

Overview and Future of Switching Solutions for Industrial Ethernet

György Kálmán, Dalimir Orfanus, and Rahil Hussain
 ABB Corporate Research Norway
 {gyorgy.kalman, dalimir.orfanus, rahil.hussain}@no.abb.com

Abstract—Industrial Ethernet is the preferred network technology in industrial green field deployments. Although the performance of these networks is superior compared to legacy fieldbuses, the complexity of possible installations is an issue in the field. Both selection and placement of active devices and efficiently running network services are causing problems in planning and also during the life of the network. This paper is giving an overview on implementation possibilities of switched network in industrial applications with respect to performance, features, logical architecture, and flexibility. It also shows the importance of time synchronization and presents current standardization work with focus on impact on industrial applications. An outlook for expected future possibilities is shown, including the use of Software Defined Networking (SDN).

Keywords—industrial Ethernet, switch, embedded, discrete, soft switch, forwarding, performance, QoS, SDN, I588, 802.1AS

I. INTRODUCTION

This paper is an extended version of our paper at ICN 2014, An Overview of Switching Solutions for Wired Industrial Ethernet [1]. Compared to the original, it contains an extended state-of-the-art analysis, a section on time synchronization, which is one of the most important features for industrial applications, extended experiments and an outlook with an introduction on how *Software Defined Networking* (SDN) could be used in an industrial setting.

Ethernet is already the dominating technology at the control and higher levels of an automation network and is expected, along with wireless Ethernet, to be the primary choice in field installations.

Because of resource constraints and *Quality of Service* (QoS) requirements, most of the automation networks are implemented as *Local Area Networks* (LANs) (Figure 1). The motivation is twofold: first, this approach leads to a simpler network configuration as there are no routing tables and there are plenty of resources to utilize without the need to coordinate with external nodes or other actors. Second, as the operation is closer to the physical layer, the QoS parameters can be held within more strict boundaries.

The paper is structured as follows: the second section provides background overview on industrial Ethernet, then current top design priorities are presented in the third: topology and fourth: time synchronization. In the fifth section, the possible switch architectures are presented including performance requirements of different applications.

Then the possible architectural solutions are explained, with discrete, embedded and soft switches as main categories.

Performance comparison is given based on our testbed measurements focusing on latency and jitter. A conclusion on possible fields of use for the discrete, embedded and soft solutions is given. As a possible evolution, an outlook on SDN for future industrial Ethernet planning and extensions towards using Layer 3 networks and wireless solutions is shown.

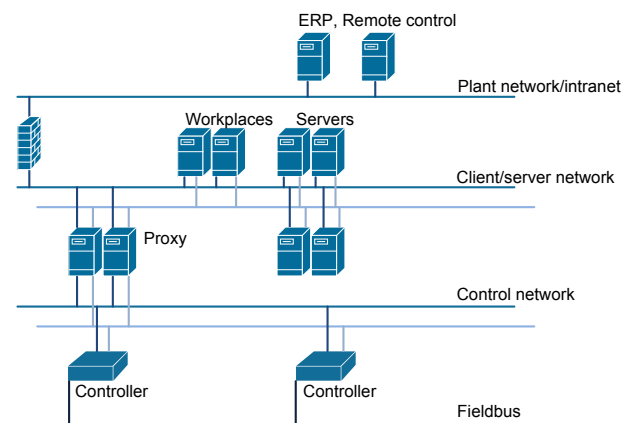


Fig. 1. An example of automation network architecture

II. INDUSTRIAL ETHERNET BACKGROUND

Industrial Ethernet enables the use of standard Ethernet devices and the IP protocol suite in automation networks [2]. By implementing industrial network based on Ethernet, vendors can create infrastructures that provide improved bandwidth, resiliency, and network security compared to fieldbus solutions (Figure 1). As an additional value, the use of already established standards lowers the risk associated with technology development.

A number of issues emerge from the fact that the Ethernet networks are replacing the fieldbuses. A heritage of the fieldbus past including design processes is the dominant cause of bus-like topologies (Figure 2) resulting in suboptimal operation of Ethernet [3]–[5].

The most challenging topology types are long chains of switches, which are often closed to form rings. While *Rapid Spanning Tree Protocol* (RSTP) was designed with loop-avoidance in mind, it is currently widely used redundancy protocol in the industry. In a ring structure, RSTP will disable one link and render the active network topology into a special tree, a line of switches.




Fig. 2. Industrial Ethernet topologies

Although a line is a valid Ethernet topology and the technology will work, the industrial network will suffer from scalability issues due to much smaller end-node count than it could be expected from office experience [6], [7]. One of the reasons is that the most industrial Ethernet protocols are using standard Ethernet as bearer, with some exceptions such as PROFINET IRT or EtherCAT, so the network operation is in practice not different from the office counterparts. However, it runs on much less optimal infrastructure with more QoS sensitive applications. The other problem is raising in the requirements of precise timing. The need for synchronous operation is apparent, e.g., in case of synchrophasor operations, where the precise sampling of the waveform requires today a GPS time source. To enhance resiliency and to lower costs, time synchronization protocols are expected to provide a stable solution allowing the operation of such services even if the network is large and without the need to place an external precise time source at each critical network point.

The very long and sparse spanning tree is having a low branching factor (the average number of child links at each switch) and can lead to excess latency and jitter [8]–[11]. As compared to office counterparts, more manually configured nature of the industrial networks is creating another drawback: in a typical network the root bridge is configured manually and this node is typically not the one in the middle of the longest chain to break it into two halves. Thus, in calculating expected QoS parameters, the worst case having one single long line of nodes is used.

Still, in the most of the industrial operations, processes and automation tasks have tolerances several magnitudes higher than the total jitter of a reasonably built industrial Ethernet network could have. One of the exceptions is motion control where deadlines can be close to the jitter and/or delay of a large network, thus giving certain planning constraints. This particular field of automation shares requirements with applications in very different fields, e.g., finance and telecommunication.

In the current applications, however, automation tasks are focused on LAN operations, thus mostly the time-critical planning stops at the LAN-WAN interface. In telecom uses WAN environments are also included, which results in, e.g., different focus in standardization.

This convergence can clearly be seen from, e.g., the standardization efforts, where in addition to traditional networking and industrial companies, telcos and silicon vendors are taking a considerable effort.

III. TOPOLOGY

In office environments, high port count switches are used to implement a high branching factor network, thus the issues associated with cascaded switches are less important [12], [13]. Also, in a typical setup, an office LAN is much earlier divided into subnetworks using firewalls and/or routers than reaching a deep spanning tree.

As an indirect result of the low branching factor and the pressure for lower costs and relatively high price of managed industrial Ethernet switches, the industrial installations are experiencing pressure from both sides: on the one, the required performance level is relatively high, especially for processing delays and jitter. On the other hand, the price of such devices on the market is relatively high in a project, because of the low port count and low general utilization of the resources. The recent trend is to include or integrate low port count switches directly into devices, i.e., devices can be interconnected into a cascade or a low branching factor tree without the need of other external devices [14]. Such topology has also practical impact on physical installation of cables: generally it is more far more expensive to lead a bundle of cables across the whole factory rather than a single cable that connects with the most or even all devices.

Moving towards switch as an integrated feature might lead to scalability problems because of increased latency and jitter on long sparse trees. This scalability issue is more apparent when integrated modules use low port count (2-3 ports, only providing enough ports for daisy-chaining). However, higher port count modules can be as flexible as the current situation with discrete devices in fieldbuses. The only limitation for the scalability of networks [15]–[17] can be when the daisy-chaining solution with few ports is combined with time-critical automation tasks.

Integrated switches are expected to lower the cost of building the network with lowering the amount of standalone devices and with better integration to the automation devices. The expected feature set is similar or equivalent to discrete units. In following sections we provide an overview of the typical embedded switch architectures and how they could be fitted into the industrial landscape.

IV. TIME SYNCHRONIZATION

One of the most important features in the current industrial deployments is the possibility and capability of precise time synchronization throughout the network.

The convergence between the industrial, telecom and finance requirements for a synchronous operation lead to that two standardization groups in IEEE are working on extending the current synchronization solutions towards additional networks and to widen their application areas.

Synchronizing time was always an important question over network applications, time references, timestamps are widely used in communication both for management, message exchange, logging and security applications. However, the granularity and the level of synchronization very much depends on the application. For less demanding tasks, e.g., process automation, a protocol with a fraction of a second precision would be sufficient. On the other end, (fine) motion control applications or power electronics would require sub microsecond level synchronization [18] to operate in a coordinated manner (Figure 3). Also in the incident management is important to be able to put the chain of events into the correct order.

Today the standard requirement for switched Ethernet is to support IEEE 1588v2 as the time synchronization protocol [19].

Another motivation for high precision time synchronization is also to filter out the non-deterministic processing delay in the active network components (switches, routers etc.). For this reason, IEEE 1588v2 supports Layer 2 time stamping of frames and different layers of procedures to enhance the precision.

The first place is in the network itself. Relay nodes in the network (switches, routers etc.) are introducing variable frame delay because they queue frames. This delay depends on the amount of traffic in the transit at a particular relay point. In IEEE 1588 the delay at a transit point can be calculated from timestamps put on ingress and egress frames.

The second level is the node level, which can be either an end node or a relay node. Timestamps are created in a way that the end node can calculate the correct time within a precision limit. In intermediate nodes, both at an ingress and egress ports the frame timestamps are updated to correct the transit time with the residence time in the node in order to keep the precision. The reason why IEEE 1588 and also other high precision synchronization protocols cannot rely solely on filtering solutions in upper layers is, that the measurement of time (here arrival and departure times) should be done as close to the physical medium as possible. Closeness to the physical world lowers uncertainties and allows for a more accurate measurement. Explicitly, IEEE 1588, if enabled, uses a hardware time stamper connected directly to the Ethernet card between the PHY and the MAC, to allow accurate measurement in the required precision range.

While this is a very efficient and good placement for a high precision stamping unit [20], the direct connection to wired Ethernet was also found to be a limitation, so in the future versions, it is expected that a hardware abstraction layer will be introduced to open for more networking technologies.

A. IEEE 1588v3

The work on the new version of the current standard Precision Time Protocol (PTP) is ongoing. The standard is expected to undergo a major overhaul compared to the v2 released in 2008 [22].

The new standard is expected to follow the example of the architecture of IEEE 802.1AS, which was originally developed




Fig. 3. Timing requirements for industrial Ethernet [21]

for Audio/Video applications but currently evolving into a much more generic protocol under the IEEE Time Sensitive Networks group. To forecast the evolution of industrial Ethernet, it is important to get an view on how IEEE 1588v3 is expected to evolve compared to v2 [23].

- *Common Management Information Base (MIB)*: The current 1588v2 landscape is somewhat disrupted as several parallel profiles exist, which are not fully compatible with each other. An example is the IEEE C37.238, the power profile, which uses a different MIB compared to other profiles, thus management of the complete time domain is more challenging.
- *A layered protocol structure*: The current PTP version uses functions which are stretching across network layers and thus limiting generic implementations of the protocol, where, e.g., the clock state machine can be reused even if the physical layer is exchanged.
- *High availability synchronization*: the protocol is planned to offer a solution for having multiple clocks in the time domain, where they can take over the master's role within specified limits.
- *Security*: IEEE 1588v3 is expected to provide a security mechanism to protect the time quality by using cryptographic functions for authenticity and integrity checking.
- *Wider scope*: The new PTP version is expected to support a wider range of networking technologies, including Wifi networks and wireless sensor networks. Support for time synchronization over IP is expected.

The most important change is to apply a proper protocol architecture where the media-dependent parts are connecting to the upper layers of the protocol through a standard interface. This enables 1588v3 to run on different networks.

Another difference w.r.t. 802.1AS is that 1588 does not require time-aware infrastructure. It can apply sophisticated filtering to reach good precision, although it would be less precise than with direct layer 2 stamping.

B. IEEE 802.1AS

The development of 802.1AS started as an audio/video technology. The intended use was to support synchronized network streams to be delivered over a LAN and requires time-aware infrastructure. The protocol was defined partially as a

strict profile (subset) of 1588v2, but also has extended the 1588 protocol so the compatibility is limited.

IEEE 802.1AS was planned already with a proper protocol layering and from the beginning it was expected to support several network technologies. It has raised the interest of industrial actors and later the audio/video focus was extended to cover any kind of service which needs high precision time synchronization.

The 802.1AS focuses on synchronized services on bridged networks and the standardization is ongoing with high effort. The most important current fields are:

- A frame preemption technique in the bridges (P802.1Qbu)
- An Ethernet traffic scheduling technique (P802.1Qbv)
- New traffic shapers
- Next version of the Stream Reservation Protocol (SRP)(P802.1Qcc)

With IEEE 802.1 TSN, the network itself will use the time that is transported. Both the Ethernet traffic scheduling functionality and time triggered traffic shapers will utilize time for correct execution.

The addition of a reservation protocol, having preemption and scheduling makes this standardization effort very important for industrial actors. This could mean that IntServ-like QoS functions could be deployed on an L2 network, which until now had to accept to have only simple filtering, priorities and very simple traffic management rules.

C. IETF TICTOC

Internet Engineering Task Force (IETF) is working on a time synchronization solution with focus on WANs as well. Their focus is on running time synchronization over large networks, e.g., running 1588 over *Multiprotocol Label Switching* (MPLS).

D. Convergence

It is expected that the two protocols, as they both are being developed by nearly the same actors under the umbrella of the same organization, will be able to at least coexist on the same network. Even more, a cooperation between the protocols would be beneficial, as there 802.1AS could take over in LANs and 1588v3 could cover all the other networks.

V. ARCHITECTURE POSSIBILITIES

With a few exceptions, only managed switches are being deployed in industrial Ethernet networks. This is a result of the different requirements, e.g., prioritization and Virtual LANs (VLANs) are rising in the industrial environment compared to the office networks.

A. Discrete switches

Unmanaged switches offer a low-price connectivity solution, where leafs are placed into the same network and the ingress traffic can be treated with the same rules independently of the port. The biggest advantage of the unmanaged devices is their simplicity. The simple silicon offers more deterministic

operation (there are no multiple user-configured rules and preferences to check), higher reliability (typically in safety deployments an unmanaged switch represents much lower failure risk), lower heat dissipation and lower price.

Unfortunately, in most cases and despite their advantages, unmanaged switches are not considered as typical requirements in industrial installations. Typical requirements are redundancy functions, traffic prioritization and extended status reports, which are not available in unmanaged devices.

Managed switches offer redundancy and loop-avoidance functions (e.g., RSTP) with logical segmentation using VLANs, remote management with *Simple Network Management Protocol* (SNMP) and troubleshooting features such as port mirroring.

There are also devices in the office networks, located between these two levels, called smart switches. They offer the most of the managed switch functions, but lack for example SNMP management. Introducing a similar class of devices into automation networks where only the necessary protocols are selected might be of interest, since having a more grained approach on switch features can lead to a more cost-effective network architecture. Also, in safety and security fields, it is beneficial to have as simple devices as possible, also in the software running on them. A switch, where a part of the protocols are being left out (e.g., VoIP-related if no VoIP system is in operation) simplifies the checking and certification tasks.




Fig. 4. Typical switch size comparison

There are few arguments against the use of discrete switches and most of them originate from the specific industrial landscape: the low branching factor, which results in a high number of low port-count switches, as shown on Figure 4. The high number of standalone switches and the rugged hardware leads to an expensive network infrastructure with most of the resources being unused.

The low port count is even more apparent in the daisy-chained field networks, where Ethernet is also expected to replace the legacy communication solutions but typically it has to utilize the same topology.

In such environments, using the typical 8-10 port managed switches is rather expensive, as even such a low port count will not be utilized in addition to the higher management effort. To overcome price pressure, excessive engineering complexity and dependency on third party devices, vendors move towards embedded solutions.

VI. EMBEDDED SWITCHES

Integrating a switch module into devices like controllers is on the agenda of automation vendors. These modules could take tasks of discrete switches in the lower levels of automation networks. The construction of these units is potentially cheaper than using a separate switch (e.g., a low-end switch fabric) and can provide a few gigabit/second of non-blocking bandwidth.

There are several important issues around the integration of devices. The first is the question of interface towards the host device. The typical architecture offers an internal interface towards the host, which is implemented as a standard, but internal, Ethernet link. This setup is analogue with the discrete switch case, only the interface connecting the host and the switch has been exchanged with the internal connection.

Switch modules, by default, only forward the traffic and all features, which are needed to implement a managed switch, have to be run on the host or the switch module has to be extended with traffic processing capabilities.

Integrated modules are expected to deliver similar performance results as their low port count discrete counterparts and also to offer the similar range of services. The cost of such a solution could be still lower as compared to the separate unit, as several components can be shared with the host, e.g., power supply, casing and user interface, while a potential risk is the software support and compatibility for mixed or brown field deployments.

If the management functions are implemented by using the CPU on the host, multicore platforms can be exploited by moving the forwarding-connected functions to one core and running the other functions on another core, even on a different operating system if needed by utilizing virtualization.

The possible drawback with these modules is that the host is still only connected with one internal port. This means that if the aggregated bandwidth use exceeds the host's bandwidth, the host has no possibility to monitor the whole network traffic. The lack of full monitoring possibility is problematic if the switch module itself cannot report hazardous traffic situations with, e.g., raising a signal when a port experiences traffic congestion. Also the implementation of traffic prioritization can be problematic if not at least parts of the task is implemented in the silicon.

It is clear, that for the industrial applications, unmanaged switches offer less than the required functionality, but on the other hand, fully featured managed switches, especially the low latency ones originally developed for core operation or finance applications are too expensive and most of their features will not be utilized.

A. Minimum acceptable service level

A non-conventional approach is to minimize the implemented features of the devices. Typical requirements state that the switches used should be managed, but the actual feature set is not defined. Currently, managed switches typically implement the whole feature set expected from a managed switch (complying to IEEE 802.1D), but in most cases, only a handful

of features are actually used and also out of these, some are enabled by using the default configuration (e.g., weighting of frames in QoS queues).

A reduction in both cost and management effort could be realized by implementing a set of minimum acceptable service-level switches. An office network approach is the smart switch device class, for example the NetGear JFS524e, where the feature set is restricted to offer easier management and cost reduction in hardware. The smart switches represent more a restricted managed switch, but in the industrial environment, an approach from the opposite direction, extending the features of an unmanaged switch might be more interesting.

The motivation to use such devices is partly supported by the reduced development and device cost, but more importantly, the special circumstances of industrial deployments are also supporting this solution. One of the more complex services of the discrete switches is connected to traffic manipulation and security functions.

The gain associated with, e.g., *Internet Group Management Protocol* (IGMP) is that it can reduce the link load by grouping the receivers of multicast streams and also to protect other devices from using resources on traffic, which they have no use for. *Protocols Multiple MAC Registration Protocols* (MMRP) and *Multiple VLAN Registration Protocol* (MVRP), are also expected to reduce traffic load in areas where, e.g., a VLAN has no clients configured.

The other group of protocols are the network operation functions, e.g., RSTP, and IEEE 802.1X using *Remote Authentication Dial In User Service* (RADIUS). These protocols are run to keep the network loop-free and ensure network integrity and to allow secure authentication of new nodes.

Although all of these protocols are useful in an average network topology, in the industry-typical long and sparse trees, their gain is reduced. For increased traffic effectiveness: because of operational safety, networks anyway have to be designed so, that they can carry the whole network traffic, so the gain offered by grouping protocols might be limited. The main problem associated with grouping protocols in the typical line topology is that the resources need to be reserved over the whole path if nodes are expected to join or leave on the ports. The traffic reduction efficiency for line topologies depend on the actual traffic type.

For example, Multiple VLAN Registration Protocol (MVRP) might cut out some VLANs to be carried on a specific path, but if new nodes are allowed to join to a segment, the bandwidth for carrying additional or all of the existing VLANs shall be possible, thus the bandwidth spared by MVRP shall be reserved. In case of multicast protocols, like MMRP can be beneficial, but in this case also, at least the bandwidth need for all multicast groups shall be reserved even if not all of the groups are transmitted.

The execution of RSTP might also be of limited use, if the switches are organized in a chain and in every case, if the main uplink is broken, the other designated port towards the other switches will be chosen. Also, the topology of these networks is very static.

A possible solution is to use a compromise: deploy as simple as possible switches where chained topologies are used and include fully-featured discrete units where a tree connection structure is used (e.g., interconnecting rings or network backbone). Thus, the discrete units can run all the grouping protocols and reduce the load introduced to the ring, but inside the ring no further optimization is done. The simple devices shall be transparent on all protocols they do not support.

VII. SOFT SWITCHES

Embedded communication solutions are now allowing the implementation of a soft switch processing traffic of several gigabit/s of traffic on low consumption System on a Chips (SoCs). These are typically combined from an embedded CPU, a set of independent network controllers and a chipset, which integrates these into one system.

The positive point with these setups is that the host is the switch: it is possible to monitor the whole traffic flow directly on the interfaces. Also, the platform can provide a good basis for feature extensions toward implementing a router, firewall or network monitoring appliances. The industry can profit here with the evolution of SoCs in the recent years, reaching very high performance and integration with low power consumption and versatile usage areas, e.g., different ARM cores with integrated analogue-digital converters, various network adapters, on dye memory, cryptographic functions and storage.

A high performance, multicore SoC can also serve as a platform for automation tasks and with the use of a multicore CPU, the communication and automation tasks could be run separated. With virtualization being also offered in this price and performance segment, it is also possible to even run tasks on separate, isolated operating systems, using dedicated cores.

The main drawback of soft switches is the absence of the dedicated switching fabric. The throughput of the platform is prone to the actual implementation of the system and network adapter interconnection, used drivers and operating systems as well. Also, the limitations of the bus system and the network interfaces are summed, which can lead to insufficient performance in low latency environments. The price tag of such a solution can be justified if the device is utilized also in other tasks not only bridging.

VIII. FEATURE COMPARISON

The reviewed architectures show that if the switching solution is chosen, the future possibilities regarding traffic management, performance and feature set are being reduced.

Discrete switches offer high performance and a long list of management features and supported protocols. Embedded switch modules are implementing switching, but protocol and management features have to be implemented by the host or by a separate CPU and they only offer statistic multiplexing towards the host if utilized bandwidth exceeds what the host interface can carry. Soft switches are in practice implementing the embedded switch scenario but without the hardware switch

module, thus while offering full access to all traffic crossing the interfaces, they also suffer from the largest delays.

From the forwarding performance side, for large port counts, discrete switches offer the best solution, since a high-speed, non-blocking backplane is a hard requirement in this area. For smaller and medium sized switches (4-16 ports), an integrated module can also be viable. For low port count even the cheaper backplane solutions can provide enough bandwidth. It is also less probable, that such a switch will be experiencing a situation where all of the ports are fully utilized.

Our measurements on the forwarding latency and throughput of switches showed marginal differences between discrete and embedded solutions while the tests executed on the soft switch platform resulted in weaker performance figures.

IX. PERFORMANCE MEASUREMENT

A. Measurement and test equipment

Our test was implemented with the use of an array of discrete managed switches. The traffic generator was a Softing Industrial Ethernet Tester (OEM Psiber LanExpert 80), which can generate traffic between its two gigabit Ethernet interfaces and was acting as traffic source and sink.

Our tests were split into two areas: one was to measure the latency between two ports of the same switch to provide a way to compare the raw performance. The second area was to show switched Ethernet behavior in a typical industrial setup, where switches are chained and the ingress and egress links are only 100Mbps while inter-switch links are 1Gbps. The initial results based on the LanExpert measurements showed no significant difference in latency or throughput between the embedded and discrete units.

To measure the latency between 100Mbps endpoints, we need preciseness ideally at the level of a bit-duration or better, which is 10ns for the Fast Ethernet. Since the LanExpert's measurement capabilities were not satisfactory for generation of the statistics and exact measurement of forwarding behavior in a cascade, we decided to use the EtherCAT network consisting of the master (a Freescale P2020 development board utilizing a dual-core PowerPC e500 CPU) and two slaves (Figure 5. EtherCAT provides service called *Distributed Clock* (DC), which can precisely synchronize clock in slaves with time resolution of 10ns and has dedicated hardware in slaves to measure network latency).




Fig. 5. Testbed setup

To assess the performance of the selected equipment and to be able to provide guidelines for network planning, we set up the following measurements.

B. Default forwarding latency

Measurement of the time it takes for the frame to traverse the switch. It is composed from store and forward latency (L_{SF}), the switch fabric latency (L_{SW}), the wire line latency (L_{WL}) and the queuing latency (L_Q) [24].

L_{SF} depends on the frame length. The results are expected to show a linear growth of the latency with the longer frames [25].




Fig. 6. Test segment setup

Our architecture related measurement scenarios deals with latency between endpoints (both with 100Mbps) of serial connected switches and without any additional interfering traffic (see Figure 6). The purpose is to see the raw latency scaling of a network built by a chain of switches. We have four scenarios, each of them consisting with 1 up to 4 switches. Switches are between themselves connected with 1Gbps link. Initial measurements showed, in accordance with the LanExpert measurements, no significant difference between the discrete and embedded units, so the testbed was created by using 4 RuggedCom RS940G switches.

C. Standalone forwarding

Latencies and throughput between two interfaces of the same switch was measured with the LanExpert device and the results showed no significant difference between the capabilities of the embedded or the discrete units.

Measurements were performed on switches, which represent a significant part of the market: RuggedCom RS940G, Hirschmann RSR30, Moxa EDS-G509, a board based on Marvell 88E6352 switch chip and a soft switch using a stock Ubuntu linux and an Intel Xeon CPU with four chipset-integrated gigabit Ethernet interfaces. As a control, a test was also executed on a Cisco SG 200 switch (approximately the same performance class as the tested industrial variants), where differences in the results were also insignificant compared to the industrials. The measured latencies of the Marvell module are marginally lower, than the discrete counterparts, which are expected to be the result of the simpler architecture, as the module in the tested form implements only an unmanaged switch.

The only considerable difference could be observed with the soft switch platform. It was not expected to hold the same latency figures but the maximal frame frequency of approximately 180kfps is low compared to the rest of the devices (Table I). Although the latency is also higher (Table II), the figures stay mostly within acceptable range for the majority of networking tasks. The low throughput observed with shorter frames in contrast, limits the specific setup's usability since it will not be able to utilize the bandwidth in case of a setup like our test segment, where two interfaces need to carry the

TABLE I
THROUGHPUT IN K FRAMES PER SECOND FOR RESPECTIVE FRAME SIZES USING 1 GBPS LINKS

Frame size	RS940G	RSR30	EDS-G509	88E6352	soft
64	1481	1485	1485	1485	179
128	840	842	842	842	178
256	452	452	452	452	178
512	234	234	234	234	166
1024	119	119	119	119	119
1280	96	96	96	96	96
1518	81	81	81	81	81

TABLE II
LATENCY IN MICROSECONDS FOR RESPECTIVE FRAME SIZES USING 1 GBPS LINKS

Frame size	RS940G	RSR30	EDS-G509	88E6352	soft
64	5	5	5	3	16
128	5	5	5	4	16
256	6	6	6	5	18
512	8	9	8	7	33
1024	12	13	12	11	114
1280	14	15	14	13	93
1518	16	17	16	15	99




Fig. 7. Scenarios 1-4

aggregated traffic. It is expected that with different operating system and driver optimizations, better performance can be achieved.

D. Scaling of latency in a chain

Our measurements using the testbed extended with a variable cascade of switches (Figure 6 show that the discrete switches scaled as expected. When no additional traffic was injected to the measured interfaces, the latency growth was linear with minor variations. Measurements using an industry-typical scenario with 100 Mbps edge links and 1 Gbps internal links were executed. Four scenarios were measured, compromising of chains of 1-4 switches (Figure 7).

Also, the histogram on Figure 8 shows the expected behavior: the longer the chain is built, the wider is the range of latencies measured. The determinism of the switching solutions can be seen on the measurements and that in low traffic installations a linear growth of latency can be expected.

The histogram of the measurements showed the expected result, with having the most step-like distribution at using one switch and a still narrow but wider distribution of frame latencies for longer switch chains.




Fig. 8. Histogram Scenarios 1-4

E. Behavior under load

As an extension, the testbed was also used to execute measurements under load, to check how the devices react on a selection of traffic situations.

Behavior on crossing traffic was tested, where the switch fabric was tested if it is in fact non-blocking. In the range of devices tested, there was no significant difference compared to the no traffic scenarios (Figure 9, which shows that the hardware is using a non-blocking architecture).




Fig. 9. Switch fabric stress test histogram

Also when jitter with regard to frame length was checked, the results were in line with the expectations with showing stable low deviations in the no traffic case (where the measurement frame was the only payload transmitted on the network) and the case where minimum length frames were sent over the network.

The traffic made from full length frames, as expected were suffering more (Figure 10), as both their transmission time and the possible waiting time is longer.




Fig. 10. Jitter with different frame sizes

The histogram (Figure 11) also shows the expected behavior, where in the no other traffic situation the roundtrip times have little variation, where at the maximum length frame measurements show a relatively high jitter, which is introduced by the long transmission time.




Fig. 11. Histogram of jitter for different frame sizes

One of the main sources of industrial skepticism against Ethernet is the non-deterministic operation of the network. While our experiments show that the expected non-deterministic transmission times do happen and that in case of long frames, the jitter is high compared to the minimal or the typical transmission times, it is important to point out, that this low jitter values only have relevance in a few applications only. The measurements showed a maximum deviance of approximately 5000 nanoseconds per hop. In more realistic traffic mixes, where the industry-typical short to medium sized frames would take the bulk of the traffic, also with the typically low bandwidth utilization, the experienced jitter would be much lower. Still, even if the worst case with maximum length frames is taken, compared to real applications, like process automation and most of manufacturing, the tolerance of the measurement or control loop is several magnitudes higher than the jitter introduced by the communication subsystem.

The insignificance of jitter is especially valid if there is a Layer 3 device in the path (e.g., firewall or a gateway), where it is expected, that device alone will introduce a higher delay than the whole section of L2 network before and after. In case operator reaction is needed, the tolerance is even higher, as humans react in fraction of a second or close to a second, so the delay introduced by the operator would be much higher than the jitter.

X. OUTLOOK

In the foreseeable future, industrial network installations will follow the chained structures, mainly because of the industry's traditional skepticism for changes and partly because of the costs associated with additional cabling.

The evolution of industrial Ethernet will proceed towards higher device integration, where first, the switches will be included in the automation controllers and other devices, and which will proceed with integrating most of the active network devices.

It is expected, that only the site level high performance devices, which have a large resource need on their own will still be independent devices, but these will most probably

be shared between the industrial and telecom fields. This sharing is opening a very promising field, *Software Defined Networking* (SDN) for the industrial market as the support functions will be available in the high-end active devices supplied by networking companies and the industrial area offers a relatively simple, static and homogeneous world.

The strongly limited industrial environment represents a good starting point for deploying networks using SDN. The range of devices is much smaller as compared to office environments and, especially in green field deployments; the network is controlled by one vendor.

Traffic on an industrial network also tends to be deterministic, having periodic communication as part of the control loops, periodic updates of the logs and infrequent best effort traffic related to end node maintenance.

SDN fits very well the wishes of industrial vendors:

- hiding network complexity, which lowers the pressure on engineers,
- allows simpler service deployment,
- possibility for central resource management,
- seamless integration of wireless technologies.

A. Network complexity

The network topology and node types provide a good basis for experimenting with SDN in an industrial setting. The used topologies are not complex, the type of nodes is limited and the typical operation is static.

So an abstraction from the physical nodes could be implemented with limited effort as there is no need for a generic solution. In a typical installation, even before using integrated switches, a handful of types would be deployed together with a large number of end nodes, which from the network's viewpoint, operate nearly the same. As the typical traffic is machine-to-machine, the dynamism of the traffic is limited compared to non-industrial scenarios.

The use of SDN would although come from a different motivation as compared to, e.g., telecom installations. While there the use of SDN would enable better dynamism and abstraction of control from forwarding to allow more flexible network utilization, the industrial scenario is more about having a centralized control over the network, including forwarding decisions (no local configuration of the integrated devices is necessary) and for implementing a type of call admission control. With admission control, the SDN controller could check if a newly deployed control loop or other traffic source could be securely fit into the current resource utilization or there is a need for more changes in the infrastructure.

Also the granularity of control would be extended, as with SDN it is possible to influence forwarding decisions on a per flow basis. This means, that the engineer will be able to globally prioritize traffic flows, have a total overview on forwarding decisions and have a central entity to ensure that the configuration of devices is actually the one what is expected.

B. Service deployment

In and SDN case, the central management entity can change the configuration and forwarding behavior of the underlying devices. This could lead to cost savings in both infrastructure, as currently over provisioning is typical in the field, or in lowering the resources needed for engineering since the traffic estimates would be made by the SDN system and not the engineers.

An additional gain would be that an SDN system could deploy a new service without disturbing the current operation, which would reduce costs related to planned downtimes. The abstraction of hardware and central intelligence leads also towards simpler integrated devices.

C. Central resource management

Currently, SNMP-based Network Management Systems (NMSs) are widely used for monitoring the health and status of large network deployments. Using SDN could also here be beneficial, as the monitoring functionality would be extended with the ability of actively changing configurations and resource allocations if needed.

D. Wireless integration

Another key field currently is the integration of wireless networks into industrial deployments. Although there is skepticism for the usage of wireless solutions both because of determinism and security, already now, a growing part of the deployments use wireless solutions. Here SDN could help with integration of wireless technologies by checking if the needs of a new service, e.g., can be satisfied with a path having one or more wireless hops or a new rule has to be deployed into the network to steer the traffic of that service on a different path.

E. Vendor-neutrality

As also Ethernet deployments are getting old enough to be part of plant upgrades or new deployments with multiple suppliers, SDNs capability of giving a vendor-neutral abstraction of the network would be a great advantage.

XI. CONCLUSION

Our review shows that with regard to forwarding performance and latency, embedded switching solutions present a competitive solution compared to discrete units. Although, the offered set of features might differ and for managed functions, either the host CPU or an additional CPU for the switching board needs to be used, the performance expectation can be the same.

Soft switching on the other hand might be problematic when using a non-real time operating system and non-optimized drivers. If the software selection would move towards these, on the other hand, the flexibility of the platform would be limited. Our measurements showed that soft switches might be too slow to be used in a chained topology, but might be applicable in cases, where additional processing is required, for example as a controller with several network interfaces.

Our conclusion is that if there are no clear requirements for traffic monitoring capabilities exceeding the bandwidth of the host-switch module link, embedded switches are a viable and effective solution for low branching factor industrial networks.

Soft switches are a viable solution for implementing routers or other network functions, where the additional latency compared to the other switching solutions is not critical as the processing of the data on higher layers will contribute to more latency and jitter as the switching.

In a broader perspective, the identified problems for industrial Ethernet switching lie more on the feature set than in the hardware performance.

REFERENCES

- [1] G. Kálmán, D. Orfanus, and R. Hussain, "An overview of switching solutions for wired industrial ethernet," in *Proceedings of the Thirteenth International Conference on Networks*, IARIA, Nice, France, 2014, pp. 131–135.
- [2] J. Kay, R. Entzminger, and D. Mazur, "Industrial ethernet- overview and best practices," in *Pulp and Paper Industry Technical Conference, Conference Record of 2014 Annual*, June 2014, pp. 18–27.
- [3] F. Cen, T. Xing, and K.-T. Wu, "Real-time performance evaluation of line topology switched ethernet," *International Journal on Automation and Computing*, vol. 5, no. 4, pp. 376–380, October 2008.
- [4] J. Jasperneite and P. Neumann, "Switched ethernet for factory communication," in *Emerging Technologies and Factory Automation, 2001. Proceedings. 2001 8th IEEE International Conference on*, October 2001, pp. 205–212, vol.1.
- [5] J.-P. Georges, N. Krommenacker, T. Divoux, and E. Rondeau, "A design process of switched ethernet architectures according to real-time application constraints," *Engineering Applications of Artificial Intelligence*, vol. 19, no. 3, pp. 335 – 344, 2006.
- [6] N. Kakanakov, M. Shopov, G. Spasov, and H. Hristev, "Performance evaluation of switched ethernet as communication media in controller networks," in *Proceedings of the 2007 International Conference on Computer Systems and Technologies*, ACM, 2007, pp. 30:1–30:6.
- [7] J.-P. Georges, T. Divoux, and E. Rondeau, "Comparison of switched ethernet architectures models," in *Emerging Technologies and Factory Automation, 2003. Proceedings. ETFA '03. IEEE Conference*, vol. 1, September 2003, pp. 375–382 vol.1.
- [8] G. Marsal, B. Denis, J.-M. Faure, and G. Frey, "Evaluation of response time in ethernet-based automation systems," in *Factory Communication Systems, 2006 IEEE International Workshop on*, 2006, pp. 95–98.
- [9] H.-T. Lim, L. Volker, and D. Herrscher, "Challenges in a future ip/ethernet-based in-car network for real-time applications," in *Design Automation Conference (DAC), 2011 48th ACM/EDAC/IEEE*, June 2011, pp. 7–12.
- [10] T. Skeie, S. Johannessen, and O. Holmeide, "Timeliness of real-time ip communication in switched industrial ethernet networks," *Industrial Informatics, IEEE Transactions on*, vol. 2, no. 1, pp. 25–39, February 2006.
- [11] J. Jasperneite, M. Schumacher, and K. Weber, "Limits of increasing the performance of industrial ethernet protocols," in *Emerging Technologies and Factory Automation, 2007. ETFA. IEEE Conference on*, September 2007, pp. 17–24.
- [12] A. Manita, F. Simonot, and Y.-Q. Song, "Multi-dimensional markov model for performance evaluation of an ethernet switch," INRIA, Tech. Rep. 4813, 2003.
- [13] K. Beretis and I. Symeonidis, "Experimental evaluation of end-to-end delay in switched Ethernet application in the automotive domain," in *SAFECOMP 2013 - Workshop CARS (2nd Workshop on Critical Automotive applications : Robustness & Safety) of the 32nd International Conference on Computer Safety, Reliability and Security*, M. ROY, Ed., Toulouse, France, 2013, p. NA.
- [14] J.-D. Decotignie, "Ethernet-based real-time and industrial communications," *Proceedings of the IEEE*, vol. 93, no. 6, pp. 1102–1117, June 2005.
- [15] A. Jacobs, J. Wernicke, S. Oral, B. Gordon, and A. George, "Experimental characterization of qos in commercial ethernet switches for statistically bounded latency in aircraft networks," in *Local Computer Networks, 2004. 29th Annual IEEE International Conference on*, November 2004, pp. 190–197.
- [16] H. Systems, "Real time services (qos) in ethernet based industrial automation networks," Hirschmann, Tech. Rep., 2000.
- [17] M. Knezic, B. Dokic, and Z. Ivanovic, "Performance evaluation of the switched ethercat networks with vlan tagging," *Serbian Journal of Electrical Engineering*, vol. 9, no. 1, pp. 33–42, 2012.
- [18] M. Kezunovic, A. Sprintson, J. Ren, and Y. Guan, "Signal processing, communication, and networking requirements for synchrophasor systems," in *Signal Processing Advances in Wireless Communications (SPAWC), 2012 IEEE 13th International Workshop on*, June 2012, pp. 464–468.
- [19] "IEEE Standard Profile for Use of IEEE 1588 Precision Time Protocol in Power System Applications," IEEE Std. C37.238-2011, 2011.
- [20] A. Mahmood, R. Exel, and T. Sauter, "Delay and jitter characterization for software-based clock synchronization over wlan using ptp," *Industrial Informatics, IEEE Transactions on*, vol. 10, no. 2, pp. 1198–1206, May 2014.
- [21] Siemens, "Profinet answers for industry," Siemens, Tech. Rep., 2010.
- [22] *1588-2008 IEEE Standard for a Precision Clock Synchronization Protocol for Networked Measurement and Control Systems*, IEEE Std., 2008.
- [23] F.-J. Goetz, "IEEE 802.1AS bt (gPTP) & IEEE 1588 v3 (PTP v3)," 2013, presentation of Siemens at IEEE 802.1 TSN WG Meeting.
- [24] S. RuggedCom, "Latency on a switched ethernet network," RuggedCom, Tech. Rep., 2008.
- [25] J.-P. Georges, T. Divoux, and E. Rondeau, "Network calculus: Application to switched real-time networking," in *Proceedings of the 5th International ICST Conference on Performance Evaluation Methodologies and Tools*, ser. VALUETOOLS '11. ICST, Brussels, Belgium, Belgium: ICST (Institute for Computer Sciences, Social-Informatics and Telecommunications Engineering), 2011, pp. 399–407.