

Time-sensitive networking for industrial automation



Pekka Varis
Senior Member of Technical Staff

Thomas Leyrer
Distinguished Member of Technical Staff

Texas Instruments

Introduction

Time-sensitive networking (TSN) is an Ethernet extension defined by the Institute of Electrical and Electronic Engineers (IEEE) designed to make Ethernet-based networks more deterministic. Industries like automotive, industrial and performance audio use real-time communication with multiple network devices and will benefit from the TSN standard.

The consumer and enterprise world of Ethernet and wireless Ethernet communication is bandwidth oriented. For example, while browsing the Internet you accept a varying amount of delay before video playback starts. Although there is a preference for quick interaction, for the average user it is acceptable if one out of 100 clicks perform an order of magnitude worse. However, if a video is bad quality or even halted the typical consumer will be frustrated.

Even infrequent delays are unacceptable in control systems such as those inside automobiles, production lines or concert halls. The most important aspects for these systems are latency and jitter or variation in the latency of control data through the network. The maximum time a packet takes to reach the destination in the system defines the communication cycle or control frequency in the network.

Table 1 gives an overview of network parameters for certain application examples. The network size and topology can be either fixed (to a certain application) or variable. The Internet is the worst-case example when it comes to the number of nodes and the route that a packet takes through the network. Latency is in the seconds range and jitter is very high when you repeat a packet transfer over the Internet.

By contrast, real-time deterministic Ethernet communication typically limits the number of devices connected to the network. Using a machine tool as an example of an embedded product, the number of motors connected with Ethernet to a single piece of control hardware is less than 100. New motor control parameters are exchanged

every 250 μ s. This fixed and pre-engineered setup requires deterministic real-time Ethernet with short cycle times and high-precision clock distribution.

Production systems in a modern factory are fully connected using real-time Ethernet. **Figure 1** shows various control systems in a production cell.

Network feature	Automotive radar	Machine tool	Professional audio	Consumer video
Scale	4 sensors	64 axes	20 speakers	1 screen
Bandwidth	1 Gb	100 Mb	100 Mb	100 Mb
Jitter	20 ns	100 ns	10 ns	100 ms
Latency	1 ms	100 μ s	10 μ s	1 s
Cycle time	10 ms	<1 ms	Stream	Burst
Time synchronized	Yes	Yes	Yes	No
Topology	Star	Line, ring	Star, line	Point to point

Table 1. Network parameters for various application examples.

A time-sensitive network is a key technology with which to connect various control systems in real time.

Although the requirements for control systems are different in terms of scale, cycle time and accuracy, they can use the same communication interface to transfer data deterministically. Many sensors and actuators are deployed inside the control system. They either connect directly to real-time Ethernet or connect to a concentrator in the real-time Ethernet network using serial point-to-point connections. TSN with its deterministic networking performance is a good fit here, in the “field” level of a manufacturing floor.

The high number of connected sensors and actuators of the industrial control system is one of the key challenges for TSN network configuration.

industry. The only challengers to its dominance have been more sophisticated technologies in the areas of determinism and quality of service such as Asynchronous transfer mode (ATM), token ring and RapidIO®. TSN is a set of roughly 12 IEEE 802 standards aimed at addressing determinism and quality of service without compromising the strengths of Ethernet such as interoperability. Most of the above mentioned TSN standards are now included in the main Ethernet specification IEEE 802.1Q-2018.

TSN is a local area network (LAN)-level solution that can work with non-TSN Ethernet, but timeliness is only guaranteed inside the TSN LAN. You can group TSN standards based on what use case it solves: a common view of time, guaranteed maximum latency, or co-existence with background or other traffic.

Like any popular standard, the TSN toolbox of standards is evolving; some of the individual standards like 802.1AS-Revision (Rev) have not yet been approved, and new alternative shapers are being introduced. Because of this on-going evolution, when choosing a solution it is important to consider the upgradability of the solution to support new or changed standards.

As shown in **Figure 2** on the following page, IEEE 802 Ethernet including the TSN features is a layer 2 or data link layer technology. Applications will require an upper-layer protocol such as UDP/IP or PROFINET above TSN (**Figure 4** on page 6).

Industry 4.0 | Production Cell

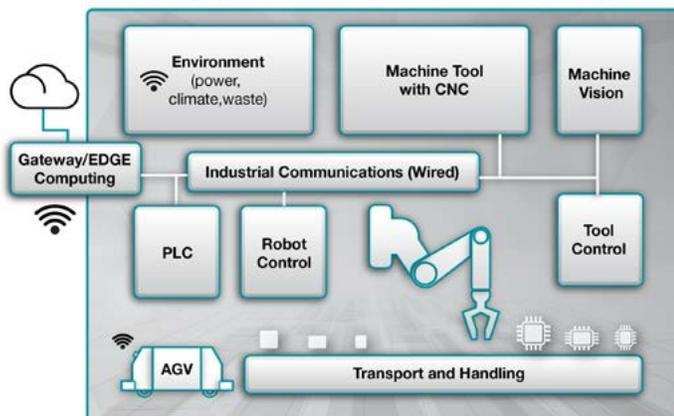


Figure 1: Industry 4.0 production cell with TSN backbone.

TSN provides the capability to classify streams and tag them with deterministic delivery through a network. In an ideal case, a converged network hosts different stream classes ranging from motion control to traffic to the cloud, in one TSN domain. But first, you should understand the basic functions of TSN before answering the question of network engineering.

Principles of TSN

IEEE 802 Ethernet combined with Internet Protocol has been a success story like few others in the tech

802.1AS-Revision – Timing and synchronization for time-sensitive applications

All devices in a network expecting deterministic packet transmission will require a common understanding of time. The clock master or masters distribute time over Ethernet packets to all devices in the network running the Best Master Clock Algorithm (BCMA).

802.1AS-Rev is a tightly defined subset or a profile of 1588v2 precision timing over packet. The additions of 802.1AS-Rev to 802.1AS add support

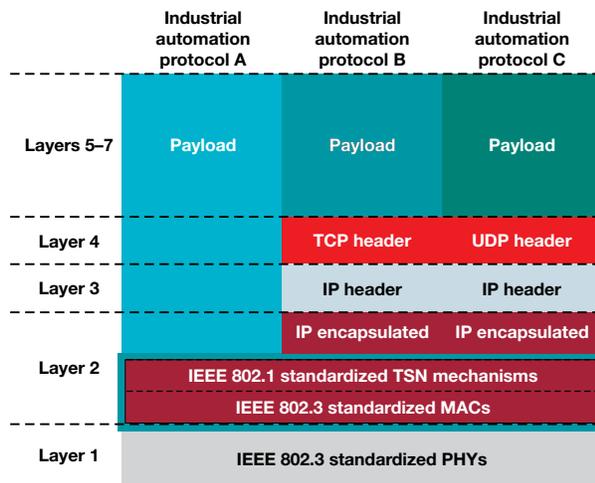


Figure 2: IEEE TSN and the communications stack.

for more than one time domain, support for two timescales (precision time protocol and arbitrary), and optional support for one-step in addition to two-step. The underlying hardware must support the time stamping of transmitted and received packets as close to the wire as possible. For one-step delay reporting, the hardware must also be able to insert a time stamp into the packet as it is being transmitted. For two-step delay reporting, the transmit time stamp is included in a follow-up packet, in some cases increasing the packet load created by time synchronization over packet.

802.1AS-Rev includes peer-to-peer line-delay measurement and bridge-delay calculation. Beyond time stamping, the rest of 802.1AS-Rev is typically implemented in software that runs on a dedicated core. TSN implementation for Texas Instruments (TI) Sitara™ processors supports 802.1AS and 802.1AS-Rev.

802.1Qbv – Enhancements for scheduled traffic (EST)

Time-aware shaper (TAS) makes switches and endpoints aware of the cycle time for real-time traffic. The IEEE802.1Q-2018 name for this feature is Enhancements for Scheduled Traffic (EST). A pergress port scheduler for packets creates a periodic window during which there is no interfering traffic.

TI's TSN implementation for Sitara processors supports TAS. TAS is mostly a hardware feature,

with a software stack configuring the hardware shaper in each bridge port and endpoint talker.

802.1Qbu – Frame preemption and 802.3br – Interspersing Express Traffic (IET)

Ethernet is a store-and-forward network. Once a packet starts to go on the wire it will block the wire from other packets until the end of the packet is reached. For example, a 100-Mbps network and a typical maximum transmission unit (MTU) packet size of 1.5 kB create a head-of-line blocking of about 120 μs (1.5 kB/100 Mbps). Higher-speed links reduce this linearly, but even 1-Gbps networks can have resulting jitter in the tens of microseconds.

To reduce head-of-line blocking issues, IEEE defined frame preemption (802.1Qbu) and the related physical layer standard interspersing express traffic (802.3br). Only express traffic can preempt, providing guaranteed latency for express traffic. The IEEE802.1Q-2018 name for this feature is Interspersed Express Traffic (IET).

Cut-through switching, together with TAS and frame pre-emption, are the basic technologies to reduce worst-case latency—even in a long daisy-chain topology network. TI's TSN implementation for Sitara processors supports cut-through switching, frame pre-emption and interspersing traffic.

802.1Qch – Cyclic queuing and forwarding

Cyclic queuing and forwarding defines completely deterministic delays for all streams. TI's TSN implementation for Sitara processors does not initially support 802.1Qch; instead, we propose using 802.1Qbv (TAS) and a fully managed network to avoid interfering traffic.

802.1CB – Frame replication and elimination for reliability

Typical Ethernet networks rely on higher-level protocols such as Transmission Control Protocol (TCP) retransmission to recover from dropped

Ethernet frames, and the Spanning Tree Protocol (STP) to construct new routes through the network. Both approaches sacrifice a nondeterministic amount of time to deliver the frame.

TSN uses redundancy to guarantee latency even in the presence of single-point failures such as cut cables or broken switches. To proactively guarantee the delivery of frames inside a LAN topology with multiple routes, 802.1CB provides redundancy by selectively duplicating frames at the sender and then discarding the duplicate at the destination. Selecting which frames to duplicate is based on the concept of streams. 802.1CB includes the definition of TSN streams based on layer-2 level headers called “null stream” and higher level called “IP stream”, which is used in other TSN features.

802.1CB is compatible with existing industrial networks where earlier redundancy protocols such as High-Availability Seamless Redundancy (HSR) and Parallel Redundancy Protocol (PRP) provided no latency impact from single-point failures.

802.1Qcc – Stream reservation protocol enhancements and performance

TSN uses three identifying labels: stream ID, stream destination address and traffic class.

- Stream ID is the media access control (MAC) source address concatenated with a 16-bit handle.
- The stream destination address is the MAC destination address concatenated with the virtual LAN (VLAN) ID (802.1Q – VLAN support). Addresses are usually locally managed or multicast addresses.
- VLAN priority bits, typically

using only one or two classes, determine the traffic class.

Stream ID is the unique identifier used by the resource management. Stream destination address and traffic class identify the data path taken.

802.1Qcc supports a centralized configuration model with a centralized user configuration (CUC), as shown in **Figure 3**. A centralized network configuration (CNC) calculates resource allocations and availability and configures the bridges.

Alternative architectures are possible: talkers and listeners talking directly to CNC, or even a fully distributed architecture. A centralized architecture and YANG-based network management protocol (YANG is a data modeling language for network configuration developed by IETF and defined in RFC 7950) like RESTCONF or NETCONF used over a standard secure networking stack like Transport Layer Security (TLS) are likely.

TSN's impact on industrial Ethernet

TSN adds real-time capabilities to standard IEEE Ethernet – capabilities that were once only available on specialized industrial field buses (also called industrial Ethernet). TSN does not remove the need for or replace the protocol used above Ethernet.

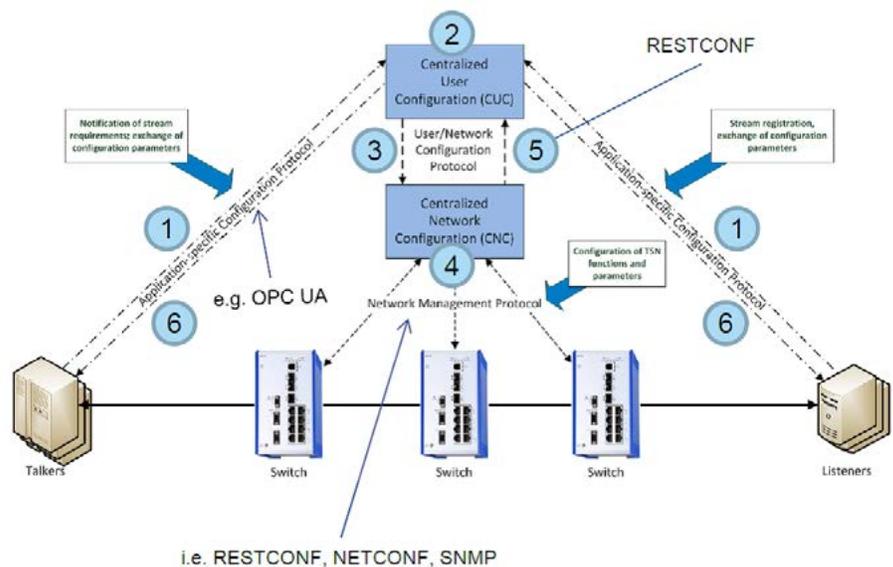


Figure 3: TSN configuration (Source: <https://www.belden.com/blog/industrial-ethernet/what-does-tsn-configuration-look-like-today-and-in-the-future>).

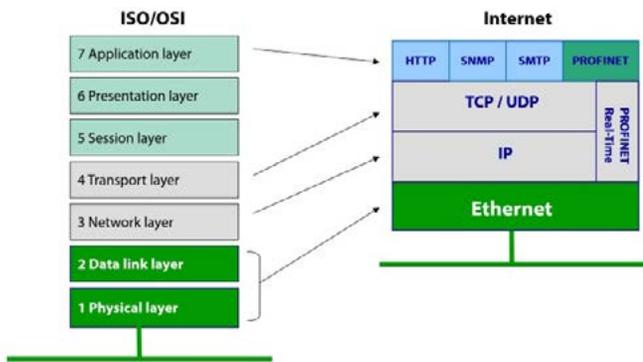


Figure 4: Industrial Ethernet protocols and Ethernet.
 (Source: <http://profinews.com/2017/07/profinet-and-tsn/>).

The interface to software is a good example. For Transmission Control Protocol/Internet Protocol (TCP/IP), a Berkeley software distribution (BSD) socket has become the standard interface with TCP/IP and networking in general, and has proven portable and scalable for a wide set of applications.

A Hypertext Transfer Protocol (HTTP) application works unmodified while reading from a local file or over the Internet. But these sockets are not necessarily relevant interfaces for a protocol that prioritizes solving worst-case latency and exposing the concept of time all the way to the application. For example, the industrial Ethernet protocol PROFINET expects TSN-enabled Ethernet to be just one data-link layer over which to run the protocol.

IEEE TSN defines layer-2 functionality and LAN-level switching, including the concept of time. What it does not define is the software interface with which to configure these hardware features. This means that management software for a switch from vendor A will need to use one application programming interface (API) to another API for vendor B.

A second (and perhaps more unique) area outside the scope of IEEE specifications centers around the concept of latency and the variation or the jitter of latency in the data-path software. As we mentioned, earlier sockets are great but do not even try to address real time or latency.

It is very likely that the API and software architecture around the data path for TSN networks will evolve

over time. The PRU-ICSS-based programmable TSN solution addresses both the software stack portion of latency and configuration and management incrementally. The TSN solution can adapt to a software architecture that requires a very specific buffering mechanism for real time, while supporting mainline Linux[®] networking in parallel. And once Linux leverages more real-time features, a programmable solution can adapt to the new software architecture. TI is working with the community to enable more determinism and an open source configuration API for TSN hardware. For example TAS hardware offload on AM65x and AM64x is supported in Linux kernel version 5.4.

Industrial Ethernet protocols like PROFINET[®] and EtherNet/IP[™] already assume the IEEE Ethernet learning bridge as the underlying switching technology. These protocols can now adapt the extension of TAS and frame pre-emption to use standard TSN hardware for industrial Ethernet. EtherNet/IP uses User Datagram Protocol (UDP) packets for data exchange. PROFINET supports a direct layer-2 buffer model for consumer and provider data supported by PRU-ICSS TSN solution. Both PROFINET and EtherNet/IP are compatible with the TSN switching layer and can benefit from the real-time enhancements.

The IEEE standardization of redundancy protocols were not included in the 802.1Q-2018 standard, and it remains to be seen whether 802.1CB will replace the redundancy protocols of PROFINET and EtherNet/IP, such as media redundancy protocol (MRP) and device level ring (DLR). The transition from existing industrial communication protocols to new standards typically spans many years. During that transition phase, old and new protocols are used concurrently in production systems.

Gateway functions that bridge between existing and new protocols will accelerate the introduction of TSN hardware for industrial communication.

Possible gateway applications include:

- Legacy field buses (DeviceNet[™], PROFIBUS[®])

to new Ethernet protocols (EtherNet/IP at TSN, PROFINET at TSN).

- Industrial Ethernet, which cannot work with TSN in the same network (EtherCAT®, SERCOS® III).
- A gateway between the field level (controller) and the control level (device).
- An IO-Link master gateway to the TSN network.
- The object linking and embedding for process control unified architecture (OPC UA) over TSN-to-cloud interfaces using standard wired and wireless Ethernet.
- A combination of protocols at the field and control levels with an uplink into the cloud.

Sitara processors can realize the gateway functions on this list. Industrial communication subsystems (PRU-ICSS), an additional gigabit switch and the ability to interface to a Wi-Fi® module mean that the Sitara processor can support up to seven communication channels from a single device. Each PRU-ICSS can implement either controller- or device-side industrial Ethernet protocols. Sitara AM65x processors scale up this support to gigabit data rate in up to three sets of PRU-ICSS. At lower power consumption point Sitara AM64x supports up to 5 TSN ports capable of gigabit Ethernet.

Figure 5 shows possible gateway options to deploy at different levels in a production system. The **AM65x Industrial Development Kit** supports six real-time gigabit Ethernet ports, an additional gigabit Ethernet port and Peripheral Component Interface Express (PCIe) generation 3 interface with support for quality of service with 4 virtual channels. AM64x evaluation board supports three real-time gigabit TSN Ethernet ports with another 2 available with an expansion board.

Production System

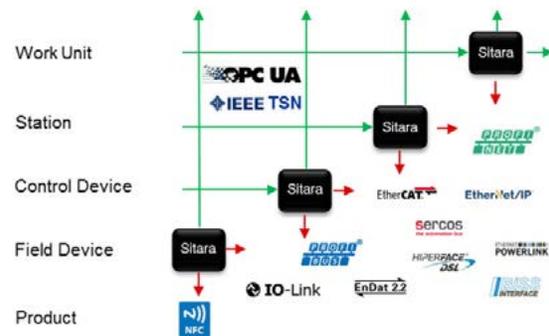


Figure 5: Industry 4.0 gateway.

Conclusion

TSN is moving in the right direction to provide a rich set of capabilities with which to stream packets through a larger industrial network with guaranteed latency. Leading industrial Ethernet organizations are adopting TSN technology and integrating it into existing engineering systems and application profiles.

Higher-layer control systems and other applications outside the factory may also work with central network configuration tools that are independent of the application. Manufacturing networks will continue to evolve with the introduction of TSN features, however traditional industrial Ethernet (PROFINET, EtherCAT, EtherNet/IP, etc) will continue to be supported by Sitara processors. Using Sitara processors with integrated TSN hardware switches can accelerate the transition from 100 Mb to 1 Gb industrial Ethernet.

Related resources

[Industry 4.0 page](#) on TI.com

[Sitara processors - Industrial communication overview](#)

Important Notice: The products and services of Texas Instruments Incorporated and its subsidiaries described herein are sold subject to TI's standard terms and conditions of sale. Customers are advised to obtain the most current and complete information about TI products and services before placing orders. TI assumes no liability for applications assistance, customer's applications or product designs, software performance, or infringement of patents. The publication of information regarding any other company's products or services does not constitute TI's approval, warranty or endorsement thereof.

Sitara is a trademark of Texas Instruments Incorporated. All other trademarks are the property of their respective owners.

IMPORTANT NOTICE AND DISCLAIMER

TI PROVIDES TECHNICAL AND RELIABILITY DATA (INCLUDING DATASHEETS), DESIGN RESOURCES (INCLUDING REFERENCE DESIGNS), APPLICATION OR OTHER DESIGN ADVICE, WEB TOOLS, SAFETY INFORMATION, AND OTHER RESOURCES "AS IS" AND WITH ALL FAULTS, AND DISCLAIMS ALL WARRANTIES, EXPRESS AND IMPLIED, INCLUDING WITHOUT LIMITATION ANY IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT OF THIRD PARTY INTELLECTUAL PROPERTY RIGHTS.

These resources are intended for skilled developers designing with TI products. You are solely responsible for (1) selecting the appropriate TI products for your application, (2) designing, validating and testing your application, and (3) ensuring your application meets applicable standards, and any other safety, security, or other requirements. These resources are subject to change without notice. TI grants you permission to use these resources only for development of an application that uses the TI products described in the resource. Other reproduction and display of these resources is prohibited. No license is granted to any other TI intellectual property right or to any third party intellectual property right. TI disclaims responsibility for, and you will fully indemnify TI and its representatives against, any claims, damages, costs, losses, and liabilities arising out of your use of these resources.

TI's products are provided subject to TI's Terms of Sale (<https://www.ti.com/legal/termsofsale.html>) or other applicable terms available either on [ti.com](https://www.ti.com) or provided in conjunction with such TI products. TI's provision of these resources does not expand or otherwise alter TI's applicable warranties or warranty disclaimers for TI products.

Mailing Address: Texas Instruments, Post Office Box 655303, Dallas, Texas 75265
Copyright © 2021, Texas Instruments Incorporated