Evaluation of a Low-Latency, High-Accuracy Intrusion Prevention System for Digital Substations

Nicola Cibin, Marijn Boringa, Ioannis Semertzis, Alexandru Ştefanov
Department of Electrical Substationable Energy
Delft University of Technology
Delft, The Netherlands
n.cibin@tudelft.nl, m.w.boringa@student.tudelft.nl, i.semertzis@tudelft.nl, a.i.stefanov@tudelft.nl

***Abstract*—The increasing digitalization of the power grid and widespread adoption of the IEC 61850 standard have raised concerns regarding the cyber security of electrical substations. In particular, the Sampled Values (SV) protocol lacks inherent security mechanisms, leaving protection and control functions vulnerable to malicious SV frames injection. To address this challenge, a method using hybrid statistical-deep learning was previously proposed to detect, prevent, and localize SV injection attacks in digital substations, while ensuring minimal latency and near-zero false positives. In this work, we evaluate and experimentally validate the real-time implementation of the proposed Intrusion Prevention System (IPS) using a hardware-in-the-loop testbed consisting of a Real-Time Digital Simulator (RTDS), Raspberry Pi implementing the intrusion prevention method, and an industrial intelligent electronic device. Multiple operating scenarios are considered to evaluate fault clearance latency, communication reliability, IPS classification performance, and IPS impact on protection schemes. Results demonstrate that the IPS introduces less than 150 µs of additional delay, achieves a near-zero false positive rate in the absence of cyber attacks, and attains an F1-score above 99% in preventing SV injection attacks, while preserving substation protection and control functionalities. These findings confirm the practicality and effectiveness of the proposed IPS for deployment in IEC 61850-compliant digital substations.**

***Keywords*—Cyber Attacks, Digital Substation, Hardware-in-the-Loop, IEC 61850 Sampled Values, Intrusion Prevention System, Statistical Analysis**

I. Introduction

In recent years, successful cyber attacks on electrical substations have demonstrated their potential to disrupt the electricity supply [1, 2]. While the digitalization of substations through the adoption of the IEC 61850 standard offers multiple benefits, including more efficient protection, automation, and control systems, it also introduces an expanded cyber attack surface. The emerging cyber security risks primarily arise from the increased integration of Information and Operational Technology (IT-OT) infrastructures.

Message authentication and encryption are not implemented to secure IEC 61850 Sampled Values (SV) data streams due to the strict latency requirement of 3 ms for power system protection schemes, high frequency of SV communications, and preference of not increasing devices' computational requirements and communication network infrastructure complexity. This allows various cyber attacks that exploit the SV protocol to send malicious current and voltage measurements into the OT communication network in digital substations, potentially causing circuit breakers tripping or disrupting the normal operation of protection schemes, e.g., flooding, spoofing, replay, high *smpCnt*, and Man-in-the-Middle (MitM) attacks [3]. Various Intrusion Detection Systems (IDS) have been proposed in the literature to detect such attacks [4]. However, the detection of SV injection attacks is insufficient to prevent the malicious opening of circuit breakers. Thus, in [3], a novel method using hybrid statistical-deep learning is proposed for the detection, prevention, and source localization of IEC 61850 SV injection attacks.

Power system operators are often reluctant to deploy Intrusion Prevention Systems (IPS) in digital substations given the risk of potential OT communication network disruptions, increased communication latencies, and required computational resources. To address these issues, this paper evaluates and experimentally validates the real-time implementation of the previously proposed IPS method in an Hardware-in-the-Loop (HIL) IEC 61850-compliant testbed consisting of a Real-Time Digital Simulator (RTDS), an industrial Intelligent Electronic Device (IED), and a Raspberry Pi 5 (RPi). Unlike the previous work that evaluated the IPS method offline using multiple datasets, the setup presented in this paper enables online, high-fidelity, real-time evaluation through an Electromagnetic Transient (EMT) simulation and interactions among systems within an operational substation.

The contributions of this paper are summarized as follows:

1) We evaluate the proposed IPS method under multiple operating conditions, including steady-state power system operation, three-phase-to-ground faults, presence or absence of SV injection cyber attacks, and combinations of these scenarios. The metrics considered for the evaluation include intrusion prevention False Positive Rate (FPR), True Positive Rate (TPR), and F1-score, and the introduced additional latency due to the processing and forwarding of the filtered SV frames.
2) We analyse the impact of the IPS on the protection schemes' normal operation. This also includes the impact on power system stability due to a potential increase in fault-clearing time and potential unintended protection scheme operations due to misclassified SV frames.
3) We discuss the implementation challenges that need to be faced when deploying an in-band IPS on computationally constrained OT devices, while ensuring low-latency and high-throughput communications. In particular, we describe how such strict requirements can be met on a general-purpose single-board computer.

The remainder of the paper is organized as follows. Section II discusses the IEC 61850 SV protocol and intrusion prevention method under evaluation. Section III describes the experimental setup, simulated scenarios, and. IPS real-time deployment. Section IV analyses the IPS performance and impact on protection schemes. Finally, Section V concludes the paper.

This work has received funding from the European Commission in the Horizon Europe ESTELAR Project with Grant Agreement No. 101192574.
Corresponding author: Alexandru Ștefanov (a.i.stefanov@tudelft.nl).

979-8-3315-8138-1

II. Intrusion Prevention in Digital Substations

A. IEC 61850 Sampled Values

IEC 61850 SV is a publisher-subscriber protocol employed for transmitting three-phase current and voltage measurements from Merging Units (MUs) to IEDs. In this scheme, MUs multicast SV frames, while IEDs selectively process only the subscribed SV data streams. Depending on the power system frequency and application requirements, the standard specifies multiple SV publishing rates. In practice, the most common configurations in digital substations are 4000 and 4800 frames per second for 50 and 60 Hz systems, respectively. The standard also defines the internal structure of SV frames, including the *smpCnt* field, a counter that increments with each transmitted frame and resets to zero at the start of every second. As a result, *smpCnt* filed values range from zero to one less than the total number of frames published per second by the MU. This mechanism introduces a degree of determinism in the expected arrival time of each SV frame identified by a given *smpCnt* value. Specifically, for a given second i, a frame publishing rate FS, and a *smpCnt* value c, the theoretical arrival time of an SV frame, denoted as $F_a^e(i,c)$, can be expressed as:

$$F_a^e(i,c) = i + {}^{c}/_{FS}, \quad \forall c \in \{0,1,\dots,FS-1\} \tag{1}$$

This determinism in the SV frames' arrival time provides the method under validation with an informative advantage, which can be exploited to discern legitimate and malicious SV frames in real-time, on a per-frame basis.

B. Intrusion Prevention Method

As described in [3], the difference between the theoretical ($F_a^e(i,c)$) and measured ($F_a^{real}(i,c)$) SV frames arrival time, defined as F_{as}, is modeled using an Exponentially Modified Gaussian (EMG) Probability Distribution Function (PDF). The parameters of this distribution are continuously estimated during operation using the Method of Moments Estimator (MME) by using the latest F_{as} values computed on the received SV frames. Then, for each newly received frame, the EMG PDF (ϕ_{EMG}) is used to compute the likelihood of the frame being legitimate (F_p), such that $F_p = \phi_{EMG}(F_{as})$. As discussed in [3], for an attacker, it is extremely difficult to estimate F_a^{real} from a device different than the targeted IED and to inject the malicious SV frames matching the expected PDF. Thus, when received by the SV subscriber, the SV frames published by the legitimate MU have a higher legitimacy likelihood than the malicious frames sent by the attacker from a compromised device. This allows the proposed method to discard the malicious frames and forward only the legitimate ones to the IED.

Given the strict timing requirements of digital substation communications, when deploying the method it is essential to minimize the latency introduced by the IPS. To this end, the method was implemented in C++ and designed around four threads being executed in parallel:

1) *Frames capture and intrusion prevention:* Monitors the Network Interface Card (NIC) for incoming SV frames belonging to the subscribed SV stream and performs intrusion prevention by verifying that: (a) *smpCnt* field value lies within *[0, FS-1]*, where *FS* is the number of SV frames published per second; (b) the elapsed time since the last accepted frame with the same *smpCnt* is higher than 500 µs; (c) the frame arrival time falls within the DoS protection interval, $\mu + \lambda \pm 5 \cdot \sigma$, where μ, σ, and λ correspond to the EMG distribution mean, standard deviation, and exponential decay, respectively; (d) and, if multiple frames with the same *smpCnt* are received by the IPS within one second, only the frame with the highest legitimacy likelihood is not discarded.
2) *Frame forwarding:* Forwards the SV frames with the highest legitimacy likelihood to the IED.
3) *EMG parameter update:* Updates the EMG PDF parameters using the MME on the measured frames' arrival time shift of the latest accepted frames.
4) *IPS logging and performance monitoring:* Continuously records IPS performance metrics, including the number of received, forwarded, and discarded frames, frames' forwarding latency, and EMG parameters values.

III. Experimental Setup

The experimental testbed comprises RTDS and an Ethernet-based substation network used for IEC 61850 communication experiments. The RTDS is interfaced to the network switch through a GTNETx2 card; one card module is configured as an SV publisher, while the other is configured to publish and subscribe to Generic Object-Oriented Substation Event (GOOSE) messages. An IED is connected to the same network switch, and, when the simulation scenario requires it, an RPi is inserted between the IED and network switch. Two additional workstations are connected to the switch: the first one is attached to a mirror port and is used to monitor the OT network traffic, whereas the second one is used to inject malicious SV frames into the OT network. The testbed configuration is depicted in Figure 1.

A. Power System Model

RTDS simulates a small power system consisting of a 13.8 kV, 60 Hz voltage source acting as an infinite bus and supplying power to a load of 5 MW. The load was connected to the source by a circuit breaker and a transmission line. A voltage and current transformer measure the three-phase voltage and current between the source and circuit breaker. The measured voltages and currents for all three phases and neutral are sent as SV frames at a rate of 80 samples per AC cycle, i.e., 4800 frames per second, to the IED. The breaker facilitates tripping of all three phases separately and is controlled by GOOSE messages sent by the IED. Under normal operation, the IED sends a GOOSE message every second as a heartbeat, thus ensuring it is working properly. During a fault, the sending rate increases to one per millisecond and then decreases gradually to 1 per second. This ensures that the tripping signal is timely received by the

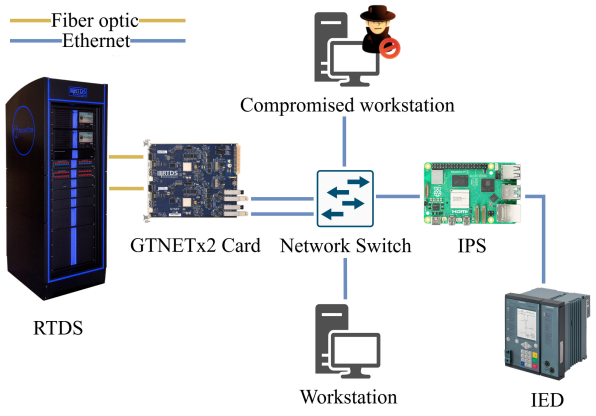

Figure 1: Experimental testbed architecture.

GOOSE subscribers. Finally, different types of faults are simulated in the middle of the transmission line, including line-to-ground and line-to-line faults. Still, for the scope of this work, the simulations are limited to three-phase-to-ground faults without loss of generalization.

B. Protection Scheme Configuration

The IED deployed on the testbed is configured for overcurrent and earth fault protection. As the current flows only from the source to the load, the overcurrent protection is non-directional. Furthermore, to achieve a fast fault clearing, two stages of overcurrent protection are configured. The first stage is described by an IEC extreme inverse operating curve, which ensures fast or slower fault clearing for high or lower currents, respectively. The second stage is configured with definite time characteristics, so it operates with the same time delay independently of the current magnitude. When the overcurrent protection is activated due to a detected fault, GOOSE messages are sent by the IED to instruct which phases of the circuit breaker need to be disconnected.

C. Intrusion Prevention System Implementation

The IPS is deployed on an RPi running RPi OS Lite. To improve determinism and reduce SV frames processing delay, the Linux kernel was recompiled with the PREEMPT_RT patch, thereby enabling fully preemptible kernel operation and reducing scheduling latency. The IPS code was compiled with GCC using high-level optimization flags to maximize instruction throughput and minimize execution overhead. The RPi is equipped with two Ethernet interfaces, one connected to the network switch and the other to the IED. Whereas the network traffic received by the first interface is forwarded to the IED by the IPS process running on the RPi, the traffic received by the second interface is directly forwarded to the first one, and then to the network switch, through a kernel-level software bridge. This bridge ensures a low processing overhead and minimizes the forwarding latency for the GOOSE traffic exchanged between the IED and RTDS.

Even though, for the scope of this work, the method is implemented in a dedicated device connected to the IED, vendors could consider implementing the method in the IED itself. This would further reduce the latency introduced by the IPS, without significantly increasing the computational load of the IED supporting the IPS.

D. Cyber Attack Scenarios

The SV injection attacks are conducted from a workstation directly connected to the network switch used to deliver the OT network traffic to the IED. Given that SV frames are sent in multicast, the attacker is able to monitor the network and sniff the legitimate SV frames both during normal and fault conditions. This allows the attacker to gain information about the SV frames' content, in particular, the legitimate source and destination MAC addresses, the SV APPID field, and the reported current and voltage measurements. Depending on the simulated attack scenario, the attacker either replays normal or fault operating conditions, while spoofing the identity of the legitimate SV publisher. The attack was practically implemented in a Python script, which aimed at executing the tcpreplay tool [5] with the correct timing to match the SV frame arrival time expected by the attacked SV subscriber. The five scenarios in which the attacker was able to inject the SV frames with the best timing are considered for the evaluation of the IPS.

TABLE I
TESTBED CONFIGURATION DURING THE SIMULATIONS OF DIFFERENT SCENARIOS.

Case	Power System State	Raspberry Pi	Cyber attack	No. of runs
N0	Steady	Not deployed	Inactive	5
N1	Fault	Not deployed	Inactive	1000
N2	Steady	Deployed	Inactive	5
N3	Fault	Deployed	Inactive	1000
N4	Fault	Deployed*	Inactive	1000
C0	Steady	Not deployed	Active	5
C1	Fault	Not deployed	Active	5
C2	Steady	Deployed	Active	5
C3	Fault	Deployed	Active	5

E. Simulated scenarios

A series of scenarios, also referred to as “cases” in the following, were designed to evaluate the proposed IPS under both steady and three-phase-to-ground fault conditions, and in the absence and presence of cyber attacks. For each scenario, power system and OT network traffic data were acquired for 60 seconds, with each scenario running a different number of times depending on the variability of the operating conditions, as presented in Table I.

Case N0 represents the baseline scenario of normal operation in the absence of the RPi and cyber attacks, serving to confirm that the IED remains idle. Case N1 uses the same configuration as case N0, but it introduces a fault in the power system; this case provides the baseline to measure the IED reaction and fault clearance delay. Case N2 reproduces normal operations but with the RPi connected and executing the IPS, ensuring that no false positives are caused by the IPS. Case N3 uses the same configuration as case N2 but introduces a fault, enabling measurement of IED response and fault clearance delay while monitoring for potential false positives. Moreover, a case named N4 is considered, in which a fault is simulated while the RPi is connected but not executing the IPS; packets are forwarded transparently by the RPi thanks to a bidirectional kernel bridge between its two NICs. This latter scenario allows to estimate the latency introduced by the RPi communication stack itself.

The cyber attack scenarios begin with case C0, which simulates normal operation without the RPi, while an SV injection attack forces the IED to malfunction. Case C1 repeats this process under fault conditions. Case C2 evaluates normal operations with the RPi connected and executing the IPS during an SV injection attack. Finally, case C3 represents a fault with both the RPi deployed and an ongoing cyber attack. These latter two cases validate that the IED, protected by the IPS, neither issues unintended control signals nor becomes unresponsive, and successfully clears the fault while maintaining low response and fault clearance latency.

IV. RESULTS AND DISCUSSION

In the following subsections, the IPS performance are evaluated considering the following metrics: 1) additional fault clearance delay introduced by the RPi and IPS; 2) FPR provided by the IPS in absence of cyber attacks; 3) FPR, TPR, and F1-score provided by the IPS during SV injection attacks; and 4) impact of the IPS on the efficacy of the protection functionalities.

A. Fault Clearance Delay

To evaluate the fault clearance delay introduced by the IPS, three experimental configurations are considered, namely, cases N1, N3, and N4. To have a reliable estimate of the introduced additional delay, the number of required simulations for each case was determined using Chebyshev's inequality applied to the sample mean:

$$\Pr\left(|\bar{X} - \mu| < {}^{k\sigma}/_{\sqrt{n}}\right) \geq 1 - {}^{1}/_{k^2} \quad (2)$$

where $\bar{X}$ is the sample mean, μ the population mean, σ the population standard deviation, n the number of samples, and k a positive constant controlling the confidence level. By performing 1000 simulations per case, Chebyshev's inequality guarantees that, given σ equal to 2.6, the difference between the estimated latency mean $\bar{X}$ and the true latency mean μ is less than 370 µs with a probability greater than 95%.

For this evaluation, the fault clearance delay is measured as the time between a phase current exceeding 110% of its nominal peak value and the first change in the GOOSE control signals received by the RTDS from the IED. This metric is obtained by analysing the comma-separated values (.csv) files exported from the RTDS simulation. As presented in Table II and Figure 2, the additional latency introduced by the RPi when forwarding the network traffic is 96 µs. Then, when the IPS is enabled, the introduced latency increases to 248 µs. These results highlight that if the IPS is directly deployed with an IED with computational resources similar to the RPi, the latency increase due to the IPS is 152 µs.

B. IPS Performance During Non-Attack Conditions

Given that cyber attacks on critical infrastructures are low-frequency, high-impact events, most of the time the IPS has to ensure no interference on the reception of legitimate SV frames by the IED in the absence of cyber attacks. The IPS performance in this scenario is evaluated by considering cases N2 and N3. Since the IPS received only SV frames from the legitimate publisher, each frame was assigned the highest legitimacy likelihood, and the IPS ensured their delivery to the IED with minimal additional delay. In these scenarios, the IPS provides an FPR of 0%, indicating that none of the 2,880,000 legitimate SV frames received during each simulation run were discarded. Moreover, the measured processing time for each frame is 99 µs, with a standard deviation of 122 µs.

C. IPS Performance During SV Injection Attacks

To assess the performance of the proposed IPS during an SV injection attack, cases C0, C1, C2, and C3 are considered. As described in Section III.D, the malicious SV frames were injected in all cases from a workstation directly connected to the network switch, while spoofing the MAC address of the legitimate SV publisher, and replaying normal and fault operating conditions in cases C0 and C2, and C1 and C3, respectively.

TABLE II
MEASURED FAULT CLEARANCE DELAY IN DIFFERENT TESTBED CONFIGURATIONS.

Case	Raspberry Pi	Fault clearance delay (mean, std.) [ms]
N1	Not deployed	(55.559, 2.635)
N4	Deployed*	(55.655, 2.675)
N3	Deployed	(55.807, 2.624)

Both cases C0 and C1 confirm that when the IED receives two conflicting SV streams, which are the legitimate and malicious ones, it stops operating. While this behavior does not have any impact on the power system during normal operation, when a fault occurs during the SV injection attack, the IED fails to clear the fault. This lack of fault clearance can have a disruptive impact on the power system, as discussed in Section IV.D.

On the contrary, in cases C2 and C3, thanks to the IPS, the IED consistently receives only the legitimate frames, and thus does not get blocked by the SV injection attack. In fact, as shown in Table III, the IPS provides near-zero FPRs and close to 100% TPRs and F1-scores in all cases and simulation runs. Moreover, even though a small fraction of legitimate frames gets misclassified as malicious and dropped, and a small fraction of malicious frame is forwarded to the IED, it does not cause any mis-operation. This is because the attacker needs multiple consecutive malicious frames to reach the IED to cause a protection scheme malfunction. Finally, to highlight the informative advantage exploited by the IPS, in Figure 3, the differences between the arrival time shifts of the legitimate and malicious SV frames for case C2 are depicted. Similar results are obtained for case C3.

Given that in cases C2 and C3 the IPS was receiving twice the number of SV frames received in cases N2 and N3, this

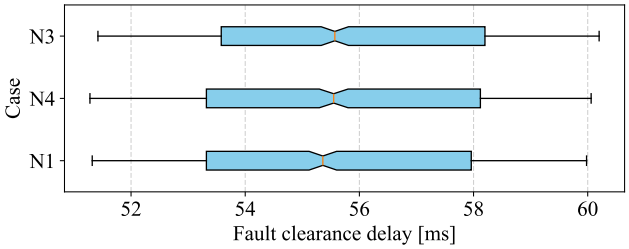

Figure 2: Measured fault clearance delay in different testbed configurations.

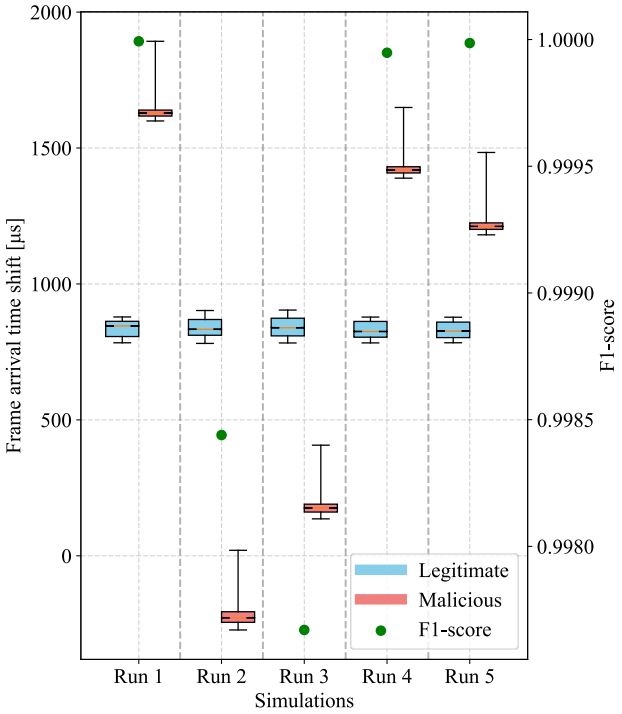

Figure 3: Legitimate and malicious SV frames arrival time shift during the different simulation runs of case C2.

TABLE III
C2 AND C3 MEAN, MIN, AND MAX FPR, TPR, F1-SCORE.

	C2			C3		
	Min.	**Avg.**	**Max.**	**Min.**	**Avg.**	**Max.**
FPR [%]	0.0	0.03	0.10	0.0	0.03	0.12
TPR [%]	99.74	99.91	99.99	99.63	99.90	99.99
F1-score [%]	99.77	99.92	99.99	99.68	99.92	99.99

caused a slight increase in SV frames processing delay. During SV injection attacks, the measured processing time for each frame is 105 µs, with a standard deviation of 178 µs. According to prior studies that validated the detection time of related IDSs for IEC 61850 SV under comparable real-time testing conditions, the SV frame processing delay observed in this setup remains significantly lower than the reported SV injection detection times of 480, 1200, 2800 µs in [6, 7, 8], respectively.

D. Impact on Protection and Control Functionalities

Figure 4 illustrates the power system behavior during one of the simulation runs for cases N1, C1, and C3. The three-phase-to-ground fault is simulated at a simulation time of 0.10 s, and it is cleared by the opening of the three-phase circuit breaker due to the tripping of the overcurrent protection relay. Since the overcurrent relay serves as the primary protection implemented in this test system, its operation occurs within the range of tens of milliseconds. As shown in Figure 4, for both cases N1 and C3, the protection relay instructs the circuit breaker to open phases A and C after 56 and 60 ms, respectively, and consequently to open phase B. After 102 and 106 ms, the fault is cleared for both cases N1 and C3, respectively, regardless of the SV injection attack being performed in case C3. Thus, due to the implemented IPS, the SV injection attack is successfully mitigated. Moreover, as presented in Table II, the additional latency introduced by the inclusion of the IPS on the RPi is 152 µs. This minimal delay is negligible and indicates that the primary protection function remains unaffected, preserving the protection coordination strategy in the digital substation. On the contrary, in case C1, the cyber attack successfully prevents the IED from reacting to the fault and from issuing the required control signal. As a result, the fault must be cleared either by the backup protection within the same substation or by a protection IED from a neighboring substation. This leads to a greater impact on power system stability. The delayed fault-clearing action may cause damage to physical equipment, while intervention by protection IEDs from connected substations could result in additional components being de-energized.

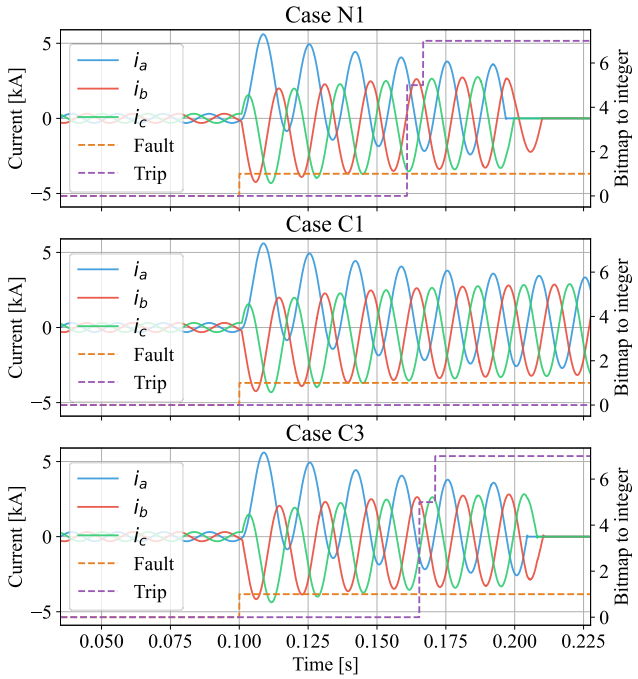

Figure 4: Power system operation during cases N1, C1, and C3. "Trip" corresponds to the bitmap to decimal value conversion of the signal sent by the IED to control the three-phase circuit breaker; each bit corresponds to one phase, with the least significant bit corresponding to phase A. "Fault" identifies when the fault occurs. The instantaneous current values are identified by i_a, i_b, and i_c.

V. CONCLUSIONS

This paper evaluated and experimentally validated the real-time implementation of a hybrid statistical-deep learning IPS designed to mitigate SV injection attacks in IEC 61850-based digital substations. Using a HIL setup comprising RTDS, an RPi, and an industrial IED, the proposed IPS was evaluated under various operating scenarios. The results demonstrated that the IPS achieves high detection accuracy with near-zero false positives, while introducing less than 150 µs of additional latency, thereby not interfering with the real-time requirements of protection and control functions. Furthermore, the system effectively prevented malicious SV frames from impacting substation operation, confirming its practicality and robustness for real-world deployment. Future work will include the deployment of the IPS directly on the industrial IED and validation of the method against more advanced attacks capable of matching the expected SV frames arrival time PDF.

REFERENCES

[1] P. Kozak, I. Klaban, and T. ˇSlajs, "Industroyer cyber-attacks on ukraine's critical infrastructure," in 2023 International Conference on Military Technologies (ICMT). IEEE, 2023, pp. 1–6.

[2] V. S. Rajkumar, A. Ştefanov, J. L. Rueda Torres, and P. Palensky,"Dynamical analysis of power system cascading failures caused by cyber attacks," IEEE Transactions on Industrial Informatics, vol. 20, no. 6, pp. 8807–8817, 2024

[3] N. Cibin, B. Mulder, H. Carstens, P. Palensky, and A. Ştefanov, 'Cyber Attacks Detection, Prevention, and Source Localization in Digital Substation Communication using Hybrid Statistical-Deep Learning', arXiv preprint arXiv:2507. 00522, 2025.

[4] J. Gaspar, T. Cruz, C.-T. Lam, and P. Simões, "Smart substation communications and cybersecurity: A comprehensive survey," IEEE Communications Surveys Tutorials, vol. 25, no. 4, pp. 2456–2493, 2023.

[5] A. Turnbull, 'Tcpreplay', 2025. [Online]. Available: https://tcpreplay.appneta.com

[6] J. Mo and H. Yang, "Sampled value attack detection for busbar differential protection based on a negative selection immune system," Journal of Modern Power Systems and Clean Energy, vol. 11, no. 2, pp. 421–433, 2023.

[7] 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.

[8] J. Hong and C.-C. Liu, "Intelligent electronic devices with collaborative intrusion detection systems," IEEE Transactions on Smart Grid, vol. 10, no. 1, pp. 271–281, Jan. 2019.