

Analysis of Smart Grid Fault Recovery Protocols

Luana Uchôa, Silvio Quincozes, Juan Lucas Vieira,
Diego Passos, Célio Albuquerque
Laboratório MídiaCom, Instituto de Computação
Universidade Federal Fluminense
Niterói/RJ, Brazil
E-mail: {luchoa, sequincozes, juanlucasvieira}@id.uff.br,
{dpassos, celio}@ic.uff.br

Daniel Mosse
Department of Computer Science
University of Pittsburgh
Pittsburgh-PA, USA
E-mail: mosse@cs.pitt.edu

Abstract—The Smart Grid is a high availability system that requires fault recovery protocols to minimize downtime. In addition, recovery, in most cases, must be made within certain temporal constraints. The IEC 62493-3 standard defines the Parallel Redundancy Protocol (PRP) and High-availability Seamless Redundancy (HSR) as protocols for dealing with link layer communication failures in electrical substations. There are variations that perform the same functionality handling packets at the IP layer, such as iPRP. Those protocols promise zero-time recovery by proactively sending duplicate packets across distinct and independent paths. In this paper, those fault recovery protocols are analyzed by means of simulations. Our results show that, in realistic setups, recovery time is never actually zero and that, under certain conditions, this time can exceed the temporal requirements that some protection applications demand for proper operation.

Index Terms—PRP, HSR, iPRP, Smart Grids, Fault Recovery Protocols, Recovery Time

I. INTRODUCTION

The growing number of communication nodes in Substation Automation Systems (SAS) has increased the complexity in the engineering of communication networks and created new challenges. This kind of system was introduced into the electric grid control decades ago when centralized systems based on serial communication began to migrate to a distributed approach [1]. The IEC 61850 standard has played an important role in that process by defining protocols for the communication between different devices operate in a substation. That includes specialized devices (*e.g.*, Intelligent Electronic Devices (IED), Merging Units (MU), protection relays), as well as measuring and monitoring devices [2].

Reliable communication is essential for collecting and disseminating data, as well for issuing alerts and reporting problems detected in the electrical system. A critical failure in the communication network could mean the non-execution of important tasks, such as the switching of circuit breakers during a short-circuit. Risks include financial losses and even harm to human life, so vulnerabilities should not be overlooked [3].

Because of that, many redundancy protocols seeking to reduce the effects of critical failures were proposed and applied to substations. The IEC 62439-3 standard [4], for

example, defines two fault-tolerant protocols, the Parallel Redundancy Protocol (PRP) and the High-availability Seamless Redundancy (HSR) [5]. Both employ redundancy strategies to improve the reliability of the communication network used in electrical systems.

Those protocols deliberately duplicate packets and send them across disjoint paths. Unless failures to both paths happen simultaneously — which is unlikely — at least one copy of the packet should be delivered to its final destination. Because of this proactive nature, recovery time is assumed to be zero and failures are not likely to stop network operation [6].

Although redundancy is applied for all transmitted packets, the delivery time still depends on the length of the paths and the conditions of each composing link. Therefore, when computing the recovery time of those methods, it is important to consider the latency difference between each path used for replication. Thus, recovery time may not be zero as expected, unless paths are precisely symmetrical in terms of performance and current load.

The main contributions of this paper are to experimentally demonstrate this fact, to quantify its negative effects in realistic scenarios, and to bring awareness to it. To this end, simulations were performed in Mininet¹ using network topologies that mimic communication networks found in real substations and different traffic patterns. Our analysis reveals that the strategy of proactively replicating packets can be interfered by many variables, as the delivery times for different paths will likely be different — even if by a small margin. Given the stringent temporal requirements of protocols used by IEC 61850 (*e.g.*, GOOSE, SV [7]), this time difference may cause important functions of the electrical system to work poorly, such as the teleprotection which needs timely exchange of GOOSE messages to avoid damages in the electrical infrastructure in case of a short circuit. On the other hand, we also experimentally demonstrate that traffic control can be used to avoid those issues, thus contributing to a more reliable management of SAS.

This paper is organized as follows. Section II presents some background on teleprotection. Section III discusses related work. We describe state-of-art protocols for fault-tolerant com-

munication with temporal restrictions in Section IV. Our experiments and analysis are presented in Section V. Section VI discusses and illustrates how traffic control techniques can be used to improve performance of those protocols. Finally, a brief conclusion is presented in Section VII.

II. BACKGROUND ON TELEPROTECTION

Current power systems must meet the challenges of expanding power generation and transmission capacity to match the energy demands of modern society. This requires investments, including the creation of new generation, transmission and distribution infrastructures. Various types of equipment are used in this type of effort, some of which stand out: generators, transformers, transmission lines, disconnect switches and circuit breakers.

Circuit breakers, for example, are particularly important for the protection of power systems as their main function is to quickly interrupt circuit faults to mitigate possible damages caused to equipment in case of short circuits. In addition to current faults, circuit breakers must be able to interrupt rated load currents, transformer and reactor magnetization currents, capacitive capacitor bank currents, and no-load currents.

Table I illustrates the approximate cost of Three-Phase Transformers (TPT) and Series Capacitor Banks (SCB), according to the Brazilian National Electric Energy Agency (ANEEL). Such components, typically employed in electrical substations, are very expensive. Therefore, ensuring their proper functioning and avoiding damages is vital to keep overall costs controlled. In particular, they should be disconnected as soon as possible to reduce the possibility of damage in case of a fault.

TABLE I: Approximate prices for some components typically found on electrical substations.

Voltage (kV)	Power (MVA)	TPT Price	SCB
230/69	100	\$ 1,673,214.91	\$ 4,270,494.05
345/69	100	\$ 2,436,678.53	\$ 4,484,018.02
500/138	100	\$ 3,008,245.10	\$ 3,925,277.74
500/345	100	\$ 3,342,494.55	\$ 4,121,541.62

The functionality of monitoring and possibly acting on the power system in order to avoid damages to components is known as *protection*. When IEDs perform protective functions by exchanging messages over a communication network, that functionality becomes a *teleprotection*. Those messages may carry measurements of electrical parameters — *e.g.*, current — or commands for devices that will act on the power system — *e.g.*, a trip message for a circuit breaker to open a circuit. As such, they are very time-critical.

The IEC 61850 standard defines a set of protocols and architectural elements for the teleprotection of power systems. In particular, the standard defines three communication protocols: Sample Value (SV), used for the transmission of measurements of the monitored electrical parameters; GOOSE (Generic Object Oriented Substation Event), used for the exchange of commands between IEDs; and MMS (Manufacturing Message

Specification), used for the communication between IEDs and the supervisory system.

Maximum acceptable delays for the transmission of protection messages may vary according to their type, although SV and GOOSE are particularly demanding. Some substation protection messages have a maximum tolerable delay between 4 ms and 12 ms [8]. For trip messages to circuit breakers, in particular, delay should not exceed 4 ms. Trips to backup circuit breakers (after breaker failure timeout) must be received in at most 12 ms. For a circuit breaker re-closing, up to 8 ms is tolerated [8].

Therefore, besides robust protection schemes, teleprotection requires a communication network with high performance and reliability. To meet these requirements, fault-tolerance should be considered both at the device and link levels. In Section IV, fault recovery protocols are discussed.

III. RELATED WORK

The growing demand for electricity by the society has forced companies of the power sector to implement increasingly more efficient solutions. One important example is the Smart Grid, which is the evolution of the traditional electrical grid, including, among other differences, a heavy usage of communication technologies for improved control [9]. Due to the critical requirements of their components, there is a necessity to increase monitoring and management of Smart Grids' communication networks. As such, Software-Defined Networking (SDN) [10]–[12] can be regarded as a widely flexible architecture able to solve many of the aforementioned issues [13], [14].

As a requirement of Smart Grids, communication must be reliable, *i.e.*, the system must be able to detect failures and respond adequately. In substations, in particular, the stringent temporal restrictions for the teleprotection messages implies a demand for very fast response to failures. As such, even though there are classical protocols that can be used for fault-recovery, *e.g.*, the Spanning Tree Protocol (STP) and the Rapid Spanning Tree Protocol (RSTP), their response time is not fast enough for this kind of application. In particular, there will always be a non-zero time required for a fault to be detected and bypassed by triggering an alternate route [5], [15], [16], making this kind of approach unfeasible for teleprotection communication networks.

The easiest way to provide the required reliability levels, then, is to use redundancy as a strategy to proactively react to failure — even if it never happens [17]. There are some recommended redundancy protocols for different types of topology, *e.g.*, [18], [19], but only PRP and HSR were defined as standard protocols by IEC 62439-3 [4]. Those two protocols' operations and their application at real-time systems were defined by Kirrmann *et.al.* in [20] and [21], respectively.

Since then, many authors attempted to better understand the applicability and limitations of those protocols. Kumar *et.al.* [22] evaluated the performance of the transmission of SV-type packets between interconnected substations, with a digital protection scheme based on a PRP-compatible topology

and common substations' equipment, *e.g.*, MU, switches, IED and Red boxes. The experiments showed that the PRP-compatible topology was adequate for SV traffic in substations. However, when large errors and delays were introduced in the communication network, the IEDs behaved poorly. Araujo *et al.* [23] analyzed the network parameters required by the energy automation protocols and proposed a memory system for detecting duplicates and redundant frames in the PRP and HSR protocols.

Regardless of their proactive response to potential failures, this work will demonstrate that the redundancy protocols used in substations (PRP and HSR) do not actually provide zero-time recovery, as widely accepted by many authors in the literature [22]–[26].

IV. FAULT RECOVERY PROTOCOLS

The IEC 61850 standard is a guide to the implementation of communication between existing devices in a SAS, however, it does not deal with redundancy in the protection architecture, which is covered by the IEC 62439-3 standard. It addresses this with two protocols PRP and HSR. Those protocols have become a benchmark for ensuring fast processing, increased availability and minimum delay in data transfer [4]. Besides the standardized protocols, there are also variants, *e.g.*, IP Parallel Redundancy Protocol (iPRP) [27] which applies redundancy strategies at application-level using the IP layer to deliver the packets.

The main purpose of these protocols is to send copies of each packet through different interfaces, causing these packets to travel through disjoint paths until they reach the destination. In absence of failure, duplicates are discarded so that only one packet will actually be delivered at the upper layers. However, if a failure occurs, as long as there is still one path working properly, one of the copies will still be delivered, regardless of any failure recovery method used in the network.

Because they are able to deliver each packet without requiring any explicit, reactive failure correction, these protocols are usually assumed to provide zero-time recovery. However, the recovery time depends on the specifics of the used network paths, such as physical asymmetries and different patterns of concurrent traffic. Given the stringent time-requirements of protocols used for the communication of devices in the IEC 61850 standard, such as GOOSE and SV, further studying that non-zero-time recovery becomes important.

In the remainder of this section, basic concepts related to PRP (Section IV-A), iPRP (Section IV-B) and HSR (Section IV-C) are presented. PRP and HSR are already applied to the communication redundancy in Smart Grids in practice, while iPRP is a proposal to demonstrate that it is possible to perform a similar function at higher layers, while still respecting the time constraints of IEC 61850.

A. PRP: Parallel Redundancy Protocol

Designed to be used within networks of the automation industry and described by the IEC 62439-3 standard, PRP does not require a specific network topology. Each generated

packet is duplicated and sent through different interfaces. PRP assumes that each of those interfaces is connected to the destination node by separate paths. If a failure occurs on one of those paths, the copy of the packet sent through the other is still expected to be successfully received by the recipient.

While it does not require a particular topology to work, the effectiveness of PRP's proactive fault recovery technique depends on the type of physical redundancy available on the network. IEC 62439-3 requires all network infrastructure between the communicating IEDs to be duplicated in order for the alternative paths used by each interface to be completely independent.

PRP is implemented by means of a Link Redundancy Entity (LRE). The LRE is a mid-layer placed between the link and physical layers, which is responsible for handling messages before sending, at the source, and upon reception, at the destination. In practice, the upper layers should not notice the duplicity of messages nor be required to send copies of the same packet. Instead, LRE performs such tasks on their behalf. At the sender, LRE duplicates each packet, adds the PRP header and then sends each copy through different interfaces. At the receiver, LRE processes incoming messages to identify whether each message represents new content or a duplicate of an already received packet. In the latter case, the new packet is discarded. Otherwise, the message is delivered for processing by the upper layers [4], [16].

Duplicate control of messages occurs basically according to the information entered and loaded into the header of the *Ethernet* frame by the LRE, prior to transmission. This information is known as Redundancy Control Trailer (RCT), a unique sequence number. The RCT can be used to identify the network source of the message since MAC addresses of both are identical. Therefore, LRE is able to identify duplicates and guarantee the uniqueness of messages.

B. iPRP: IP Parallel Redundancy Protocol

PRP works well when its requirements are met. Nevertheless, it is not feasible for long-distance networks which are part of the emerging Smart Grid solutions. On large-scale networks — *e.g.*, those that provide communication for an entire distribution grid — routers are often required. That prevents the usage of PRP, which runs over Ethernet.

Motivated by this issue, iPRP [27] was designed based on the original PRP, but on top of IP. This variant does not duplicate the entire traffic as its precursor does. Instead, only time-critical UDP traffic is duplicated. iPRP is commonly installed on end devices that have multiple interfaces, such as streaming and receiving devices. Experiments by Popovic *et al.* [27] suggest that it is possible to use redundancy at the application layer, even for time-critical traffic.

C. HSR: High-availability Seamless Redundancy

The HSR protocol appeared as an alternative to PRP, with the objective of fulfilling the zero-time recovery, but without the need to duplicate the entire physical infrastructure of the network. This is possible thanks to the use of ring topologies,

which are capable of providing redundant communication without the need for additional devices [4], [16].

In an HSR topology, nodes typically exchange information through full-duplex connections. As in PRP, nodes send redundant messages — *i.e.*, duplicates of each transmitted packet. Unlike PRP, in which the duplicates are transmitted through distinct networks, in the HSR protocol they are transmitted in the same network, but in opposite directions of the ring. Similarly to PRP, LRE is responsible for guaranteeing that only the first copy of a message that reaches the receiver is delivered to the upper layers. Once received, LRE removes the message's HSR header and delivers it in its usual Ethernet format.

V. ZERO-TIME RECOVERY ANALYSIS

In Section IV, we presented protocols that promise zero-time recovery when a failure occurs. We argue that claim must be questioned since the packets' duplicates are transmitted by different paths, which may have different performance. In particular, it is necessary to analyze the concurrent traffic that exists in each path, as well as possible asymmetries.

We present an analysis based on experiments performed over a controlled environment, using the Mininet simulator to represent an electrical system containing IEDs and switches for communication. Mininet was running on a Ubuntu workstation with a Intel Core i7-8700 3.2 GHz processor and 16 GB of RAM. Experiments were performed using three different topologies (see Figure 1), representing typical scenarios in which PRP, iPRP and HSR could be applied. Each experiment is repeated using the same settings for all three protocols.

The three topologies used in this evaluation were designed based on the IEC 62439-3 standard. Both topologies 1 and 2 support the usage of PRP and iPRP by having two segregated LANs through which each duplicate of a packet can be sent. In particular, for Topology 2 we assumed two asymmetrical LANs to highlight that the PRP specification does not requires them to be identical. Then, to avoid forwarding loops, we used STP in LAN A.

Topology 3 is a ring network, supporting the execution of HSR. We used the *brctl* tool to allow IEDs to receive a packet through one network interface and forward it through the other. Notice that one can also treat the two directions of the ring as separate LANs, which allowed us to also execute iPRP on Topology 3.

The main performance metrics evaluated in this analysis are end-to-end delay and packet loss rates — both evaluated for each utilized path. To measure that, we implemented a simplified version of each protocol (without the LRE layer, for example) so we could log and process raw transmission and reception times, as well as delivered and dropped duplicates. Despite the simplifications of our implementations, each protocol adds headers of realistic sizes: 6 bytes for PRP messages, 6 bytes for HSR messages and 48 bytes for iPRP messages.

A unique identifier is associated with each packet and the receiving interface uniquely identifies which LAN/direction




Fig. 1: The three topologies were based on IEC 62439. Topologies 1 and 2 support PRP and iPRP, while Topology 3 supports HSR as well.

of the ring was used for that duplicate. That allowed us to calculate the the delay for the duplicate of each packet. We further defined a *recovery time* metric for a given packet as the difference between the arriving times of each duplicate of that packet. The idea is that, if the path with the lower delay fails at the time that packet is transmitted, then this difference characterizes the additional time the protocol required for recovering from the failure.

Each topology was set up as to mimic the real conditions of a substation as close to reality as possible. One of these conditions is the distance between devices, which may cause some difference in the delivery time of the packets. When setting up each topology, we estimated that shorter links would have a propagation delay of $1 \mu\text{s}$, while longer ones would have delays of $1.5 \mu\text{s}$. Taking Topology 1 as an example, the propagation delay of the link connecting IED1 to SW1 is $1 \mu\text{s}$, whereas the delay of the link between IED1 and SW2 is $1.5 \mu\text{s}$. This same criterion was used for all links in Topologies 1 and 2. For the ring topology, Topology 3, we assumed equal distances between IEDs, resulting in a $1 \mu\text{s}$ delay for all links. All links are assumed to operate at 10 Mb/s.

In a SAS, critical messages (*e.g.* voltage and current samples and circuit breaker commands) must be delivered in a short period of time while non-critical services — such as file transfers — have a longer delivery threshold. However, the messages of the latter might travel in the same links of the prior, corresponding to concurrent traffic.

To reproduce a realistic scenario of a substation, we injected SV packets with a size of 142 bytes, which is equivalent to critical service messages. For each experiment, 10,000 (ten thousand) packets are sent from IED1 to IED6, always sending two packets simultaneously over different network interfaces.

Concurrent TCP/UDP traffic both operate at full link capacity, *i.e.*, 10 Mb/s, and starts from IED2 to IED4, when in topologies 1 and 2, and from IED2 to IED3, in Topology 3. In this way, concurrent traffic will always share the link with one of the packets that are being sent redundantly, except for topology 1, which has dedicated links between the IEDs and switches.

A. Experiments and Results

For each topology, two types of experiments were performed: with and without concurrent traffic. In the first case, the only traffic source in the network is that of the packets of the evaluated protocol (PRP, HSR or iPRP, depending on the case). In the second case, we have used the iperf tool [28] to generate TCP/UDP packets, with the purpose of simulating concurrent traffic in one of the LANs/directions of the ring.

Table II shows the average recovery time for each evaluated protocol considering all topologies and all interference scenarios. As the recovery times measured without any interference show, there is always a delay between the delivery of each duplicate of a packet. With no concurrent traffic, however, one expects very low recovery times, since the only sources of asymmetry are the possible differences in links' propagation delays and — in case of Topology 3 — the difference in the number of hops. Since the evaluated topologies are fairly small, recovery time is limited. Nevertheless, even in such an ideal situation, it is not possible to state that the recovery time is effectively zero.

On the other hand, the experiments confirm that concurrent traffic can disturb the stability of the evaluated protocols, causing long delays. The only exceptions are the experiments performed at Topology 1, in which concurrent traffic does

TABLE II: Average Recovery Time (ms)

	Topology 1		Topology 2		Topology 3	
	PRP	iPRP	PRP	iPRP	HSR	iPRP
w/o Interference	0.025	0.041	0.027	0.041	0.029	0.028
TCP 10 Mb/s	0.023	0.038	6.951	6.483	5.813	6.548
UDP 10 Mb/s	0.029	0.016	5.772	5.257	5.359	5.399

not increase recovery time, because of the direct connection between every IED and all switches. Clearly, this provides an ideal scenario for those protocols. However, one may question whether such a topology is scalable due to possible limitations in the number of ports available in a single switch. In the other topologies, on the other hand, we can see the concurrent traffic causes a significant increase in recovery time. On average, the recovery time with TCP interference is slightly longer when compared to UDP interference.

Figure 2 shows the boxplot of the recovery times with and without interference for PRP — Topologies 1 and 2 — and HSR — Topology 3. Notice that the recovery time is considerably higher in the scenarios with interference traffic, as well as the variability².




Fig. 2: Average recovery time for the three topologies. Topologies 1 and 2 represent PRP results, while Topology 3 represents HSR results.

Figure 3 shows a time series of the behavior of iPRP in Topology 2 with the TCP concurrent traffic. The blue line shows the recovery time for each packet. One can notice that the average recovery time — the constant green line — exceeds 5 ms. However, some individual packets experimented a recovery time greater than 10 ms. It is also possible to see that there are packet losses — marked with a red “x” — as a consequence of the congestion caused by TCP. In fact, the behavior of the recovery time closely mimics the expected tooth-saw behavior of TCP’s congestion window under congestion avoidance. However, the increases are not

²Observe that time is shown in log scale in the vertical axis. Thus, although the upper and lower limits — as well as the first and third quartiles — are graphically closer to the median, variance is actually larger in the scenarios with traffic.





Fig. 3: Recovery time in a sample of one thousand packets. The green constant line shows the average recovery time of the analyzed sample. Losses are marked with a red “x”.


strictly linear, with some momentary decreases followed by sharp increases. That can be explained by TCP New Reno’s nature of transmitting in bursts followed by relatively long periods of silence. Not coincidentally, the losses of iPRP packets are generally associated with peaks of the recovery time.

To illustrate the relationship between TCP’s transmission rates and the spikes in recovery time, Figure 4 shows in more detail both quantities during the interval of from 100 ms to 250 ms in this same experiment. Figure 4 (a) shows the transmission rate for the TCP flow while Figure 4 (b) shows the recovery time in that same time span. In Figure 4 (a), it can be observed, for example, that TCP’s transmission rate drops to 8 Mb/s shortly after 150 ms. This reduction corresponds to the natural behavior of TCP’s congestion control mechanism. By comparing the two graphs, it is possible to notice that the recovery time is strongly correlated with the transmission rate of TCP, although there is a slight offset between TCP’s actions and the corresponding effect on recovery time (consider, for example, the sharp valleys in both time series after 150 ms). This offset can be explained by the delay between the moment a burst of packets is transmitted at the TCP source until those packets arrive at the bottleneck link’s queue, causing congestion.

While the bursty nature of TCP traffic contributes to the increase in recovery time, Figure 5 illustrates that the same may occur with UDP. It shows the distribution of the recovery times obtained with PRP in Topology 2 with and without interference caused by UDP. Without interference, recovery time is not only shorter, but also very predictable (*i.e.*, it has a lower variance). Interference, however, creates a bimodal distribution, in which a little under 5% of the samples experience less than 0.5 ms of recovery time, while for the other 95% it oscillated between 4 ms and 7.5 ms. Thus, interference not only shifted the average recovery time, but it also increased



(a) TCP’s transmission rate.



(b) iPRP’s recovery time in the same time span shown in (a).

Fig. 4: Detailed view of the experiment shown in Figure 3, where (a) shows the transmission rate of the TCP interference traffic, while (b) shows the recovery time, both in the same interval of the experiment.

its variance.

These results show that the recovery times of the evaluated fault recovery protocols are not zero, as commonly assumed for the PRP, iPRP, and HSR protocols. Also, in certain cases, recovery times exceed the requirements of certain time-critical messages of the IEC 61850 standard, jeopardizing teleprotection schemes and, consequently, expensive infrastructure. Therefore, factors such as the network topology and currently active network flows must be assessed to ensure the communication network can deliver the necessary Quality of Service.

VI. TRAFFIC CONTROL

One of the main tasks of network management is to ensure the proper operation of the network. This requires the identification and classification of the information that is transmitted over the network. In the particular case of teleprotection, data sent via GOOSE and SV messages should be prioritized, considering that failing to deliver such packets on time may cause problems in the overall functioning of protection tasks. Therefore, awareness regarding the priority relationships




Fig. 5: Histogram of the recovery time for PRP with and without interference caused by UDP in Topology 2.

between different data streams and the application of traffic control mechanisms are essential duties to maintain proper network performance.

As shown in Section V, background traffic or interference can increase the delay for critical message delivery. Thus, an alternative to circumvent this problem is to use traffic control mechanisms, prioritizing messages that have strict temporal restrictions.

To simulate network traffic control, we repeated a portion of the experiments described in Section V, but now assigning different traffic priorities to each flow. The prioritization rules were implemented using the *Traffic Control* (TC) tool available on Linux systems [29]. This tool is used to scale and control existing traffic on the network. Following the execution of the experiments, we have analyzed the results to validate our hypothesis that it is possible to mitigate the impact caused by interference traffic on the IEC 61850 time-critical flows.

A. Recovery Time Analysis with Traffic Control

The experiment using Topology 2 and iPRP was repeated using a priority queue so that the most critical traffic — *i.e.*, iPRP packets — has the highest priority in contrast to the non-critical ones. Thus, at any bottleneck link, packets of any other type of traffic are not transmitted until all iPRP packets queued there are served.

Figure 6 shows that the recovery time when using priority-based packet scheduling is shorter than the value measured in the previous experiments without prioritization. Compared to the results presented in Figure 3, we can see that the recovery time has dropped significantly, with values that are comparable to those of the scenario without interference.

However, this setting also has some undesirable side effects. Due to the prioritization, non-critical traffic can starve if the prioritized packets are transmitted at high rates. This makes the network operation dependent on priority traffic, jeopardizing other applications and services that are also required in the network. To avoid that, new experiments were performed using a Token Bucket to control the bandwidth assigned to the priority class.




Fig. 6: Histogram showing the recovery time of the iPRP protocol using packet scheduling with priority queue.




Fig. 7: Recovery time in a sample of one thousand packets. The sample was taken from the iPRP experiment with simultaneous TCP traffic in Topology 2.

Figure 7 shows that with this configuration recovery time can be kept at acceptable levels even in the face of interfering traffic. We utilized a priority queue with TBF to limit bandwidth utilization. Additionally, the competing application will still be able to work — albeit with a lower throughput.

VII. CONCLUSION

Fault-recovery strategies are necessary for fulfilling the time-critical requirements of teleprotection scenarios. Those strategies must be able to reestablish the communication promptly in the face of failure, and, if possible, to allow applications to continue their execution without negative effects.

PRP and HSR were proposed by the IEC 62439-3 standard to provide network redundancy, as a fault-recovery strategy. Both proactively send redundant packets across distinct paths so that, in case of a failure in one of the paths, a copy of the packet will still reach the destination. Thus, both PRP and HSR are assumed to provide zero-time recovery.

As shown in this paper, however, the actual recovery times of HSR and PRP — and its iPRP variant — are not zero,

because, in practice, there are sources of asymmetry between the paths. Those asymmetry sources range from small differences in propagation times due to the physical lengths of links to unbalanced levels of concurrent traffic. In particular, concurrent traffic can cause large increases in recovery time, which may reach unacceptable levels considering the strict temporal requirements of the messages in certain protocols such as GOOSE and SV.

This paper has also shown that there are viable solutions to the problem of delayed packet delivery when there is concurrent network traffic in the teleprotection scenario. The simplest solution is to use traffic control policies giving higher priority to flows that have time constraints. This type of configuration can also be performed using an SDN controller, a paradigm that has been considered for improving the management of communication networks in teleprotection scenarios.

Regardless of the implementation method, experimental results confirm that using traffic differentiation contributes considerably to the reduction of the recovery time of the PRP, iPRP, and HSR protocols. While the recovery time for such protocols remains non-zero, the effect of network interference is mitigated, making those times acceptable for protection applications such as GOOSE and SV.

ACKNOWLEDGEMENTS

The authors would like to thank CAPES, CNPq, FAPERJ, TAESA and P&D ANEEL (PD-07130-0053/2018) for their financial support.

REFERENCES

- [1] R. E. Mackiewicz, "Overview of IEC 61850 and benefits," in *Power Systems Conference and Exposition, 2006. PSCE'06. 2006 IEEE PES. IEEE*, 2006, pp. 623–630.
- [2] J. Tan and W. Luan, "Iec 61850 based substation automation system architecture design," in *2011 IEEE Power and Energy Society General Meeting*, July 2011, pp. 1–6.
- [3] L. Wisniewski, "Communication reliability," in *New methods to engineer and seamlessly reconfigure time triggered Ethernet based systems during runtime based on the PROFINET IRT example*. Springer, 2017, pp. 105–123.
- [4] S. Kumar, N. Das, and S. Islam, "High performance communication redundancy in a digital substation based on IEC 62439-3 with a station bus configuration," in *Power Engineering Conference (AUPEC), 2015 Australasian Universities*. IEEE, 2015, pp. 1–5.
- [5] J. Araujo, J. Lazaro, A. Astarloa, A. Zuloaga, and A. Garcia, "PRP and HSR version 1 (IEC 62439-2), improvements and a prototype implementation," in *Industrial Electronics Society, IECON 2013-39th Annual Conference of the IEEE*. IEEE, 2013, pp. 4410–4415.
- [6] J. A. Araujo, J. Lázaro, A. Astarloa, A. Zuloaga, and A. García, "Prp and hsr version 1 (iec 62439-3 ed.2), improvements and a prototype implementation," in *IECON 2013 - 39th Annual Conference of the IEEE Industrial Electronics Society*, Nov 2013, pp. 4410–4415.
- [7] Technical Committee on Power systems management and associated information exchange, "IEC 61850 - communication networks and systems in substations," International Electrotechnical Commission (IEC), Tech. Rep., 2002– 2013.
- [8] "IEEE standard communication delivery time performance requirements for electric power substation automation," <https://standards.ieee.org/content/ieee-standards/en/standard/1646-2004.html>, accessed in: 19-09-2019.
- [9] G. Rolim, D. Passos, C. Albuquerque, and I. Moraes, "Moskou: A heuristic for data aggregator positioning in smart grids," *IEEE Transactions on Smart Grid*, vol. 9, no. 6, pp. 6206–6213, Nov 2018.
- [10] M. H. Rehmani, A. Davy, B. Jennings, and C. Assi, "Software defined networks based smart grid communication: A comprehensive survey," *IEEE Communications Surveys Tutorials*, pp. 1–1, 2019.
- [11] X. Dong, H. Lin, R. Tan, R. K. Iyer, and Z. Kalbarczyk, "Software-defined networking for smart grid resilience: Opportunities and challenges," in *Proceedings of the 1st ACM Workshop on Cyber-Physical System Security*, ser. CPSS '15. New York, NY, USA: ACM, 2015, pp. 61–68. [Online]. Available: <http://doi.acm.org/10.1145/2732198.2732203>
- [12] N. Dorsch, F. Kurtz, H. Georg, C. Hägerling, and C. Wietfeld, "Software-defined networking for smart grid communications: Applications, challenges and advantages," in *2014 IEEE International Conference on Smart Grid Communications (SmartGridComm)*, Nov 2014, pp. 422–427.
- [13] Y. Cui, S. Xiao, C. Liao, I. Stojmenovic, and M. Li, "Data centers as software defined networks: Traffic redundancy elimination with wireless cards at routers," *IEEE Journal on Selected Areas in Communications*, vol. 31, no. 12, pp. 2658–2672, December 2013.
- [14] H. I. Kobo, A. M. Abu-Mahfouz, and G. P. Hancke, "A survey on software-defined wireless sensor networks: Challenges and design requirements," *IEEE Access*, vol. 5, pp. 1872–1899, 2017.
- [15] M. Goraj and R. Harada, "Migration paths for IEC 61850 substation communication networks towards superb redundancy based on hybrid PRP and HSR topologies," in *11th IET International Conference on Developments in Power Systems Protection (DPSP 2012)*. IET, 2012.
- [16] J. Araujo, J. Lazaro, A. Astarloa, A. Zuloaga, and A. Garcia, "High availability automation networks: PRP and HSR ring implementations," in *Industrial Electronics (ISIE), 2012 IEEE International Symposium on*. IEEE, 2012, pp. 1197–1202.
- [17] M. Médard and S. S. Lumetta, *Network Reliability and Fault Tolerance*. American Cancer Society, 2003. [Online]. Available: <https://onlinelibrary.wiley.com/doi/abs/10.1002/0471219282.eot281>
- [18] G. Mirjalily, F. A. Sigari, and R. Saadat, "Best multiple spanning tree in metro ethernet networks," in *2009 Second International Conference on Computer and Electrical Engineering*, vol. 2, Dec 2009, pp. 117–121.
- [19] T. Li, B. Cole, P. Morton, and D. Li, "Cisco hot standby router protocol (hsrp)," RFC Editor, United States, Tech. Rep., 1998.
- [20] H. Kirmann, K. Weber, O. Kleineberg, and H. Weibel, "Hsr: Zero recovery time and low-cost redundancy for industrial ethernet (high availability seamless redundancy, iec 62439-3)," in *2009 IEEE Conference on Emerging Technologies Factory Automation*, Sep. 2009, pp. 1–4.
- [21] H. Kirmann, M. Hansson, and P. Muri, "Iec 62439 prp: Bumpless recovery for highly available, hard real-time industrial networks," in *2007 IEEE Conference on Emerging Technologies and Factory Automation (EFTA 2007)*, Sep. 2007, pp. 1396–1399.
- [22] S. Kumar, N. Das, and S. Islam, "Performance evaluation of two interconnected high voltage utility substations using prp topology based on iec 62439-3," in *2017 Australasian Universities Power Engineering Conference (AUPEC)*, Nov 2017, pp. 1–5.
- [23] J. A. Araujo, J. Lazaro, A. Astarloa, A. Zuloaga, and J. I. Garate, "Prp and hsr for high availability networks in power utility automation: A method for redundant frames discarding," *IEEE Transactions on Smart Grid*, vol. 6, no. 5, pp. 2325–2332, Sep. 2015.
- [24] M. Urbina, T. Acosta, J. Lázaro, A. Astarloa, and U. Bidarte, "Smart Sensor: SoC architecture for the Industrial Internet of Things," *IEEE Internet of Things Journal*, vol. 6, no. 4, pp. 6567–6577, Aug. 2019.
- [25] A. Leal and J. F. Botero, "Defining a reliable network topology in software-defined power substations," *IEEE Access*, vol. 7, pp. 14 323–14 339, 2019.
- [26] P. Jafary, A. Supponen, M. Salmenperä, and S. Repo, "Analyzing reliability of the communication for secure and highly available goose-based logic selectivity," *Security and Communication Networks*, vol. 2019, 2019.
- [27] M. Popovic, M. Mohiuddin, D. Tomozei, and J. L. Boudec, "iprp—the parallel redundancy protocol for ip networks: Protocol design and operation," *IEEE Transactions on Industrial Informatics*, vol. 12, no. 5, pp. 1842–1854, Oct 2016.
- [28] "Iperf - The TCP/UDP bandwidth measurement tool," <https://iperf.fr/>, accessed in: 20-09-2019.
- [29] "Linux Advanced Routing & Traffic Control HOWTO," <https://lartc.org/lartc.html#LARTC.QDISC>, accessed in: 20-09-2019.