

An Overview of Switching Solutions for Wired Industrial Ethernet

György Kálmán, Dalimir Orfanus, 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. Active devices interconnecting the nodes are switches. In an industrial deployment, nodes are typically located on the same Layer 2 network. Routers or firewalls are almost exclusively used at the network edges. The current trend on engineering of industrial devices is to include an embedded Ethernet switch, instead of using discrete units. This paper is giving an overview on switch implementation possibilities, with respect to performance, features, logical architecture and flexibility.

Index Terms—industrial Ethernet, switch, embedded, discrete, soft switch, forwarding, performance, QoS

I. INTRODUCTION

Ethernet is already the dominating technology on the control and higher levels of an automation network and is expected to spread also into the field networks.

Because of resource constraints and Quality of Service (QoS) requirements, most of the automation networks are implemented as Local Area Networks (LANs). Although separating firewalls or routers are used between the automation network and the company network or the internet, inside the system, the network is typically interconnected on layer 2, by switches, as shown on figure 1.

The paper is structured as follows: the second section provides an background overview on industrial Ethernet, focusing on topologies and QoS. 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.

II. INDUSTRIAL ETHERNET BACKGROUND

Industrial Ethernet enables the use of standard Ethernet devices and the IP protocol suite in automation networks. By implementing networks based on Ethernet, vendors can create infrastructures, which 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 arise from the fact that the Ethernet networks are replacing the fieldbuses. A heritage of the fieldbus past is the dominant use of bus-like topologies (figure 2 resulting in suboptimal operation of Ethernet [1], [6].

The most challenging topology type are long chains of switches (figure 2, which are often closed to rings. While

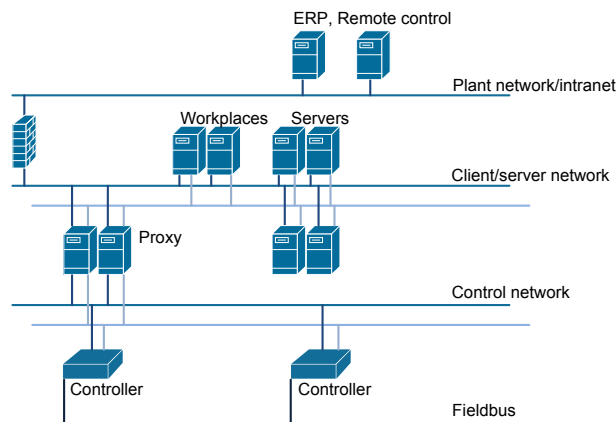


Fig. 1. An example of automation network architecture

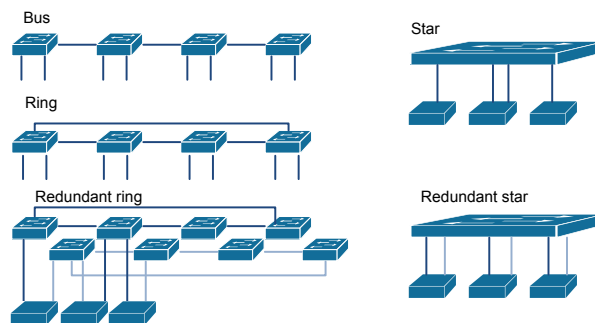


Fig. 2. Industrial Ethernet topologies

Rapid Spanning Tree Protocol (RSTP) was designed with loop-avoidance in mind, it is a widely used redundancy protocol in industry. In a ring structure, RSTP will disable one link and render the network into a special tree, a line of switches.

Although a line is a valid Ethernet topology and the technology will work, the industrial network will suffer from scalability issues at a much smaller end-node count than it could be expected from office experience [2]. The very long and sparse spanning tree (in practice only a single path) is having a low branching factor and can lead to excess latency and jitter [3], [7], [8].

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 [4], [5].

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 recent trend is to include low port count switches into the devices, so that devices can be interconnected into a cascade or a low branching factor tree without the use of external switches. This trend also shows that latency and jitter from the long sparse trees will be persistent in coming years and in combination with time-critical automation tasks, might be a limitation for the scalability of networks [9], [11].

Integrated switches are expected to lower the cost of building the network but without a penalty in features or QoS. In the following sections, we provide an overview of the typical embedded switch architectures and how they could be fitted into the industrial landscape.

III. ARCHITECTURE POSSIBILITIES

With a few exceptions, only managed switches are being deployed in industrial Ethernet networks. This is a result of the different requirements arising in the industrial environment compared to 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. In a direct comparison, unmanaged switches fail to meet the redundancy and logical segmentation requirements of the upper network layers of the automation network architecture.

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

There are devices in the office networks, located between these two levels, called smart switches. They offer the most of the managed switch functions, but lack, e.g., SNMP management. Introducing a similar class of devices into automation networks might be of interest, since having a more grained approach on switch features can lead to a more cost-effective network architecture.




Fig. 3. 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 3. The high number of standalone switches and the rugged hardware leads to a high per port price.

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 these number of ports 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.

IV. EMBEDDED SWITCHES

Integrating a switch module into devices like controllers is on the agenda of automation vendors. These modules could take the tasks of discrete switches in the lower levels of automation networks. The construction of these units is potentially cheaper than using a separate switch, as, e.g., a low-end switch fabric, 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 analog 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 in the host.

Another solution might be with an additional cost, to include a CPU on the switch module and implement the managed functions on the board, thus in practice be an independent managed switch inside the housing of the host.

These 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. 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 an other one, even on a different operating system if needed.

The possible drawback with these modules is that the host is still only connected with one internal port, which means that the host has no possibility to monitor the whole network traffic, if the aggregated bandwidth use exceeds the host's bandwidth.

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 802.3D), 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 are 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 Group Registration Protocol (MGRP), 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.

V. 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.

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.

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 chipset, and used driver and operating system as well. Also the limitations of the bus system and the network interfaces are summed, which can lead to insufficient performance in a 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.

VI. 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 embedded solution can also be viable, as 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.

VII. 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 behaviour 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 behaviour in a cascade, we decided to use the EtherCAT network consisting of the master (P2020 board) and two slaves (figure 4). 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. 4. 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 wireline latency (L_{WL}) and the queuing latency (L_Q) [10].

L_{SF} depends on the frame length. The results are expected to show a linear growth of the latency with the longer frames [12].

Our architecture related measurement scenarios deals with latency between endpoints (both with 100Mbps) of serial connected switches and without any additional interfering traffic




Fig. 5. Test segment setup

(see Figure 5). 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 is expected to be the result of the simpler architecture, as the module in the tested form implements only an unmanaged switch.

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

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 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. 6. Scenarios 1-4




Fig. 7. Histogram 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 5 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 6).

Also the histogram on figure 7 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 more wide distribution of frame latencies for longer switch chains.

VIII. 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.

REFERENCES

- [1] F. Cen, T. Xing and K. Wu, Real-time Performance Evaluation of Line Topology Switched Ethernet, International Journal on Automation and Computing, October 2008, pages 376–380.
- [2] N. Kakanakov, M. Shopov, G. Spasov and H. Hristev, Performance Evaluation of Switched Ethernet as Communication Media in Controller Networks, International Conference on Computer Systems and Technologies - CompSysTech 2007, Article No. 30.
- [3] G. Marsal, B. Denis, J. Faure and G. Frey, Evaluation of Response Time in Ethernet-based Automation Systems, Conference on Emerging Technologies and Factory Automation 2008, pages 380–387.
- [4] A. Manita, F. Simonot and Y. Song, Multi-Dimensional Markov Model for Performance Evaluation of an Ethernet Switch, Research report at INRIA, no. 4813, 2003.
- [5] K. Beretis, I. Symeonidis, Experimental evaluation of end-to-end delay in switched Ethernet application in the automotive domain, International Conference on Computer Safety, Reliability and Security, Toulouse 2013.
- [6] J. Jasperneite, P. Neumann, Switched Ethernet for Factory Communication, Conference on Emerging Technologies and Factory Automation 2001, pages 205–212.
- [7] H. Lim, L. Völker and D. Herrscher, Challenges in a Future IP/Ethernet-based In-Car Network for Real-Time Applications, Design Automation Conference 2011, pages 7–12.
- [8] T. Skeie, S. Johannessen and O. Holmeide, Timeliness of Real-Time IP Communication in Switched Industrial Ethernet Networks, IEEE Transactions on Industrial Informatics Vol.2, No.1, February 2006, pages 25-39.
- [9] A. Jacobs, J. Wernicke, S. Oral, B. Gordon and A.D. George, Experimental Characterization of QoS in Commercial Ethernet Switches for Statistically Bounded Latency in Aircraft Networks, Report, High Performance Networking Group, University of Florida
- [10] RuggedCom, Latency on a Switched Ethernet Network, 2008.
- [11] Hirschmann, Real Time Services (QoS) In Ethernet Based Industrial Automation Networks, 2000.
- [12] J. Georges, T. Divoux and E. Rondeau, Network Calculus: Application to switched real-time networking, ICST 2011, pages 399–407.