θ_{med} , and Θ_{low} respectively, when cross traffic is added. This is an increase by 44% on average for all three streams. Therefore, we characterize the SPQ approach as *soft* QoS, as it is not possible to enforce strict guarantees. This can be explained by the fact that lower priority frames can be scheduled and selected for transmission when a higher priority frame arrives at the DuT. In this case, the lower priority frame is not interrupted and the higher priority frame must wait until the lower priority frame is transmitted, which for full regular Ethernet frames with 1522 B takes around 12.3 μ s in the worst case. This effect can be observed in Figure 11a: without cross traffic, the highest priority stream Θ_{high} is delayed from 4894 ns (minimum value) to at most 18 540 ns, a growth by about one full Ethernet frame transmission.

D. Hard Quality-of-Service with Time-aware Traffic Shaping

1) *Overview*: In a second step towards providing QoS, we measured the one-way packet latency for TAS. Figure 12 shows measurements for the GCL configurations in Table V-D1 and Table VI. We strive to gain insights for the different TAS parameters settings. Therefore, we systematically created four different GCL configurations and, where appropriate, we added a guard band to protect the highest priority stream. We considered stream set Θ , and we added cross traffic in *all* twelve measurements. Furthermore, we varied the slot size (the time a gate of the respective queue is open): a value of 1 corresponds to our unit of 15 μ s which in theory should allow transmission of one full Ethernet frame. Since we estimated that frames are not sent out at the sender with the highest precision, we wanted to investigate further

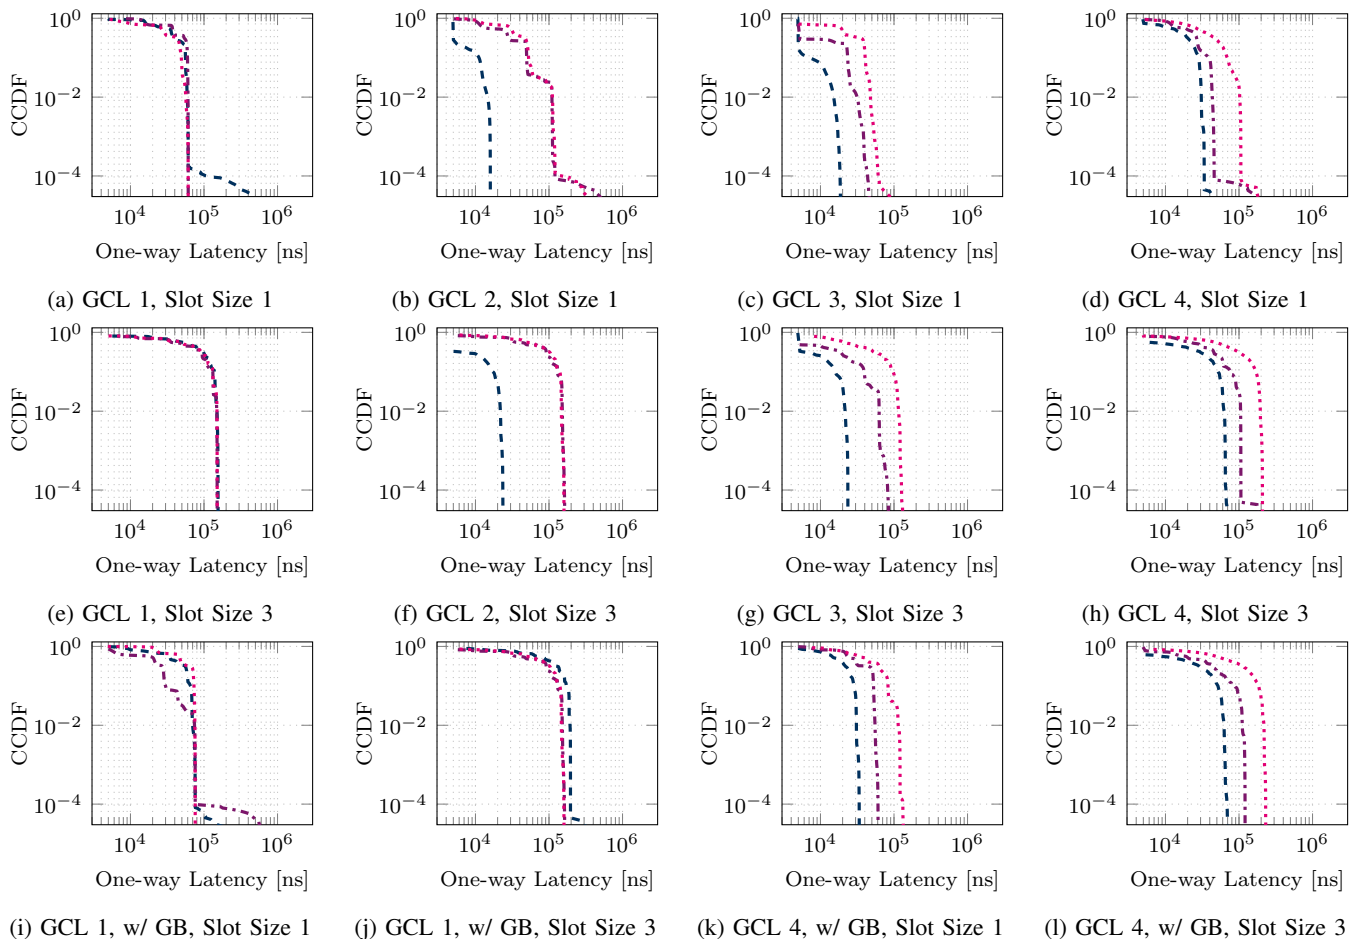


Figure 12: Measurement of stream set Θ with enabled Time-aware Shaper (TAS) at the DuT. We tested with multiple generic Gate Control List (GCL) configurations (see Table VI) and slot sizes in order to examine the TAS effects. The slot size refers to the time unit a gate of the corresponding queue is opened—a slot size of 1 equals $15\mu\text{s}$ (see row 1), and 3 corresponds to $45\mu\text{s}$ (see row 2). We conducted measurements without (w/o) and with (w/) enabled Guard Band (GB) for protecting the highest priority (see the last row).

the effect of increased slot size. Therefore, we also considered $45\mu\text{s}$ which corresponds to a slot size of 3.

2) *GCL Configuration 1*: We observe in Figures 12a, 12b, 12c, and 12d the fundamental effects of the GCL settings: Figure 12a has a round-robin-like configuration in which every stream priority gets the same amount of time. The goal of this configuration is to allocate equal bandwidth to all streams, without accounting for the different stream characteristics. However, we observe that the blue dashed curve has an increased tail-latency. This is due to the higher number of sent frames than in the other two streams (see Table IV). If the GCL configuration is composed of only short slots for each single stream, then streams with higher frame frequency are more affected by blocked time slots than other streams—a frame from the previous cycle which was buffered at the switch now uses the next cycle. Therefore the probability of outliers increases. The measurements with stream set Ψ and stream set Ω in Figure 15e and Figure 15i (see Appendix) clearly demonstrate that this GCL configuration imposes less delay to sporadic cyclic streams, and higher delays to bursty streams with higher data rates. Also, the comparison of Figure 15e and

Figure 15i demonstrates the negative effect of the short slots.

3) *GCL Configuration 2*: Configuration 2, which is considered Figure 12b, allows the highest priority to always transmit data (the gate is never closed), but other streams must wait for their window (similarly to configuration 1). Stream Θ_{high} is always preferred and should experience the lowest delay. In Figure 12b, we observe the effects in a much lower one-way latency for the highest priority stream Θ_{high} (compared to Figure 12a), however with a discontinuous gradient which indicates a relatively high PDV. Due to the small slots, there are collisions where the tails of forwarded streams block the next time slot. Increasing the slot size, see Figure 12f, solves this issue for high-priority traffic; however, increases the overall latency of the lower priority streams. The effects on stream set Ψ and stream set Ω are similar: the prioritization for the high priority stream is guaranteed, and a larger slot size decrease the PDV (see Figures 15 and 16).

4) *GCL Configuration 3*: Figure 12c considers configuration 3 which gives more transmission opportunities to higher priority streams, striving to isolate the stream priorities. For stream set Θ and stream set Ω this works in the expected




Figure 13: Boxplot of one-way delay measurements for different GCL configurations with a window size of 1 for stream set Θ . The x-axis labeling refers to the six GCL configurations in Table VI. This boxplot allows a direct comparison between all combinations with insights into particular statistical values, such as outliers, which are of high importance for TSN.

manner (see Figure 13 and Figure 16, respectively). In stream set Ψ , the large and bursty second priority stream has an effect on the high priority stream: Due to the short $15\mu\text{s}$ slots, the stream Ψ_{med} can block the high priority stream from being transmitted. Increasing the slot size can be a solution, but reduces the bandwidth for the low priority, high-bursty stream Ψ_{low} , as shown in Figure 15k (vs. Figure 15g). This can even increase the delay of stream Ψ_{low} , when the frame rate is higher than the GCL cycle time, in which case the DuT needs to buffer more and more frames.

5) *GCL Configuration 4*: Configuration 4 is similar to configuration 3, but distributes the transmission opportunities over time (the shape resembles a tree). The target was to provide a priority isolation between the streams, respecting their stream characteristics. The measurements confirm this behavior except for stream set Ω with a short slot size, see Figure 16h. The overall one-way delay increases since there are more slots in which the corresponding stream is not allowed to send and forced to wait. Nevertheless, we observe a clear distinction between the stream priorities.

6) *Comparison of GCL Configurations*: Figure 13 facilitates the comparison of the GCL configurations in one plot. Figures 12a, 12b, 12c, 12d, 12i, and 12k are combined in Figure 13. We observe that configuration 3 has the least outliers and clearest differentiation between the stream priorities. This is to be expected, since configuration 3 increases the number of transmission opportunities linearly with the

priority. Configuration 4 is similar to configuration 3 and achieves similar one-way latencies, although the overall one-way latency increases compared to configuration 3. It is important to note that the y-axis uses logarithmic scaling. Therefore, a larger box in the box plots indicates a vastly increased latency variation. Furthermore, the logarithmically scaled boxplot reveals outliers more clearly. Especially in configurations 1, 1 with guard band, 2, and 4 without guard band, there are (numerous) outliers.

In summary, for a service with clear priority distinction, the measurement results indicate that a configuration that is similar to the third configuration is recommended. If only the latency and PDV of the highest priority stream is of importance, then configuration 2 can perform slightly better, while degrading the performance of the remaining stream more significantly.

7) *Slot Size*: With increased slot size (see second row in Figure 12), the overall delay increases since the gates are opened and closed for a longer period and frames from another queue have to wait longer for the next send window. This can be seen, e.g., in Figure 12d and 12h, where the average latency increases from $18.4\mu\text{s}$ to $34.0\mu\text{s}$ for stream Θ_{high} . On the other hand, the tail-latency can be slightly decreased due to a more relaxed time frame where the gate is open (see Figure 12d). In general, we observe from all four measurements with an increased slot size that the shapes of the CCDF curves are *smoother*. That means the PDV is distributed over a longer period, since the slot size increases and therefore the slots are more contiguous than with a higher frequent GCL.

The measurements with different slot sizes show that the GCL slots need to be aligned with the maximum burst size of the streams. Otherwise, the streams can block each other, which cannot be resolved with guard bands in all cases. In conclusion, our setting with a slot size of one is not suitable for every stream set, but clearly demonstrates the importance of this slot size parameter.

8) *Guard Band*: Theoretically, a guard band should reduce the effect, that the tail of long frames blocks the next time slot. In the bottom row in Figure 12, we investigate the effects of an added guard band. The guard band should protect the highest priority stream Θ_{high} from interference from the lower priority streams (Θ_{med} and Θ_{low}) and is configured to always have a slot size of 1, regardless of the slot length for the test data, since the guard band should only cover the worst case: A full lower-priority Ethernet frame is enqueued for transmission right before the gate closes. If this occurs at the end of a cycle (usually the highest priority gets the first transmission window in a cycle), this could interfere and delay the next high priority frame. In the four GCL configurations, it is only necessary to add guard bands for configurations 1 and 4, since the gate for the highest priority is always kept open in the other two configurations.

We observe from Figure 12i a similar behavior as in Figure 12a, the guard band has only a minor effect. However, now the second highest priority exhibits lower one-way delays in the range below 99%. An additional guard band combined with short slots has a negative effect, see Figure 12i, because the guard band reduces the available bandwidth for the streams.

The guard band usage for GCL configuration 1 shows

even a negative impact on the priority isolation, because the guard band forces the high priority stream to wait. On the other hand, we can now detect outliers below 99.99% for stream Θ_{med} which can be explained by the too short slot size (see same behavior in Figure 12a). With an added guard band, the average latency increases by 16% compared to the test without guard band.

If the slot size increases to 3, the behavior resembles Figure 12e; however, with one exception: the total average latency of the highest priority stream increases by 24%, or $42.8\ \mu\text{s}$ for the decade from 99% to 99.9%. This can be explained by the fact of the high frequency of the stream Θ_{high} , which is in the same region as the GCL cycle time. Sometimes, the transmission is delayed and therefore misses its time slot and must wait until the next one.

For configuration 4, we observe less changes. The latency increases in the range below 50% for Figure 12k for all three streams. The mean latency for all streams increases on the other hand, because the time for accessing the channel again increases as well. Although, this effect is less visible for stream Θ_{high} , since there are the most send opportunities reserved for this stream in the GCL configuration 4. In Figure 12l, we observe only small deviations to Figure 12h, but we can detect a similar effect as for the same configuration with a slot size of 1 (shift of latency caused by the additional guard band period).

For all configurations applies: a negative effect of the guard band is the bandwidth reduction for all other streams. However, the proportion of the guard band in relation to the overall cycle time is configurable, i.e., by extending the slot sizes the channel utilization can be enhanced.

9) *Traffic Burstiness*: Another stream characteristic is the traffic burstiness: stream set Θ and stream set Ω are mostly cyclic, whereas the Video VBR stream of stream set Ψ has a bursty behavior, stream Video CBR is slightly bursty. Where TAS 1 provide an equal channel bandwidth of $250\ \text{Mbit s}^{-1}$ for each stream, the bandwidth slightly changes if a guard-band is applied. For small slot-size each streams has still $200\ \text{Mbit s}^{-1}$ and for large slotsize there are $230\ \text{Mbit s}^{-1}$ available. The stream set Θ have a fixed data rate of $7.6\ \text{Mbit s}^{-1}$, $5.1\ \text{Mbit s}^{-1}$ and $3.0\ \text{Mbit s}^{-1}$. So without bursts the data streams are not affected by congestion. If the TAS has a slot size of 1, then a burst of two frames will be most likely (and a burst of three or more frames will certainly) not be forwarded immediately. The influence of frame burstiness can be observed by comparing the first and second row in Figure 15 (see GCL configurations). We observe that the stream Ψ_{low} is strongly delayed because the GCL configuration does not respect the stream particularities. Also, increasing the GCL slot size to 3 does not provide enough bandwidth for stream Ψ_{low} . In comparison, the performance of the slightly bursty stream Ω_{low} is improved by providing more bandwidth with a slot size of 3, see second and third row of Figure 16.

10) *Summary*: Comparing subsection VI-C and subsection VI-D does not reveal consistent advantages of traffic shaping and hard scheduling. The TAS does not always perform better, especially for the average latency. For the TAS,




Figure 14: Measurement with TAS applied at sender and intermediate switch to synchronize sender packet transmission with switch gate states to avoid queuing delay, for stream set Θ . Ideally, Figure 9a and this figure should be identical, implying that an optimized and distributed TAS can achieve the same latencies as dedicated communication paths.

we observe an upper (a guaranteed) limit of latency, but the PDV varies often widely. However, the explanation for this delay variation is simple: our senders are not synchronized with the gates on the switch. The measurement runs were started randomly (compared to the cycle of the switch). Also, we have discussed in subsection V-B that our tools for sending data depend on the host machine's precision. Therefore, frames arrive at the switch not necessarily when the corresponding gate is open. The frames are then enqueued and must wait until the gate opens again. This queuing is not efficient and can significantly increase the one-way delay.

E. Time-aware Traffic Shaping with Synchronized Senders

The TAS scheduling can be used for stream priority isolation preventing the streams influencing each other negatively. But this feature comes at a price: reduced bandwidth and additional delay for all streams. An optimally configured TAS setup should have all components configured so as to avoid collisions. If a data source is sending traffic already periodically, the source can be synchronized to the network.

In this section, we examine for the first time a strategy to achieve near minimal latency and PDV in a measurement testbed using Commercial-of-the-Shelf hardware. As described in subsection IV-C3, the Linux kernel provides tools to use a software TAS. A software TAS can align (time-synchronize) the sender nodes with the switch since the network testbed nodes are already time-synchronized via PTP with the DuT. The suggested technology uses buffering on the sender to shape the packet traffic even before it is transmitted by the sender node. Essentially, we *distribute* the GCL configuration on the switch over the network nodes with some adjustments for path delay and other system inaccuracies [30, Sec. V-D-1]. This is realistic as in real systems, such as in robotics systems,

the cycle of the robot nodes can be adjusted to be time-synchronized with the time-aware network architecture. The described approach dramatically reduces the latency and PDV, even on generic systems, such as the used COTS hardware in our TSN-FlexTest testbed.

In Figure 14, we show that it is possible to synchronize a generic application running on COTS hardware with the TAS by utilizing a scheduling on the transmitter side itself. With this configuration the packet delay caused by queuing in the network can be reduced to a minimum, compared to the baseline measurements in Figure 9. Figure 14 shows the CCDF of the one-way packet latency with the time synchronization distributed to the sender nodes. From a comparison with Figure 9a we conclude with the synchronized senders using a distributed TAS, we achieve essentially the same low one-way latencies as in our baseline measurement (when only a single stream was transmitted in Figure 9a), but now with added cross traffic of 2 Gbit s^{-1} transmitted at the same time in Figure 14. There is close to no PDV for the streams from stream set Θ while the DuT is forced to drop at least half of the incoming data packets. At the same time, the one-way latency remains at a minimum, i.e., 4968 ns and a standard deviation of 40 ns for stream Θ_{high} .

This approach with synchronized senders using a distributed TAS allows all types of interference, such as high-bandwidth file sharing, on the same physical connection while providing exceptional performance for critical services. However, there is a downside when using a software TAS: it is not possible to achieve very short cycles compared to the hard TAS in the DuT. The time a gate was opened was therefore increased by a factor of 10 while still providing the same throughput.

VII. CONCLUSION

We introduced TSN-FlexTest, a flexible highly-precise TSN measurement testbed, for evaluating TSN features. Following a comprehensive review of the available hardware and software, TSN-FlexTest is designed with Commercial-of-the-Shelf (COTS) hardware and open-source software to foster further studies, allowing lost-cost TSN testbed measurements. The validation of the TSN-FlexTest testbed highlighted the following features: The underlying PTP clock synchronization provides nanosecond precision, which enables our high-precision TSN measurements. The flexible cyclic traffic generator empowers researchers to reproduce various stream characteristics without the necessity of obtaining expensive devices, e.g., to simulate multi-modal feedback in the context of the Tactile Internet.

We conducted extensive evaluation studies with our TSN-FlexTest testbed with multiple stream sets. We measured KPIs for widely used Quality-of-Service (QoS) configurations, including *net-neutral* transmissions, Strict Priority Queuing, and the usage of a Time-aware Shaper (TAS). We found that the TAS can provide an upper bounded one-way delay, although the Packet Delay Variation (PDV) may vary significantly in certain cases, especially for sporadic data streams, e.g., video traffic. However, adjusting the GCL in terms of slot size and inserting guard bands can mitigate the PDV to some extent.

We identified the sender node behavior as the root cause of the high PDV: randomly time-shifted data transmissions by sender nodes can introduce high delays due to possibly closed gate states at a switch. We conducted measurements with sending nodes synchronized to the state of the TSN switch. By leveraging solely COTS hardware and open-source software in the TSN-FlexTest testbed, we have thus been able to achieve a level of QoS that is comparable to a dedicated link while accommodated multiple parallel streams, with a more than 200% over-saturation of the link. We provide the source code of the TSN-FlexTest testbed publicly available at <https://github.com/5GCampus/tsn-testbed>.

A. Lessons Learned

Our experience with the TSN-FlexTest testbed setup provided valuable lessons about tools and issues that can be helpful in future research and development:

- To obtain reliable measurement results, all power-saving techniques on all COTS hardware components should be deactivated. This mostly entails adjusting the CPU clock. Since we leveraged generic CPUs and a general-purpose operating system (Ubuntu Linux), it is furthermore recommended to increase process isolation. Under Linux, it is possible to configure the scheduler to not use certain CPU cores and further to assign processes/threads to these cores. This increased process isolation improves the measurement precision by reducing interrupts from co-running services on the same machine blocking the actual measurement.
- Drivers may support fewer features than the official data sheet reports. For example, the documented *one-step timestamping* feature of the used Intel NIC had to be manually patched into the `igb` driver. The patch is available in the <https://github.com/5GCampus/tsn-testbed> repository.
- It is recommended to use the latest software: Particularly, the employed `tcp_replay` tool behaves differently below version 4.3.4, where the frame transmission times become longer. At the point of time when the measurements were conducted, the Ubuntu repository did not offer the already patched version. More information on the cycle time evaluation can be found in subsection V-B.
- We encountered hardware problems while stress-testing with traffic generators, such as *MoonGen*. We conducted the measurements with two other TSN switches. During high traffic loads and PTP long-term measurement, the TSN switches became unstable, with one device irreversibly breaking and needing to be returned to the manufacturer. Our assumption is that not all features of the switches are designed to be run under all conditions.
- Our results confirmed previous observations (see subsection IV-C2) that non-DPDK-based software packet generators cannot saturate a link for all frame sizes.
- Handling large data sets with high precision requires several iterations with enhancements for the software used for evaluation. We want to highlight one very specific aspect of the measurement process, that can lead

to deviations in the results. A Unix epoch timestamp with nanosecond resolution, as it is employed in the TSN-FlexTest testbed, should not be stored in the *float* data type (see Python PEP410 issue for more information); there is a loss of precision after 194 days of absolute time. This means the time cannot be stored anymore with a 1 ns resolution if the “time” is greater than July 14, 1970. In order to maintain correctness inside the evaluation, the more precise data type of `Decimal()` must be chosen. This is specific to the *Python* language. However, this can occur in other programming languages as well and especially *Python* is very well suited for evaluation because of the rich tools for statistical data analysis.

B. Future Work

The proposed TSN-FlexTest testbed expands opportunities for investigating new methods and for evaluating existing methods through measurements on a real hardware testbed. We proceed to summarize some potential future research directions. A first step could be to investigate the effectiveness of other TSN standards, such as Frame Preemption (FP), Frame Replication and Elimination for Reliability (FRER), and Stream Reservation Protocol (SRP). Future work could also consider a wider set of traffic profiles. Furthermore, the TSN-FlexTest testbed topology could be extended to multiple switches and a larger number of sending and receiving nodes in future work.

Evaluating commercial industry-grade hardware components on the market can reveal additional information about their behaviors and performance levels and may provide insights for academic researchers. For example, using a time-coordinated CPU in the testbed could potentially improve the precision of the TSN network, which could be investigated further.

Also, the configuration can have a significant impact on the TSN testbed performance. Consequently, one important future work direction is to design a framework for determining proper pre-configuration for static topologies or reconfiguring parameters according to changes in traffic profile and network resources over time.

Finally, integrating TSN with wireless technologies can significantly improve flexibility and mobility, which can be advantageous for a wide range of new use cases [5], [119]. Integrating the 5G System (5GS) with a TSN network is considered in recent 3GPP standards [120] in an architecture where the 5GS acts as a virtual TSN bridge. The TSN-FlexTest testbed can be employed to develop a 5GS in accordance with recent 3GPP standards to enable measurement based performance evaluations for a 5G-TSN network.

APPENDIX A MEASUREMENTS WITH REAL-WORLD DATA SETS

This appendix provides an overview of the conducted measurements with stream set Ψ and stream set Ω , following the configurations described in section VI. The purpose of these

measurements is to gain insights into TSN behaviors for *real-world* data sets. As described in subsection IV-D1, real data sources behave differently compared to our generic data set. Figure 15 and Figure 16 provide an overview of how stream characteristics, e.g., non-periodic and bursty traffic, can influence the one-way delay and PDV. Figure 17 and Figure 18 allow additional insights for the cycle time and the PDV, respectively.

ACKNOWLEDGMENT

Funded in part by the German Federal Ministry for Economic Affairs and Climate Action (BMWK) projects “TICCTEC” – grant 01MC22007A, “5G-OPERA” – grant 01MJ22008A, and “stic5G” – grant 01MJ22018C, and by the German Research Foundation (DFG, Deutsche Forschungsgemeinschaft) as part of Germany’s Excellence Strategy – EXC 2050/1 – Project ID 390696704 – Cluster of Excellence “Centre for Tactile Internet with Human-in-the-Loop” (CeTI) of Technische Universität Dresden. We also would like to thank the open-source community.

REFERENCES

- [1] S. Senk, M. Ulbricht, J. Acevedo, G. T. Nguyen, P. Seeling, and F. H. P. Fitzek, “Flexible measurement testbed for evaluating Time-Sensitive Networking in industrial automation applications,” in *Proc. IEEE 8th Int. Conf. on Network Softwarization (NetSoft)*, 2022, pp. 402–410.
- [2] R. Behraves, D. Harutyunyan, E. Coronado, and R. Riggio, “Time-sensitive mobile user association and SFC placement in MEC-enabled 5G networks,” *IEEE Transactions on Network and Service Management*, vol. 18, no. 3, pp. 3006–3020, 2021.
- [3] J. Falk, H. Geppert, F. Dürr, S. Bhowmik, and K. Rothermel, “Dynamic QoS-aware traffic planning for time-triggered flows in the real-time data plane,” *IEEE Transactions on Network and Service Management*, vol. 19, no. 2, pp. 1807–1825, 2022.
- [4] V. Gavriluț, A. Pruski, and M. S. Berger, “Constructive or optimized: An overview of strategies to design networks for time-critical applications,” *ACM Computing Surveys (CSUR)*, vol. 55, no. 3, pp. 62:1–62:35, 2022.
- [5] T. Hoeschele, C. Dietzel, D. Kopp, F. H. Fitzek, and M. Reisslein, “Importance of internet exchange point (IXP) infrastructure for 5G: Estimating the impact of 5G use cases,” *Telecommunications Policy*, vol. 45, no. 3, p. Art. no. 102091, 2021.
- [6] A. Sacco, F. Esposito, and G. Marchetto, “Restoring application traffic of latency-sensitive networked systems using adversarial autoencoders,” *IEEE TNSM*, vol. 19, no. 3, pp. 2521–2535, 2022.
- [7] S. Senk, M. Ulbricht, I. Tsokalo, J. Rischke, S.-C. Li, S. Speidel, G. T. Nguyen, P. Seeling, and F. H. P. Fitzek, “Healing hands: The tactile internet in future tele-healthcare,” *Sensors*, vol. 22, no. 4, p. Art. no. 1404, 2022.
- [8] Y. Seol, D. Hyeon, J. Min, M. Kim, and J. Paek, “Timely survey of time-sensitive networking: Past and future directions,” *IEEE Access*, vol. 9, pp. 142 506–142 527, 2021.
- [9] Y. Huang, S. Wang, T. Huang, and Y. Liu, “Cycle-based time-sensitive and deterministic networks: Architecture, challenges, and open issues,” *IEEE Communications Magazine*, vol. 60, no. 6, pp. 81–87, 2022.
- [10] V. Balasubramanian, M. Aloqaily, and M. Reisslein, “An SDN architecture for time sensitive industrial IoT,” *Computer Networks*, vol. 186, p. Art. no. 107739, 2021.
- [11] —, “Fed-TSN: Joint failure probability-based federated learning for fault-tolerant time-sensitive networks,” *IEEE Transactions on Network and Service Management*, vol. 20, no. 2, pp. 1470–1486, 2023.
- [12] N. G. Nayak, F. Dürr, and K. Rothermel, “Incremental flow scheduling and routing in time-sensitive software-defined networks,” *IEEE Transactions on Industrial Informatics*, vol. 14, no. 5, pp. 2066–2075, 2017.
- [13] D. Thiele, R. Ernst, and J. Diemer, “Formal worst-case timing analysis of Ethernet TSN’s time-aware and peristaltic shapers,” in *Proc. IEEE Vehicular Netw. Conf. (VNC)*, 2015, pp. 251–258.

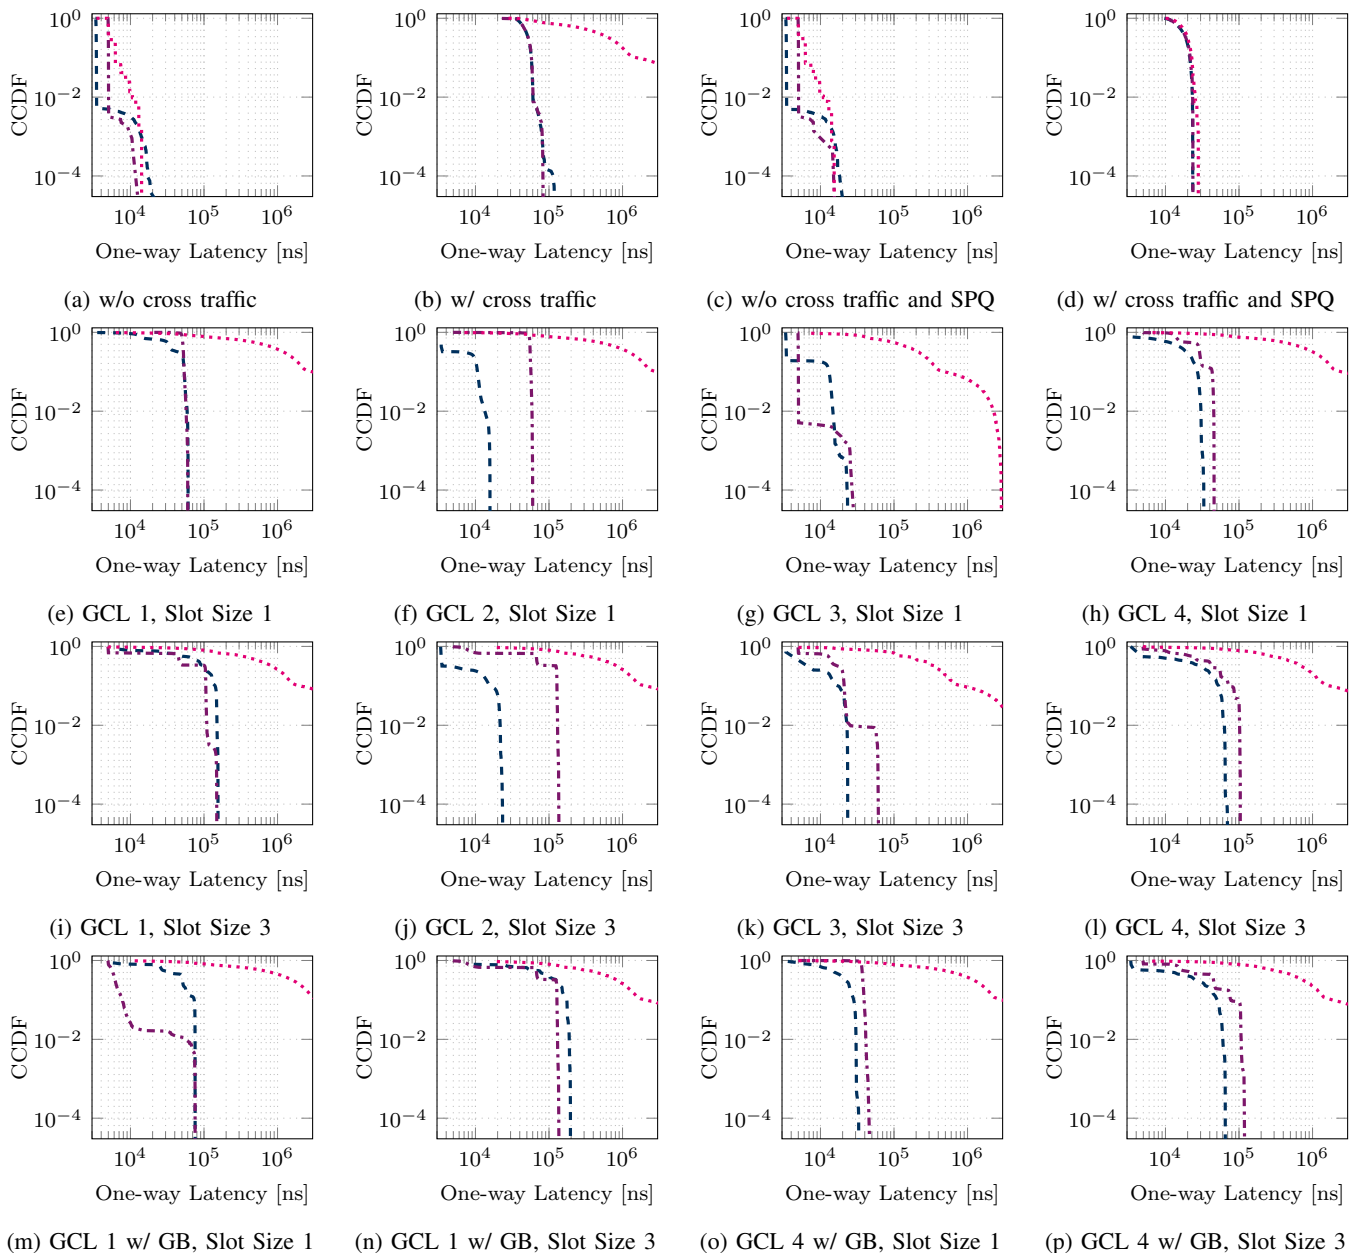


Figure 15: Measurement of stream set Ψ for configurations in section VI: We measured the streams with equal priority (Figures a and b), with enabled SPQ (Figures c and d), and with GCL configurations described in subsection VI-D (see Figures e to p). We conducted measurements without (w/o) and with (w/) cross traffic and enabled Guard Band (GB) for protecting the highest priority traffic. The GCL slot size was varied between 1 and 3 (see subsection VI-D).

- [14] J. Migge, J. Villanueva, N. Navet, and M. Boyer, “Insights on the performance and configuration of AVB and TSN in automotive Ethernet networks,” in *Proc. ERTS*, Jan. 2018, pp. 1–14.
- [15] F. He, L. Zhao, and E. Li, “Impact analysis of flow shaping in Ethernet-AVB/TSN and AFDX from network calculus and simulation perspective,” *Sensors*, vol. 17, no. 5, 2017.
- [16] W. Guo, Y. Huang, J. Shi, Z. Hou, and Y. Yang, “A formal method for evaluating the performance of TSN traffic shapers using UPPAAL,” in *Proc. IEEE LCN*, 2021, pp. 241–248.
- [17] *OMNeT++ Discrete Event Simulator*, July 2022. [Online]. Available: <https://omnetpp.org/>
- [18] *ns-3 Network Simulator*, June 2022. [Online]. Available: <https://www.nsnam.org/>
- [19] M. Ulbricht, J. Acevedo, S. Krdoyan, and F. H. Fitzek, “Emulation vs. reality: Hardware and software co-design in emulated and implemented Time-Sensitive Networks,” in *Proc. European Wireless*, Verona, Italy, Nov. 2021.
- [20] “An instant virtual network on your laptop (or other pc),” Oct. 11, 2022. [Online]. Available: <http://mininet.org/>
- [21] P. Emmerich, S. Gallenmüller, D. Raumer, F. Wohlfart, and G. Carle, “MoonGen: A scriptable high-speed packet generator,” in *Proc. ACM Internet Measurement Conf.*, 2015, p. 275–287.
- [22] R. Kundel, F. Siegmund, J. Blendin, A. Rizk, and B. Koldehofe, “P4STA: High performance packet timestamping with programmable packet processors,” in *Proc. IEEE/IFIP NOMS*, 2020, pp. 1–9.
- [23] M. Kassouf, L. Dupont, J. Béland, and A. Fadlallah, “Performance of the Precision Time Protocol for clock synchronisation in smart grid applications,” *Trans. on Emerging Telecommunications Technologies*, vol. 24, no. 5, pp. 476–485, 2013.
- [24] B. Ferencz and T. Kovácskóczy, “Hardware assisted COTS IEEE 1588




Figure 16: Measurement of stream set Ω for configurations in section VI: We measured the streams with equal priority (Figures a and b), with enabled SPQ (Figures c and d), and with GCL configurations described in subsection VI-D (see Figures e to p). We conducted measurements without (w/o) and with (w/) cross traffic and enabled Guard Band (GB) for protecting the highest priority traffic. The slot size of the GCL was varied between 1 and 3 (see subsection VI-D).

- solution for x86 Linux and its performance evaluation,” in *Proc. IEEE Int. Symp. on Prec. Clock Synch. for Measurement, Control and Commun. (ISPCS)*, 2013, pp. 47–52.
- [25] E. Kyriakakis, J. Sparsø, and M. Schoeberl, “Hardware assisted clock synchronization with the IEEE 1588-2008 Precision Time Protocol,” in *Proc. ACM RTNS*, 2018, pp. 51–60.
- [26] S. Sudhakaran, K. Montgomery, M. Kashef, D. Cavalcanti, and R. Candell, “Wireless time sensitive networking for industrial collaborative robotic workcells,” in *Proc. 17th IEEE Int. Conf. on Factory Communication Systems (WFCS)*, 2021, pp. 91–94.
- [27] T. Agarwal, P. Niknejad, F. Rahmani, M. Barzegaran, and L. Vanfretti, “A time-sensitive networking-enabled synchronized three-phase and phasor measurement-based monitoring system for microgrids,” *IET Cyber-Phys. Sys.: Theory & Appl.*, vol. 6, no. 1, pp. 1–11, 2021.
- [28] P. Kehl, J. Ansari, M. H. Jafari, P. Becker, J. Sachs, N. König, A. Göppert, and R. H. Schmitt, “Prototype of 5G integrated with TSN for edge-controlled mobile robotics,” *Electronics*, vol. 11, no. 11, 2022.
- [29] M. Böhm, J. Ohms, and D. Wermser, “Multi-domain time-sensitive networks - an east-westbound protocol for dynamic TSN-stream configuration across domains,” in *Proc. 24th IEEE Int. Conf. on Emerging Technologies and Factory Automation (ETFA)*, 2019, pp. 1363–1366.
- [30] M. Bosk, F. Rezabek, K. Holzinger, A. G. Marino, F. Fons, A. A. Kane, J. Ott, and G. Carle, “Methodology and infrastructure for TSN-based reproducible network experiments,” *IEEE Access*, in print, pp. 1–1, 2022.
- [31] L. Zhao, P. Pop, and S. Steinhorst, “Quantitative performance comparison of various traffic shapers in time-sensitive networking,” *IEEE Trans. Network and Service Management*, in print, pp. 1–1, 2022.
- [32] B. Caruso, L. Leonardi, L. L. Bello, and G. Patti, “Experimental assessment of TSN support in heterogeneous platforms with virtualization for

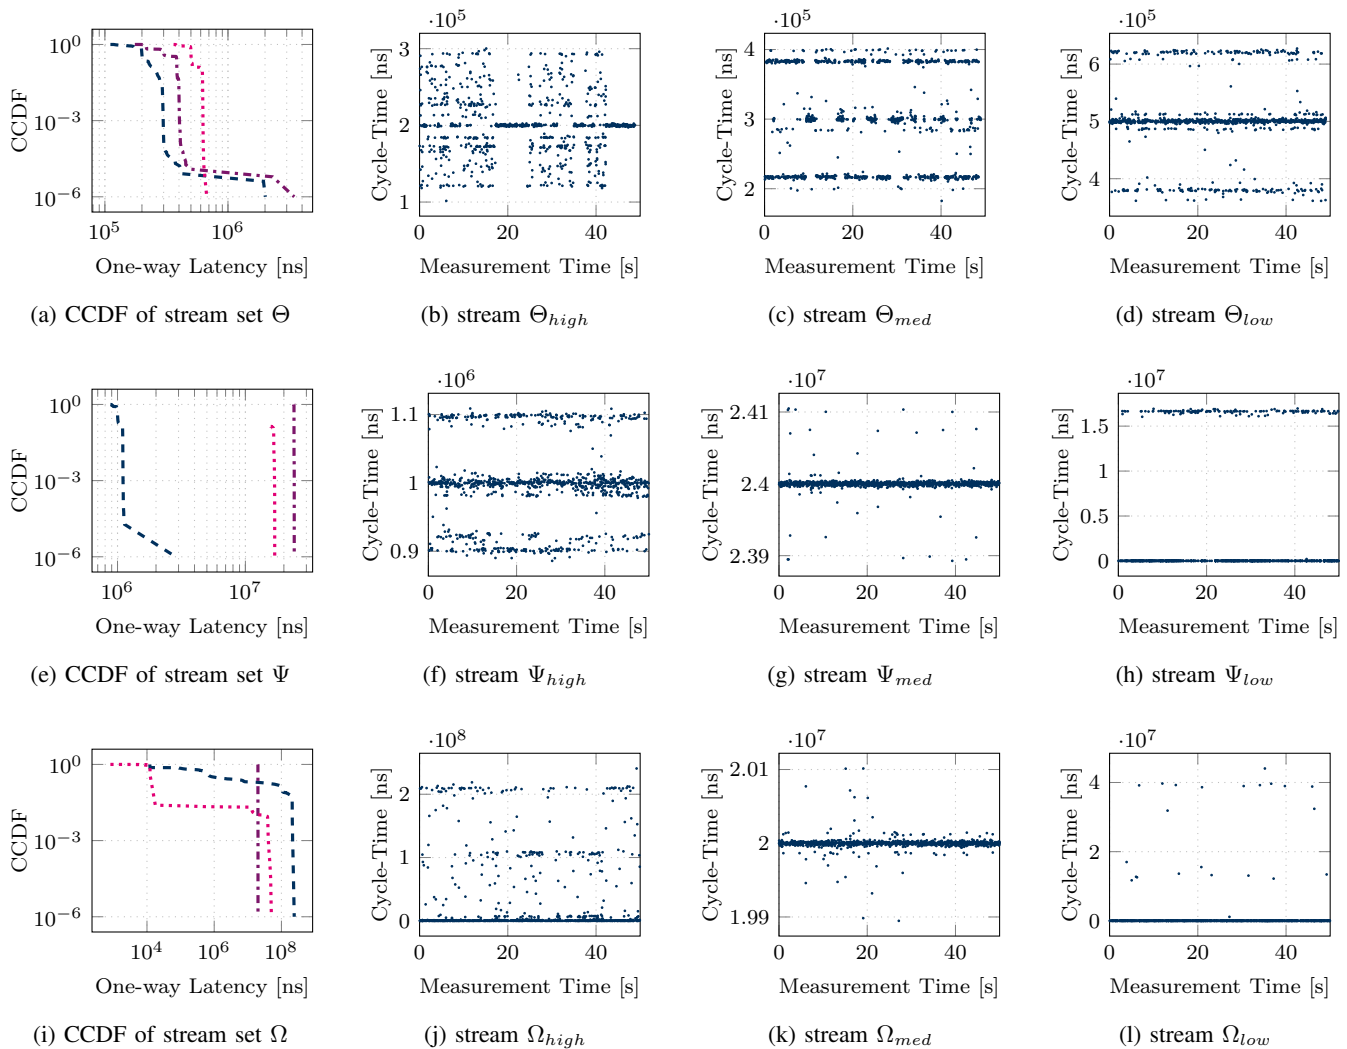


Figure 17: The figure illustrates the different cycle times of the three stream sets. In each row, the first figure shows the CCDF for a comparison of all three streams within a stream set, and then the three streams individually with a plot showing the behavior over time. Since the streams especially in stream set Ψ and Ω have rather varying characteristics, the y-axis is not fixed, otherwise, outliers may not be easily visible.

- automotive applications,” in *Proc. IEEE AEIT AUTOMOTIVE*, 2021, pp. 1–5.
- [33] S. Chouksey, H. S. Satheesh, and J. Åkerberg, “An experimental study of TSN-NonTSN coexistence,” in *Proc. IEEE 11th Ann. Computing and Commun. Workshop and Conf. (CCWC)*, 2021, pp. 1577–1584.
- [34] I.-H. Liu, L. Y. Chang, J. S. Li, and C. G. Liu, “Industrial control system cybersecurity testbed with TSN feature,” in *Proc. Int. Conf. on Artificial Life and Robotics (ICAROB)*, 2022, pp. 291–294.
- [35] M. M. Tajiki, S. H. G. Petroudi, S. Salsano, S. Uhlig, and I. Castro, “Optimal estimation of link delays based on end-to-end active measurements,” *IEEE Transactions on Network and Service Management*, vol. 18, no. 4, pp. 4730–4743, 2021.
- [36] W. Quan, W. Fu, J. Yan, and Z. Sun, “OpenTSN: An open-source project for time-sensitive networking system development,” *CCF Transactions on Networking*, vol. 3, no. 1, pp. 51–65, 2020.
- [37] A. Heideker and C. Kamienski, “Network queuing assessment: A method to detect bottlenecks in service function chaining,” *IEEE TNSM, in print*, pp. 1–1, 2022.
- [38] I. Kunze, J. R uth, and O. Hohlfeld, “Congestion control in the wild—investigating content provider fairness,” *IEEE Transactions on Network and Service Management*, vol. 17, no. 2, pp. 1224–1238, 2020.
- [39] V. Maglogiannis, D. Naudts, S. Hadiwardoyo, D. van den Akker, J. Marquez-Barja, and I. Moerman, “Experimental V2X evaluation for C-V2X and ITS-G5 technologies in a real-life highway environment,” *IEEE TNSM*, vol. 19, no. 2, pp. 1521–1538, 2022.
- [40] B. Martini, M. Gharbaoui, D. Adami, P. Castoldi, and S. Giordano, “Experimenting SDN and cloud orchestration in virtualized testing facilities: Performance results and comparison,” *IEEE Transactions on Network and Service Management*, vol. 16, no. 3, pp. 965–979, 2019.
- [41] N. Makris, C. Zarafetas, A. Valantasis, and T. Korakis, “Service orchestration over wireless network slices: Testbed setup and integration,” *IEEE TNSM*, vol. 18, no. 1, pp. 482–497, 2021.
- [42] O. S. Penaherrera-Pulla, C. Baena, S. Fortes, E. Baena, and R. Barco, “KQI assessment of VR services: A case study on 360-Video over 4G and 5G,” *IEEE TNSM, in print*, pp. 1–1, 2022.
- [43] K. Polachan, J. Pal, C. Singh, and P. T. V., “Assessing quality of control in tactile cyber-physical systems,” *IEEE Transactions on Network and Service Management, in print*, pp. 1–1, 2022.
- [44] S. Troia, M. Mazzara, M. Savi, L. M. M. Zorello, and G. Maier, “Resilience of delay-sensitive services with transport-layer monitoring in SD-WAN,” *IEEE Transactions on Network and Service Management*, vol. 19, no. 3, pp. 2652–2663, 2022.
- [45] N. Deric, A. Varasteh, A. Van Bempten, A. Blenk, and W. Kellerer, “Enabling SDN hypervisor provisioning through accurate CPU utilization prediction,” *IEEE Transactions on Network and Service Management*, vol. 18, no. 2, pp. 1360–1374, 2021.
- [46] M.-R. Fida, A. F. Ocampo, and A. Elmokashfi, “Measuring and localising congestion in mobile broadband networks,” *IEEE Transactions on Network and Service Management*, vol. 19, no. 1, pp. 366–380, 2022.
- [47] M. A. Imtiaz, D. Starobinski, and A. Trachtenberg, “Empirical com-

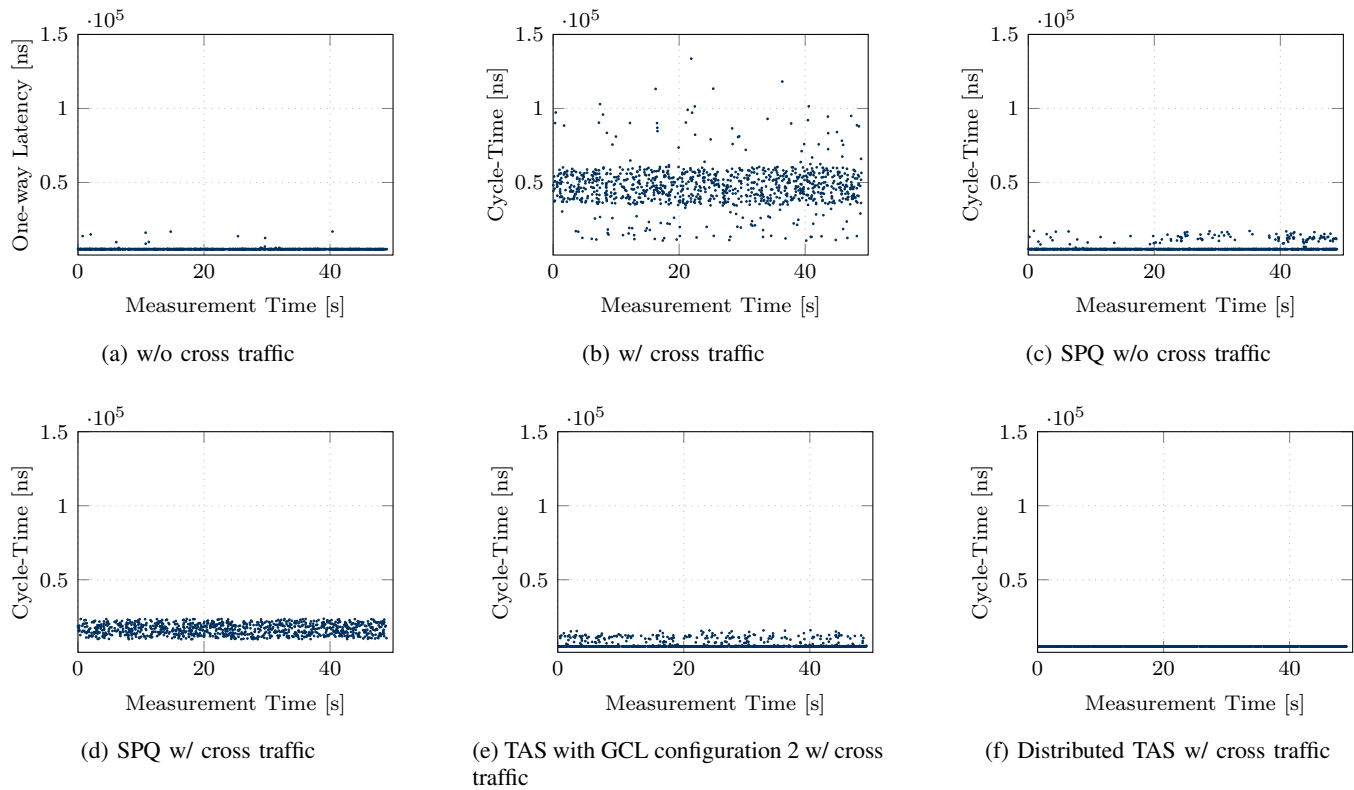


Figure 18: The figure shows the measurement for stream Θ_{high} for different applied QoS features: no applied QoS, SPQ, TAS, and distributed TAS. Figure 10, 11, and 12 are using the CCDF to visualize the measurement results. Here, the plot elucidates the behavior over time. The effects of additional cross traffic is clearly noticeable in the respective configurations. Furthermore, we can observe the advantages of all applied QoS features. However, the plots highlight the superiority of the TAS, in particular with the distributed approach.

- parison of block relay protocols,” *IEEE Transactions on Network and Service Management*, in print, pp. 1–1, 2022.
- [48] O. Michel, J. Sonchack, G. Cusack, M. Nazari, E. Keller, and J. M. Smith, “Software packet-level network analytics at cloud scale,” *IEEE TNSM*, vol. 18, no. 1, pp. 597–610, 2021.
- [49] “IEEE Standard for Local and Metropolitan Area Networks—Timing and Synchronization for Time-Sensitive Applications,” pp. 1–421, 2020.
- [50] “The Linux PTP Project,” 2022. [Online]. Available: <https://linuxptp.sourceforge.net/>
- [51] Intel, *Intel Ethernet Controller I210*, January 2021.
- [52] InnoRoute GmbH, *TrustNode networkdriver git*, 2019. [Online]. Available: <https://github.com/InnoRoute/packages/tree/master/TrustNodeDriver/src>
- [53] “IEEE Standard for a Precision Clock Synchronization Protocol for Networked Measurement and Control Systems,” pp. 1–154, 2002.
- [54] “IEEE Standard for a Precision Clock Synchronization Protocol for Networked Measurement and Control Systems,” pp. 1–269, 2008.
- [55] “IEEE Standard for a Precision Clock Synchronization Protocol for Networked Measurement and Control Systems,” pp. 1–499, 2020.
- [56] A. Nasrallah, A. S. Thyagaturu, Z. Alharbi, C. Wang, X. Shao, M. Reisslein, and H. Elbakoury, “Ultra-low latency (ULL) networks: The IEEE TSN and IETF DetNet standards and related 5G ULL research,” *IEEE COMST*, vol. 21, no. 1, pp. 88–145, 2019.
- [57] “IEEE Standard for Local and metropolitan area networks – Bridges and Bridged Networks – Amendment 26: Frame Preemption,” pp. 1–52, 2016.
- [58] Y. Huang, S. Wang, X. Zhang, T. Huang, and Y. Liu, “Flexible cyclic queuing and forwarding for time-sensitive software-defined networks,” *IEEE Trans. Network and Service Management*, in print, pp. 1–1, 2022.
- [59] E. Li, F. He, Q. Li, and H. Xiong, “Bandwidth allocation of stream-reservation traffic in TSN,” *IEEE Trans. Network and Service Management*, vol. 19, no. 1, pp. 741–755, 2022.
- [60] A. Nasrallah, A. S. Thyagaturu, Z. Alharbi, C. Wang, X. Shao, M. Reisslein, and H. Elbakoury, “Performance comparison of IEEE 802.1 TSN time aware shaper (TAS) and asynchronous traffic shaper (ATS),” *IEEE Access*, vol. 7, pp. 44 165–44 181, 2019.
- [61] “IEEE Standard for Local and metropolitan area networks– Virtual Bridged Local Area Networks Amendment 12: Forwarding and Queuing Enhancements for Time-Sensitive Streams,” pp. 1–72, 2010.
- [62] “IEEE Standard for Local and metropolitan area networks – Bridges and Bridged Networks - Amendment 25: Enhancements for Scheduled Traffic,” pp. 1–57, 2016.
- [63] M. Kim, D. Hyeon, and J. Paek, “eTAS: Enhanced time-aware shaper for supporting nonisochronous emergency traffic in time-sensitive networks,” *IEEE IoTJ*, vol. 9, no. 13, pp. 10 480–10 491, 2022.
- [64] “IEEE/ISO/IEC International Standard - Information technology - Telecommunications and information exchange between systems - Local and metropolitan area networks - Specific requirements - Part 1Q: Bridges and bridged networks - AMENDMENT 7: Cyclic queuing and forwarding,” pp. 1–34, 2019.
- [65] J. Yan, W. Quan, X. Jiang, and Z. Sun, “Injection time planning: Making CQF practical in time-sensitive networking,” in *Proc. IEEE INFOCOM*, 2020, pp. 616–625.
- [66] “IEEE Standard for Local and Metropolitan Area Networks–Bridges and Bridged Networks - Amendment 34:Asynchronous Traffic Shaping,” pp. 1–151, 2020.
- [67] M. Máté, C. Simon, and M. Maliosz, “Asynchronous time-aware shaper for time-sensitive networking,” *Journal of Network and Systems Management*, vol. 30, no. 4, p. 76, 2022.
- [68] J. Prados-Garzon, L. Chinchilla-Romero, P. Ameigeiras, P. Muñoz, and J. M. Lopez-Soler, “Asynchronous time-sensitive networking for industrial networks,” in *Proc. EuCNC/6G Summit*, 2021, pp. 130–135.
- [69] A. Ademaj, D. Puffer, D. Bruckner, G. Ditzel, L. Leurs, M.-P. Stanica, P. Didier, R. Hummen, R. Blair, and T. Enzinger, “Time sensitive networks for flexible manufacturing testbed characterization and mapping of converged traffic types,” Website, 2019. [Online].

- Available: https://www.iiconsortium.org/pdf/IIC_TSN_Testbed_Char_Mapping_of_Converged_Traffic_Types_Whitepaper_20180328.pdf
- [70] “Profiles Part 1: Fieldbus profiles,” pp. 1–716, 2019.
- [71] S. Brooks and E. Uludag, “Time-sensitive networking: From theory to implementation in industrial automation,” Website. [Online]. Available: <https://www.tttech.com/wp-content/uploads/TSN-in-industrial-automation.pdf>
- [72] C. Simon, M. Maliosz, and M. Mate, “Design aspects of low-latency services with time-sensitive networking,” *IEEE Communications Standards Magazine*, vol. 2, no. 2, pp. 48–54, 2018.
- [73] N. Finn. Current state of ieee 802.1 time-sensitive networking task group. [Online]. Available: <https://iee802.org/1/files/public/docs2014/tsn-nfinn-IEEE-TSN-Status-for-IETF-1114-v01.pdf>
- [74] —, “Introduction to time-sensitive networking,” *IEEE Communications Standards Magazine*, vol. 2, no. 2, pp. 22–28, 2018.
- [75] L. Maile, K.-S. Hielscher, and R. German, “Network calculus results for TSN: An introduction,” in *Proc. IEEE ICTC*, 2020, pp. 131–140.
- [76] L. Thomas, A. Mifdaoui, and J.-Y. L. Boudec, “Worst-case delay bounds in time-sensitive networks with packet replication and elimination,” *IEEE/ACM Trans. Netw.*, in print, pp. 1–15, 2022.
- [77] M. H. Farzaneh and A. Knoll, “Time-sensitive networking (TSN): An experimental setup,” in *Proc. IEEE VNC*, 2017, pp. 23–26.
- [78] *INET Framework*, July 2022. [Online]. Available: <https://inet.omnetpp.org/>
- [79] P. Heise, F. Geyer, and R. Obermaisser, “TSimNet: An industrial time sensitive networking simulation framework based on OMNeT++,” in *Proc. IFIP NTMS*, 2016, pp. 1–5.
- [80] T. Steinbach, H. D. Kenfack, F. Korf, and T. C. Schmidt, “An extension of the OMNeT++ INET framework for simulating real-time ethernet with high accuracy,” in *Proc. 4th Int. ICST Conf. on Simulation Tools and Techniques*, 2011, p. 375–382.
- [81] J. Specht and S. Samii, “Urgency-based scheduler for time-sensitive switched ethernet networks,” in *Proc. 28th Euromicro Conf. on Real-Time Systems (ECRTS)*, 2016, pp. 75–85.
- [82] J. Falk, D. Hellmanns, B. Carabelli, N. Nayak, F. Dürr, S. Kehrer, and K. Rothermel, “NeSTiNg: Simulating IEEE Time-sensitive Networking (TSN) in OMNeT++,” in *Proc. NetSys*, 2019, pp. 1–8.
- [83] J. Jiang, Y. Li, S. H. Hong, M. Yu, A. Xu, and M. Wei, “A simulation model for Time-sensitive Networking (TSN) with experimental validation,” in *Proc. IEEE ETFA*, 2019, pp. 153–160.
- [84] G. A. Rodrigo, “Simulation of Instability in Time-Sensitive Networks with Regulators and Imperfect Clocks,” Master’s thesis, Universitat Politècnica de Catalunya, Barcelona, Spain, 2020.
- [85] D. Krummacker and L. Wendling, “TSN simulation: Time-aware shaper implemented in ns-3,” in *Proc. Workshop on Next Generation Networks and Applications*, 12 2020.
- [86] L. Linguaglossa, et al., “Survey of performance acceleration techniques for network function virtualization,” *Proc. IEEE*, vol. 107, no. 4, pp. 746–764, 2019.
- [87] P. Shantharama, A. S. Thyagaturu, and M. Reisslein, “Hardware-accelerated platforms and infrastructures for network functions: A survey of enabling technologies and research studies,” *IEEE Access*, vol. 8, pp. 132021–132085, 2020.
- [88] A.S. Thyagaturu, et al., “Operating systems and hypervisors for network functions: A survey of enabling technologies and research studies,” *IEEE Access*, vol. 10, pp. 79825–79873, 2022.
- [89] A. Liatifis, P. Sarigiannidis, V. Argyriou, and T. Lagkas, “Advancing SDN: from OpenFlow to P4, a survey,” *ACM Computing Surveys (CSUR)*, in print, 2022.
- [90] R. Kundel, F. Siegmund, R. Hark, A. Rizk, and B. Koldehofe, “Network testing utilizing programmable network hardware,” *IEEE Communications Magazine*, vol. 60, no. 2, pp. 12–17, 2022.
- [91] T. M. Runge, A. Beifuß, and B. E. Wolfinger, “Low latency network traffic processing with commodity hardware,” in *Proc. IEEE SPECTS*, 2015, pp. 1–8.
- [92] A. Beifuß, T. M. Runge, D. Raumer, P. Emmerich, B. E. Wolfinger, and G. Carle, “Building a low latency Linux software router,” in *Proc. 28th Int. Teletraffic Congress (ITC 28)*, vol. 01, 2016, pp. 35–43.
- [93] “IEEE Standard for Local and Metropolitan Area Networks—Bridges and Bridged Networks – Amendment 31: Stream Reservation Protocol (SRP) Enhancements and Performance Improvements,” pp. 1–208, 2018.
- [94] M.-T. Thi, S. Ben Hadj Said, and M. Boc, “SDN-based management solution for time synchronization in TSN networks,” in *Proc. 25th IEEE ETFA*, 2020, pp. 361–368.
- [95] F. Rezabek, et al., “EnGINE: Developing a flexible research infrastructure for reliable and scalable intra-vehicular TSN networks,” in *Proc. IEEE CNSM*, 2021, pp. 530–536.
- [96] A. Garbugli, A. Bujari, and P. Bellavista, “End-to-end QoS management in self-configuring TSN networks,” in *Proc. 17th IEEE Int. Conf. on Factory Communication Systems (WFCS)*, 2021, pp. 131–134.
- [97] F. Groß, T. Steinbach, F. Korf, T. C. Schmidt, and B. Schwarz, “A hardware/software co-design approach for ethernet controllers to support time-triggered traffic in the upcoming IEEE TSN standards,” in *Proc. IEEE ICCE-Berlin*, 2014, pp. 9–13.
- [98] J. Coleman, S. Almalih, A. Slota, and Y.-H. Lee, “Emerging COTS architecture support for real-time TSN Ethernet,” in *Proc. 34th ACM/SIGAPP Symp. on Applied Computing (SAC)*, 2019, p. 258–265.
- [99] M. Vlk, Z. Hanzálek, K. Brejchová, S. Tang, S. Bhattacharjee, and S. Fu, “Enhancing schedulability and throughput of time-triggered traffic in IEEE 802.1Qbv time-sensitive networks,” *IEEE Transactions on Communications*, vol. 68, no. 11, pp. 7023–7038, 2020.
- [100] N. G. Nayak, F. Duerr, and K. Rothermel, “Routing algorithms for IEEE802.1Qbv networks,” *ACM Sigbed Review*, vol. 15, no. 3, pp. 13–18, 2018.
- [101] G. Miranda, E. Municio, J. Haxhibeqiri, D. F. Macedo, J. Hoebcke, I. Moerman, and J. M. Marquez-Barja, “Time-sensitive networking experimentation on open testbeds,” in *Proc. IEEE INFOCOM (INFOCOM WKSHPS)*, 2022, pp. 1–6.
- [102] N. Beheshti, Y. Ganjali, M. Ghobadi, N. McKeown, J. Naous, and G. Salmon, “Performing time-sensitive network experiments,” in *Proc. ACM/IEEE ANCS*, 2008, p. 127–128.
- [103] A. Pruski and M. Berger, “Design considerations for high-performance time sensitive networking switches,” in *Proc. 10th Int. Conf. on Networks of the Future (NoF)*, 2019, pp. 114–117.
- [104] Z. Li, H. Wan, Y. Deng, X. Zhao, Y. Gao, X. Song, and M. Gu, “Time-triggered switch-memory-switch architecture for time-sensitive networking switches,” *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, vol. 39, no. 1, pp. 185–198, 2020.
- [105] “Introduction to Intel Time Coordinated Computing,” 2022. [Online]. Available: <https://www.intel.com/content/www/us/en/developer/videos/time-coordinated-computing-introduction.html>
- [106] C. Pfefferle, F. Wiedner, and C. Schwarzenberg, “IEEE 802.1 Qcr asynchronous traffic shaping with Linux traffic control,” in *TUM Network Architectures and Services*, doi: 10.2313/NET-2022-01-1_03, 2021, pp. 11–14.
- [107] J. Rischke, P. Sossalla, S. Itting, F. H. Fitzek, and M. Reisslein, “5G campus networks: A first measurement study,” *IEEE Access*, vol. 9, pp. 121786–121803, 2021.
- [108] J. Rischke, C. Vielhaus, P. Sossalla, S. Itting, G. T. Nguyen, and F. H. Fitzek, “Empirical study of 5G downlink & uplink scheduling and its effects on latency,” in *Proc. IEEE WoWMoM*, 2022, pp. 11–19.
- [109] Intel, *Intel 82599 10 GbE Controller Datasheet*, December 2010.
- [110] P. Emmerich, S. Gallenmüller, G. Antichi, A. W. Moore, and G. Carle, “Mind the gap—a comparison of software packet generators,” in *Proc. ACM/IEEE ANCS*, 2017, pp. 191–203.
- [111] S. M. Wong, “An evaluation of software-based traffic generators using docker,” 2018.
- [112] A. Kothuru, S. P. Nooka, and R. Liu, “Application of audible sound signals for tool wear monitoring using machine learning techniques in end milling,” *The International Journal of Advanced Manufacturing Technology*, vol. 95, no. 9, pp. 3797–3808, 2018.
- [113] *Spot Automate sensing and inspection, capture limitless data, and explore without boundaries.*, July 2022. [Online]. Available: <https://www.bostondynamics.com/products/spot>
- [114] M. Ulbricht, J. Acevedo, S. Krdoyan, and F. H. Fitzek, “Precise fruits: Hardware supported Time-Synchronisation on the RaspberryPI,” in *Proc. IEEE SmartNets*, Sep. 2021, pp. 1–6.
- [115] C. Liß, M. Ulbricht, U. F. Zia, and H. Müller, “Architecture of a synchronized low-latency network node targeted to research and education,” in *Proc. IEEE HPSR*, 2017, pp. 1–7.
- [116] F. Fitzek, S.-C. Li, S. Speidel, T. Strufe, M. Simsek, and M. Reisslein, *Tactile Internet: With Human-in-the-Loop*. Academic Press, London, UK, 2021.
- [117] N. Promwongsa, A. Ebrahimzadeh, D. Naboulsi, S. Kianpisheh, F. Belqasmi, R. Glitho, N. Crespi, and O. Alfandi, “A comprehensive survey of the tactile internet: State-of-the-art and research directions,” *IEEE COMST*, vol. 23, no. 1, pp. 472–523, 2020.
- [118] Z. Xiang, F. Gabriel, E. Urbano, G. T. Nguyen, M. Reisslein, and F. H. Fitzek, “Reducing latency in virtual machines: Enabling tactile internet for human-machine co-working,” *IEEE JSAC*, vol. 37, no. 5, pp. 1098–1116, 2019.

- [119] J. Navarro-Ortiz, P. Romero-Díaz, S. Sendra, P. Ameigeiras, J. J. Ramos-Munoz, and J. M. Lopez-Soler, "A survey on 5G usage scenarios and traffic models," *IEEE COMST*, vol. 22, no. 2, pp. 905–929, 2020.

- [120] 3GPP, "System architecture for the 5G System (5GS)," 3rd Generation Partnership Project (3GPP), Technical Specification (TS) 23.501, June 2022, version 17.5.0.



Marian Ulbricht is a PhD student at Technische Universität Dresden. He studied at the Hochschule für Telekommunikation Leipzig (HfTL) for Bachelor and Master in communication technology. The focus of the studies was on embedded systems and micro-controller programming. Since 2015, he is with the InnoRoute GmbH in Munich as software developer and project engineer with focus on TSN and network node design.



Stefan Senk studied at Technische Universität Dresden where he received his degree "Diplom-Ingenieur" (Dipl.-Ing.) in electrical engineering in July 2019. Since late 2019, he is a Ph.D. student at "Deutsche Telekom Chair of Communication Networks" at TU Dresden. He is currently working on 5G non-public networks with interest on deterministic communication. His main research interest is Time-sensitive Networking (TSN) and the focus towards human-machine-collaboration.



Hosein K. Nazari received his B.Sc. in Information Technology Engineering in 2018 from Institute for Advanced Studies in Basic Sciences (IASBS) in Zanjan, Iran. In 2021, he graduated with an M.Sc. in Computer Science from IASBS. He is currently pursuing a Ph.D. in Electrical and Computer Engineering Department of Technische Universität Dresden, Germany. His research interests include Network coding, IoT, and Time-sensitive networking.



How-Hang Liu How-Hang Liu received his B.Sc. in Communication Engineering in 2014 from National Central University (NCU) in Taoyuan, Taiwan. In 2016, he graduated with an M.Sc. in Communication Engineering from National Taiwan University (NTU). He is pursuing a Ph.D. in Electrical and Computer Engineering Department of Technische Universität Dresden, Germany. His research mainly focuses on Time-sensitive networking.



Martin Reisslein (S'96-M'98-SM'03-F'14) received the Ph.D. in systems engineering from the University of Pennsylvania, Philadelphia, PA, USA, in 1998. He is currently a Professor with the School of Electrical, Computer, and Energy Engineering, Arizona State University (ASU), Tempe, AZ, USA. He is currently an Associate Editor for *IEEE Access*, *IEEE Transactions on Education*, and *IEEE Transactions on Mobile Computing*.



Giang T. Nguyen is currently an Assistant Professor, heading the Haptic Communication Systems research group at the Cluster of Excellence for Tactile Internet with Human-in-the-Loop (CeTI) and Faculty of Electrical and Computer Engineering, TU Dresden, Germany. He received a Ph.D. degree in Computer Science from TU Dresden in 2016. His research interests include network softwarization, in-network computing, and distributed systems, aiming at networked systems' low latency, flexibility, and resilience to facilitate haptic communication.



Frank H.P. Fitzek is a Professor and head of the Deutsche Telekom Chair of Communication Networks at the Technische Universität Dresden, coordinating the 5G Lab Germany. He is the spokesman of the DFG Cluster of Excellence CeTI. He received his diploma (Dipl.-Ing.) degree in electrical engineering from the University of Technology Rheinisch-Westfälische Technische Hochschule (RWTH) Aachen, Germany, in 1997 and his Ph.D. (Dr.-Ing.) in electrical engineering from the Technical University Berlin, Germany in 2002 and became Adjunct Professor at the University of Ferrara, Italy in the same year. In 2003 he joined Aalborg University as Associate Professor and later became Professor. In 2005 he won the YRP award for the work on MIMO MDC and received the Young Elite Researcher Award of Denmark. He was selected to receive the NOKIA Champion Award several times in a row from 2007 to 2011. In 2008 he was awarded the Nokia Achievement Award for his work on cooperative networks. In 2011 he received the SAPERE AUDE research grant from the Danish government and in 2012 he received the Vodafone Innovation prize. In 2015 he was awarded the honorary degree Doctor Honoris Causa from Budapest University of Technology and Economics (BUTE).