RESILIENT 5G TECHNOLOGIES OPTIMIZED FOR POWER GRID PROTECTION SOLUTIONS USING IEC 61850 TIME-CRITICAL COMMUNICATIONS

Ana Aleixo[1*], Rogério D Paulo[2], Rui D Jorge[1], Alberto Rodrigues[3], Carlos Arantes[2], José Cabaça[4], Pedro M Neves[4]

[1]Protection, Automation and Control Division, Efacec, Oeiras, Portugal
[2]Protection, Automation and Control Division, Efacec, Porto, Portugal
[3]Innovation and Development/Products and Technology Division, Efacec, Porto, Portugal
[4]Innovation and Technology Strategy, Altice Labs, Aveiro, Portugal
**ana.aleixo@efacec.com*

Keywords: 5G, POWER SYSTEM PROTECTION, PMU, R-GOOSE, RELIABILITY

Abstract

This paper reports part of the ongoing work being carried out in the scope of the SliceNet project, a vertical business-driven H2020 project in which 5G technologies and services are being developed with the specific goal of serving a heterogeneous set of vertical applications.

Efacec and Altice Labs are contributing with a vertical business use case focused on communication-based protection and automation solutions for medium voltage (MV) energy distribution networks. The use case includes the implementation and testing of several distributed algorithms, including high-speed selective blocking using IEC 61850 Routable-Generic Object Oriented Substation Events (R-GOOSE), high-speed post-fault service restoration using R-GOOSE, and differential protection using Phasor Measurement Unit (PMU) data transmitted over R-GOOSE.

The paper reports on the first stage of lab-based trials that are currently being carried out in the use case testbed, consisting mainly of R-GOOSE communication resilience and high-speed selective blocking over 5G. A small set of the lab tests aim at providing early results for a future viability assessment of the implementation of a differential protection solution over 5G. The results from phase-one trials will be disclosed and analysed, and the plan for the ensuing use case trial stages will be subsequently presented.

1 Introduction

Wide-area protection schemes for energy distribution networks are currently restricted by the lack of a low-latency communication network that provides the required levels of reliability and availability at a moderate cost. Access to a dedicated communication infrastructure requires a considerable investment, and presently available public networks do not fully comply with the quality of service requirements of many wide-area time-critical applications.

Emerging fifth generation cellular network technologies (5G) are expected to introduce a wide range of features and capabilities, such as ultra-Reliable Low-Latency Communication (uRLLC), that will promote the implementation and deployment of new communication-based protection solutions for the power grid [1].

The H2020 SliceNet project [2] aims at developing a 5G framework suited for diverse vertical business applications with diverging requirements [3]. The work described throughout this paper relates to one of the vertical business use cases that is being implemented for validating the 5G technologies developed and integrated in the project. The presented use case refers to communication-based protection and automation solutions, and includes communication-intensive applications such as high-speed selective blocking and differential protection using synchrophasor data transmitted over IEC 61850 R-GOOSE [4-5].

Effective high-speed protection coordination delivering improvements in the order of hundreds of milliseconds over traditional methods requires uninterrupted access to a deterministic communication infrastructure that ensures latencies of few milliseconds. Such is the case for selective blocking, which is dependent on the prompt transmission of blocking signals between remotely located protection devices.

Communication network latency, reliability, and availability are key aspects of the 5G technologies being developed and integrated in the SliceNet framework. These include highly customizable network slices with optimized resources for vertical business applications capable of spanning across multiple network domains, and a cognition engine for communication network fault prediction and mitigation.

Having access to a 5G network provides a significant opportunity for deploying distributed protection applications grid-wide, which are currently only deployed in limited sites where an adequate physical communication infrastructure is available, typically to supply specific critical loads. Differential protection solutions for MV power distribution networks fall within the referred set of applications, which is not technically viable using currently available public cellular networks, as it relies on publishing and subscribing continuous streams of power system quantities at sub-cyclic rates.

2 Communication-Based Power Grid Protection Solutions

The use case includes three power system protection (or protection-related) applications that rely on IEC 61850 horizontal communications: high-speed selective blocking, high-speed post-fault service restoration, and synchrophasor-based differential protection. Although all cited solutions rely on device-to-device (D2D) communication using unicast R-GOOSE over User Datagram Protocol (UDP), the first and third applications are the more demanding in terms of communication requirements and, consequently, are the focus of the phase-one trials.

2.1 High-speed selective blocking

The goal of the algorithm is to ensure high-speed logical selectivity between protection relays integrating radial topologies in energy distribution networks. The solution requires that several protection relays placed along a distribution feeder, often in remote locations, coordinate with each other.

In the event of a fault, all the Intelligent Electronic Devices (IEDs) that detect the fault immediately send blocking messages to all upstream devices. IEDs that receive blocking indications will promptly reset the corresponding protection elements, securing selectivity through high-speed protection coordination. Figure 1 depicts an example of the R-GOOSE messages exchanged by the IEDs when a fault is detected in the feeder.

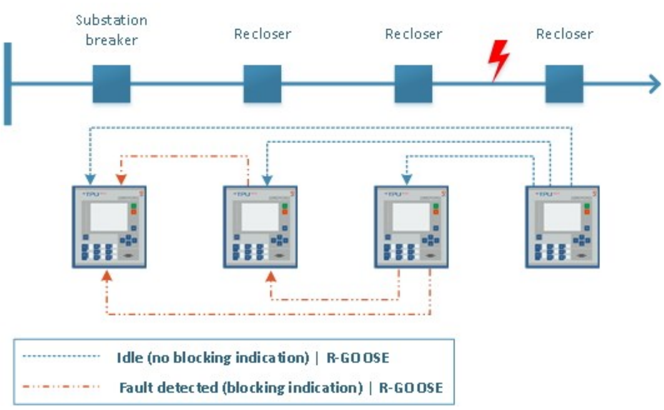

Figure 1 R-GOOSE data flows for selective blocking.

2.1.1 Communication requirements: Selectivity is only assured if the blocking messages reach the upstream devices before their protection elements trip. This is only possible with a highly reliable communication infrastructure that guarantees continuous availability with deterministic low-latency levels.

Table 1 includes a set of requirements specified for event-driven communications used in high-speed selective blocking. End-to-end (E2E) latencies higher than the values indicated as ideal may be tolerated by the application but will impact the effectiveness of the solution.

Table 1 Event-driven communication requirements.

Requirements	Values
Availability/reliability	99.999 %
Ideal E2E latency	< 10 ms
Maximum E2E latency	40 ms
Jitter	< 10 ms

2.2 Synchrophasor-based differential protection

The presented solution consists of a line differential protection algorithm that compares locally measured current levels with the values obtained from devices installed in remote downstream locations, using synchrophasor data transmitted via R-GOOSE.

The solution includes main and backup protection schemes, which require the continuous transmission of up to three streams of phasor data (the PMU data flows are represented in Figure 2). PMU data is published at ¼ power network cycle rates (*i.e.*, 5 ms @ 50 Hz) as sets of three synchrophasors which include magnitude, angle, quality, and timestamp of the three phase currents.

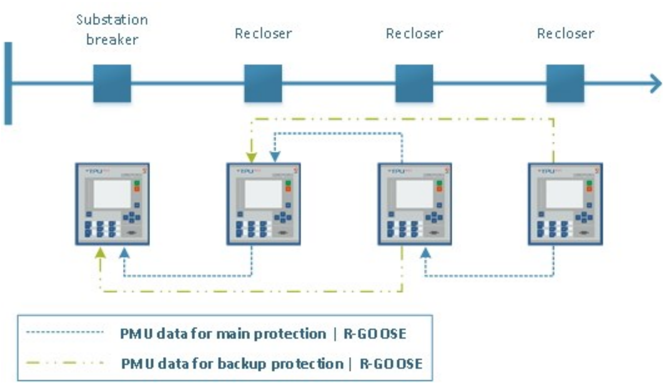

Figure 2 PMU data flows for differential protection.

2.2.1 Communication requirements: Although the solution is robust to short delays and sporadic single packet losses, PMU data should not be delayed by more than ½ power network cycle. Packets that are delayed for longer than 30 ms are treated as lost packets and not processed by the application.

Table 2 PMU data communication requirements.

Requirements	Values
Availability/reliability	99.999 %
Ideal E2E latency	< 5 ms
Maximum E2E latency	30 ms
Jitter	< 5 ms

Table 2 includes communication requirements specific for the differential protection application. E2E latencies higher than the values indicated as ideal may be tolerated by the application but will impact the effectiveness of the solution.

3 Resilient 5G Technologies

One of the main goals of 5G is to increase the reliability and availability of public cellular networks. Communication resilience is being tackled in the SliceNet project by implementing and integrating a set of vertical-oriented technologies and services that rest on network slicing and slice management.

3.1 Network slicing

Network slicing is used for allocating a set of network resources with predefined characteristics, optimized for specific applications, services or customers. The network slicing concept [6] is composed by three layers: the service instance layer, which represents the supported end-user or business services; the network slice instance layer, which consists of a logical network, composed by a set of network functions and resources, instantiated with the purpose of providing the necessary characteristics required by the service instances; the resource layer, composed by physical resources (*e.g.*, computation power, storage, radio access) and the logical partition or grouping of those resources.

This architectural feature maximizes the flexibility and scalability of 5G networks by allowing network operators to dynamically deploy multiple isolated “virtual networks” on their infrastructure. Slice isolation will enhance both security and availability, since the data being transmitted through a slice is not accessible from the remainder of the network and should not be affected by poor Quality of Service (QoS) on neighbouring slices.

3.2 Cognitive fault prediction and mitigation

The project also encompasses cognition-based optimization algorithms that resort to machine-learning (ML) to improve QoS and Quality of Experience (QoE).

One of the implemented solutions consists of a network fault prediction and mitigation algorithm that uses a ML engine to predict imminent communication failures and tries to prevent them, if possible. If a communication network fault is considered inevitable, the algorithm can attempt to mitigate it by, for instance, temporarily allocating more resources to a slice.

4 Test Environment

Integration and validation tests are presently being held in the use case testbed, which is set up across two different sites, in Aveiro, Portugal. One of the sites which hosts the 5G network management and the other hosts the 5G network infrastructure and the power grid IEDs and test sets.

Undergoing tests are part of the phase-one trials and aim at ensuring the integration of the vertical business application prototypes in the SliceNet system. Phase-one trials include preliminary testing of high-speed selective blocking and communication-centred tests of R-GOOSE over 5G, comprising both event-driven communication and the transmission of synchrophasor data.

Upcoming trial phases will focus on the validation of the developed 5G technologies, with the main goal of proving their suitability for the proposed power grid protection and automation solutions.

4.1 Testbed

The use case testbed, represented in Figure 3, includes the hardware and software elements necessary for the vertical-business, Digital Services Provider (DSP) and Network Services Provider (NSP) actors to perform their actions and play their roles in the final validation and demonstration trials. In the diagram, the vertical is represented in purple, the DSP and NSP1 correspond to the same entity, represented in yellow, and the NSP2 is represented in blue.

The testbed is currently equipped with four Efacec protection relays, all of which are connected to the SliceNet 5G network. Three are field IEDs connected to Long Term Evolution (LTE) modems; the other is a substation IED, connected to a physical Local Area Network (LAN), which in turn is also connected to the 5G network. All IEDs connect to second LAN, not connected to the 5G network, which is used for management and Precision Time Protocol (PTP) synchronization.

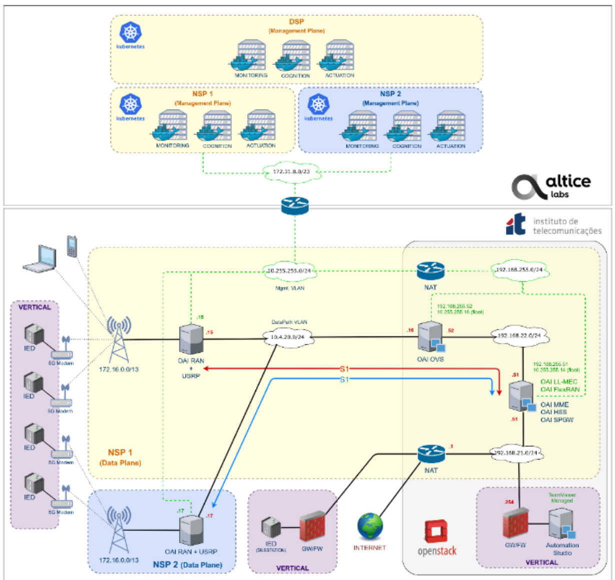

Figure 3 Use case prototyping and validating testbed.

4.2 Test conditions and metrics

During the phase-one integration trials, the 5G communication resilience is being evaluated through the analysis of the quality of the R-GOOSE flows under different communication network conditions:

- Ideal conditions (good signal strength; no additional traffic);
- Degraded or congested communication network;
- Isolated network slice with ideal conditions; degraded or congested network in a neighbouring slice.

The following metrics are being used for evaluating the quality of the R-GOOSE communication:

- GOOSE failure indications and accumulated duration;
- Number of lost GOOSE packets;
- Application latency and jitter;
- Communication latency and jitter.

4.2.1 Test environment limitations: The test environment hosting the SliceNet system has some constraints, enforced by the project scope, which pose some implications in the overall test results: the user equipment (UE) connected to the IEDs are commercial LTE modems (and not 5G); all the IEDs are in the same geographical location, and are presently connected to the same cellular base station, through identical UE, thus limiting the possibility of studying the impact of more complex aspects of cellular radio networks (*e.g.*, channel asymmetry) in smart grid Protection, Automation, and Control (PAC) applications.

5 Test Results

The presented test results were gathered during the integration phase, which is currently being undertaken. Although the SliceNet system is not yet fully integrated and not all network supervision and management modules and tools are presently deployed, it was nevertheless possible to perform preliminary trials in the testbed that already show some promising results. The phase-one trials were able to provide a preliminary assessment of the advantages and benefits that may be obtained for PAC systems by having access to a 5G communication network.

5.1 Latency and jitter assessment

Albeit not being the main focus of the phase-one trials, particularly in light of the referred test environment limitations, carrying out a first assessment of the latency levels that are achievable in the use case testbed provides a fundamental basis for the subsequent tests.

Two distinct latency metrics were considered relevant: communication latency and application latency (the diagram in Figure 4 represents the two measurement paths). The corresponding jitter values were considered relevant as well.

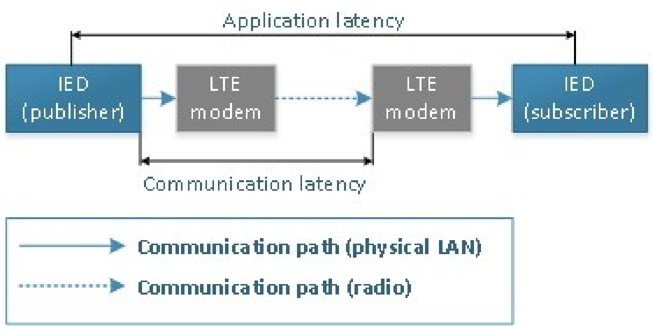

Figure 4 Latency measurement paths.

Communication latency was measured using Internet Control Message Protocol (ICMP) ping requests to one of the modems, using an equipment connected to one of the other modems. The values presented in Figure 5 correspond to ½ of the round-trip ping response times. The average communication latency and jitter observed throughout the tests were of approximately 31.7 ms and 2.3 ms, respectively.

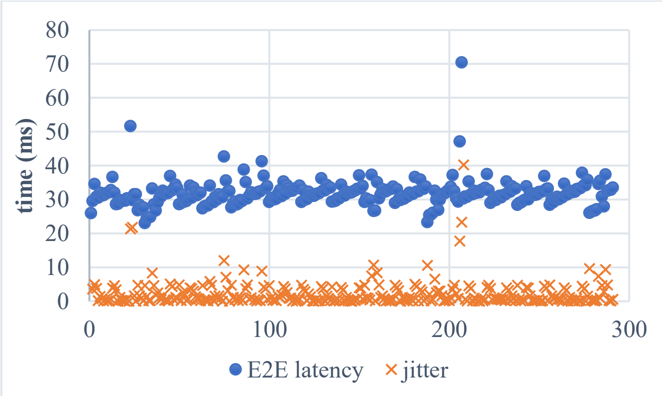

Figure 5 Communication latency and jitter measurements.

Application latency was measured by exchanging a R-GOOSE message between two IEDs (one set as publisher, the other as subscriber) with a dataset containing a single point status (SPS) [7] that periodically changed value. The values presented in Figure 6 correspond to the time elapsed since the SPS value was modified by the publisher until that value was processed and logged by the subscriber. The average application latency and jitter observed throughout the tests were of approximately 45.9 ms and 13.5 ms, respectively.

Table 3 includes average, minimum and maximum values for the observed communication and application latency and jitter. The observed maximum values correspond to sporadic occurrences, that will be further analyzed and addressed during the ongoing integration phase. Even considering the current limitations, the average latency and jitter values are within the range required by the selective blocking and post-fault service restoration algorithms, and not far from the levels required for the synchrophasor-based differential protection.

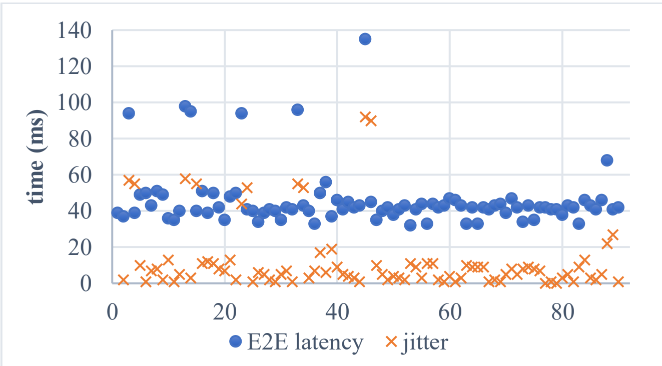

Figure 6 Application latency and jitter measurements.

Table 3 Latency and jitter measurement statistics.

Measurements (ms)	Average	Max.	Min.
Communication latency	31.73	70.45	23.02
Communication jitter	2.26	40.18	0.01
Application latency	45.89	135	32
Application jitter	13.53	92	1

5.2 Communication network resilience trials

The preliminary assessment of the 5G framework resilience was made by evaluating the R-GOOSE message exchange between two IEDs under different communication network conditions, with and without network slices:

- **Test #1**: No slices were deployed. The communication network had no additional traffic.
- **Test #2**: No slices were deployed. The communication network was congested with additional traffic up to 90 % of its capacity.
- **Test #3**: No slices were deployed. The communication network was congested with additional traffic up to 100 % of its capacity.
- **Test #4**: Two slices were deployed: one slice was used by the IEDs, the other was congested up to 90 % of its capacity.
- **Test #5**: Two slices were deployed: one slice was used by the IEDs, the other was congested up to 100 % of its capacity.

Each test lasted 10 minutes, in which an IED published an R-GOOSE message with a heartbeat of 250 ms that changed value every 10 seconds. The second IED subscribed this R-GOOSE message and logged the GOOSE statistics displayed in Table 4.

Table 4 GOOSE statistics registered during communication network resilience trials.

Test	Lost packets	GOOSE failures	Acc. failure duration (s)	Availability (%)
#1	8	1	0.219	99.96
#2	134	26	5.659	99.06
#3	107	25	360.825	39.86
#4	25	2	0.435	99.93
#5	16	4	0.279	99.95

The effect of a network congestion on the R-GOOSE time-critical communication is evident in the results displayed in Table 4 for tests #2 and #3, where additional traffic occupied up to 90 % and 100 % of the network capacity, respectively. Both tests registered an elevated number of lost packets and GOOSE failures, and the network was unavailable for long periods of time. With the network congested to 100 % of its capacity, the R-GOOSE communication failed completely after approximately 4 minutes and did not recover for the extent of the test.

When the IEDs communicate through an isolated slice, however, they are virtually unaffected by excessive traffic in neighboring slices, as can be inferred by the results of tests #4 and #5, where communication availability levels are similar to the levels observed in ideal conditions.

These tests were repeated with phasor data published using R-GOOSE. The PMU data was packed into a GOOSE dataset containing instantaneous magnitude and phase angle values, quality, and timestamp for the three phase currents, and was published at a cadence of 5 ms.

Similarly to what could be ascertained from the previous tests, high network congestions in a network with no slices led to a complete PMU data communication failure a few minutes into the test. An identical scenario with the time-critical communication and the traffic congestion running in two distinct isolated slices presented satisfactory results, similar to the ones observed while transmitting the PMU data in a communication network in optimal conditions, with no additional traffic.

5.3 Application tests

Phase-one trials also included prefatory tests on the high-speed selective blocking solution, with the aim of validating the communication network resilience in an application context.

The protection coordination tests were performed using three IEDs connected to the network through LTE modems, integrating the scheme represented in Figure 7. All PAC devices were set to trip in 80 ms for phase currents with a magnitude higher than 200 A (primary value).

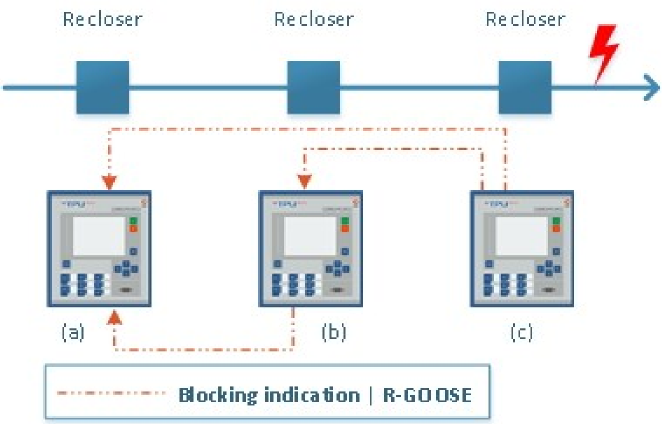

Figure 7 High-speed selective blocking test scenario.

The test consisted on the simulation of a single-phase fault downstream from all the devices. As the fault is detected by all devices, each device sends blocking indications to upstream devices. In order to guarantee selectivity, upstream devices must receive and process the blocking indications in less than the parameterized tripping time of 80 ms.

The simulation was replayed multiple times in different communication network scenarios: (1) no additional traffic; no slices; (2) network congested up to 90 % of its capacity; no slices; (3) two network slices: one specifically for R-GOOSE communication, the other one congested up to 90 % of its capacity.

The high-speed selective blocking solution ran as expected in scenario (1), where all blocking indications were received and processed in time. When simulated over a congested network, however, the blocking signals often failed to arrive on time to one or to both upstream IEDs, compromising the system selectivity. When the simulation was replayed using network slices, it was possible to ascertain that the blocking indications were received in a timely basis, thus ensuring selectivity.

Figure 8, Figure 9, and Figure 10 are disturbance records registered by the three IEDs during a simulation of the latter scenario, where it is possible to observe the protection elements of the upstream devices dropout upon the reception of the blocking signals from the downstream devices.

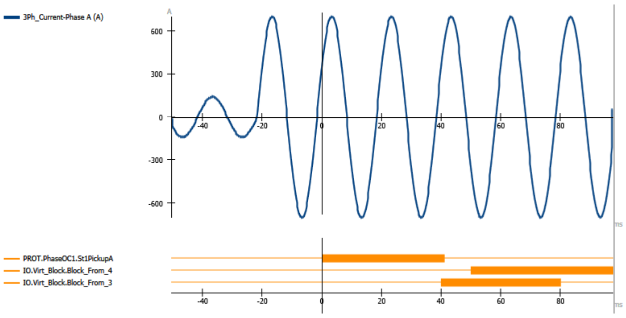

Figure 8 Disturbance record registered by IED (a).

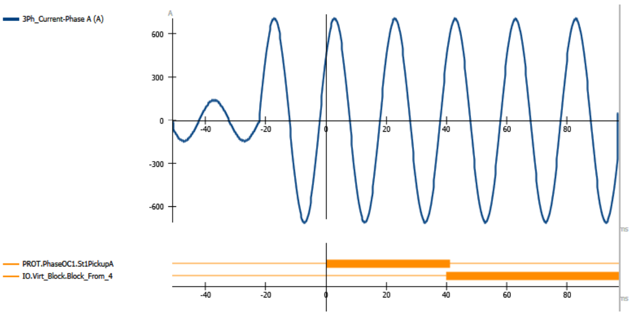

Figure 9 Disturbance record registered by IED (b).

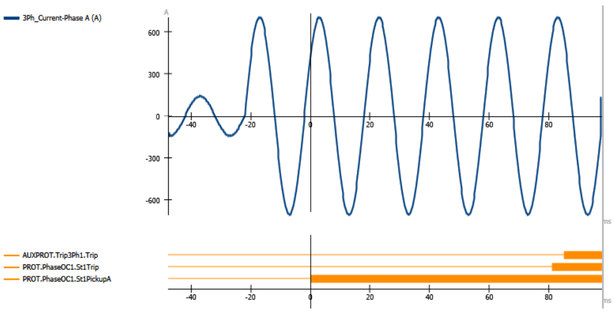

Figure 10 Disturbance record registered by IED (c).

6 Conclusion

There is an emerging need for a cost-effective uRLLC communication infrastructure capable of providing an adequate foundation for time-critical wide-area Smart Grid applications. 5G technologies will undoubtedly play a crucial role in the development of new algorithms and solutions and will likely be one of the key elements for the constantly evolving Smart Grid.

A relevant part of the 5G framework and services which are currently being developed and integrated in the scope of the SliceNet project is targeted at creating communication network mechanisms with the specific purpose of improving reliability and resilience for vertical businesses, such as the energy sector.

The phase-one trials were able to provide a set of interesting conclusions. The advantages of network slicing are already evident – the outcome of the tests performed with PAC devices using a dedicated slice proved that it was possible to maintain high levels of QoS even when the network was congested or degraded outside that slice. Test results confirm that network slice and, in particular, network slice isolation increases communication determinism, which is a key requirement for virtually any distributed power system protection solution, as is the case of protection coordination algorithms or solutions that rely on PMU data.

The SliceNet project is ongoing and close to reaching its final stage, which will incorporate further improvements, including cognition-based algorithms specifically aimed at improving communication network resilience, even across multiple domains. The PAC algorithms that compose the Smart Grid use case will undergo further trials as the 5G framework and services evolve. These trials will constitute the basis for an in-depth analysis of the viability of communication-demanding solutions and of the foreseeable gains that can be obtained from using 5G communications for wide-area PAC solutions.

7 Acknowledgements

This work has received funding from the European Union's H2020 program, under grant agreement No. 761913. The authors wish to thank all SliceNet partners for their support.

8 References

[1] Thrybom, L., Landernäs, K., Valtari, J., et al.: '5G and Energy', 5G-PPP, 2015.

[2] '5G-PPP SliceNet Project, "End-to-End Cognitive Network Slicing and Slice Management Framework in Virtualised Multi-Domain, Multi-Tenant 5G Networks' (H2020-ICT-2016-2/761913)', https://slicenet.eu/, accessed 1 November 2019.

[3] Wang, Q., Alcaraz-Calero, J., Ricart-Sanchez, R., et al.: 'Enable Advanced QoS-Aware Network Slicing in 5G Networks for Slice-Based Media Use Cases', IEEE Transactions on Broadcasting, 2019, 65, (2), pp 444-453.

[4] Apostolov, A.: 'R-GOOSE: what it is and its application in distribution automation', Proc. CIRED, Glasgow, UK, 2017.

[5] IEC, 'IEC TR 61850-90-5 Ed.1: Communication networks and systems for power utility automation – Part 90-5: Use of IEC 61850 to transmit synchrophasor information according to IEEE C37.118', 2012.

[6] NGMN Alliance, 'Description of Network Slicing Concept', NGMN, 2016.

[7] IEC, 'IEC 61850-7-3 Ed.2: Communication networks and systems for power utility automation – Part 7-3: Basic communication structure – Common data classes', 2010.