QoS - Application of TSN to EtherNet/IP Networks

Jordon Woods
Director, Deterministic Ethernet Technology Group
Analog Devices, Inc.

Steve Zuponcic
Technology Manager
Rockwell Automation

Abstract

Deterministic Ethernet is already widely used in industrial control systems. Many Industrial protocols solve the problem of determinism over Ethernet using proprietary layer 2 solutions. The new IEEE 802.1 TSN standards are aimed at the same class of problems encountered in industrial control and promise to supplant proprietary solutions in favor of a standards-based approach. EtherNet/IP has always relied upon standard, commercially available Ethernet technologies to deliver deterministic performance and is therefore well-positioned to leverage these emerging standards.

In this paper we discuss specific use cases and examine how the emerging TSN standards can be applied to EtherNet/IP networks to provide improved determinism and performance. We will also contrast the results expected in a TSN-based network with those of the technologies currently in use.

Keywords

Industrial Ethernet, Industrial Control, Motion Control, EtherNet/IP, CIP Motion, CIP Sync, Time-Sensitive Networking, Audio-Video Bridging, AVB, AVB Gen2, software-defined networking, SDN, deterministic Ethernet.

Definition of terms (optional)

Bridge: A system that includes Media Access Control (MAC) Bridge or Virtual Local Area Network (VLAN) Bridge component functionality and that supports a claim of conformance to Clause 5 of IEEE Std 802.1Q-2014 for system behavior. Note: "Bridge" is often used interchangeably with "switch" in TSN discussions.

Centralized Network Configuration (CNC): A centralized component that configures network resources on behalf of TSN applications (users).

Centralized User Configuration (CUC): A centralized component that discovers and configures application (user) resources in end stations. The CUC exchanges information with the CNC to configure TSN features on behalf of its end stations.

Credit Based Shaper: A prioritization mechanism that operates on one or more outgoing queues in a bridge and guarantees "fairness" on the wire by providing lower priority traffic access to the network when higher priority traffic have exhausted their "credits" during a given frame of time.

End station: A device attached to a local area network (LAN) or metropolitan area network (MAN), which acts as a source of, and/or destination for, traffic carried on the LAN or MAN.

Grandmaster: The time-aware system that contains the best clock, as determined by the best master clock algorithm (BMCA), in the generalized precision time protocol (gPTP) domain.

Listener: The end station that is the destination, receiver, or consumer of a stream.

Stream: A unidirectional flow of data from a Talker to one or more Listeners.

Synchronized time: The synchronized time of an event is the time of that event relative to the grandmaster.

Synchronized time-aware systems: Two time-aware systems are synchronized to a specified uncertainty if they have the same epoch and their measurements of the time of a single event at an arbitrary time differ by no more than that uncertainty.

Talker: The end station that is the source or producer of a stream.

Traffic class: A classification used to expedite transmission of frames generated by critical or time-sensitive services. Traffic classes are numbered from 0 to N-1, where N is the number of outbound queues associated with a given Bridge Port, $1 <= N <= 8$, and each traffic class has a one-to-one correspondence with a specific outbound queue for that port. Traffic class 0 corresponds to non-expedited traffic; nonzero traffic classes correspond to expedited classes of traffic. For a given priority associated with a frame and a given number of traffic classes, a fixed mapping determines which traffic class will be assigned to the frame.

Deterministic Ethernet Overview

Deterministic Ethernet refers to an extended set of capabilities that allow standard Ethernet to be used in real-time, mission critical applications such as Factory Automation, Process Control, and Automobile Networks.

Ethernet has traditionally been a "best-effort" network. To allow Ethernet to be deployed in mission critical applications, it is necessary to add specific features including time synchronization, scheduled traffic, ingress policing, seamless redundancy and others. These features allow network designers to ensure that certain classes of traffic can be delivered on time, every time throughout the entire network topology.

Deterministic Ethernet started in the factory automation market, where large OEMs defined their own methods for adding these capabilities to Ethernet and this has spawned a number of "open standards" that are used in Industrial Ethernet today. ODVA has deployed deterministic Ethernet networks, with time-synchronization and quality of service for critical traffic, for many years.

The IEEE standards organization is currently working to add Time Sensitive Networking (TSN) features to standard 802.1 and 802.3 Ethernet to provide deterministic performance. When this work is completed, it will become practical to deploy standard TCP/IP Ethernet (with the TSN extensions) in real-time, mission critical applications. However, it is unlikely that the existing industrial protocols will be replaced by TSN. It is more likely that they will be adapted to make use of the inherent deterministic capabilities provided by TSN.

The goal behind the new TSN standards is to achieve a truly converged network where all classes of traffic can seamlessly coexist. This would allow mission critical real-time traffic to coexist on the same

network as traditional QoS prioritized traffic and best-effort traffic. A major driving force behind the development of these new TSN standards is the emerging Automotive Ethernet market. However, Industrial Automation market share many common requirements with automotive control applications and can therefore leverage the standard technologies and economies of scale inherent in this very large market.

It is important point to note that the new IEEE standards are based on some of the same basic techniques that have been used in the Industrial Ethernet protocols for many years. However, many organizations have relied upon proprietary layer 2 techniques to achieve determinism. In contrast, ODVA and EtherNet/IP have relied exclusively on widely available, ubiquitous standards from IEEE and other organizations. For this reason, ODVA is uniquely poised to successfully leverage the emerging TSN standards.

TSN and the Distributed Motion and Time Synchronization SIG
ODVA has tasked the Distributed Motion and Time Synchronization SIG with evaluating how to merge these TSN standards into the existing CIP standards. It is a common misconception that TSN is a single standard. In reality, TSN is a set of new standards and enhancements to exiting standards. In other words, TSN is a basket of new Ethernet features including:

  • Time Synchronization
    In the context of TSN, Time synchronization refers to proposed modifications to the existing IEEE P802.1AS standard that was defined for Audio Video Bridging. To understand the proposed changes, first we should describe IEEE 802.1AS.

    IEEE802.1AS is the Audio-Video Bridging (AVB) profile of the IEEE1588 Precision Time Protocol. IEEE1588, and therefore IEEE802.1AS uses a master-slave protocol to synchronize real-time clocks in the nodes of a distributed system that communicates using a network. In simple terms, PTP ensures that every node on the network knows what time it is. It does not specify what a given node does with its knowledge of time. The AVB profile (IEEE802.1AS), has features tailored to the "plug and play requirements" of AVB components. It does not use "transparent clocks" to compensate for bridge latency because transparent clocks violate IEEE802 layering conventions. Instead, each node accepts time information from the best available master clock and produces a slave clock, in a manner similar to a boundary clock, to compensate for latency. By nature, this approach is peer-based. This approach is not compatible with end-to-end transparent clock defined as default IEEE-1588 profile and used in CIP Sync. For this reason, time-bridging mechanisms need to be developed to allow existing technologies and installations to migrate into the newer TSN domains. By developing time gateways, brownfield installations can be included in the larger TSN eco-system and benefit from many of the new features and capabilities that TSN has to offer.

    The End-to-End TC function does not require all nodes to be time aware which is important in brownfield installations. While time accuracy will be compromised if there are non-time-aware nodes, a common understanding of time is maintained among the time-aware nodes. In contrast, the Peer-to-Peer mechanism utilized by IEEE802.1AS requires every node be time-aware. Obviously, this requirement is not practical for brownfield uses cases.
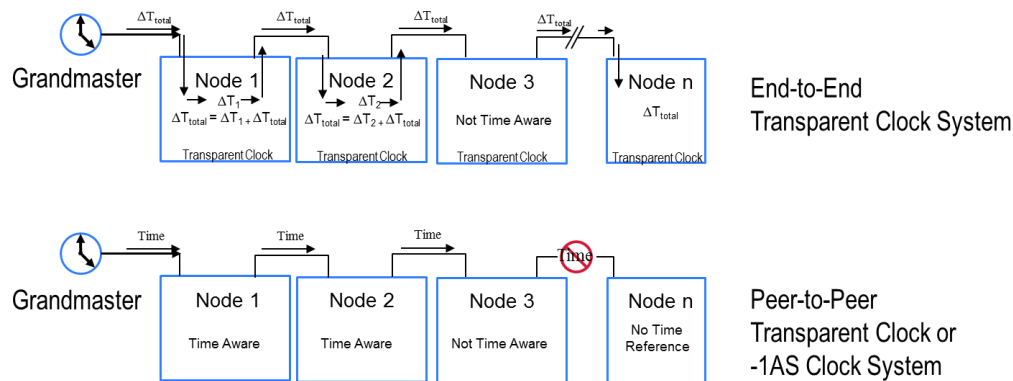
Figure 1 - Peer-to-Peer and End-to-End Clock systems

IEEE802.1AS-REV introduces new features needed for time-sensitive applications. These features include the ability to support a multiple time domains to allow rapid switchover, should a Grandmaster fail, and a more precise measurement of time.

While time synchronization is a foundational component of most TSN features (AVB, ingress policing, scheduled traffic), none of these features require that specific PTP profile be utilized. For instance, ingress policing requires an understanding of time that can come from the default IEEE1588v2 profile, the IEEE802.1AS profile, or some other PTP profile.

- Scheduled Traffic
  Applications implementing control loops over Ethernet (factory automation, robotics, automotive control systems) require delivery of control data at precise times with minimal latency and jitter. Existing priority mechanisms provide for prioritization of traffic or guarantees of bandwidth (AVB), but the time of delivery is unpredictable. IEEE802.1Qbv supports this requirement by divided Ethernet traffic into different classes thus ensuring that, at specific times, only one traffic class (or set of traffic classes) has access to the network. This division, in effect, creates a protected "channel" that is used by that traffic class alone. For clarity, this traffic class *is* scheduled traffic in all cases, and scheduled traffic is given the highest priority on the wire.

As shown below IEEE802.1Qbv accomplishes this division by introducing time-aware "transmission gates". These gates are used to enable separate transmission queues. The Qbv shaper provides a time-based circular schedule which opens and closes the transmission gate at specific times. Note that once a gate is open, existing IEEE802.1Q transmission selection algorithms function normally. A given queue may operate on strict priority (best effort) basis while another may use the credit-based shaper defined for AVB applications. However, it is worth noting that IEEE provides no guidelines regarding which QoS/traffic-shaping mechanism to use in a given situation. Such uses cases are application-specific. It will fall to industry bodies such as Avnu and ODVA to develop TSN profiles for these use cases. This task is vitally important in providing system models and behaviors that are consistent across applications and even within the same application. Without consistent models and behaviors, one CNC may configure a network in one way – while another CNC may configure the same network in a different way.
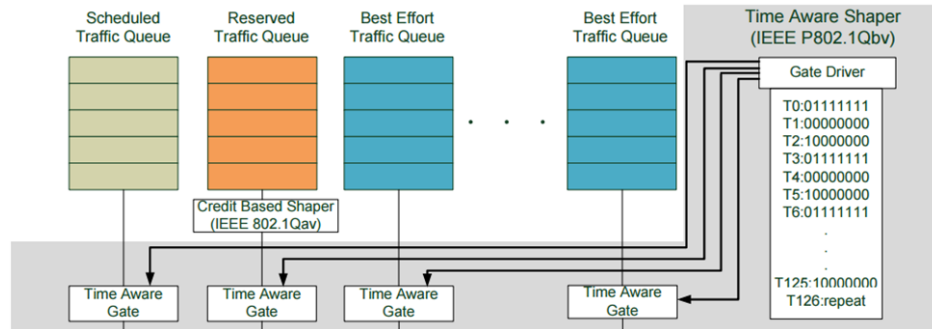
Figure 2 - IEEE802.1Qbv Queuing Structure

- Preemption
  Today's infrastructure components are designed to complete transmission of an entire packet before the next packet can be transmitted, if the initial packet has already been placed in the egress port. Therefore, a 1500 byte packet of lower priority can "hold off" a higher priority packet (~120 us at 100Mb/s wire speed). Preemption (also called Interspersing Express Traffic per IEEE802.3br/IEEE 802.1Qbu) defines a mechanism that allows the switch to stop a transmission in mid-stream in order to allow a higher priority packet to move through the system. Note that only one level of traffic is defined as preemptive.
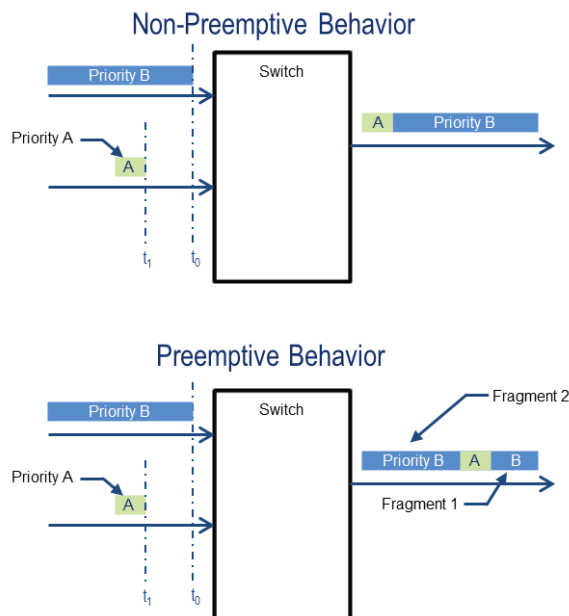
Figure 3 - Preemption

- Seamless Redundancy
  Frame Replication and Elimination for Reliability (IEEE P802.1CB)
  To ensure robust and reliable communication, control systems must be tolerant to packet loss due to congestion, link failures, cable breakage and other faults. To minimize the impact of such faults, P802.1CB aims to send duplicate copies of critical traffic across disjoint paths in the network. If both frames reach their destination, the duplicate copy is discarded. If one copy fails to

reach its destination, the duplicate message is still received, effectively providing seamless redundancy.

To minimize network congestion, packet replication can be selected based upon address/traffic class and path information. Likewise, duplicate frame elimination can be based upon address/traffic class and timing. In other words, only critical traffic need be replicated. Best effort and other traffic tolerant to congestion loss can still be transmitted normally.

- Ingress Policing
Ingress policing generically refers to methods used to prevent traffic overload conditions (e.g., Distributed Denial of Service or DDoS or erroneous delivery) from affecting the receiving node or port. These methods may be used to protect against software bugs on endpoints or switches/bridges but also against hostile devices or attacks. P802.1Qci proposes to provide filtering on a per stream (traffic class) basis by providing an input gate for each stream. These gates would be responsible for passing or blocking a given stream or streams based upon a policing function. Supported functions include a time window (only allowing streams to pass at a certain time), only allowing specific streams to pass on specific ingress ports, a maximum burst size, a leaky bucket algorithm and others. Note: a leaky bucket algorithm checks that streams conform to defined limits on bandwidth and burstiness (a measure of the unevenness or variations in the traffic flow). Thought of another way: each talker has a contract with a respective listener (excess bandwidth, burst sizes, packet sizes, misuse of labels, etc.). The input gate serves to enforce that contract.

- Centralized Configuration
IEEE P802.1Qcc (Stream Reservation Protocol (SRP) Enhancements & Performance Improvements) enhances the existing Stream Reservation Protocol with the addition of a User Network Interface (UNI) which allows for a centralized network configuration (CNC) entity. This CNC can then provide a centralized means for performing network calculus, scheduling and other configuration via a remote management protocol such as NETCONF or RESTCONF. A Centralized User Configuration (CUC) entity communicates to the CNC via a standard API. The CUC may be used to discover end stations, retrieve end station capabilities and user requirements, and configure TSN features in end stations.
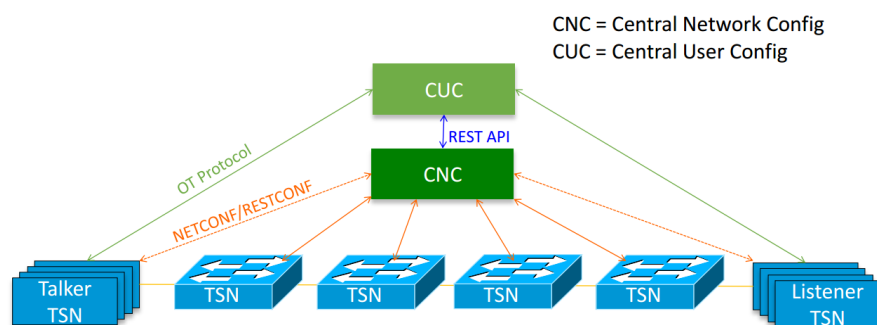

Figure 4 - Network utilization of Centralized Configuration

Industrial Control Overview
Modern industrial networks combine the disciplines of both information technology (IT) and operational technology (OT) to meet the requirements of industrial applications   The applications served by the industrial control sector are sophisticated processes requiring complex and modularized designs that often require modification and augmentation during runtime. These applications require high performance, determinism, and predictability; their infrastructures demand highly robust and reliable designs both in hardware durability and in resiliency features to support high-availability needs.

The networking legacy in the industrial control market is long established with roots that extend to multiple industrial technology standards as well as to multiple vendor-specific, proprietary technologies. The industrial installed base that has settled on these technologies is a conservative community of manufacturers that move carefully and slowly from one generation of technology to the next. Solutions need to be proven and ROI is carefully calculated before investments are made. The move from any current position is always done through migration and evolution in order to protect installed assets and to maximize profitability for any given change in architecture.

Prior to the advent of TSN technologies, there hasn't been much motivation to drive these differing industrial solutions to a more unified approach. The TSN conversation, however, has driven an awareness that future networking solutions will require a holistic approach and a true system perspective. The future network requires a comprehensive answer that includes all elements - infrastructure and end stations alike – to be included in the final solution. Moreover, in order to solve the complete Industrial Internet of Things (IIoT) problem statement, this comprehensive solution will require the collaboration of the existing, differing, technologies and standards to morph toward solutions that can allow their inclusion in this new eco-system. If any specific technology, vendor, or product does not participate in this new paradigm, there is no way for "the system" to accommodate or plan for the traffic that needs to be managed to or from these components. The result is that these components will either not be properly served in the overall design – or that they will interfere with the overall design in a non-productive fashion.

Industrial Use Case
Consider the following representative architecture of an industrial control application and associated network design:
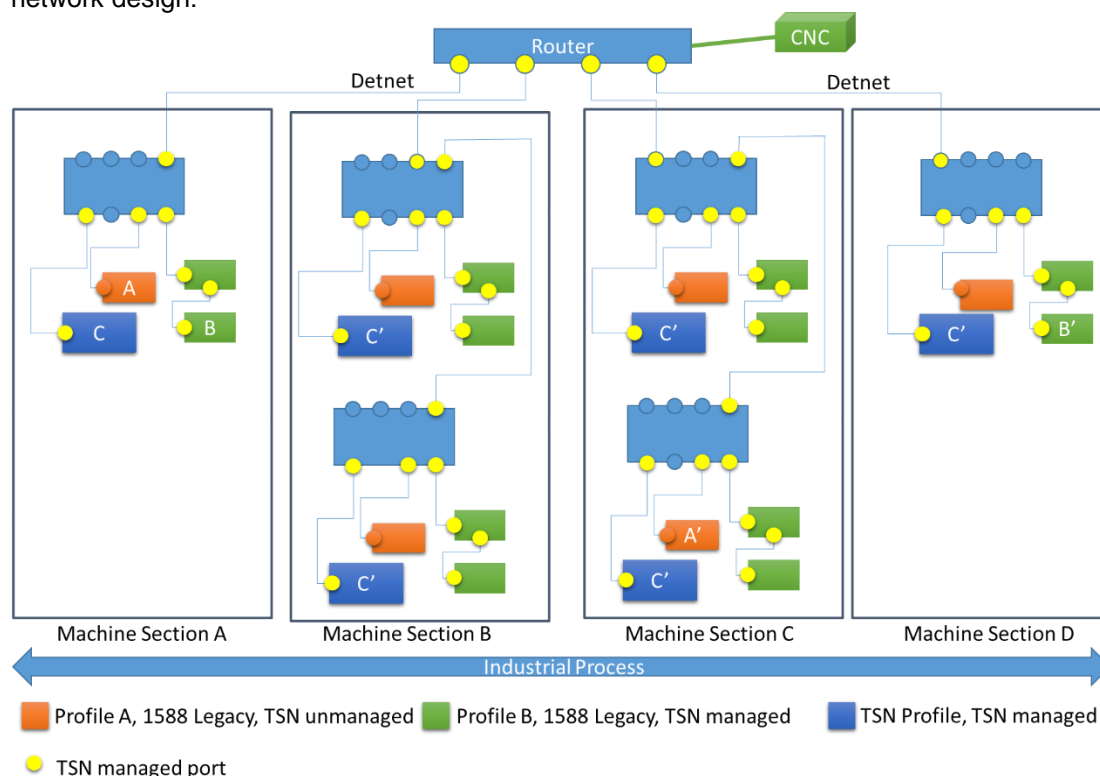

Figure 5 – Industrial Use Case

In this system, an industrial process is controlled by a single machine which consists of four different sections of machinery. Each section of machinery is delivered to the process via different OEMs, each of which have different specialties in the parts of the process that they control. The end user has 7 manufacturing sites around the globe. There are 15 of these machines per site, and each machine IP addressing scheme is identical to the other machines in the same manufacturing facility, which matches the other addressing schemes of the other facilities. Moreover, each machine section is a subnet with a unique VLAN so that the equipment can be constructed modularly and functional organization is achieved. All sections of the machine are synchronized and coordinated to produce final product and relevant events are timestamped so that data on the manufacturing floor can be correlated against data in the MES system and data from the supply chain. The entire manufacturing facility uses the same understanding of absolute time and all events are related back to the common notion of "wall clock time." Additionally, this use case acknowledges the pre-existence of products and technologies that utilize other 1588 based solutions. In this case, component A in Machine Section A may communicate with component A' in Machine Section C; Component B in Section A may communicate with component B' in Machine Section D. Finally, component C in Machine Section A produces data to be consumed by C' components in sections B, C, and D. The implementation of time bridges or time gateway mechanisms via the switches provides a migration path for an extended community of existing products and technologies to be included in the wider TSN value proposition. In this use case, components A and A', and B and B' would require time gateway translations.

Scope and timing
It's important to note that the use case given above is typical for a very wide range of industrial applications. This use case illustrates the need for solutions at layer 2 (switching), layer 3 (routing) and for time bridging functions.

Today, the IEEE-802.1 TSN Workgroup focuses on layer 2 in the architecture and this is where the majority of TSN work is being defined. Here, there is work complete enough to begin implementation in a system or in a product. Additionally, AVnu has defined compliance specifications sufficient enough to certify some products for this level of operation. This version of the "Theory of Operation" document focuses on the technology in this part of the architecture and will expand its scope in future versions to include other areas as those technologies become defined.

Layer 3 (routing) functionality is currently being designed by the Internet Engineering Task Force (IETF) and it essentially extends the capabilities of the 802.1 TSN specifications into layer 3. This work, often called "Detnet", (deterministic network) is not yet complete, at the time of this writing.

Time bridging, also referred to as time gateway capability, is shown in the layer 2 switches and provides the mechanism for migrating legacy technologies into a TSN system. This functionality has not been identified as a required work item for any standards communities. Individual suppliers could develop these bridge functions as solutions for the market.

Finally, it shouldn't go unnoticed that a CNC device is shown in this diagram as the centralized configuration tool that is used to manage the network and ensure proper management of networking traffic and streams for this given application. This is consistent with the concepts discussed above around Figure 4 where CNC and CUC relationships are described.

Effect of TSN features on Control Applications

To examine the effects TSN can have on control applications, let us consider the model depicted below:
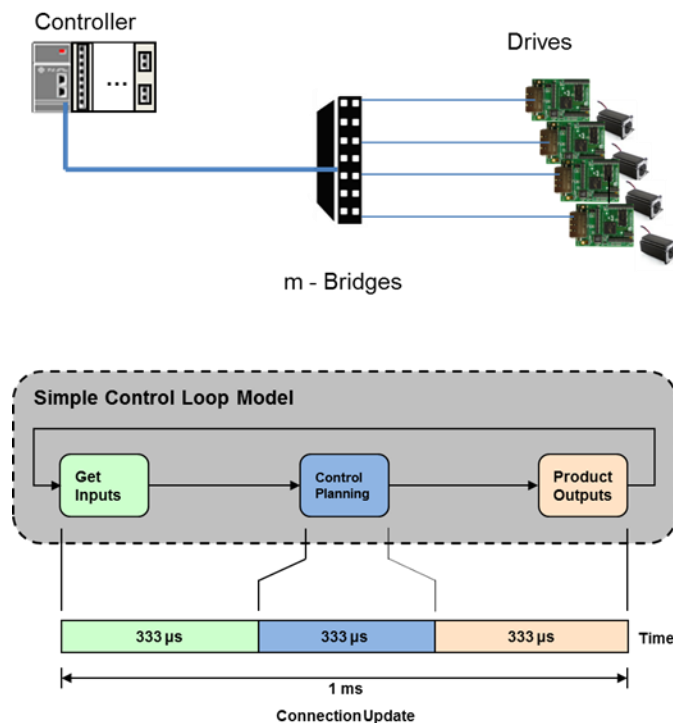

Figure 6 - A Simple Control Model

This model consists of a controller and drives connected through an infrastructure consisting of "m" switch hops. Input data must arrive at Controller before the end of the input interval. Planned data outputs should be transmitted before the end of the planner interval. Output data must arrive at the drive before the end of the output interval. For the sake of simplicity, these intervals are divided in thirds at an update rate of 1 mS. Further the model assumes:

- All network elements are time-aware
- Standard QoS/priority throughout.
- Cut-through switching (cut-through latency ~2usec @ 100 Mbs; ~1usec @ 1 Gbs)
- Some control of traffic volume and the size of interfering traffic on the network

This analysis will focus on a part of the problem associated with network performance. Variables such as controller/drive latency are assumed to be understood and constant. A simple, mathematical model describing such a system is given by:

Max Axis = 1 + {1/3 * Connection Update Period – (Drive Transmission Delay + (m + 1) * Ethernet Transmission Time + m * Switch Latency + NIC Packet Processing Delay + Bus Interface Delay)}/NIC Packet Processing Delay

- *(Where m = # of hops)*
- *Drive Transmission Delay:* We'll assume all drives have outputs queued prior to transmission, so this is contribution is small with respect to other operands, effectively 0 *usec*
- *Assume update packets are fairly small(124 bytes), so Ethernet Transmission Time is (124+20)*80ns/byte = 11.52 usec (at 100 Mbs)*
- *Switch Latency = (interfering packet size+20)*80ns/byte*
- *NIC Packet Processing Delay – There are techniques to ensure the network is the bottleneck (e.g. 2 cycle processing): 11.5 usec for 100 Mbs, 1.15 for Gigabit.*
- *Bus Interface Delay: has a lot to do with the overall system architecture. Delay could go effectively to 0 (given good bus structure, DMA/ etc.). We'll assume 0 for this analysis.*

Ideally, we'd like all of the drives to transmit their output data simultaneously. In this way, the link between the controller and bridge is optimally utilized. As shown below, performance is strongly influenced by interfering traffic and thus, the number of hops. In practice, control systems will engineer the network to limit the size of interfering packets (the example below assumes 100 Mbs link speed with both 1500 and 500 bytes of interfering traffic).
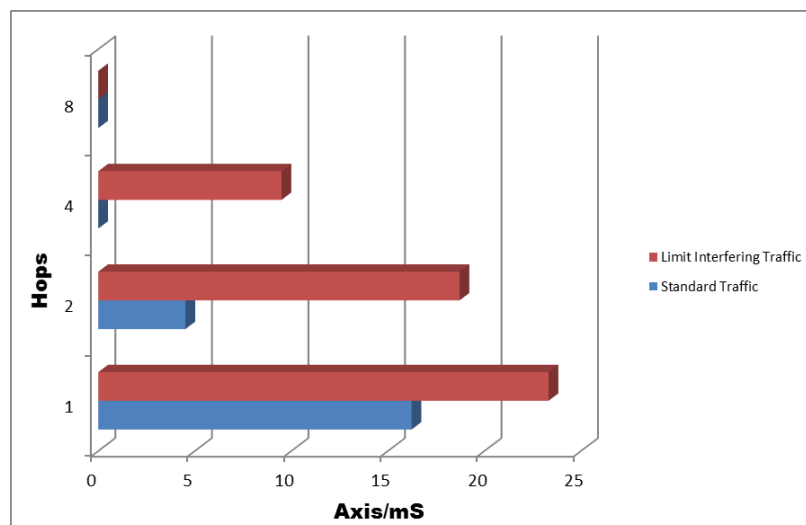

Figure 7 - 100 Mbs Baseline

Preemption offers a standard, unified means to limit the effects of interfering traffic. With a maximum fragment size of 64, the maximum interfering frame size is:

(2*64 bytes)-1 =127 bytes

Thus, Preemption simplifies the problem of isolating the control network from interfering traffic. The network designer still needs to ensure that other traffic of the same priority is not present on the wire or that bandwidth is sufficient to deal with all such traffic.
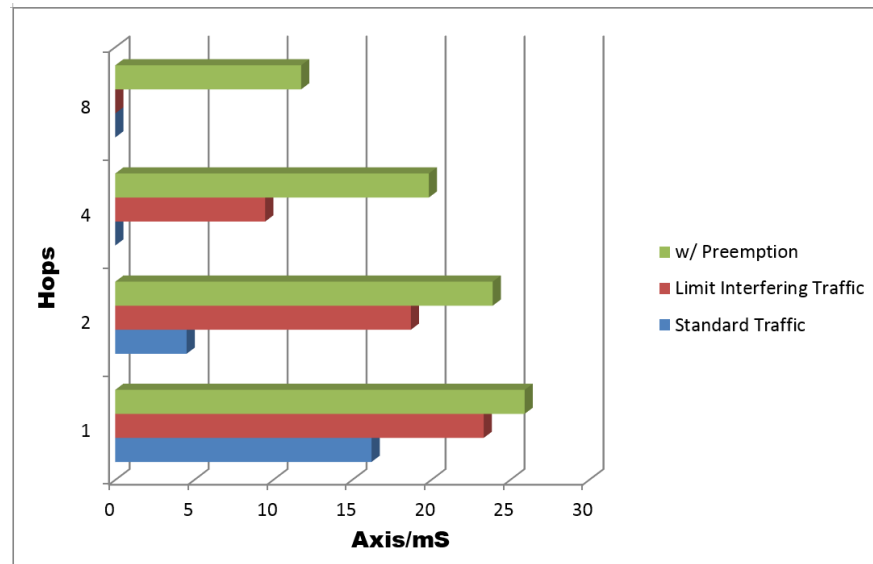

Figure 8 - 100 Mbs w/ Preemption

Utilization of a line topology and scheduled traffic can further minimize effect of interfering traffic. As shown below, the schedule of drives can be individually adjusted to compensate for drive transmission delay, transmission time and switch latency.
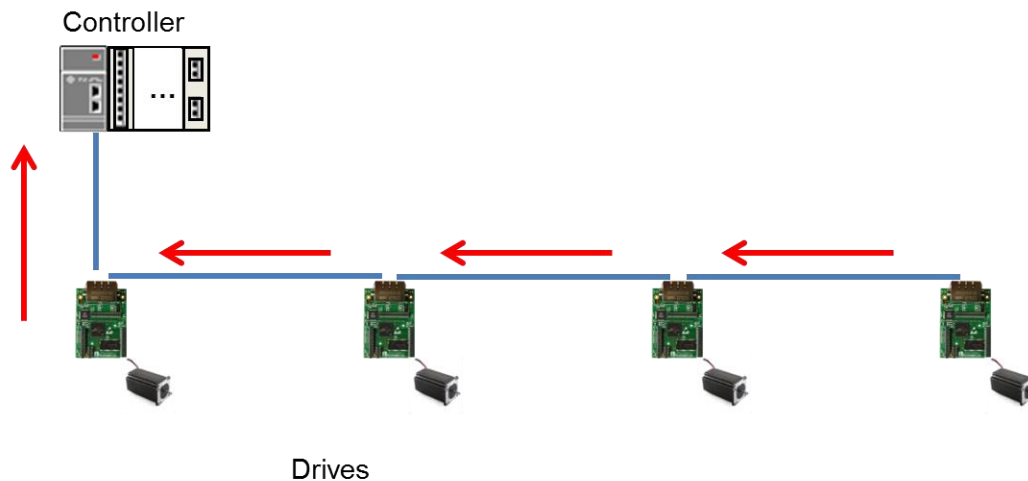
Figure 9 - Use of Line Topology to Minimize the Effects of Switch Latency

Thus, the effects due to switch latency are minimized and the effects of interfering traffic are of less consideration than the ability of the controller to process incoming packets. This example still assumes a somewhat isolated network (i.e. there is no other traffic of the same class which might interfere with control packets).

As shown below, Gigabit transmission speeds further reduce the effects of interfering traffic. Note that the benefit from scheduled traffic across hops is much less significant.
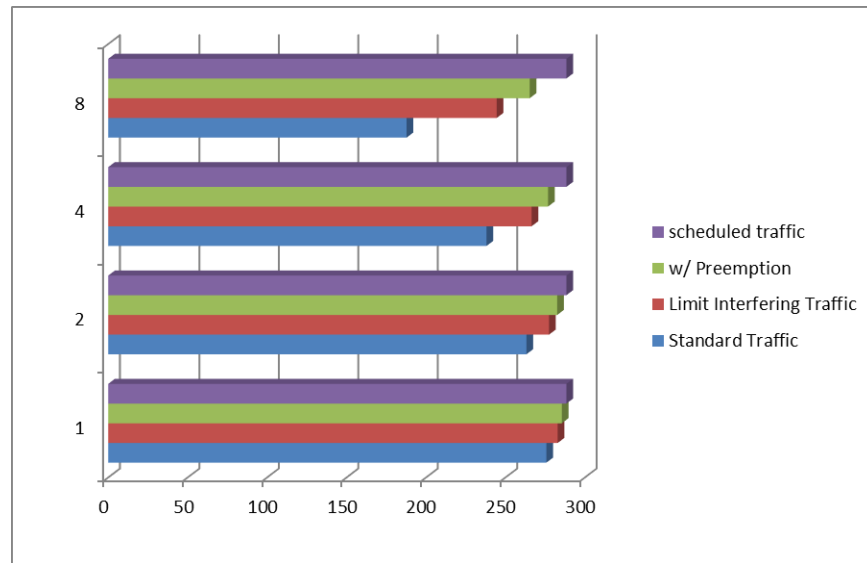
Figure 10 - The effects of Gigabit Link Speeds

Models vs. Reality

It is worth noting that what is shown in this analysis is a simple model intended to illustrate the effects of TSN on an Ethernet-based control solution. There are many potential points of optimization in a complex, real-world system. Assumptions made herein are for the purpose of discussion, not to suggest design approaches or solutions.

> *You will never strike oil by drilling through the map! BUT: this does not, in any way, diminish the value of a map! (Solomon Golomb: Mathematical models – Uses and limitations. Aeronautical Journal 1968)*
> • Source: Dr Edward Lee (UC Berkley) TSNA'15 – The Internet of Important Things

For instance, the model used herein assumes a very conservative 11.5 microseconds for "NIC Packet Processing Delay". If one assumes that the system bottleneck is the network, rather than controller packet processing, then Gigabit speeds show drastically improved results.
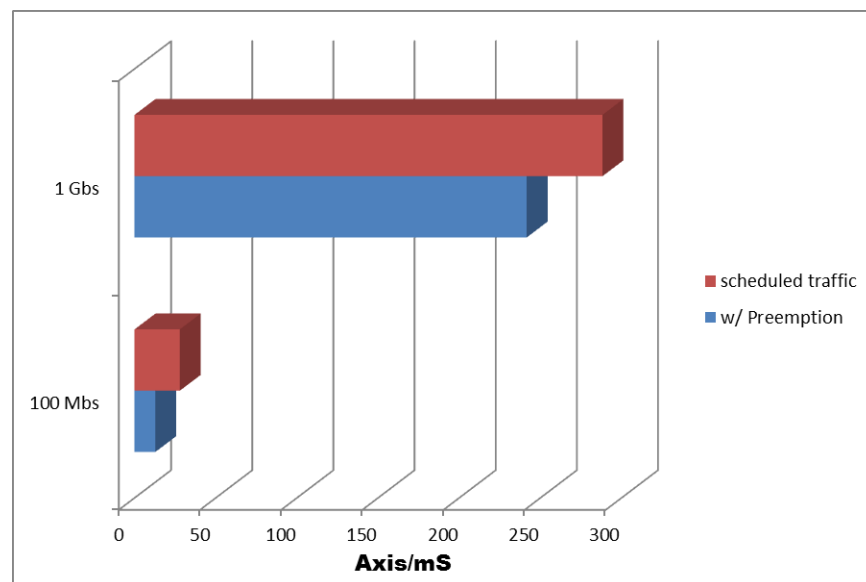

Figure 11 - Improved Gigabit Performance

The Distributed Motion and Time Synchronization SIG uses a more sophisticated model for modeling motion problems. This tool is proprietary to ODVA and is therefore not included as part of this paper, but results from that model validates the results contained herein.
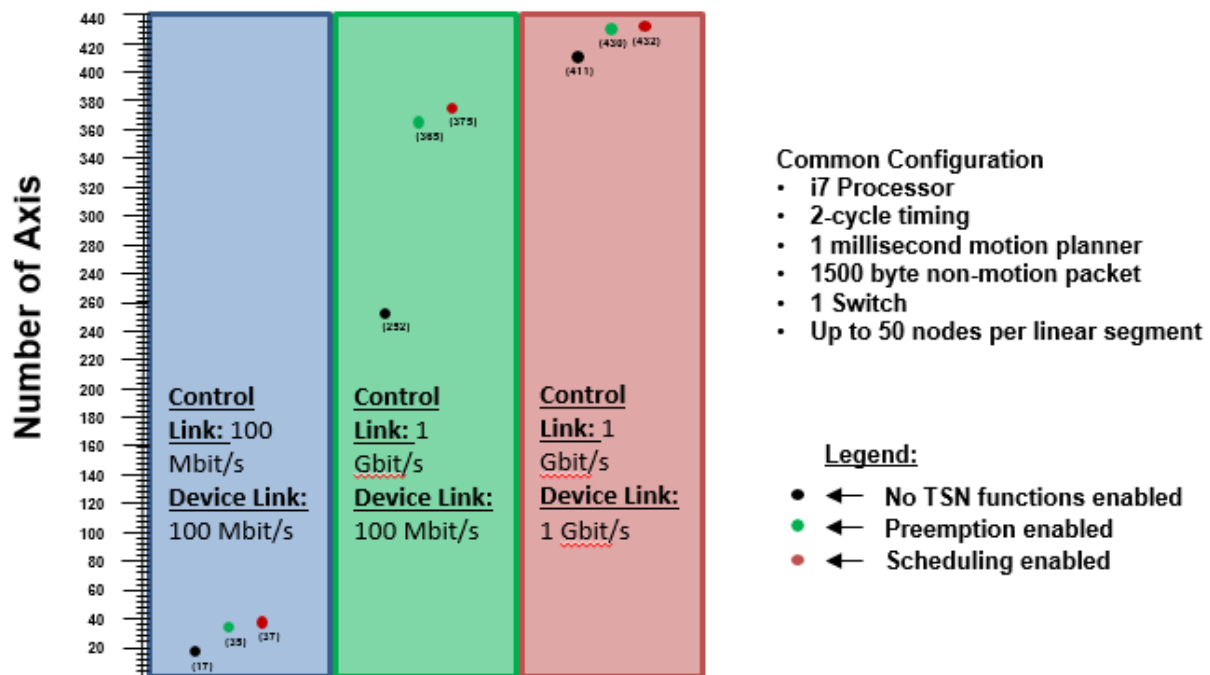
**Figure 12 - Impact of Wire Speed vs. TSN Functionality per the CIP Motion Model
(Distributed Linear Segment Topology)**

Wire speed contributes to the majority of throughput and performance on the wire. TSN functions allow for better "packing" of data at any given wire speed. Preemption offers a standard, unified means to limit the effects of interfering traffic control systems. Analysis suggests many control applications could be served with Gigabit Ethernet and preemption without the additional complexity of scheduled traffic.

Network Convergence and Centralized Configuration

Given these results, one can reasonably ask whether industrial applications need to consider scheduled traffic. Certainly, the data would suggest that the benefits of scheduled traffic for control applications are not as dramatic as one might expect, particularly at Gigabit speeds. However, it is important to remember that, as deployment of Ethernet in Industrial Automation grows data access from Industrial networks become more desirable. As the Enterprise and Automation networks become more integrated, data flows from various applications must converge.

The centralized configuration model proposed in IEEE802.1Qcc provides a practical approach for achieving this vision of convergence. Such convergence requires publication of network traffic requirements to and from disparate and possibly conflicting traffic sources. Scheduled traffic is one of the tools that the CNC can use manage and resolve these conflicts.

Referring to the use case outlined in "**Error! Reference source not found.**igure 5 – Industrial Use Case", we see that a centralized network control entity can be used to configure individual subsystems and that a CNC entity can be used to integrate these various pre-configured subsystems. In this sense, TSN is closely aligned with Software Defined Networking (SDN) concepts. Software Defined Networking or SDN is an approach to computer networking that allows network administrators to manage network services through abstraction of lower-level functionality. This translates to providing a system level view of network configuration and management vs. the traditional component level view.
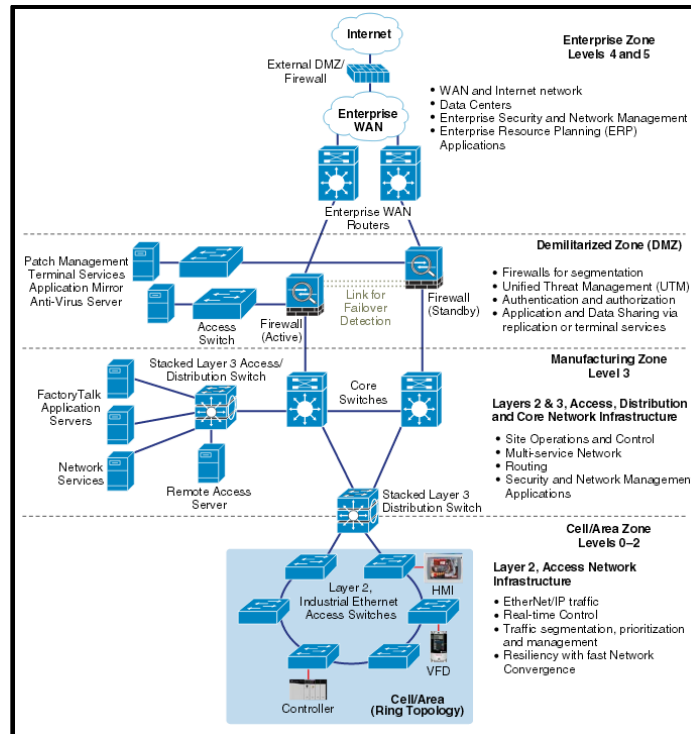
Figure 13 - Concept of tomorrow's configuration view

This approach also lends itself to network modeling prior to network commission. Given a known topology and set of traffic requirements/constraints, the CNC can accurately predict network performance.

Conclusions

TSN technologies offer a scalable, predictable approach to deterministic networking. Because Ethernet/IP products have always relied upon standardized technologies, ODVA is in an excellent position to leverage these emerging standards. However, significant challenges remain. The integration of various PTP profile and the convergence of EtherNet/IP traffic with a scheduled TSN network are chief among these challenges.

References (optional)

1. Browne, B. Hotchkiss, M. Weingartner, T. "Real-Time Performance of Industrial Ethernet in Field Devices." November 7, 2011, http://www.innovasic.com/technology/PriorityChannel.
2. Alsup, D. Weingartner, T. "What is PriorityChannel™?" May 14, 2012, http://www.innovasic.com/technology/PriorityChannel.
3. Chaffee, Mark. "CIP Motion Implementation Considerations." Proc. of ODVA 2009 Conference & 13th Annual Meeting, Howey-in-the-Hills, Florida USA.
4. Frithjoh Klasen, Jürgen Jaskolla, "Profinet at VW de Mexico", Profinet International – North America General Assembly Meeting Presentation, Scottsdale, AZ, October 2011
5. Lee, Edward A. "Absolutely Positively on Time: What Would It Take?" *Computer* 38.1 (2005): 85-87.
6. Lee, Edward A. "The Internet of Important Things?" TSNA 2015, Santa Clara, CA USA.
7. Prettyjohns, Keith, and Alsup, David. "Enhancing Determinism, Fault-Tolerance and Reliability in EtherNet/IP Systems." Proc. of ODVA 2009 Conference & 13th Annual Meeting, Howey-in-the-Hills, Florida USA.
8. Woods, Jordon. "Minimizing Contributions to System Delay from NIC Packet Processing in CIP Motion and Other Real-Time Control Applications." Proc. of ODVA 2011 Conference & 14th Annual Meeting, Phoenix, Az USA.

****