Characterizing the Tradeoff between Fault Tolerance and Cost of Redundant TSN Networks

Mateu Jover, Manuel Barranco, Josep Naranjo, Julián Proenza, Alberto Ballesteros
Departament de Ciències Matemàtiques i Informàtica, Universitat de les Illes Balears, Spain
mateu.jover@uib.es, manuel.barranco@uib.es, josep.naranjo1@estudiant.uib.cat, julian.proenza@uib.es, a.ballesteros@uib.es

***Abstract*—New emerging Distributed Control Systems (DCSs), like Substation Automation Systems (SASs) of Smart Power Grids, raise new requirements on their underlying control networks. To meet these new requirements, both Industry and Academia are promoting the Time-Sensitive Networking (TSN) Ethernet standards. In particular, TSN includes mechanisms to exchange information simultaneously through several paths of practically any spatially redundant network topology. This topological flexibility can offer a better balance between fault tolerance (FT) and redundancy cost (extra number of components) than classical Industrial Ethernets. However, the mentioned TSN mechanisms may also increase the cost in terms of extra latency and jitter, which could jeopardize real-time communications. In this paper we show our ongoing work to experimentally assess this extra latency and jitter and, thus, characterize the benefits of TSN in terms of balance between FT and cost.**

***Index Terms*—fault tolerance, network redundancy, TSN**

I. Introduction

Many industrial automation systems that emerge from the integration of *Operation Technologies* (OT) with *Information Technologies* (IT) are supported by *Distributed Control Systems* (DCSs). A DCS is made up of several computing nodes that coordinate with each other, exchanging information through what is called a *control network*. OT-IT systems impose new requirements on DCSs, such as network convergence and advanced fault tolerance (FT) with low-moderate cost [1]. To satisfy these new requirements, both Industry and Academia are intensively promoting a new set of Ethernet standards known as *Time-Sensitive Networking* (TSN).

Some OT-IT systems that will benefit from TSN are the new energy production/distribution automation systems that are expected to compose the *Smart Power Grid* [1]. In particular the reliability of *Substation Automation Systems* (SASs) is key for the smart grid, as they transform electricity between transmission lines and between those and local distribution systems. Thus, we are specially interested in providing highly reliable TSN *Substation Communication Networks* (SCNs).

The most common mean to achieve communication reliability is *Fault Tolerance* (FT) based on *redundancy*. Specifically, the two main *Industrial Ethernet* protocols that have been classically used to provide SCNs with FT are the *High-availability Seamless Ring* (HSR) and the *Parallel Redundancy Protocol* (PRP) specified in IEC 62439-3. Basically, both of them provide FT by means of a *spatially redundant network*, i.e. a network that includes redundant (multiple different) paths between each pair of critical end nodes (nodes for short).

The potential FT benefits of TSN over HSR and PRP have been discussed to some extent in the literature, e.g. [2] [1] [3]. One important advantage of TSN is its flexibility from the point of view of the network topology. In HSR the redundant topology is restricted to a ring or multi-rings; which limits the potential number of redundant paths. As regards PRP, it allows using any topology. However, PRP relies on the use of two independent networks, which must have the same latency and thus almost an identical topology. Thus in PRP there are two alternatives to provide a redundant path between a pair of nodes. One is to consider a non-redundant network and then duplicate it, which may not be cost-effective (in number of components). The other one is to consider an inherently redundant network and, then, duplicate it; which is even less cost-effective. In contrast to HSR and PRP, TSN allows implementing any redundant topology and does not require separate networks. This topology flexibility of TSN opens room for exploring multiple topology (and routing) alternatives to provide nodes with redundant paths within a single network, thereby achieving a potential better balance between FT and extra number of components.

The general objective of the present work is to propose guidelines to provide redundant topologies, for TSN-based SCNs, with an adequate balance between FT and cost. Note that TSN switches (bridges according to TSN terminology), and TSN *Network Interface Cards* that nodes use to communicate, are more complex than typical Ethernet devices, they are not consolidated in the market, and their firmware can be updated as they include new/reviewed TSN standards. Thus TSN devices are priced much higher than those of other Ethernets and, hence, redundant TSN networks are hardly price-effective even if they do not include many components.

However, note that part of the complexity of TSN bridges lies in the mechanisms they include to support network redundancy. These mechanisms can introduce *latency* and *jitter* in the communications, which may prevent deadlines of *hard real-time traffic* from being met. Thus, while waiting for TSN prices to drop, it is necessary to assess the cost in terms of the extra latency and jitter that the TSN bridges that support network redundancy may introduce in communications.

In this paper we show our first results in this direction by experimentally comparing the FT, latency and jitter of two of the most simple redundant topologies than can be respectively implemented with PRP and TSN using a similar number of components, namely a PRP network with *dual-line topology*, and a TSN network with *ladder topology*.

II. NETWORKS' BASICS

Figure 1 depicts the two spatially redundant topologies we consider. The PRP dual-line consists of two independent parallel identical networks with a line topology. Each network (line) is formed by a succession of N full-duplex Ethernet switches. As concerns the TSN ladder, it is a single network whose topology results from including two lines (top and bottom) and then connect each pair of full-duplex switches (one from the top and one from the bottom) along the lines. We start by considering just $N = 2$ switches per line, and just one transmitter node and one receiver node. These nodes represent the two ones located at the opposite extremities of the network, which allows us to measure the latency and jitter considering the longest paths within each topology.

To understand the FT of each topology, it is necessary to understand how frames are transmitted and forwarded. Generally speaking, when a transmitter wants to exchange a piece of data, encapsulates and transmits it in a frame. Each frame the transmitter generates to transmit a new piece of data, e.g. a new sampled sensor value, is called a *frame edition*.

In PRP each frame edition has an associated *sequence number*. After generating a frame edition, the PRP transmitter creates two *frame replicas* (copies) of that frame and inserts the frame edition's sequence number in both of them. This operation is known as *duplication*. The transmitter simultaneously transmits each replica through a different network (each replica in a different line in this case) and, then, each replica is forwarded by the corresponding switches. Note that in PRP the switches are agnostic about the PRP protocol; they just forward frames as in a regular switched-Ethernet network.

In absence of faults each frame replica is expected to traverse its corresponding line and reach the receiver. The receiver then uses the sequence number of both replicas to identify that they are copies of the same frame edition and, then, to discard one of them. This operation is known as *de-duplication*. Given a frame edition, from a topological perspective, PRP dual-line can tolerate faults effecting several links and switches as long as one of the lines is non-faulty. Since it is enough that each line suffers from a single fault for the frame edition to be lost, it can be said that in the worst case PRP dual-line guarantees tolerance to 1 fault.

In TSN the transmission and forwarding over a spatially redundant network is carried out by means of a set of mechanisms standardized as IEEE 802.1CB *Frame Replication and Elimination* (FRER). As in PRP, each frame edition has a sequence number, and the transmitter inserts that number within each one of the frame replicas. Unlike PRP, the FRER mechanisms in the transmitter can generate as many frame replicas (of a given frame edition) as network interfaces the transmitter has. Anyway, in the case of the ladder topology, FRER generates just two replicas. Another difference is that TSN bridges are aware of the replication, so that they include FRER mechanisms. Each one of the bridges both de-duplicates and duplicates. Specifically, when a bridge receives the first replica of a given frame edition, it accepts that replica and

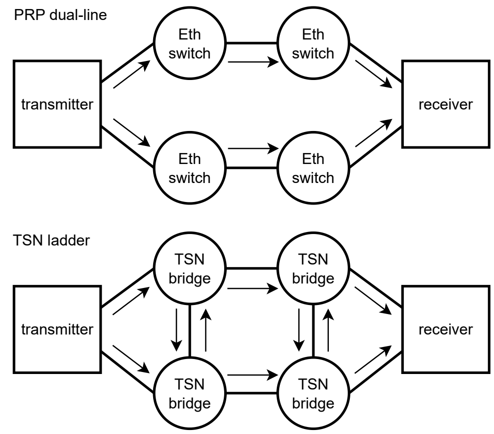

Fig. 1. Topologies

discards any later replica of that edition. Then, the bridge duplicates the accepted replica by generating as many replicas (copies) as ports it forwards the frame through.

The choice of forwarding ports for each bridge depends on the topology and the routes configured on that topology. To fully exploit the redundancy of our ladder topology, we configured each bridge to forward each frame edition through its *horizontal* link and its *vertical* link, as depicted in Figure 1. From the topological point of view, our TSN ladder can tolerate faults effecting several links and bridges as long as there is at least one non-faulty path. However, there are combinations in which 2 faults are enough to lose all the replicas of a given frame edition. Specifically the fault combinations are the following: (1) two vertically adjacent bridges; (2) two vertically adjacent horizontal links; (3) two obliquely adjacent bridges; (4) an obliquely adjacent bridge and horizontal link, e.g. 1st bridge of the bottom line and 2nd horizontal link (the central one) of the top line.

III. TESTBED

In order to experimentally test the FT and measure the latency and jitter of the just-described networks, we built both of them using the same testbed. The testbed is basically composed of 4 FPGA-based full-duplex Ethernet switches, 2 nodes (one transmitter and one receiver), a set of 1 Gbit/s Ethernet cables, and a *Network analyzer* (NA).

The switches are the *SoCe 1G MTSN* bridges [4] v22.3, each with 4 TSN Ethernet ports supporting 10/100/1000 Mbps. The bridges can be configured to work as off-the-shelf switches to implement our PRP network, or as TSN bridges to build our TSN network. In this later case the bridges can be set up to implement the main TSN standards, e.g. the *traffic shapers* 802.1Q(bv, av), the *filtering and policing mechanisms* 802.1Qci, clock synchronization 802.1AS, etc. However, we want to measure the latency and jitter introduced by the TSN mechanisms bridges use to (de-)duplicate and forward over a redundant network, independently of any specific application, scheduling or policing. Thus, we configured our TSN bridges to just include FRER 802.1CB, so that they basically de-duplicate, duplicate and forward frames.

Each node is built using hardware for embedded devices; specifically, a Jetway JBC373F38-525-B barebone, which includes an Intel Atom executing Ubuntu 16.04 SO, and 4 regular (non-TSN) Ethernet *Network interface cards* (NICs). As explained later, we measure the latency and jitter introduced by the communication subsystem, i.e. by the bridges and links. Thus we need neither a real-time SO, nor TSN NICs.

The purpose of the transmitter node is to generate and transmit a given number of dummy data frames, to emulate the periodic transmission of a piece of data, e.g. a periodically sampled sensor value. Please recall that in both, PRP and FRER, each replica includes the sequence number of the frame edition of which it is a copy. In PRP the sequence number is inserted in the payload within which it is known as the *Redundancy Control Trailer* (RCT); in FRER it is included after the *Ethertype* field within the so called *Redundancy Tag* (R-Tag). In order to generate the frame editions and transmit the replicas with the corresponding sequence number (using the appropriate frame format), we implemented the transmitter application as a script that relies on a set of libraries. The script includes parameters to specify aspects such as: the number of frames to transmit, the size of their payload, their priority and the inter-transmission time. The script uses the libraries to basically: (1) build up a PRP / FRER frame; (2) create two replicas of that frame; and (3) request the SO for the transmission of both replicas, each through a specific NIC.

The role of the receiver node is to log the frame replicas received at each one of its NICs. This allows us to assess the FT capabilities of the network by checking that the frame replicas that are logged for each NIC are those that were expected according to the redundancy of the topology. To log the replicas we implemented, for each NIC, a pair of threads, a buffer and a log file. When a given NIC receives a frame replica, one of the corresponding threads writes the relevant information of that replica at the buffer. Eventually the other corresponding thread transfers this information to the log file dedicated to that NIC. The relevant information that is logged includes aspects like the sequence number and the frame size.

As regards the *Network analyzer* (NA), it is a testing tool that allows us to collect the necessary information to measure the latency and jitter caused by the communication subsystem on the frame replicas. Specifically, the NA is the *RELY-TSN-LAB* [5], which has two measurement interfaces. Any of these interfaces can be connected to a given link to log and take a timestamp of each frame that traverses that link. The two NA measurement interfaces use the same clock. Thus to calculate the latency of a frame in a given path, e.g. as shown in Figure 2, we log and take a timestamp of that frame when it traverses the first link and when it traverses the last one. Then we subtract the first timestamp from the second. As for calculating the jitter in a path, we repeatedly measure its latency and then subtract the lowest measurement from the highest one.

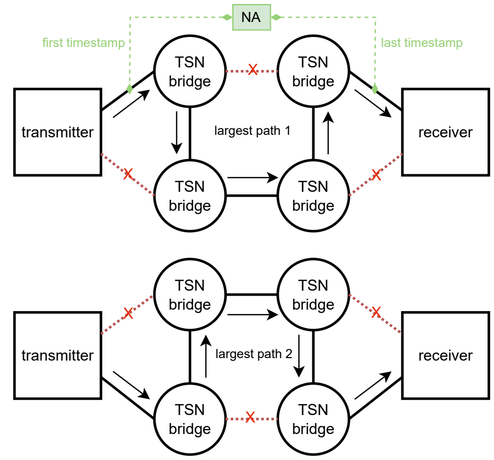

Fig. 2. Degraded TSN ladder

IV. Experimental results

We carried out different sets of experiments to assess the FT, the latency and the jitter. All experiments where conducted at 1 Gbit/s and using links of around 1.5 m each (so that the propagation time on the links is negligible).

In the 1st set of experiments we corroborated that, in absence of faults, the receiver receives exactly two replicas of each frame edition (one per NIC) in both, PRP and TSN.

The 2nd set consisted of several fault-injection experiments in the TSN network, aimed to corroborate the superior FT of the TSN ladder topology. In each experiment we manually unplug a given combination of links of the ladder topology and, then, we triggered the transmission of 1000 frame editions, one every 60 ms. In each experiment we set a specific payload size and frame priority. Specifically, we considered payloads of either 64 or 1420 bytes, and priorities 2 and 5. We successfully observed that the TSN receiver always receives at least one replica of each frame edition as long as there is a non-faulty path from the transmitter. In particular, we corroborated that the receiver receives exactly one replica of each frame edition, also when the only non-faulty path is one of the two largest symmetrical paths that may remain in a ladder topology after a series of link failures. These two paths are shown in Figure 2; as can be seen, for a frame replica to reach the receiver, it has to switch from one line to the other repeatedly (following a zigzag pattern). We call the resulting topology shown in Figure 2 as *degraded TSN ladder*; moreover, we refer to these two paths as *largest path 1* and *largest path 2*.

The rest of the experiments were aimed to assess the negative impact, on the latency and jitter, of TSN due to the use of FRER within its bridges to (de-)duplicate and forward. For this purpose we considered 3 topological scenarios, namely: (1) the *fault-free PRP dual-line*, (2) the *fault-free TSN ladder*, and (3) the *degraded TSN ladder*.

Note that when the PRP and the TSN networks are not affected by any fault, i.e. topological scenarios (1) and (2), the first replica (of any given frame edition) that reaches the receiver is either the one that crosses the top path (top line), or the one that crosses the bottom path (bottom line). Thus,

TABLE I
RESULTS

	mean (ns)	median (ns)	std (ns)	jitter (ns)
(1) fault-free PRP dual-line				
top path	28187.41	28188.00	12.96	72.00
bottom path	28173.18	28170.00	13.94	81.00
(2) fault-free TSN ladder				
	mean (ns)	median (ns)	std (ns)	jitter (ns)
top path	28296.42	28296.00	13.51	81.00
bottom path	28282.36	28287.00	12.77	81.00
(3) degraded TSN ladder				
	mean (ns)	median (ns)	std (ns)	jitter (ns)
largest path 1	55473.71	55476.00	17.08	108.00
largest path 2	55475.00	55476.00	17.71	117.00

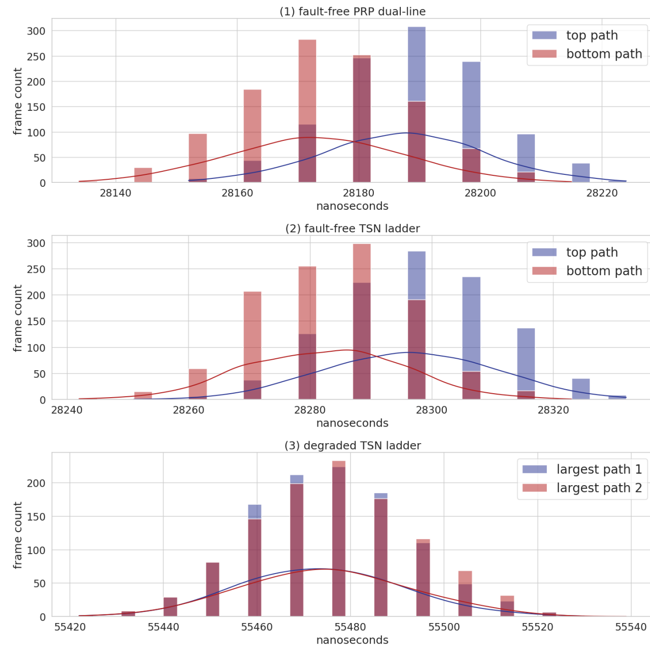

Fig. 3. Histograms of the results

to compare PRP and TSN in absence of faults, we measured the latency and jitter in these two paths in both scenarios.

As regards scenario (3), note that the largest possible path within the ladder topology can be either the largest path 1 or the largest path 2. Thus, to assess the latency and jitter when the TSN ladder topology is degraded to its most, we measured them on both paths.

For all 3 scenarios we considered small and large frames, i.e. with payload sizes of 64 and 1420 bytes respectively. We observed that the results do not noticeably depend on the frame size, thus, due to page limitations, Figure 3 and Table I only show the results for large frames.

In absence of faults, i.e. in scenario (1) and (2), we observe that the latency of the top path is slightly higher than that of the bottom path. Anyway, this difference, around ≤ 20ns, can be considered negligible and caused by differences in the manufacturing and assembly of the components that constitute the bridges. In fact, we observe almost no difference between the two largest paths of scenario (3).

Results of (1) and (2) also show that, in absence of faults, TSN (FRER) bridges in the ladder introduce an extra latency of around 109ns (54ns per bridge) in both the top and the bottom paths. This extra latency per bridge is very small; equivalent to the transmission time of 54 bits at 1Gbit/s. Moreover, results show that TSN bridges do not increase the jitter, as the jitter difference between (1) and (2) is 0 or lower than the measurement precision of the NA (≈ 12ns).

Results of scenario (3) show that the latency of the largest paths is almost twice than that of the top/bottom ones. This was expected since the largest paths have twice the hops of the top/bottom paths. Interestingly, results show that, compared to the top/bottom paths, the largest paths increase the jitter by around 31ns in average. Since scenarios (1) and (2) show that TSN bridges do not introduce extra jitter compared to PRP, this jitter increase is not due to FRER but to the fact that the largest paths have a greater number of bridges than the top and bottom paths. Anyway, this increase is so low that, even if we considered more bridges (up to the maximum number of hops acceptable by TSN clock synchronization, i.e. 7), the jitter increase would not be very noticeable.

V. Conclusions and Future Work

Compared to classic Industrial Ethernets, TSN allows implementing almost any redundant topology. This increases the potential redundant topology alternatives that can be explored to find a good tradeoff between fault tolerance (FT) and the monetary cost associated with extra components. It is evident that currently TSN components are very expensive, and thus a redundant TSN network is hardly price-effective. However, while waiting for TSN prices to drop, it is still necessary to characterize the cost in terms of the extra latency and jitter that the higher complexity of TSN bridges that support network redundancy may introduce in communications.

In this paper we describe our first results in this direction, by comparing two simple redundant networks: a PRP dual-line and a TSN ladder. Results are promising since they show that TSN bridges do not noticeably increase neither the latency nor the jitter. Thus, we plan to study other redundant topologies and additional TSN FT mechanisms that can be included in the bridges, e.g. error-containment.

Acknowledgment

This paper is part of project PID2021-124348OB-I00 funded by MCIN/AEI/10.13039/501100011033/ERDF, EU. Mateu Jover was supported by the Conselleria d'Educació i Universitats del Govern de les Illes Balears under the contract FPU2023-007-C.

References

[1] International Electrotechnical Commission (IEC), "Communication networks and systems for power utility automation – Part 90-13: Deterministic networking technologies", IEC TR 61850-90-13:2021-02(en), Switzerland, 2021.

[2] I. Alvarez et al., "Fault Tolerance in Highly-Reliable Ethernet-based Industrial Systems", Proceedings of the IEEE, 107(6):977-1010, June 2019.

[3] M. Jover et al., "Opportunities and Specific Plans for Migrating from PRP to TSN in Substation Automation Systems", 2023 IEEE 28th Int. Conf. on Emerging Tech. and Factory Automation (ETFA'23).

[4] SoCe, "1G MTSN – Multiport TSN Switch IP Core", https://soc-e.com/mtsn-multiport-tsn-switch-ip-core/

[5] RELYUM, "Time-Sensitive Networking Testing Tool, RELY-TSN-LAB, User Guide", https://www.relyum.com/web/rely-tsn-bridge-3/