Experimental study of performance and vulnerabilities of IEC 61850 process bus communications on HSR networks

Stéphane Mocanu, Jean-Marc Thiriet

► To cite this version:

Stéphane Mocanu, Jean-Marc Thiriet. Experimental study of performance and vulnerabilities of IEC 61850 process bus communications on HSR networks. EuroS&PW 2020 - IEEE European Symposium on Security and Privacy Workshops, Sep 2020, Gênes, Italy. pp.584-593, 10.1109/EuroSPW51379.2020.00085 . hal-02921495

HAL Id: hal-02921495

<https://hal.science/hal-02921495>

Submitted on 25 Aug 2020

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Experimental study of performance and vulnerabilities of IEC 61850 process bus communications on HSR networks

Stéphane Mocanu

*Laboratoire d'Informatique de Grenoble
Univ. Grenoble Alpes, CNRS, Inria, Grenoble INP
Grenoble France
stephane.mocanu@imag.fr*

Jean-Marc Thiriet

*GIPSA-Lab
Univ. Grenoble Alpes
Grenoble France
jean-marc.thiriet@univ-grenoble-alpes.fr*

Abstract—Modern power-network communications are based on the IEC 61850 series standards. In this paper we investigate the real-time performance, the vulnerabilities and the attack scenarios on the sensor level communication networks, more precisely on the Sampled Measured Value (SMV) protocol. There are two main contributions of our work. First, we evaluate statistically the measured real-time performance of the communication network. The second contribution is the description, implementation and experimental validation of the attacks on SMV protocol targeting electrical protection functions.

Index Terms—IEC 61850, Process Bus, Sampled Measured Value, High-availability Seamless Redundancy networks

1. Introduction

IEC 61850 standards collection was intended to be a universal specification for the design and operation of intelligent power grids. Although focusing on the communication part, the standards are far more general. Such that, it specifies the electrical protection and control function models, device configuration languages, physical process data model and even electromagnetic compatibility and environmental requirements. The standard collection is intended to answer to two specific electrical domain needs. The first one is the distributed control and protection. Due to the complexity, versatility and interdependence of electrical networks, control and protection functions are also complex. Given that the power grid is a very large size physical process, even considered at transformation substation level, control and protection function have to be distributed such that communication between controllers and protection relays is paramount. The second need is interoperability. Industrial communication is, historically, a proprietary world. Each manufacturer will tend to support only his own communication protocol stack such that interoperability is a real issue. IEC 61850 aims to provide a universal protocol stack (actually 3 protocols) for power systems communication and a distributed control and protection functions modeling framework.

Our work is concerned with the analysis of one of the protocols from the IEC 61850 collection: Sampled

Measured Value (SMV). This Ethernet based protocol is intended to be used for the transmission of sampled current and voltage data from sensors to protection relays. One of the characteristic of SMV traffic is its intensity. Data is sampled on sensors at 4kHz in 50Hz electrical networks then a frame is sent every 250 μ s for each measurement point. Due to the criticality of the traffic a highly available network is commonly used. There are two solutions which are supported currently by the commercial device : High-availability Seamless Redundancy (HSR) and Parallel Redundancy Protocol (PRP). In this paper we focus on HSR as it is the less expensive solution on the market and therefore popular for small size electrical substations.

Our main contributions are : an experimental real-time characterization of the traffic in HSR networks, especially the jitter and a complete description, implementation and experimental validation of the attacks on SMV protocols. Proofs of the feasibility of the attacks were performed on real hardware connected to a simulated process in a hardware-in-the-loop setup. Eventually we analyze attack detection and possible mitigation measures.

The paper is organized as follows, in Section 2 we provide the IEC 61850 minimal background focusing on the SMV protocol and HSR networks. Section 3 is dedicated to the test system architecture, SMV traffic calculation and jitter measurement. In section 4 we describe attack implementation and experimental results. Some comments on intrusion detection and mitigation measures are provided in section 5. Section 6 will review related work and comment out the positioning of our results. We conclude with a summary of the work and comments on our further research in section 7.

Data sets from our experiments are available online at http://gics-hil.gforge.inria.fr/datasets_hsr/.

2. IEC 61850 and SMV

2.1. Protection functions and communication

There are two key concepts in IEC 61850 which are relevant for our work: electrical functions model and the communication stack. The electrical functions (protection, control, measurement ...) model in IEC 61850 introduce the concept of distributed functions. That means that an electrical function is composed by several standard

This work is supported by the French National Research Agency in the framework of the "Investissements d'avenir" program (ANR-15-IDEX-02) and SACADE project (ANR-16-ASTR-0023).

elementary procedures called Logical Nodes (LN) which exchange data in order to achieve the process control goal. A LN is a piece of software or hardware that accomplishes a specific basic functionality like signal acquisition and conditioning, metrology, arithmetic, etc.

In order to make the concept clear, let us consider one of the simplest protection functions. In Figure 1a we consider a simple feeder overcurrent protection. The feeder is an electrical distribution line which connects clients to the transformer substation. Internal substation electric line is called a busbar.

Figure 1. Simple protection function (a) and LN decomposition (b).

The objective of the protection function is to isolate the feeder from the busbar if an electrical fault occurs on the feeder. A current transformer is used as a sensor to measure the current on the feeder. The most common protection function is the overcurrent protection which compares the sensor measure with a given threshold and will send an opening command (a trip signal) to a circuit breaker if the threshold is crossed.

A possible decomposition of the protection function in standard LNs is as shown in Figure 1b: a LN “current transformer” (TCTR) provides current measures to a time overcurrent protection (PTOC). In case of a fault, PTOC will activate a trip conditioning LN (PTRC) which will issue the trip command to the circuit breaker LN (XCBR).

Communication requirements between LNs, including timing requirements, are provided by Piece of Information for COMMunication (PICOM). Several thousand PICOMs are specified for all the allowed LN interconnections [1].

The LN/PICOM specification does not impose a particular implementation on the actual hardware. A protection function does not need to be implemented on a single physical device. The LN used by a protection function may be distributed on several physical devices. Depending on the implementation the PICOMs will be mapped to inter-process communication on a single device or to network flows. Physical devices, which are called Intelligent Electronic Devices (IED), are very variate in terms of available sensor and actuator hardware interfaces and available LNs. Most commercial IEDs can be customized. In general, an IED will support LNs corresponding to one or several protection functions (like overcurrent, thermal, distance protection) and a variable number of sensor/actuator interfaces and or network interface. The combination corresponds to a typical electrical application (like transformer or generator or motor protection). An IED dedicated to measurement (i.e., no protection function LNs, only sensor/actuator interfaces) is called a Stand-Alone Measurement Unit (SAMU).

In Figure 2 two possible implementations of the simple protection function are shown: centralized on a versatile

IED or distributed using a SAMU with breaker control capabilities and a protection IED.

Figure 2. Two possible implementations: single IED or distributed.

Modern electrical grids heavily rely on distributed implementations and there are many reasons for this choice (see, for example, [2] for details). One of them is that network communication made the sensor information available to many consumers without extra sensor wiring and allow the deployment of many new applications requiring power grid data such as power quality or metering. Another reason is that smart-grids have to optimize production, distribution, consumption, energy storage, energy mix and increase availability of the electrical network. This requires more and more information available so, clearly, the data networks are a key element for the electrical network performance.

2.2. IEC 61850 communication protocols

Although serial-line communication is supported by IEC 61850, we will focus uniquely on Ethernet-based communication. IEC 61850 specifies three mappings to three communication protocols which correspond to the three main types of data flows in a power grid: local regular sampled values of sensors, local event transmission and remote control-room communication (supervisory control). Thus, two local Ethernet non-IP and one TCP/IP communication are defined. The TCP/IP-based protocol is Manufacturing Message Specification (MMS) [3] and the Ethernet non-IP protocols are Generic Object-Oriented Substation Event (GOOSE) [4] and Sampled Measured Values (SMV) [5]. Three logical levels of communication are present: remote TCP/IP between Supervisory Control and Data Acquisition (SCADA) and Protection IED [4], local (bay) communication between IEDs and process bus communication between SAMUs and IEDs. Figure 3 shows the respective communication levels and protocols. GOOSE protocol is designed for grid event propagation between IEDs like transmission of a trip event. SMV is designed for the periodic broadcasting of sensor measures by the SAMU.

The three communication levels do not necessarily correspond to three different networks although the standard recommends network segmentation and isolation as the three types of traffic have different requirements: low volume weak real-time for MMS (around 10kbps per IED traffic and 100 to 1000ms response time), average traffic and hard real-time for GOOSE (1kbps per flow with bursts at 1Mbps in case of events and 3ms response time) and heavy traffic hard real-time for Sample Values (SV) (around 5Mbps per measurement point and 4ms response time). Typical Ethernet connection available on a commercial IED is Fast Ethernet 100Mbps. This is mainly due to the limited resources available on IED.

Figure 3. Communication levels and protocols specified by IEC 61850.

Due to the criticality of data flows in a power grid, the use of high availability networks like Rapid Spanning Tree Protocol (RSTP), Parallel Redundancy Protocol (PRP) or High-availability Seamless Redundancy (HSR) is recommended. High quality IEDs are often equipped with up to three redundant interfaces for process, bay and SCADA communications. Some of the SAMUs available on the market cannot be configured otherwise than for a redundant network. When flows with different real-time requirements share the same network, the use of VLAN is also highly recommended.

Although there is no “best choice” of a solution between RSTP, PRP and HSR, due to the long recovery time, RSTP is not suitable for bay and process communication (GOOSE and SMV). In practice, HSR seems to be more suitable for process bus communication (SMV to IED) as it does not need supplementary interconnection equipment and therefore is cheaper for simple network topologies, while PRP is most suitable for station communication (IED to IED) while it allows complex topologies.

We adopted the HSR solution for our experimental part. We briefly present HSR as the choice of the network support has an impact on attack implementation.

2.3. HSR networks

HSR is one of the two protocols described in IEC 62439-3 [6]. The network topology is a ring built by the interconnection of nodes having two ports operated in parallel (Doubly Attached Node with HSR protocol – DANH). A source will duplicate every frame, add an identification tag “A” or “B” and send each copy of the frame on an Ethernet port. A destination will receive the two copies, remove the tag keep the first arrived and destroy the second. All nodes are forwarding frames from one port to another except if they already sent the same frame into the same direction. A Singly Attached Node (SAN) may participate to the ring if they connect to a REDundancy Box (RedBox). Several SAN may connect to the ring via the same RedBox and a switch. The RedBox is periodically broadcasting into the ring all the MAC addresses seen on the external interface. Then the RedBox may be used to interconnect the HSR ring with a regular LAN. A basic HSR ring configuration is presented in Figure 4.

Figure 4. HSR ring including DANH nodes and a RedBox.

Note that HSR does not specify a real-time traffic scheduling algorithm, determinism is not guaranteed by design. For optical rings a less than 50ns jitter is expected for each converter.

3. Experimental Workbench

Access to real substation automation network is difficult due to their critical mission and industrial secret. Moreover, as we intend to test attacks scenarios it is clearly impossible to experiment on a real electrical substation. Using the available material of our SCADA cybersecurity lab G-ICS [7], we built a small substation automation prototype connected to a simulated electrical grid using our hardware-in-the-loop simulation system [8].

3.1. Use-case and topology

The test workbench is presented in Figure 5. Four IEDs are used: a protection IED, a control IED and two SAMUs. The four IEDs participate to an optical HSR ring on multi-mode 100 Mbps fiber optic. The two SAMUs are active DANH while the two IEDs use a passive DANH extension card (only receiving and forwarding frames). A RedBox is also connected to the HSR ring. It will be mainly used for external traffic observation and attacks. The two IEDs have also two other redundant network interfaces: one that is used for real time communication with other IED in a PRP network and a second one for the traffic with the control room in an RSTP network. The full workbench includes two other protection IEDs and a SCADA, but for the scope of this paper only the four participants to the HSR ring are relevant. In accordance with the classification of the IEC 61850 standard, this use-case is relevant for a small distribution substation (1 to 5 IEDs).

The considered use-case is similar to the reference topology 7.3.2.3.9 (Process bus as a single ring) and the case studies from IEC 61850-90-4 technical report [9].

The protection IED is a distant protection. It will use measurement points (current and voltage) and estimate the location of a fault based on line impedance measurement or current measurement. In our setup the distant protection will use the measure point of one of the SAMUs and pick-up on overcurrent. The control IED is not implementing a protection function. It is used for interlocking of protection IED. That means, it is checking that the joint state of the protection IED and circuit breakers is coherent and change in the state of breakers will not

FastEthernet. Simple arithmetic shows that at most 20 SMV flows may be transmitted in a 100Mbps HSR ring. Actually, this is overestimated while it is not considering the non-real time traffic that may be present. This includes the supervision frames sent by all active nodes plus any other sporadic messages.

3.3. Experimental results

We used two different traffic measurement points: one on the external RedBox interface a second one using a network tap inside the HSR ring on one of the directions. The measurements on the two points were performed independently. At the time when the experiments were performed (2018 and 2019) no FastEthernet optical tap was available on the market. We then used two fiber/copper converters (Allied Telesis DMC 100/LC) and a copper network tap (Dualcomm DCSW-1005). Observed traffic characteristics in the two points were quite the same. We'll present below only the measurements through the external interface of the RedBox in order to avoid discussing the possible injected jitter by the converters and tap. On the other hand, this measure will show the characteristics of traffic received by a protection relay or power meter connected to the RedBox.

We perform the following measurement experiments.

3.3.1. One single SMV flow without time synchronization. This simple experiment will check the basic performance of SMV transmission. We calculated the statistic distribution of interarrival times of the SMV frames normalized by the nominal interarrival length (250 μ s). Results are displayed in Figure 8 for 100 bins. As frames are not timestamped the classical jitter calculation algorithms as the one in RFC 3550 [13] do not apply. We use the interarrival times variation as jitter estimator.

Figure 8. Histogram of normalized interarrival times for a single SMV flow.

For a perfectly stepped SMV flow all values would fall into the bin around 1 (i.e. 250 μ s between frames). The mean value of the normalized interarrival time is 0.999974 for a data set of over 100000 frames (around 27 seconds traffic capture) and that is quite perfect. But clearly there is an important spread of the values. There are 11 (out of 108509) values between 2.5 and 3.3 (i.e. an interarrival time between 625 and 825 μ s) not displayed in Figure 8

while the height of the corresponding bins is too small to be visible on the figure. The shortest expected normalized interarrival value is 0.052 (as a SMV frame duration is less than 13 μ s). However, 0.13% of the measured intervals are less than 0.052 (see the first bin in the Figure 8 histogram) and 22 measures are zero which is obviously the result of the limited WinPcap timestamping precision². We tested several configurations for WinPcap time source but, for the moment, we did not find a satisfactory solution to improve WinPcap precision so this point has to be further clarified in the future. Due to this biased timestamp we consider our measures as a worst-case lower bound of the SMV jitter.

There are several points of interest to be discussed concerning this experimental real-time performance assessment. The first one is to understand the source of this important variability of interarrival times and its consequences on the automation functions. The standard variation of the sample is 0.217 and we confirmed it on several traffic captures. That means, the expected frame interarrival interval is between 196 and 304 μ s. For a little less than 25% of the frames the interarrivals time is outside these limits. It is difficult to explain this variability as the frames are not timestamped by the SAMU. The inspection of data sets showed that traffic other than SMV on the external interface of the RedBox represents less than 0.1% of the capture and are only short ARP frames. Inside the HSR ring some HSR supervision frames are present, but this is also an insignificant traffic as there is a short (70 bytes) frame per DANH every two seconds. We assume that the interarrival times variability reflects the limit of the real-time performance of the SAMU network card. Adding time synchronization did not change the performance in a significant way (less than 1%). This is not surprising while our measurement concerns a short time performance evaluation (less than 30 seconds) while the time synchronization will correct the long-term internal clock drift. The fact that adding the 1PPS time synchronization does not improve the short term measured jitter simply means that the internal clock of the SAMU has a very small drift that will not manifest on short duration.

The crucial question is how the jitter will affect the protection functions behavior. During our measurement sessions two IEDs (one control and one protection) were present, both of them using the measures in the SMV flow. There was no alert raised by the IED concerning the quality of the samples. This is due to the fact that the protection and control functions do not use a 4kHz sample rate. According to IEC 61850-5, protection and control functions requires only 480 samples/s, i.e., only 1/8 samples from the actual SMV flow. That has an important impact on the evaluation of the jitter. The same measures of the jitter variation, taking only one out of eight samples from the same dataset, show a standard deviation of interarrival duration reduced 10 times. The interarrival duration is between 211 and 300 μ s (0.845 to 1.198 normalized intervals). The corresponding histogram is shown in Figure 9. We kept the same interval and number of bins as in Figure 8 to allow the comparison between the two histograms.

2. <https://wiki.wireshark.org/Timestamps>

Figure 9. Histogram of normalized interarrival times as seen by protection and control IED (one out of eight samples).

Metering units, on the other hand, are using all the samples from the SMV flow, but do not have strict transmission timing requirements according with IEC 61850-5. Then, if the internal sensor sampling of the SAMU is correctly timed the transmission jitter will not affect power quality measurement. We assume that metering units just “trust” that the sampling rate is deterministic disregarding the jitter of the arrival times.

3.3.2. Two SMV flows sharing the same HSR ring.

We are interested now in the interaction of several SMV flows sharing the same HSR ring and the effect on the jitter. As only two SAMUs are available in our lab we set up an experiment with only two SMV flows. The flows have identical characteristics (frame size, sampling rate and priority). The results are showing that the interaction between flows is very important. The performances of the two flows are identical. One of the statistics is represented in Figure 10. Although the average value of the interarrival duration is still very close to $250\mu\text{s}$, jitter is greatly increased. Even with only two flows, the standard deviation is multiplied by two (0.462 versus 0.217 for a single flow) and the maximal interarrival interval is greater than 1.3ms (5.3 normalized value). Over 2% of an interarrival times in each flow are superior to 1ms (normalized value of four) and 10% are very short (note the high bin around 0 in Figure 10).

Figure 10. Histogram of normalized interarrival times for a SMV flow in a two flow measurement experiment.

The SMV flow seen by protection and control IED (one out of eight samples) has also degraded performance but is still quite correct (0.003 standard deviation with interarrival times between 184 and $328\mu\text{s}$ (0.737 to 1.305 normalized values). The corresponding histogram is presented in Figure 11. Measures were performed with 34000 frames per flow (more than eight seconds traffic capture).

Figure 11. Histogram of normalized interarrival times for a SMV flow in a two flow measurement experiment.

We conclude this section with the remark that the lack of real-time traffic scheduling mechanism in HSR has an important impact on the jitter performance. Further experiments with several SAMU (and SMV flows) sharing the same HSR ring need to be conducted in order to establish if, in heavily loaded networks, the jitter increases due to the cumulative effect of the SMV flows might occur in the degradation of the electrical protection function. To our knowledge, no such public study was conducted until today.

4. Attacks on process bus

We consider two attack scenarios on SMV flows. The first one consists in injection of false measure data into the process bus. The second one is a quantitative attack (i.e. a network flood).

4.1. False data injection

As there is no authentication mechanism on SMV flows and frames are only identified by `smpCnt` counter field (see Figure 7), obviously, the protocol is vulnerable and false data injection is, in theory, simple to implement. The target of the attack would be to trigger a protection function and therefore a trip order to a circuit breaker occurring in a partial disconnection of a part of the power grid.

Practically, inserting false SMV data means that the attacker has to sniff the legal traffic, read `smpCnt`, then insert a frame with the `smpCnt` incremented before the legal frame with the same `smpCnt` value arrives. As frames with `smpCnt` less than or equal to the `smnCnt` of the previously received one are ignored, according to IEC 61850, the vulnerability seems easy to exploit.

Although the attack seems straightforward, some characteristics of the protection functions behavior and SMV flow characteristics have to be considered.

The corrupted frames must be injected inside the HSR network. That means that there are two possibilities to inject data: either through the external interface of a RedBox or directly into the HSR ring. Second case is more difficult as it will require a specially programmed DANH. The attacker device has to support HSR and to counterfeited the HSR identifications: HSR sequence number and lane. Otherwise, the frame may be rejected by the HSR layer of the target. Timing is also an issue while frames arrive every 250µs in an optical ring. The entire attack chain: optical/electrical conversion, frame reading, decoding, counters modification, transmission, electrical/optical conversion has to be performed in less than 250µs. As this attack is difficult to implement we chose to inject the false frames through the external interface of a RedBox. Then, we do not need special hardware and we do not have to counterfeit HSR frame fields.

We conduct the experiment on a computer with 2.90GHz i7-4910 processor. A first attack attempt through the RedBox showed that the computing time is not fast enough to inject a SMV frame with incremented counter before the arrival of the legal frame. But, the counter verification is not strict. IEC 61850 allows several frames to be lost so it is enough to choose a counter value big enough with regard to the last legal frame counter value. Then, the false frame is accepted by the subscriber and all the frames with intermediate counter values are rejected. As the frame counter is reset every second, after the injection of a single false SMV frame the measurement system recovers in at most one second. On another hand, electrical measures are naturally subject to noise. Protection functions will not pick up on a single sample superior to the maximal current. Often, protection functions are timed with variable timeouts ranging from millisecond to seconds. It follows that a single frame or even a very short sequence is not enough to trigger a protection function and, subsequently, a trip signal.

In our experiment we inject a false SMV flow through the external interface of the RedBox. We use the spoofed MAC address of one of the SAMUs as sender. Attack program waits for the reset of the smpCnt of the legal SMV sequence then generates a false data flow with a SMV counter incremented of 1000. The flow is not real-time stepped, but simply sequenced with usleep system call function. The attack is implemented on a Linux computer with the free libIEC61850³ development stack and libpcap-dev library. Code and capture dataset are publicly available at http://gics-hil.gforge.inria.fr/datasets_hsr/.

The attack successfully triggers the protection function on the protection IED which generates a trip signal.

An even more harmful attack is developed on the same basis. Instead of a false measure superior to the overcurrent, we inject a flat zero value sample measure. This triggers the supervision function of the protection IED. As mentioned in section 3.1 supervision functions survey the internal status of the IED and also the sensors status and block the protection function and therefore the IED will be out-of-service. The default timing of supervision function on our IED is 10 seconds and, indeed, after 10s of transmission of a false SMV flow the supervision function triggers and blocks the IED. This second attack on the

supervision function is more harmful while the IED cannot recover automatically after the attack contrary to the attack on the protection function. The IED is blocked by the supervision function and a manual reset is necessary.

4.2. Quantitative attacks (network flood)

Eventually, we test the resilience of the communication system in case of quantitative attacks. We inject a large number of frames thru the external interface of the RedBox and we observe the traffic inside the HSR ring using the copper network tap and the two optical/electric Ethernet converters. Although the type of injected Ethernet frames is not important we inject SMV traffic to the multicast addresses subscribed by the two IEDs. The idea is that we will test both the HSR infrastructure and the reaction of the IED when a large quantity of sensor data is received.

The HSR ring collapses at a global load approximately of 76Mbps measured in Wireshark at Layer 2. That means about 81Mbps at Layer 1 (including preamble and FCS bytes). The two IEDs signal a failure of the measurement points.

The exact moment when the HSR ring starts failing can be seen if we trace the HSR supervision frames throughput. HSR supervision frames are 66 bytes long in Wireshark and normally each SAMU will send one frame every 2s while the RedBox will send one frame every 2s if no SAN is attached on the external interface and two frames every 2s if a SAN is attached. That means a steady traffic of 990bps if no SAN attached and 1320bps when a SAN is attached. Then, the attack timing can be followed on the HSR supervision traffic graph (Figure 12). At point A the attacker connects to the RedBox which starts sending two supervision frames instead of one. At point B HSR ring starts falling and some supervision frames are lost. At point C the attack stops and HSR ring recovers. Then, the attacker disconnects from the RedBox (D) and the ring comes back to the initial state (E). Note that, as the graph is a 10s moving average, the real events occurred several seconds earlier than marked on the graph.

Figure 12. Evolution of the HSR supervision traffic during the quantitative attack (10s moving average).

The total traffic graph is presented in Figure 13. Due to the important difference in volume between supervision traffic and total traffic the scales in Figures 12 and 13 are not the same. The approximate position of points B and C are indicated. Points A, D and E cannot be identified on the Figure 13 while the corresponding variation is too small compared to the SMV traffic.

The exact moment when the HSR ring starts falling may be identified on the traffic capture. Inspection of the dataset shows that at some point some Ethernet frames of unknown type are transmitted into the network (Figure

3. <http://libiec61850.com/libiec61850/>

Figure 13. Traffic evolution during the quantitative attack (1s moving average).

14). A byte inspection of the frames shows that they are actually malformed SMV frames results of a probable overflow of one of the internal buffers of the RedBox. The visible “data bytes” 0x75ab63 are part of the HSR tag, 0x88b4 is the SMV Ethertype and so on. The first malformed frame in the data set identifies the moment when the HSR ring is overflowed.

Figure 14. Malformed frames transmitted by the RedBox during attack.

Finally, in Figure 15 we present the impact of the attack on the legitimate SMV traffic. The graphic displays the throughput in frames/second of one of the two legitimate SMV flows. We can remark that up to 20% of the traffic is lost at some moments during the attack. This explains the failures of the measurement points signaled by the subscribers (IEDs).

Figure 15. Effect of the quantitative attack on the legitimate SMV flow.

We conclude this section with a positive remark: despite the fact that the quantitative attack was successful and the protection function was affected, the HSR ring recovered fast after the attack and none of the network nodes (including the RedBox) needed to be manually restarted. Only the failure error of the measurement points on the IED had to be acknowledged. Note also that the attacks are exploiting protocol vulnerabilities, they are not specific to a manufacturer implementation.

5. Detection and mitigation

This section describes researches on work. We have a detection specification and reaction methodology, but hardware implementation is not yet available.

Both false data injection and quantitative attacks can be detected by measurement of the traffic inside the HSR ring. An interesting point is that early detection of attacks may be achieved with HSR supervision traffic survey. That may help detecting an attacker that connects to the

external interface of a RedBox. A skilled attacker may avoid this detection if he connects using the MAC address of an already connected device.

Mitigation is an important point as the consequences of an attack are potentially critical. Even if detection is possible, the key is the reaction time. A protection function may be triggered by an attack in several tens of millisecond or even in some milliseconds. Supervisory functions may be triggered in several seconds but this is a very short time interval for a human reaction. Therefore, we consider the automatic response in case of attack.

Electrical protection applications are designed to isolate or reconfigure a part of the electrical network in case of an electrical fault. On modern communication systems, network faults are also handled using high availability network topology. Cyberattacks are not yet included in the reaction mechanism in commercial devices even if there are several references in the literacy (see Section 6). Of course, an attack is not a fault, so, extending an electrical protection application to consider cyberattacks in the reaction loop is not immediate.

Our proposal is to use an Intrusion Detection System (IDS) alert which identifies the attack (false data or network flood attack) as an input to the electrical protection function. Then, the electrical protection application will issue the adequate reconfiguration control in order to keep the power grid operational or, at least, safe. Usually, a safe configuration for the power grid exists and the protection and control IED can activate the circuit breakers and switch the grid to safe state. For example, if one of the measurement points is detected as compromised by the IDS, given the state of the other measurement points, the topology and the load of the grid, the estimator (control room) can figure out the electrical state of the compromised measurement point and reconfigure the protection application. The main issue is not how to reconfigure the protection application, which is purely an electrical engineering problem, but how to raise the alert signal from the IDS to the IED in an effective way and in a short time. The best solution would be to implement a host IDS based on traffic measurement directly on the HSR interface of the IED. Unfortunately, given the high processor load and the limited resources of the IED this is not possible for the time being.

Then, a network IDS has to be used. The difficulty is that the IDS has to survey both lanes of the HSR ring via network taps so it has to support HSR frame format. Communication with IED has to be implemented via an IEC 61850 protocol. As an alert is an event, the adequate protocol seems to be GOOSE messaging on the bay network. In a previous work we specified such a 61850 intrusion detection function [14], [15]. The alert event transmitted by the IDS via the GOOSE messaging on the bay network may be directly used by IEDs to inhibit trip messages and to trigger reconfiguration of the electrical protection activation in real-time. Although, usually, the bay network is physically separated from the process bus, it can also be attacked and the GOOSE protocol is also vulnerable to false data injection [16]. In that case, a last resort alarm system can be used to send an alert directly to the control room via a secure protocol. This last solution has a less performance as the station bus used to communicate with the SCADA is not real-time.

The complete reaction loop is represented in Figure 16.

Figure 16. Reconfiguration loop following a process bus cyber-attack.

The implementation of such an IDS presents two issues: first, at least four network interfaces are needed as the IDS has to listen the two HSR lanes and communicate on the bay and station networks. Secondly, it has to be a real-time device as traffic measurement and bandwidth estimation have to be accurately handled. For the moment, we still benchmark hardware platforms for the implementation.

6. Related work

In this paper we tackle two different problems: traffic measurement in HSR networks and attacks on the process bus. Traffic measurement in SCADA systems and electrical grid communication network were conducted mostly to find regularity properties of the traffic (as periodic sensor reading or actuator writing) [17] or to analyze complex flow configurations and characterize the endpoints and flow duration [18] in search for traffic patterns useful for security. HSR network performance was studied in [19] and [20] using a simulated network in OPNET. Propagation delay in HSR networks was studied in [21] using a special prototype of RedBox and timestamped SMV which are not available in 61850-9-2-le. Real-time performance was studied in HSR networks for GOOSE protocol [22], but only the propagation delay was studied as this is the important performance parameter for GOOSE. Probably, the closer study to our work is [23], where the global protection function is experimentally studied with measures transported via an HSR ring, but the properties of the traffic itself are not considered. With respect to the existent state of the art, we consider that our study of jitter variation in HSR rings is a novelty.

Concerning the attacks and detection in 61850 networks, the literacy is very rich. Specifically on false data injection in power-grids, a recent survey [24] is listing model-based and model-free learning methods. All these algorithms are based on the prediction of the evolution of measures. They build a mathematical model of the dynamical evolution of the electrical network and they detect intrusions based on the deviations of the received SMV values from the predicted evolution. All the exposed

methods are completely ignoring traffic properties. As far as we know no traffic measurement methods for detection were proposed until now. Flooding attacks on smart-grid networks where considered in several papers like [25] on generic network topology. In [26] authors consider explicit flooding of SMV networks and detection based on protocol fields. No particular process network topology is considered. Unfortunately, detailed descriptions of these attacks are not available. To our knowledge our work is the only one to provide the detailed description of the attacks, the implementation and the data sets.

Reaction to attack in smart-grids is a relatively recent topic. An early resilient approach is presented in [27]. A reaction mechanism for DNP3 networks is proposed in [28]. Another approach based on a backup system was proposed in [29]. A very complex cyber-physical framework, combining network and electrical protection reaction, is presented in [30]. More generally, reaction in SCADA systems is considered in [31], [32]. Our proposed reaction mechanism acts at protection function level and, therefore, is different from the previously cited approaches. We consider our approach as a field-level technical solution which may be integrated in high-level conceptual models.

7. Conclusions and future work

In this paper we present a first exploratory study of the traffic properties of the IEC 61850 process bus communications on HSR rings. We obtain an experimental characterization of jitter and influence of interaction of two SMV flows on the jitter. The main finding is that the lack of real-time scheduling in HSR has a real impact on traffic jitter and this impact is more important when several SMV flows share the HSR ring. The results will be refined in a future research with the development of a higher performance and less intrusive traffic measurement device based on passive Keysight Flex Tap⁴ and simultaneous measures on the two lanes of the HSR ring. New experiments will include the comparison of the jitter on the two lanes and also an extended study of the interference between multiple SMV flows as we have recently acquired five more SAMUs.

In the present research we set up several attack scenarios based on false data injection and Ethernet flood. A positive finding is that the field devices (IEDs and SAMUs) and the RedBox recover well after the Ethernet flood. The main network vulnerability is the lack of protection mechanisms in HSR RedBox against external interface flooding. The conclusion is that IEC 62439-3 needs to be completed with cybersecurity clauses.

On the detection and reaction part, the main effort will focus on a hardware implementation on FPGA of the IDS and an experimental study of the reaction loop performance for bay network and system network alert transmission.

Another interesting research direction is the cross-domain alert correlation. Indeed, it would be interesting to correlate our traffic measurement-based alerts with model-based IDS alerts in the literacy. As the model-based IDS are prone to false positive our approach may help false

4. <https://www.ixiacom.com/products/network-taps/>

positive reduction, detailed attack scenario reconstruction and attack vector identification.

Eventually, we will conduct a similar study on other high availability alternative solutions in particular on Parallel Redundancy Protocol (PRP). As most of the Red-Boxes and SAMUs can be configured as HSR or PRP DAN, the reconfiguration of the test bench is simple.

Acknowledgment

The test benchmark was built using material acquired via Siemens Academic Program.

References

- [1] "IEC international standard - communication networks and systems for power utility automation - part 5: Communication requirements for functions and device models," *IEC 61850-5 Edition 2.0 2013-01*, pp. 1–306, 2013.
- [2] N. Higgins, V. Vyatkin, N. C. Nair, and K. Schwarz, "Distributed power system automation with iec 61850, iec 61499, and intelligent control," *IEEE Transactions on Systems, Man, and Cybernetics, Part C (Applications and Reviews)*, vol. 41, no. 1, pp. 81–92, 2011.
- [3] "ISO 9506 industrial automation systems — manufacturing message specification — part 1: Service definition and part 2: Protocol specification," *ISO 9506-1 and 2:2003*, 2003.
- [4] "IEC international standard - communication networks and systems for power utility automation - - part 8-1: Specific communication service mapping (SCSM) - mappings to MMS (ISO 9506-1 and ISO 9506-2) and to ISO/IEC 8802-3," *IEC 61850-5 Edition 2.1 2020-02*, pp. 1–677, 2020.
- [5] "IEC international standard - communication networks and systems for power utility automation - - part 9-2: Specific communication service mapping (SCSM) - sampled values over ISO/IEC 8802-3," *IEC 61850-5 Edition 2.0 2011-12*, pp. 1–677, 2011.
- [6] "IEC international standard industrial communication networks - high availability automation networks - part 3: Parallel redundancy protocol (PRP) and high-availability seamless redundancy (HSR)," *IEC 62439-3 Edition 3.0 2016-03*, pp. 1–354, 2011.
- [7] M. Kabir-Querrec, S. Mocanu, J.-M. Thiriet, and E. Savary, "A Test bed dedicated to the Study of Vulnerabilities in IEC 61850 Power Utility Automation Networks," in *21st IEEE Emerging Technologies and Factory Automation*, Berlin, Germany, Sep. 2016.
- [8] S. Mocanu, M. Puys, and P.-H. Thevenon, "An Open-Source Hardware-In-The-Loop Virtualization System for Cybersecurity Studies of SCADA Systems," in *C&esar 2019 - Virtualization and Cybersecurity*, Rennes, France, Nov. 2019, pp. 1–16.
- [9] "IEC international standard - communication networks and systems for power utility automation - part 90-4: Network engineering guidelines," *IEC 61850-90-4 Edition 1.0 2013-08*, pp. 1–268, 2013.
- [10] "IRIG serial time code formats," *IRIG Standard 200-04, Range Commanders Council, U.S. Army White Sands Missile Range, New Mexico 88002-5110.*, pp. 1–73, 2004.
- [11] "X.680-X.693 : Information technology - abstract syntax notation one (ASN.1) and ASN.1 encoding rules," 10 2015.
- [12] "Implementation guideline for digital interface to instrument transformers using IEC 61850-9-2," *UCA International Users Group*, pp. 1–31, 2004.
- [13] H. Schulzrinne, S. L. Casner, R. Frederick, and V. Jacobson, "RTP: A Transport Protocol for Real-Time Applications," RFC 3550, Jul. 2003.
- [14] M. Kabir-Querrec, S. Mocanu, J.-M. Thiriet, and E. Savary, "Power Utility Automation Cybersecurity: IEC 61850 Specification of an Intrusion Detection Function," in *25th European Safety and Reliability Conference (ESREL 2015)*. Zürich, Switzerland: CRC Press, Sep. 2015.
- [15] J. Hoyos, M. Dehus, and T. X. Brown, "Exploiting the goose protocol: A practical attack on cyber-infrastructure," in *2012 IEEE Globecom Workshops*, 2012, pp. 1508–1513.
- [16] M. Kabir-Querrec, S. Mocanu, P. Bellemain, J.-M. Thiriet, and E. Savary, "Corrupted GOOSE Detectors: Anomaly Detection in Power Utility Real-Time Ethernet Communications," in *GreHack 2015*. Grenoble, France: Verimag, Nov. 2015.
- [17] R. R. R. Barbosa, R. Sadre, and A. Pras, "A first look into scada network traffic," in *2012 IEEE Network Operations and Management Symposium*, 2012, pp. 518–521.
- [18] K. Mai, X. Qin, N. O. Silva, and A. A. Cárdenas, "IEC 60870-5-104 network characterization of a large-scale operational power grid," in *2019 IEEE Security and Privacy Workshops, SP Workshops 2019, San Francisco, CA, USA, May 19-23, 2019*. IEEE, 2019, pp. 236–241.
- [19] L. Xu, H. Li, and L. Chen, "Modeling and performance analysis of data flow for hsr and prp under fault conditions," in *2018 IEEE Power Energy Society General Meeting (PESGM)*, 2018, pp. 1–5.
- [20] S. Kumar, N. Das, and S. Islam, "Implementing prp and hsr schemes in a hv substation based on iec62439-3," in *2018 Condition Monitoring and Diagnosis (CMD)*, 2018, pp. 1–5.
- [21] J. Liu, Y. Li, X. Li, H. Lyu, G. Yang, and J. Wen, "Design and implementation of delay measurement in prp and hsr redbox," in *2019 IEEE 2nd International Conference on Electronics Technology (ICET)*, 2019, pp. 45–50.
- [22] M. Hosny Tawfeek Essa and P. Crossley, "Goose performance assessment on an iec 61850 redundant network," *The Journal of Engineering*, vol. 2018, no. 15, pp. 841–845, 2018.
- [23] V. Leitloff, P. Brun, S. de Langle, B. Ilas, R. Darmony, M. Jobert, C. Bertheau, P. Ferret, M. Boucherit, G. Duverbecq, J. P. Cayuela, and R. Bouchet, "Testing of iec 61850 based functional protection chain using non-conventional instrument transformers and samu," in *13th International Conference on Development in Power System Protection 2016 (DPSP)*, 2016, pp. 1–6.
- [24] A. S. Musleh, G. Chen, and Z. Y. Dong, "A survey on the detection algorithms for false data injection attacks in smart grids," *IEEE Transactions on Smart Grid*, vol. 11, no. 3, pp. 2218–2234, 2020.
- [25] F. Zhang, M. Mahler, and Q. Li, "Flooding attacks against secure time-critical communications in the power grid," in *2017 IEEE International Conference on Smart Grid Communications (Smart-GridComm)*, 2017, pp. 449–454.
- [26] M. El Hariri, E. Harmon, T. Youssef, M. Saleh, H. Habib, and O. Mohammed, "The iec 61850 sampled measured values protocol: Analysis, threat identification, and feasibility of using nn forecasters to detect spoofed packets," *Energies*, vol. 12, no. 19, 2019.
- [27] B. X. Zhu, "Resilient control and intrusion detection for scada systems," Ph.D. dissertation, UC Berkeley, May 2014.
- [28] J. Bai, S. Hariri, and Y. Al-Nashif, "A network protection framework for dnp3 over tcp/ip protocol," *Proceedings of IEEE/ACS International Conference on Computer Systems and Applications, AICCSA*, vol. 2015, pp. 9–15, 03 2015.
- [29] A. Babay, J. Schultz, T. Tantillo, S. Beckley, E. Jordan, K. Ruddell, K. Jordan, and Y. Amir, "Deploying intrusion-tolerant scada for the power grid," in *2019 49th IEEE/IFIP Int. Conf. on Dependable Systems and Networks (DSN)*, June 2019.
- [30] Y. Lopes, N. C. Fernandes, D. C. Muchaluat-Saade, and K. Obraczka, "ARES: An autonomic and resilient framework for smart grids," in *2017 IFIP/IEEE Symposium on Integrated Network and Service Management (IM)*, 2017, pp. 222–229.
- [31] B. A. Baalbaki, Y. Al-Nashif, S. Hariri, and D. Kelly, "Auto-nomic critical infrastructure protection (acip) system," in *2013 ACS International Conference on Computer Systems and Applications (AICCSA)*, 2013, pp. 1–4.
- [32] Q. Chen and S. Abdelwahed, "Towards realizing self-protecting scada systems," in *Proceedings of the 9th Annual Cyber and Information Security Research Conference*, ser. CISR '14. New York, NY, USA: Association for Computing Machinery, 2014, p. 105–108.