Update from IEEE: On-going work in Audio-Video-Bridging, Precise Time Synchronization and Time-Sensitive Networking

Norm Finn, Cisco Fellow, Cisco Systems
Rudy Klecka, Principal Engineer, Cisco Systems
Oliver Kleineberg, Program Manager, Hirschmann Automation & Control

Presented at the ODVA
2014 Industry Conference & 16th Annual Meeting
March 11-13, 2014
Phoenix, Arizona, USA

Abstract:

Real-Time networks have been successfully utilized in the Industrial Automation application space for several years. More recently, these networks have moved from fieldbus to Ethernet technology. However, these Real-Time networks are often not interoperable. This severely hampers market adoption. Recent developments in IEEE 802 and IEEE 1588 express the promise of reduced latency and jitter for IEEE standardized Ethernet. This paper provides an overview of what the IEEE is currently working on. It also shows how the new technology may tie into the existing Industrial Automation application space and into CIP™ / EtherNet/IP™ in particular.

Keywords:

Ethernet, Time-Sensitive Networks, Real-Time Ethernet, EtherNet/IP, IEEE 802, IEEE 1588

Definition of terms (optional):

ASIC	-	Application-Specific Integrated Circuit
AVB	-	Audio and Video Bridging
CIP	-	Common Industrial Protocol
CSMA/CD	-	Carrier Sense Multiple Access with Collision Detection
IEC	-	International Electrotechnical Commission
IEEE	-	Institute of Electrical and Electronics Engineers
IP	-	Internet Protocol
ISO	-	International Organization for Standardization
OSI	-	Open Systems Interconnection
PRP	-	Parallel Redundancy Protocol
PTP	-	Precision Time Protocol (colloquial name for IEEE 1588)
RSTP	-	Rapid Spanning Tree Protocol
TSN	-	Time-Sensitive Networks

1. Introduction and Motivation

1.1 Dependable Ethernet with Real-Time properties in mission-critical applications

Since its inception about 30 years ago, the development of Ethernet has been a remarkable success story. From its humble beginnings as a technology developing as a “spin-off” from efforts to connect the different branches of the University of Hawaii to Honolulu through a wireless network, Ethernet is now the pre-dominant wire-based communication technology worldwide. It has permeated into office and corporate enterprise networks, it has conquered home networks, server farms and carrier-grade networks. Today, Ethernet is even orbiting the Earth aboard the International Space Station.

Over a decade ago, Industrial Automation first took notice of Ethernet. Until that time, Industrial Automation had based its communication networks mostly on technology that was collectively referred to as “fieldbuses”. While fieldbuses were generally limited in regards to the offered bandwidth, they made up this disadvantage with excellent determinism in the time domain. This was a pre-requisite for the newly developing field of distributed control and the requirement to keep communication reaction times low for tight control loops. With the bus arbitration scheme of CSMA/CD, Ethernet was considered non-deterministic and not suitable for this task. The broad introduction of multi-port Ethernet bridges or switches, as they are colloquially called, and complete adoption of full-duplex links paved the way for the adoption of Ethernet into Industrial Automation. While the challenge of possible non-deterministic packet loss was only moved from collisions on the wire to the queue storage inside the switches, with the advent of wire speed switches the level of determinism achieved was finally up to par with many Industrial Automation requirements.

Today, ruggedized Industrial Ethernet switches satisfy the harsh environmental and communication requirements of many automation applications. And these applications are no longer restricted to the factory shop floor: Industrial Automation has broadened its scope into many other mission-critical environments, e.g. Power Transmission and Distribution, traffic control systems or medical applications. One of the most advanced Ethernet networks today is used to control the Large Hadron Collider particle accelerator of CERN near Geneva [1]. Recently, car manufacturers have started to investigate Ethernet technology for use in vehicle and first models from different manufacturers utilizing Ethernet are now reaching the market.

While Ethernet has reached a substantial level of maturity, its development is far from over. As the requirements on end to end latency, jitter, access protection, bandwidth reservation, resilience and time synchronization, especially from in vehicle networks and industrial automation, rises due to the increased requirements from applications, the major success factor of Ethernet is unveiled: flexibility.

Because of the inherent design of Ethernet to be adaptable to different requirements like physical media or speed, Ethernet was always able to scale effortlessly with increased application demands. Supported by the main standardization body of IEEE 802, Ethernet was always able to provide what industry and customers required. With the recent developments in Audio and Video Bridging, Time-sensitive Networks and precise time synchronization, IEEE 802 standardized Ethernet is about to take the next step into the most demanding networks for safety-critical industrial automation and control as well as automotive in vehicle control networks.

These developments will have a profound impact on the services Ethernet as an ISO-OSI Layer 2 technology can offer for technologies based on it. The recent developments in IEEE 802 and IEEE 1588 draft standards will have a substantial impact on CIP technology in the years to come and technologies like EtherNet/IP will be the main beneficiary.

1.2 Paper outline

The paper is organized in the following fashion: Chapter 1 describes the motivation for this paper and leads into the description of the application space. Chapter 2 continues and broadens the application space introduction and gives an overview of the current state of the art technology and how this technology applies to CIP. Chapter 3 then describes the work currently done at IEEE working groups to improve the state of the art and, subsequently, the technology CIP utilizes. Chapter 4 concludes the paper with a summary of how CIP can profit in the future and what steps need to be taken.

2. Ethernet in Automation

2.1 Current state of the technology

Today, the vast majority of Ethernet devices used in Industrial Automation are standard Ethernet end devices and bridges according to IEEE 802.3 and IEEE 802.1. The main difference between these devices and e.g. the switches used in office environments is the ability to withstand adverse environmental conditions. Depending on the use case for an automation switch, it can be hardened against electromagnetic interference, shock resistant or ingress protected with rating IP67 for outdoor use e.g. in deserts or on oil platforms.

When Industrial Automation first considered using Ethernet for communication, already several major vendors and/or vendor organizations had established their own, proprietary fieldbus solutions. Many of these solutions were standardized in IEC 61158 [2] and IEC 61784 [3]. These solutions not only covered the communication aspect of automation, but also device data modeling, monitoring, administration and configuration. Many of these technologies were, at that time, adapted from fieldbus to Ethernet usage and many technologies were developed further into large industrial consortia consisting of many vendors standardizing the technology for general usage. From this fact, today’s diversified landscape of Industrial Ethernet architectures developed. The underlying technology of an Industrial Ethernet specification is usually based on Ethernet according to IEEE 802.3 and bridging according to IEEE 802.1. At some point, depending on the technology design, specific protocol technology is “grafted” on top of IEEE 802. This point at which standardized IEEE 802 technology stops and protocol-specific technology starts greatly differs: While e.g. Profinet I/Oand EtherNet/IP mainly use the full spectrum of ISO-OSI Layer 1 and Layer 2 technology offered by IEEE 802 and EtherNet/IP extends this to Layers 3 and 4, EtherCat only uses the Physical Layer specification of IEEE 802.3 and implements a distributed shift register on the higher communication layers.

The main reason why not all vendors solely rely on IEEE 802 technology are increasing real-time requirements. IEEE 802 technology, due to its inherent design, is only able to provide a certain level of real-time behavior. The requirements fulfilled today by technologies like SERCOS III, EtherCat or Profinet IRT are–far beyond what standardized IEEE 802 technology can achieve. This, however, comes at a price: The specific real-time technologies are usually not compatible and not interoperable between each other. While being individually technically sound, this islanding of non-IEEE technologies and markets has severely impacted the adoption rate of real-time Ethernet in the market and feasibility from a sustainability point of view.

This is where future IEEE 1588 and IEEE 802 technology comes into play. It offers the promise of a fully standardized Layer 2 Real-Time Ethernet which offers full interoperability, while still offering industrial organizations like ODVA the opportunity to maintain and further develop their proven automation technology with their upper layer protocol (i.e. CIP). IEEE 802.1 AVB and TSN technology is the way out of the deadend street that is manufacturer-specific real-time Ethernet technology on ISO-OSI layers 1 and 2. The following section will illustrate how IEEE 802 and IEEE 1588 technology fit into the CIP model, based on how Ethernet technology is utilized even today.

2.2 IEEE 802 and 1588 technology in relation to EtherNet/IP

Both the current and future IEEE 802 technology is visible in Figure 1 in the blocks “Ethernet CSMA/CD” and “Ethernet Physical Layer”. IEEE 1588 is visible in “Ethernet CSMA/CD” and “Internet Protocol”.

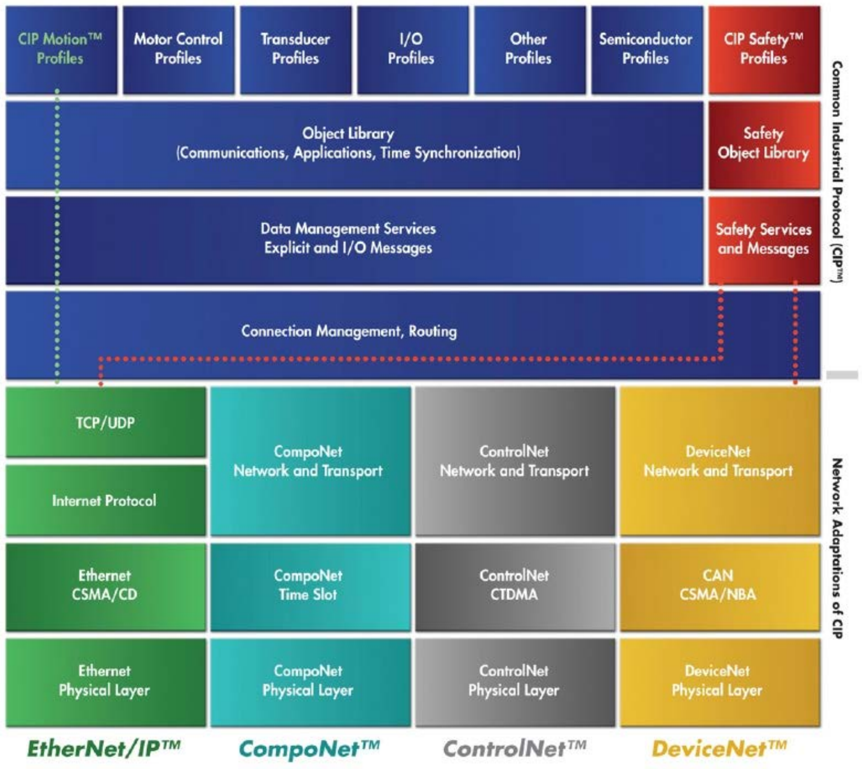

Figure 1 – CIP architecture overview (Source: ODVA)

With this, it is clearly visible that the new technologies developed by IEEE 802 and IEEE 1588 directly augment basic technology that the Common Industrial Protocol utilizes. Therefore, as soon as this new IEEE 802 and 1588 technology is phased in, CIP can immediately take advantage of its new functionality. However, because of the modular approach of IEEE 802 and 1588 technology, CIP is not forced to immediately adapt as soon as the specifications are finished and the technology is introduced but can phase in improvements step by step and at a timeframe of their choice. This is the inherent advantage of this layered approach by CIP and the systemic and modular approach of IEEE in their technology space: Adoption can be introduced as improvements over the existing technology while not invalidating existing and proven concepts – a fact that is of high importance in the Industrial Automation application space.

However, the improvements to the basic communication technology IEEE is working on at the moment are so substantial that, once the technology is market-ready, quick adoption is anticipated. With CIP and EtherNet/IP being based completely on IEEE 802 and IEEE 1588 technology, the improvements made to the basic technology can be utilized with justifiable effort. It is anticipated that the main efforts in making this available to CIP is e.g. the introduction of new EtherNet/IP specification enhancements to allow access from the CIP level to communication stream, scheduler and redundancy/topology status, control and configuration. Because of its layered technology setup, CIP will be in a unique situation to be able to adopt and make the most out of the future technology improvements at their disposal. These technology improvements are described in more detail in the following section of the paper.

3. Technology Improvements in IEEE 802 and IEEE 1588

3.1 Real-Time optimizations: Scheduling

A major improvement in the current TSN standards over the “Generation 1” AVB is the introduction of “Time-Triggered” queues with P802.1Qbv “Enhancements for Scheduled Traffic” [4]. The main point of this enhancement is to rectify issues with the AVB credit based shaper where there could be packet loss due to congestion. Although there are multiple methods of shaping currently under discussion in the TSN group, the consensus is to standardize mechanisms that will be able to guarantee bandwidth and latency requirements necessary for industrial control loops and also guarantee zero packet loss due to congestion.

There are several types of time-based shapers currently under consideration in TSN, any or all of which might be acceptable. As with the Generation 1 AVB Credit-Based Shaper, these new mechanisms are designed for implementation alongside the standard eight IEEE 802.1Q Class of Service priorities but they will be handled at an elevated priority level. Considering that the industrial space has such a wide range of applications, where trade-offs between minimizing latency, being more dynamic, ease of configuration, etc. span the range, it is conceivable that multiple methods might be employed – even in the same network over the same wires. Therefore, it is likely that more than one method of shaping traffic will ultimately be standardized.

Another key criterion of the industrial space is the need for guaranteed latencies even in a “converged” network where both high priority control and best effort traffic reside on the same wire. In such networks it is often the case that the best effort traffic can interfere with the control loop traffic even when proper prioritization is performed, as even the highest priority traffic might have to wait behind at least one lower priority packet at each hop. To help mitigate this problem the IEEE is active in defining a method for preemption with IEEE 802.1Qbu “Frame Preemption” [5] and, in combination with the physical layer 802.3 group, with P802.3br [6] “Interspersing Express Traffic Task Force”. The combination of those two efforts will allow lower latency in some networks as well as provide for better throughput of lower priority traffic in others.

3.2 High availability – Improvements for deterministic redundancy

One of the major requirements for industrial control networks always has been the need for increased resilience over e.g. home or office networks. While a sudden and unplanned loss of network connectivity in the office space usually results in a major inconvenience and possibly in productivity loss, the effects on industrial control networks can be more severe. Therefore, industrial control networks have always been designed with redundancy in mind to prevent network outages by utilizing redundant links as well as redundant networks and redundancy administration protocols. Because of the redundancy dilemma – an Ethernet network that implements redundancy e.g. through additional links is not possible because of the network’s inherent broadcast characteristic and the resulting loops – redundancy control protocols have to be utilized.

The range of redundancy control protocols used spans from the IEEE 802.1-specified general purpose Rapid Spanning Tree Protocol (RSTP) to industrial-specific ring protocols such as the Media Redundancy Protocol [7] (MRP) and the Device Level Ring [8] (DLR). These protocols share a common attribute: the network reconfiguration time in case of a fault. As industrial control loops utilizing the network can only sustain a certain amount of “dead time” due to connectivity loss, the specific upper bound reconfiguration time of a protocol needs to be known a priori. But there are also applications that cannot tolerate any reconfiguration time at all. To address this, in the recent past, seamless redundancy protocols have been specified that do not experience any reconfiguration time at all. One example of this technology is the Parallel Redundancy Protocol [9] (PRP), which uses a Dual LAN approach. In the recent past, ODVA has recently specified DLR and has adopted PRP into the EtherNet/IP specification to address these sophisticated application needs.

With the introduction of standardized real-time networks for automation and control, the necessity for redundancy arises again. Therefore, IEEE 802.1 builds on existing technology and experiences from industrial control protocols and integrates redundancy into the whole concept from the very start of the technology specification. Redundancy is no longer an “afterthought” as it has been to a certain degree in the past, but part of the systemic approach itself.
The IEEE 802.1 has two active projects that address the two important aspects of redundancy application: IEEE P802.1Qca – Path Control and Reservation – deals with communication path management for TSN streams. IEEE P802.1CB specifies frame replication and elimination to enable redundant datagram transmission and reception.

With IEEE P802.1Qca, past concerns from the industrial control community and from ODVA members in specific have been addressed. In the past, IEEE 802.1 and especially the Shortest Path Bridging protocol (IEEE 802.1aq) has been criticized for being not feasible for use in engineered networks because it lacks the capability for fixed configured communication paths through the network topology. IEEE P802.1Qca allows for fully engineered and redundant communication paths, even if these paths are not shortest paths through the topology. Further, Qca also allows to exempt these paths from network reconfiguration when a fault in the topology occurs – allowing for 100% deterministic behavior.

This ability to fully engineer the network topology and behavior is augmented by IEEE P802.1CB, which allows to specify frame replication and elimination for redundant transmission in bridges and end stations. When combining

IEEE P802.1Qca and IEEE P802.1CB, multiple resilient paths that exhibit no reconfiguration time can be engineered through any arbitrary network topology – specifically designed for use with TSN real-time streams. This systemic approach allows TSN to address both the most demanding latency and resilience requirements at the same time.

3.3 Time synchronization: The new revision of IEEE 1588

There is new work in the area of timing in both IEEE 1588 and in the TSN PTP profile named 802.1AS. The base 1588 standard work is underway in the P1588 Working Group with several areas likely of interest to ODVA, namely improvements in security, robustness, elimination of layer violations, and a goal to reduce "Profile Proliferation".

In the TSN group, there has been significant work in under the P802.1ASbt PAR with a specific charter to address Industrial Control. There have been discussions in the area of supporting "brownfield" industrial applications [10] including discussion of supporting End-to-End delay mechanisms as well as support for non-time aware and/or non-TSN elements in the path. This should make support of CIP Sync in TSN networks a much more seamless process.

4 The opportunities for CIP

4.1 Standardized L2 transport and redundancy mechanisms augment EtherNet/IP

After having reviewed the current state of the technology future for IEEE 802.3 Ethernet, IEEE 802.1 Bridging and IEEE 1588 time synchronization, the question arises: What does CIP and EtherNet/IP gain from all these technical advances? Is this a technology that can be safely ignored, relying on established principles? Or is it to be fully embraced and utilized to the full extent as soon as possible? Having already briefly touched the subject in section 2.2. of this paper, and with a market space as diverse as industrial and mission-critical automation, there is no universal answer to this question .

.

Fundamentally, the new IEEE network technologies bring major improvements in end to end latency and jitter in the time domain in combination with resilience against faults. This makes the new technology destined for usage in the more demanding application spaces, e.g. safety and/or motion control. Also, new application spaces become feasible for Ethernet, e.g. drive-by-wire or fly-by-wire in automotive or avionics applications. More new application fields are sure to follow.

Proliferation of the new IEEE technologies is aided by the basic fact that, as this is IEEE standardized technology, it will become available in many Application-Specific Integrated Circuits (ASICs) and software stacks in the future. Thus, it will be relatively easy and financially feasible to integrate into new products, especially if high volume chips can be re-used that will be developed for automotive Ethernet usage. In addition to this, all these products will be interoperable and work on a single ISO-OSI Layer 2 network. Integration of the new management entities into CIP will be on the same level of effort as integrating other protocols. This has been proven feasible in the past by e.g. integrating the RSTP and the PRP management objects into the CIP specification.

Another aspect is the seamless transition to higher network speeds. While automation networks traditionally have higher requirements towards low jitter end-to-end latency and less towards available bandwidth, this is currently in a state of change. New applications like visual goods inspection and quality assurance, video surveillance and increased data transmission needs of service-oriented architectures raise the bar for bandwidth requirements in the industrial/mission-critical application space. All those services need to run over the same network. In the future, the ability to scale network bandwidth seamlessly while still guaranteeing minimum latency for those application that require it will be the hallmark for TSN enabled automation networks.

While adoption of the technology will obviously start in the demanding application spaces where TSN will augment existing CIP/Sync mechanisms, it will not stop there. The new IEEE technologies have the potential to raise the baseline for what is expected of a layer 2 network. Where today a layer 2 network is expected to work with bridges over specific network media at specific transmission speeds, the baseline features of the future layer 2 Ethernet network will be much higher. This enables and nurtures further development in the specific application spaces. CIP and especially EtherNet/IP will see a strong boost in market feasibility and technology excellence as the technology baseline evolves.

5 Outlook and summary

EtherNet/IP and CIP are in the unique position to having abstracted from the basic layer 2 transmission technologies. This foresightful architectural approach now finds fruition in the fact that new TSN and 1588 technology can be integrated and utilized with ease into the existing architecture – with the exact feature set and individual speed of integration that the market requires. After the widespread introduction of the full duplex bridge/switch, this is the largest and most important evolution of Ethernet technology for the automation space yet – and ODVA with the CIP technology is able to immediately and fully reap the benefits. While the specification process is already ongoing in IEEE, the process is not yet finished. IEEE 802 and 1588 greatly value input, especially from knowledgeable individuals and companies in the targeted application space. Please consider contributing at a future IEEE 802 or IEEE 1588 meeting – it is most likely to have a fundamental impact on the future of ODVA and EtherNet/IP.

References:

[1] The White Rabbit Project: https://espace.cern.ch/be-dep/CO/ICALEPCS%202009/1158%20%20The%20White%20Rabbit%20Project/TUC004_FINAL.pdf

[2] IEC 61158:2010 – Industrial Communication Networks – Fieldbus Specifications

[3] IEC 61784 – Industrial Communication Networks – Profiles

[4] IEEE 802.1 Project "Enhancements for Scheduled Traffic" http://standards.ieee.org/develop/project/802.1Qbv.html

[5] IEEE 802.1 Project "Frame Preemption" http://www.ieee802.org/1/pages/802.1bu.html

[6] IEEE 802.3 PAR "Interspersing Express Traffic Task Force" http://www.ieee802.org/3/DMLT/P802_3br_PAR_030913.pdf

[7] IEC 62439-2:2010 - Media Redundancy Protocol (MRP)

[8] The CIP Networks Library Vol. 2 – EtherNet/IP Adaptation of CIP
Section 9-5 Device Level Ring Protocol
also
IEC 61158-4-2 :2010 Industrial communication networks – Fieldbus specifications
Part 4-2: Data-link layer protocol specification – Type 2 elements
Section 10 Device Level Ring (DLR) protocol

[9] IEC 62439-3:2012 – Parallel Redundancy Protocol (PRP) and High-availability Seamless Redundancy (HSR)

[10] Dan Sexton presentation to IEEE plenary, TSN Group, Dallas, Nov. 2013 http://www.ieee802.org/1/files/public/docs2013/new-tsn-dsexton-adoption-feedback-1113-v01.pdf

****