



Time-Sensitive Networking: From Theory to Implementation in Industrial Automation



Authors

Simon Brooks

TTTech
Vienna, Austria
simon.brooks@tttech.com

Ecehan Uludag

Intel® Corporation
Programmable Solutions Group
San Jose, USA
ecehan.uludag@intel.com

Introduction

Time-sensitive networking (TSN) is set to reshape the industrial communication landscape and lay the foundation for the convergence of Information Technology (IT) and Industrial Operations Technology (OT). By bringing industrial-grade robustness and reliability to Ethernet, TSN offers an IEEE standard communication technology that enables interoperability between standard-conformant industrial devices from any vendor. TSN also removes the need for physical separation of critical and non-critical communication networks, thus allowing open data exchange between operations and enterprise —a concept at the heart of the Industrial Internet of Things (IIoT).

TSN can be implemented in industrial devices as a combination of the switched endpoints and switches and requires associated software. Implementation in an SoC FPGA that integrates both the processor and the FPGA subsystems into a single device, provides increased flexibility and ensures a tight integration of TSN hardware and associated embedded software. As TSN is composed of a set of IEEE 802.1 standard extensions, the re-configurability feature inherent in FPGAs also offers a vital advantage when implementing TSN. When new standards are added or if the standards are altered, FPGAs can be quickly re-configured to support all TSN standards without costly hardware development work.

At the network system level, TSN supports deterministic communication based on network schedules that are distributed to devices via standard configuration interfaces. TTTech offers a portfolio of products that make TSN not only easy to implement at the device level but also simple to use at the system level. This end-to-end approach results in a uniquely integrated TSN implementation with more flexibility and higher performance.

This white paper introduces the features and benefits of TSN technology, as well as showing how FPGA customization and configurability can be used to develop optimized TSN devices and systems for your application.

Table of Contents

- Introduction 1
- The Challenge of Connectivity in Industrial Applications 2
- Benefits of TSN..... 2
 - IEEE Standard2
 - Converged Network.....2
- Benefits of Using an SoC FPGA-Based TSN Implementation... 3
- TSN Features..... 3
- OPC UA over TSN..... 4
- Importance of System-Level Thinking in TSN 4
- TSN Customization 5
 - FPGA Customization5
 - Linux Customization6
- TSN Configuration..... 6
 - Remote Configuration6
 - Network Scheduling6
 - Network Scheduling with OPC UA6
- Summary 6
- Appendix A 7

The Challenge of Connectivity in Industrial Applications

Today, there are multiple variants of Industrial Ethernet protocols available on the market. In most cases, the Industrial Ethernet protocol selected for use in industrial devices differs from vendor to vendor, which means that devices are only compatible with other equipment from the same vendor using the same protocol. This is known as manufacturer lock-in. It forces customers to either buy all industrial equipment from one vendor (even if it is not the most performant or price efficient), or to overcome the considerable challenge of integrating equipment from multiple vendors and by implementing protocol conversion gateways between the various Industrial Ethernet protocols. Both options have led to unnecessary expense and the limitation of innovation on the factory floor over many years, causing industrial automation architectures to be hierarchical, purpose-built, and inflexible.

However, this structure is now undergoing a dramatic change with the advent of the IIoT and Industry 4.0. demands for full automation and greater insights in manufacturing are pushing industrial automation architectures to become more interoperable, flexible, and seamless in nature. In these increasingly converged architectures, real-time connectivity is essential for executing critical processes, as well as for collecting and analyzing data from machines. TSN offers the real-time connectivity capabilities that match or exceed that current Industrial Ethernet protocols but with the added flexibility of the IEEE standards. Similar to the the enterprise world, TSN Ethernet can therefore be the common communication protocol that connects industrial equipment from different vendors, simultaneously fulfilling the challenging requirements of industrial applications.

Benefits of TSN

IEEE Standard

The overwhelming strength of TSN is its status as an open and standard technology, unaffiliated to any organization or company. For an industrial automation market that has struggled for many years with multiple incompatible proprietary communication protocols, TSN brings a number of benefits:

- TSN guarantees compatibility at the network level between devices from different vendors. This gives customers much greater choice of devices for their system, avoiding vendor lock-in and enabling connectivity across systems.
- As TSN is part of the Ethernet standard family, it naturally scales with Ethernet, which means that the technology will not be limited in terms of bandwidth or other performance criteria. New nodes can be easily added to the network and discovered via standard protocols.

- TSN can be used for communication between machines as well as from machines to enterprise systems. Communication between mission-critical TSN-based systems and existing non-critical Ethernet-based systems can be achieved without any modifications to the non-critical network infrastructure.

Overall system costs are significantly reduced when based on standard technology. Consumer choice and competition result in lower device prices. Research, development, and maintenance costs are all driven down when customers can focus on one standard technology rather than a number of proprietary solutions.

Converged Network

TSN enables the convergence of networks and systems that were previously kept separate for reasons of operational integrity, real-time performance, safety or security. Breaking down communication barriers between critical and non-critical systems is a foundational concept of the IIoT and Industry 4.0.

- Different traffic classes can coexist on the network with no impact on higher criticality level traffic from traffic with lower priority. TSN allows critical messages to be sent over the same communication line as all other Ethernet traffic, without disturbance or delay.
- Even the most challenging motion control and safety applications can be converged on Ethernet using TSN. End-to-end latency is guaranteed even under heavy traffic load and standard mechanisms can be used to accelerate message transport for high-speed communication.
- Convergence makes accessing data from industrial systems easier. With more systems on the same network, the task of gathering data from disparate sources is simplified. Data from industrial systems can be sent to enterprise systems over standard Ethernet without the need for gateways.
- New data streams can be added to the network without the risk of disturbing existing traffic and without the need to reconfigure the entire network.
- Higher layer protocols can be combined with TSN, as the technology is implemented entirely at the data link layer (OSI model layer 2). This includes the open and standard Open Platform Communications Unified Architecture (OPC UA) protocol, among others.

Overall system costs are significantly reduced by the convergence of different traffic classes on a single network infrastructure. Hardware and maintenance costs are driven down by the need for fewer devices and cables for networking.

Benefits of Using an SoC FPGA-Based TSN Implementation

Due to its broad scope and wide range of functions, the benefits of TSN can be best realized in SoC FPGA-based devices. For implementing TSN, FPGAs offer a significant advantage over other silicon (e.g. ASIC) due to their inherent feature of re-programmability. By using Intel® FPGAs, industrial equipment manufacturers can be the first to the market, increasing their return of investment and making their products IIoT ready.

Intel FPGAs can be re-programmed to implement new standards in any configuration. When new TSN standards are added or standards are altered, Intel FPGAs can be quickly re-configured to ensure that the latest TSN functionality is supported in the device and end users are benefiting from TSN features to the fullest extent.

There are multiple reasons for implementing TSN on FPGAs for use in industrial systems:

- **Re-programmability:** FPGA-based designs allow industrial equipment manufacturers to support TSN standards as well as legacy Industrial Ethernet protocols. To cope with the changing workloads and evolving standards, the ability to re-program devices is a critical factor. This re-programmability feature enables you to increase the efficiency and expand the capabilities of devices and is the key for product innovation.
- **Workload consolidation and acceleration:** The growth in network traffic is creating data transfer, management, timing, and scaling challenges. Workload consolidation at the edge is essential for Industry 4.0 factories. FPGAs are fundamental for the Ethernet protocol acceleration. By offloading the workloads, accelerating the protocol connectivity and data exchange, FPGAs help to achieve higher performance.
- **I/O flexibility:** FPGAs allow for TSN implementation along with other Industrial Ethernet protocols on one device.
- **Functional safety and security:** As TSN connects previously unconnected critical systems, functional safety and security should also be taken into account. Intel FPGAs, tools, and IP are certified to IEC61508 safety standards. Security solutions such as secure boot, device authentication to the network, secure communication channel for data exchange over the

TSN Features

TSN describes a set of features that have been added to standard Ethernet. The features are defined and published in a number of IEEE 802.1 standard extensions that address topics, such as timing, synchronization, forwarding, queuing, seamless redundancy, and stream reservation. These individual features extend the functionality and Quality of Service (QoS) of Ethernet to enable guaranteed message transmission through switched networks, providing the inherent robustness, reliability, and determinism required for an industrial communication technology.

IEEE 802.1 STANDARDS	FEATURES
802.1AS	Time Synchronization
802.1Qbv	Scheduled Traffic
802.1CB	Seamless Redundancy, Stream Identification
802.1Qcc	SRP Enhancements
802.1Qbu	Frame Preemption
802.1Qci	Filtering and Policing
802.1Qca	Path Control and Reservation
802.1Qch	Cyclic Queuing and Forwarding
802.1Qcr	Asynchronous Traffic Shaping
802.1Qcp	YANG Model for Bridging
802.1Qcw	YANG Model for Qbv, Qbu, Qci
802.1CBcv	YANG model for CB

For a detailed description of each standard, refer to appendix A.

Table 1. IEEE 802.1 TSN Standards.

The key TSN features that provide guaranteed message delivery timing are time synchronization and traffic scheduling. They are addressed by the 802.1AS and 802.1Qbv standards respectively. All devices participating in the TSN network are synchronized to a global time and are aware of a network schedule that dictates when prioritized messages will be forwarded from each switch. TSN makes use of multiple queues per port at the egress of the switch, where messages are held until a gate opens (at a time slot specified by the schedule) to release queued messages for transmission. The timed release of messages ensures that delays in the network can be deterministically predicted and managed. This allows for the convergence of critical traffic and non-critical traffic on the same network.

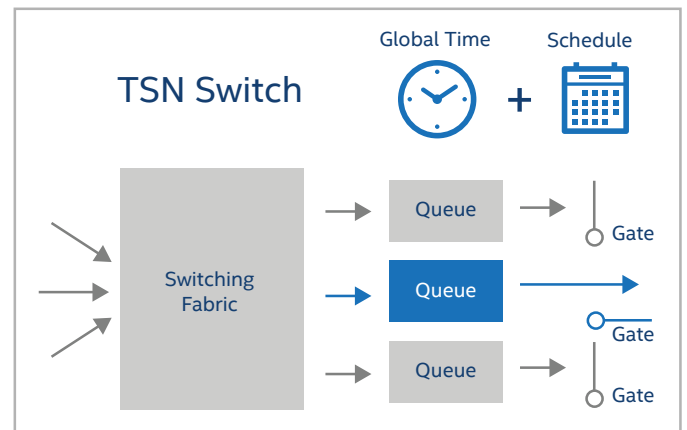


Figure 1. Representation of Scheduled Message Forwarding in TSN As Defined in the IEEE 802.1AS and 802.1Qbv Ethernet Standards

Via the 802.1CB standard, TSN provides a standard method for achieving seamless redundancy for industrial communication over Ethernet. The feature allows for the simultaneous transmission of duplicate message copies across different paths in the network. The first message copy to be received in time is processed whilst the other copies are discarded. This adds another layer of determinism to the delivery of critical messages in converged networks. The preemption feature defined in the TSN 802.1Qbu standard can be used to increase the efficiency of bandwidth use for non-critical messages. In highly converged networks, it could be the case that large low priority frames are delayed by higher priority traffic on the network and dropped. Preemption enables transmission of large frames to be interrupted, sent in smaller fragments and reassembled at the next link. This maximizes bandwidth utilization for all traffic types on the TSN network. Another important benefit of message preemption is the reduction of transmission latency for so-called Express traffic, which can preempt regular Ethernet packets. Especially on lower-speed networks (For example, 10 or 100 megabits per second (Mbps)) carrying large regular Ethernet packets up to 1,500 bytes and more, the latency reduction for Express traffic can be useful for building converged networks.

A crucial feature of TSN is support for open, vendor-independent network configuration. This is achieved through the standardization in IEEE of YANG models for various TSN standards. These can be configured over the NETCONF protocol using encoding formats such as XML or JSON. YANG models for bridging, traffic scheduling, frame preemption, seamless redundancy, and policing ensure that configuration of key TSN features is done according to standard methods. TSN networks can thus be composed of any standard compliant device from any vendor and be configured by any standard compliant network configuration software.

OPC UA over TSN

Of the many higher layer industrial communication protocols that could be combined with TSN, the most logical choice is OPC UA. Much like TSN, OPC UA is an open, standard technology that is vendor independent and useful for a wide range of industrial applications. The combination of OPC UA and TSN therefore provides a complete open, standard, and interoperable solution that fulfils a majority of industrial communication requirements.

By representing data in a uniform way, OPC UA enables interoperability between devices that could not previously share data and gives you new insight into a wealth of information. For this reason, it has been adopted and integrated into products by all of the major industrial automation vendors. OPC UA was originally limited to a client or server architecture, however the recently released publish/subscribe (PubSub) extension now enables multicast communication. In combination with TSN, OPC UA PubSub allows data to be sent with precise timing and thus be used for real-time industrial applications.

Importance of System-Level Thinking in TSN

TSN standards address a wide range of functions and their implementation can be similarly broad, encompassing switch IP core, embedded software, standard interfaces, routing algorithms and configuration tools. To ensure the highest levels of TSN performance, a system level solution is required that takes each element into account and provides a seamless interface between them. TTTech takes this approach with its Deterministic Ethernet product portfolio and reference designs for Intel SoC FPGA-based hardware.

TTTech's TSN products are:

- Designed to work seamlessly from end to end
- Standards based and interoperable
- Customizable and configurable

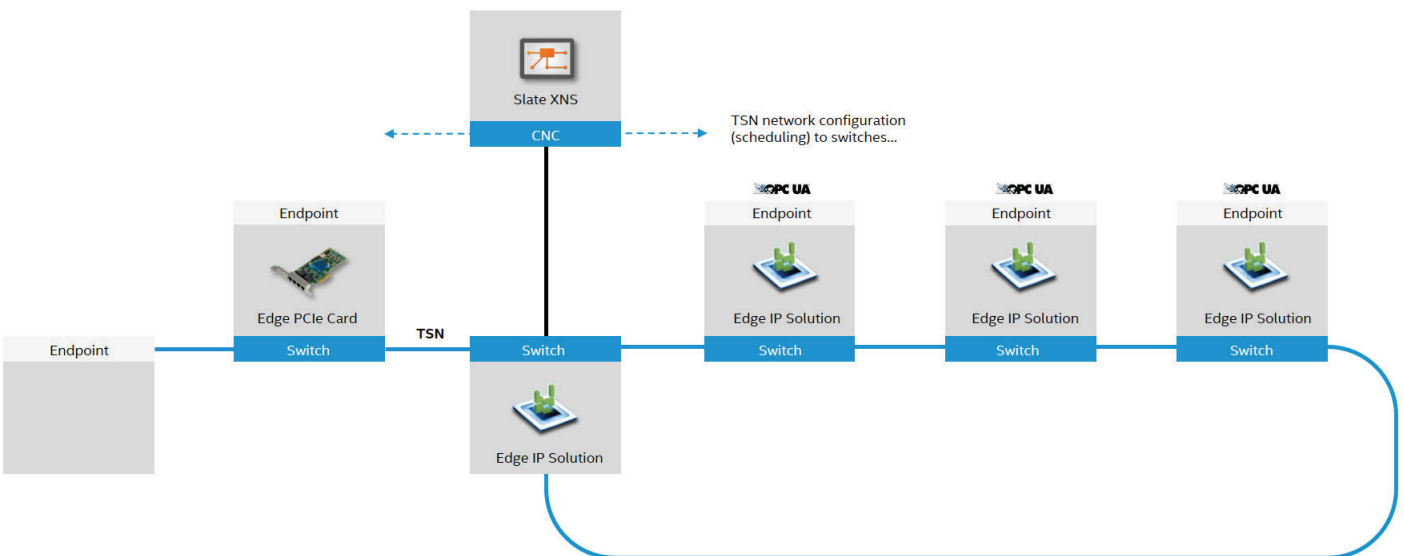


Figure 2. System-Level Design with TTTech TSN Products (Available Now)

^{DE}IP Solution Edge is an IP core and an embedded software package providing TSN switched endpoint functionality for FPGA-based devices. Up to five 10/100/1000 Mbits/s switch ports are available supporting TSN standard features including time synchronization, traffic scheduling, frame preemption, and seamless redundancy and policing. Embedded software includes a NETCONF server, 802.1AS stack, YANG module support, and switch drivers for Linux* operating system.

^{DE}PCIe Card Edge is a network interface card that integrates a defined feature set and software design for ^{DE}IP Solution Edge in an out-of-the-box hardware based on Cyclone® V SoC. The card can be used to quickly add switched endpoint functionality to existing devices with PCI Express* (PCIe*) interfaces.

Slate is a family of network configuration products that can be used to model topologies, create schedules, and deploy configurations for TSN networks. Slate software products include a NETCONF client, REST server, YANG module support, and TTTech's powerful scheduling engine. Because Slate supports standard interfaces, it can be used to schedule and configure any TSN device that is standard compliant. The Slate XNS variant includes a graphical user interface (GUI) for inputting system parameters. In future, the Slate YNS variant will include a PubSub broker for gathering system parameters from OPC UA.

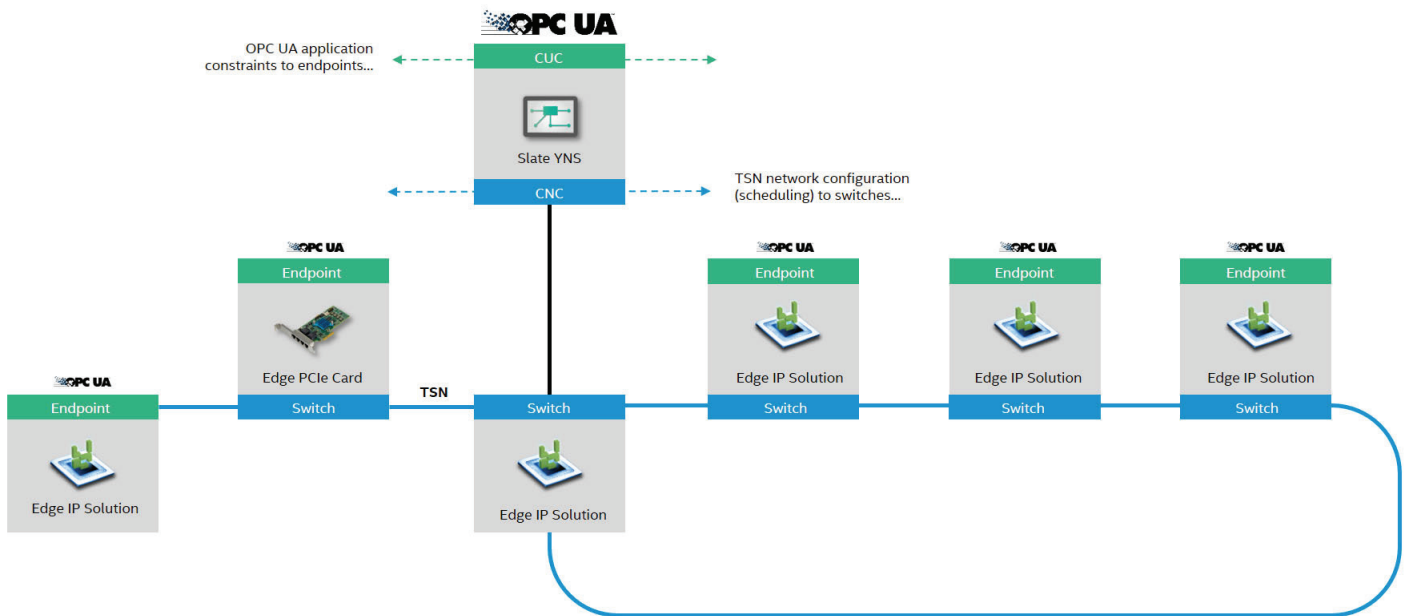


Figure 3. System-Level Design with TTTech TSN Products (Available in 2019 and Beyond)

TSN Customization

FPGA Customization

^{DE}IP Solution Edge enables you to customize TSN and other switch features to create an optimized FPGA integration. An Intel Platform Designer component has been developed for ^{DE}IP Solution Edge that makes it easy to integrate TSN IP features in Intel FPGAs. The component allows you to customize various switch functions in the Intel Quartus® Prime software, saving significant time and effort in the FPGA design process by automatically generating interconnect logic to connect IP functions and subsystems.

The IP can be customized to enable or disable switch features, scale switch features, and choose from various media-independent interfaces. IP features can be customized to fit a specific FPGA device's resource capacity and device requirements. Depending on these constraints, parameters can be set to provide fewer ports or a limited subset of TSN features.

For example, the tool can be used to define:

- Ports that can be used to forward messages that are time stamped according to the synchronized time (IEEE 802.1AS)
- Ports that can be used to forward critical messages according to the network schedule (IEEE 802.1Qbv)
- Ports that can be used to forward preemptable non-critical frames (IEEE 802.1Qbu)
- The number of priority queues with up to eight queues per port (IEEE 802.1Q)
- Ports that can be used to forward messages cut-through. For example, forward messages before the whole frame has been received
- Media-independent interfaces (MII, GMII, RMII, RGMII, SGMII) to the physical layer for each port

Linux Customization

DEIP Solution Edge enables you to customize a software image for their device. Included with the product is a Yocto Project*-build system that can be used to create tailored Linux images for embedded devices. Yocto Project allows you to define the embedded CPU architecture and add application-relevant Linux features. Standard tools in Yocto Project can be used to configure various switch functions, such as virtual local area networks (VLANs), forwarding, and port states.

DEIP Solution Edge is delivered with reference Linux architectures for Intel FPGA-based evaluation hardware that can also be used as a foundation for software customization. The Linux switch driver has been optimized for Kernel 4.9 LTSI, which gives long-term stability for implementing industrial devices.

TSN Configuration

Remote Configuration

After a switch device is set up, it usually needs to be configured as part of a wider network of devices. Remote configurability offers you another layer of flexibility when setting up TSN networks. DEIP Solution Edge enables you to remotely configure TSN and other switch features in networked devices. The product supports a NETCONF network configuration protocol server and YANG is used as a data modelling language over NETCONF. The YANG models supported in the DEIP Solution Edge can be used to configure various switch functions, such as VLANs, forwarding, and port states.

- IEEE 802.1Qcp YANG Model for Bridging enables communication of basic bridging configuration data between bridges in combination with NETCONF
- IEEE 802.1Qcw YANG Model for Qbv, Qbu, Qci extends capabilities of 802.1Qcp to communication of scheduling, preemption, and policing configuration data between bridges in combination with NETCONF
- IEEE 802.1CBcv YANG Model for CB extends the capabilities of 802.1Qcp to communication of redundancy configuration data between bridges in combination with NETCONF

Network Scheduling

The network schedule is the mechanism in TSN that triggers the forwarding of queued messages at specific timepoints and ensures that there are no unexpected delays or collisions in the network. Scheduling takes the latency, jitter, and buffer requirements of applications into account and delivers guaranteed deterministic communication of data across the network.

The IEEE 802.1Qbv standard specifies the method for defining the transmission time of a message from an end node and the forwarding times from each subsequent switch so that the message reaches a receiver within a specified time window. Each device in the network is synchronized to the same clock using the IEEE 802.1AS protocol.

Although a TSN network schedule using the 802.1Qbv standard could be written by hand, this soon becomes impractical for even slightly complex networks and for networks where devices are added incrementally. For this reason, TTTech offers the Slate network configuration software that takes application parameters as input, calculates a matching TSN network schedule, and creates configuration data, ready for deployment to the network switches via standard NETCONF/YANG interfaces. The Slate software solves the complex problem and removes the burden of network calculus and testing from network engineers.

The Slate XNS software provides a GUI that is used to model network topologies and input system parameters. In future releases, network topologies will be discovered automatically, and system parameters will be provided directly from the OPC UA application, streamlining the design process even more.

Network Scheduling with OPC UA

Combining OPC UA with TSN not only brings benefits of open, standard data exchange in real time. It also enables a standard method for configuring TSN networks online and in a dynamic way. This does not require you to input any system parameters for the scheduler as these are all taken from the OPC UA application parameters within each device. A broker mechanism as defined by the OPC Foundation provides an interface between OPC UA applications and TSN scheduling software.

TTTech is developing Slate YNS, a network configuration software combining a powerful scheduling engine with an OPC UA PubSub broker. This will translate OPC UA device parameters into TSN schedule requirements and then feedback the schedule to the devices once it is calculated. The Slate YNS software can be hosted on an industrial PC, smart switch, or server.

Summary

TSN comprises a set of IEEE Ethernet standards that address a range of use cases and scale from individual components up to large complex networks. Due to this broad scope, it is important that TSN implementations are planned and executed with a comprehensive system view, rather than focusing only on one feature or component. To fully deliver the benefits of TSN which derive from openness and convergence, it is also important that TSN implementations are designed to be standard compliant and offer standard interfaces.

For these reasons, TTTech offers a product portfolio that covers all elements of TSN implementation that is designed in accordance to the IEEE standards. This ensures that TTTech products are interoperable with other standard compliant TSN products from other vendors when integrated into wider systems. In combination with Intel FPGA technology, TTTech products deliver the flexibility through customization and configurability to develop optimized TSN devices and systems for your industrial application.

Appendix A

802.1AS

Timing and synchronization are vital mechanisms for achieving deterministic communication. 802.1AS is a profile of the IEEE 1588 PTP synchronization protocol that enables synchronization compatibility between different TSN devices. This lays the foundation for the scheduling of traffic through each participating network device. 802.1ASrev is also being defined to add support for fault tolerance and multiple active synchronization masters.

802.1Qbv

Scheduling of traffic is a core concept in TSN. Based on the shared global time provided by 802.1AS, a schedule is created and distributed between participating network devices. 802.1Qbv defines the mechanisms for controlling the flow of queued traffic through gates at the egress of a TSN switch. The transmission of messages from these queues is executed during scheduled time windows. Other queues will typically be blocked from transmission during these time windows, therefore removing the chance of scheduled traffic being impeded by non-scheduled traffic. This means that the delay through each switch is deterministic and that message latency through a network of TSN-enabled components can be guaranteed.

802.1Qbu

While the 802.1Qbv mechanisms protect critical messages against interference from other network traffic, it does not necessarily result in optimal bandwidth usage or minimal communication latency. Where these factors are important, the pre-emption mechanism defined in 802.1Qbu can be used. 802.1Qbu allows the transmission of standard Ethernet or jumbo frames to be interrupted in order to allow the transmission of high-priority frames, and then resumed afterwards without discarding the previously transmitted piece of the interrupted message.

802.1CB

Redundancy management implemented in 802.1CB follows similar approaches known from High-Availability Seamless Redundancy (HSR) - IEC 62439-3 Clause 5 and Parallel Redundancy Protocol (PRP) - IEC 62439-3 Clause 4. To increase availability, redundant copies of the same messages are communicated in parallel over disjoint paths through the network. The 802.1Qca standard for Path Control and Reservation defines how such paths can be set up. The redundancy management mechanism then combines these redundant messages to generate a single stream of information to the receiver(s).

802.1Qcc

The enhancements to Stream Reservation Protocol (802.1Qat) include support for more streams, configurable stream reservation classes and streams, better description of stream characteristics, support for Layer 3 streaming, deterministic stream reservation convergence, and User Network Interface (UNI) for routing and reservations. 802.1Qcc supports offline and/or online configuration of TSN network scheduling.

802.1Qci

Protects against faulty and/or malicious endpoints and switches. Isolates faults to specific regions in the network. It works at the ingress of the switch (forwarding engine) in order to protect the outgoing queues from flooding with frames.

802.1Qca

Protects against faulty and/or malicious endpoints and switches. Isolates faults to specific regions in the network.

802.1Qch

Defines cycles for forwarding queued traffic using 802.1Qci to assign buffers and 802.1Qbv to shape traffic.

IEEE 802.1Qcr

Provides bounded latency and jitter (lower performance levels) without time synchronization.



[†] Tests measure performance of components on a particular test, in specific systems. Differences in hardware, software, or configuration will affect actual performance. Consult other sources of information to evaluate performance as you consider your purchase. For more complete information about performance and benchmark results, visit www.intel.com/benchmarks.

OpenCL and the OpenCL logo are trademarks of Apple Inc. used by permission by Khronos.

© Intel Corporation. All rights reserved. Intel, the Intel logo, the Intel Inside mark and logo, Altera, Arria and Stratix words and logos are trademarks of Intel Corporation or its subsidiaries in the U.S. and/or other countries. Intel reserves the right to make changes to any products and services at any time without notice. Intel assumes no responsibility or liability arising out of the application or use of any information, product, or service described herein except as expressly agreed to in writing by Intel. Intel customers are advised to obtain the latest version of device specifications before relying on any published information and before placing orders for products or services. Other marks and brands may be claimed as the property of others.