Virtual PLC in industrial edge platform: performance evaluation of supervision and control communication

Massimiliano Gaffurini, *Member, IEEE*, Paolo Bellagente, *Member, IEEE,* Alessandro. Depari, *Member, IEEE*,

Alessandra Flammini, *Fellow, Member, IEEE*, Emiliano Sisinni, *Member, IEEE*, Paolo Ferrari, *Member, IEEE*

Abstract*— Edge computing allows for data processing at reduced latency since the computational power is moved close to the data sources. Traditionally, edge computing has been often used in industrial scenarios for implementing gateways between the OT (operational technology) worlds and the IT (cloud) world. Recently, big manufacturers of industrial PLC (programmable logic controller) started promoting the use of containerized virtual PLC hosted inside edge computing platforms. They foresee an innovative integration of container based applications, including automation control, with all the data centric services and application already available for edge ecosystems. Even if a clear advantage from the scalability and maintainability could be expected, would virtual PLCs meet the stringent requirements of industrial automation? This paper is part of a multistage research work, and, as a first step, it is focused on the evaluation of the performance of virtual PLC when exchanging data with other machines, controllers, supervisors, and data acquisition systems in a machine-to-machine scenario. After a brief overview of the involved technology, the design of a methodology for comparing real PLC and virtual PLC is described. Then, performance metrics, and an experimental setup for the evaluation of existing devices are defined taking care of the sources of uncertainty. The effectiveness of the proposed methodology is demonstrated considering a real use case. Through the use of the suggested methodology, important insights of the use case are revealed: for instance, the considered virtual PLC could work as fast as a real PLC with minimum communication latency in the order of 3 ms but, currently, there is a random delay with an average of 50ms whose source has been identified to be the IP stack implementation of the virtual PLC. Finally, the proposed methodology allows for the creation and the validation of analytical models of the use case.

***Index Terms*—PLC, C2C, SCADA, container-based virtualization, M2M.**

I. Introduction

The edge computing paradigm is rapidly evolving and it has been adopted in many scenarios, since edge computing can reduce latency compared to cloud computing [1].

This paragraph of the first footnote will contain the date on which you submitted your paper for review, which is populated by IEEE.
This study is within the MICS (Made in Italy – Circular and Sustainable) Extended Partnership and received funding from Next-GenerationEU (Italian PNRR – M4 C2, Invest 1.3 – D.D. 1551.11-10-2022, PE00000004).
(Corresponding author: Massimiliano Gaffurini)
Massimiliano Gaffurini, Paolo Bellagente, Member, IEEE, Alessandro. Depari, Member, IEEE, Alessandra Flammini, Senior Member, IEEE, Emiliano Sisinni, Member, IEEE, Paolo Ferrari, Member, IEEE are with the Dept. of Information Engineering, University of Brescia, 25123 Brescia, Italy. (e-mail: {massimiliano.gaffurini; paolo.bellagente; alessandro.depari; alessandra.flammini; emiliano.sisinni; paolo.ferrari}@unibs.it).

The industrial automation is now following this trend. The first version of virtualized Programmable Logic Controller (PLC) are appearing [2]. Traditional PLCs have custom firmware running on proprietary hardware, with the aim of ensuring real-time and availability. More recently SoftPLCs have control software running on standard PC hardware and RTOS (Real-Time Operating System), with the aim to reduce cost, assure portability, and provide multiple vendor sources. The newest approach proposes virtual PLCs that are the containerized version of PLC firmware: they can be executed on any platform that support containers, assuring easy maintainable, lightly-virtualized, solutions with full independency from both the hardware and the operating system. Moreover, the container based automation approach allows for microservice architectures, enabling new features like scalability, observability, traceability, accountability. In other words, the industrial machine can (independently from hardware) run exactly the required/licensed/verified services needed to produce the desired product together with its up-to-date/certified metadata (necessary for accounting the service). Maintenance and update of applications is centrally managed assuring the security and integrity of the whole system [3],[4],[5].

All the previously listed advantages are clear to machine builders that currently use traditional PLC, but an underlined question remains: what is the performance of virtual PLCs compared to real PLCs? As a matter of facts, the automation experts from operation technology (OT) field are obsessed by real-time constraints and they perfectly know that a new fancy controller that fails control deadlines would result in a useless solution (i.e. usually industrial applications cannot tolerate jitter and high latency [6]).

A. Objectives

Considering the described situation, the goal of the project this work belongs to is to provide: a methodology, an experimental procedure, and a set of metrics to evaluate performance of the communication and data exchange of PLC and virtual PLC.

Since the PLCs are placed at the center of the automation stack (also known as CIM automation pyramid [7]), they have two types of data exchange: i) they are connected to other machines, Supervisory Control and Data Acquisition Systems (SCADA) [8], and controllers for the supervision and coordination of the production line; ii) and they are connected

to sensors and actuators in order to perform their own control actions. Both aspects have been well investigated in literature [9] with respect to traditional PLCs, but a general lack of research works on containerized virtual PLCs has been noted.

In order to clearly present and discuss results, the project has been organized with two parts, and this first paper will deal only with the evaluation of the Machine-to-Machine (M2M) data exchange between PLCs or between PLC and SCADA.

The main contributions of this paper are:

- The definition of a methodology to compare performance of PLC and virtual PLC from the point of view of the flow of data between machines (or supervisors).
- The definition of an experimental setup with associate experimental procedure.
- The definition of metrics to compare performance.
- The application of the proposed methodology to a real industrial use case, demonstrating its usefulness for modeling the system, drawing conclusions, and suggesting improvements to real-time behavior.

In the following, after the overview of the involved technologies and of the existing literature, the proposed methodology is introduced and the use case is discussed. Finally, conclusions are reported.

II. Overview of Technology

It is important to point out the context of this work and the involved technologies that are used in the rest of the paper.

A. Classical SCADA and PLCs-based architecture

A classic industrial system based on SCADA and PLC devices is shown in Fig. 1 (a), it combines software and hardware components to supervise, coordinate, and control industrial processes.

The SCADA system serves as a centralized control system that collects, monitors, and analyzes data from multiple remote locations within the industrial environment. It consists of a supervisory computer, human-machine interface (HMI), and communication infrastructure. The PLC, on the other hand, is a specialized computer-based controller that performs strict real-time control functions within the industrial processes; it talks with field devices, sensors and actuators. SCADA retrieves soft real-time data from the PLCs, providing centralized view of the whole production line.

Communication between SCADA and PLCs relies on M2M or on C2C (Controller to Controller [10]) protocols for sending commands and configuration parameters (to PLCs), and production related information (to the SCADA).

B. Virtualization-based architecture

As described in [11], virtualization and containerization systems are speeding up the digital transformation of manufacturing. The rapid growth of virtualization technologies has opened new possibilities for industrial applications. Real devices often require specialized and costly hardware, making them less flexible and scalable. In contrast, virtual devices can leverage general purpose hardware, which is more affordable, easily scalable and there is also an environmental aspect [12].

With the advent of virtualization techniques in the industry, the traditional architecture shown in Fig. 1 (a) is still valid at the topological and communication level, but the single components implementation can be different.

In Fig. 1 (b), it is shown an example of an architecture based on virtual environments. The components are the same of the traditional approach but there are virtual environments where PLCs, sensors and SCADA can be virtualized. The protocols for M2M communication and sensor communication remain the same; they can be implemented directly in the virtual SCADA, in the virtual PLC, and also separately (as a microservice).

However, the adoption of virtual devices necessitates a thorough evaluation and comparison to determine their suitability for specific industrial use cases (e.g. real-time constraints). While real devices are bare-metal, so the performance are related to the hardware characteristics, the performance of virtual devices depends on many aspects such as: i) virtualization technique, ii) operating system, iii) hypervisor.

It's possible to define several virtualization techniques, the main are:

- Full virtualization: this technique provides a high level of isolation and allows running multiple operating systems simultaneously on a single physical machine.
- Containerization (Light Virtualization): containers offer a lightweight form of virtualization where the host operating system kernel is shared among multiple containers [13].

Each technique offers different levels of isolation, resource allocation, deployment systems and flexibility. For this reason, it is necessary to pay special attention to the implementation of the virtual device.

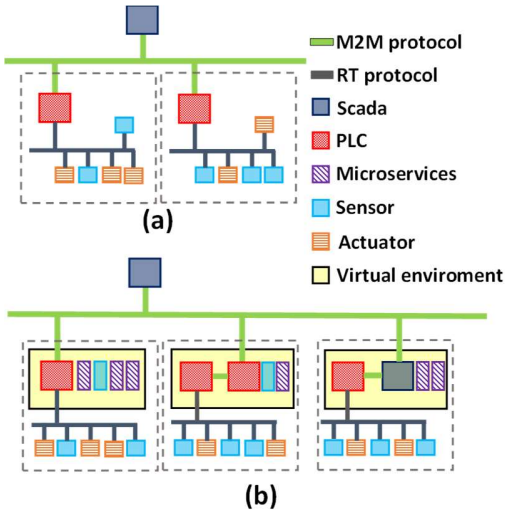

Fig. 1. (a) Classical automation architecture based on PLC and SCADA devices. (b) Virtualization applied to automation: old hardware devices are mapped to software services running inside virtual environments.

III. Related Works and Research Objectives

In literature several works on the evaluation of custom virtual

PLCs and microservices-based architectures can be found. Software-defined Automation solutions are analyzed by D. Javier Perez et al. in [14], where they compared virtualized SoftPLC to a SoftPLC without hypervisor concluding that the virtual PLC can deliver similar performance in terms of switching time while having an increased period jitter.

J. Mellado et al in [15] proposed a containerized IoT-PLC (not fully IEC 61131 compliant) running in a Raspberry Pi 4B board, they evaluated a four tanks control system scenario with a wireless communication systems, obtaining latencies suitable for control applications if process variables change slowly.

T. Cruz et al in [16] proposed a virtual PLC (vPLC) that presents a convergent approach by virtualizing and co-hosting isolated PLC devices on the same physical equipment. This convergence consolidates distributed I/O on a networked I/O fabric, resembling the integration seen in datacenter architectures. Evaluation results indicate the feasibility of vPLC from a systems virtualization perspective, especially on ×86 platforms with room for improvement.

W. Dai et al in [17] designed an orchestration method and deployment procedures, IEC 61499 compliant, based on microservice for industrial edge applications. A combined cloud and edge approach is described in [18].

M. Sollfrank et al [19] evaluated lightweight virtualization system for distributed and time-sensitive applications in industrial automation; they conclude that Docker containers can meet soft-real-time constraints of automation applications.

L. Catuogno et al. in [20] proposed a methodology for the measurements of the computational resources used by a specific container.

However, differently to the works discussed above, the goal of this research work is not to (propose and then) evaluate "custom" container based automation architectures. On the contrary, it is to design a methodology for the evaluation of existing architectures, with special focus on the data exchange performance of commercially available solutions.

L. Rosa et al in [21] developed a framework comprising a basic virtual PLC running in a Docker container, equipped with an OPC-UA middleware for IT and OT communications. For OT communications, a custom TSN-based OPC-UA configuration was utilised. The framework was evaluated on a practical testbed, which consisted of two edge nodes and an industrial network switch. The researchers concluded that the test environment demonstrated that the framework has low overhead, enhances determinism, and still maintains all of the benefits of virtualization.

In details, the scope of this first paper is to propose a method for evaluating communication performance at supervision level (M2M, C2C, SCADA). Operatively, this paper includes also the discussion of a use case, where the M2M data exchange between virtual PLCs (virtualized with different light virtualization techniques) will be compared with the reference performance of their "equivalent" real PLCs.

IV. The Proposed Methodology

The proposed setup for testing is illustrated in Fig. 2 (a). Inside the architecture under test, the first step is to identify the two partners, referred to as C (Client PLC) and E (Edge PLC), that connect to each other using the M2M (or C2C or SCADA) protocol. The second step is to identify the physical network they use to exchange data. As a matter of fact, in order to assess the network latencies, as well as the communication stack delays of C and E, a physical network access, called T, is needed. By means of T, all the relevant data packets can be captured and analyzed.

The third step is to isolate, in the M2M communication between C and E, the transaction type to be evaluated. For instance, in Fig. 2 (b), it is shown the case of a transaction of type "Request and Response". This case is very common in many supervision protocols. The method requires that for the two partners (E and C), and for T, a timestamp is taken and permanently saved for every event related to the identified transaction.

The last step of the method is to assure that transactions are not time correlated. For this reason, a suitable randomization of the Request must be introduced.

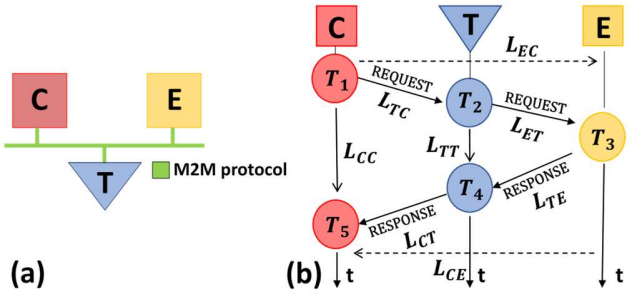

Fig. 2. Proposed methodology: (a) measurement setup, (b) exchange data diagram.

A. The metrics

The metrics of the proposed methodology are defined, without lack of generality, in the case of transaction of type "Request and Response". As a matter of fact, the other type of transaction is the "Publish" where one partner emits a message without being asked for. In practical systems "Publish with Acknowledge" and "Publish without Acknowledge" styles are possible, and the approach proposed here can deal also with them, as described after the metrics definitions.

The interaction between C and E is shown in Fig. 2 (b): i) at time T1 is generated the Request; ii) at time T2 the Request is seen on the network via T; iii) at time T3 the Request is read from E; iv) at time T4 the Response is visible on the network; v) at time T5 the Response is read by C.

The following latencies are defined and evaluated:

- $L_{CC} = T_5 - T_1$, the Request-Response round trip time;
- $L_{TT} = T_4 - T_2$, latency introduced by the elaboration of the request from the communication stack of E and the subsequent step of sending the Response;
- $L_{EC} = T_3 - T_1$, the latency from the generation of the Request to the reception of the Request;

- $L_{CE} = T_5 - T_3$, the latency from the generation of the Response to its reading;
- $L_{TC} = T_2 - T_1$, the Request traverse latency from C to the bus;
- $L_{ET} = T_3 - T_2$, the Request traverse latency from the network to E, including the communication stack of E;
- $L_{TE} = T_4 - T_3$, the Response traverse latency from E to the network;
- $L_{CT} = T_5 - T_4$, the Response traverse latency from E to the network, including the communication stack of C.

For systems that use "Publish with Acknowledge" transaction the metrics are the same since the publish message coincides with the request and the acknowledge message is equivalent to the response.

For systems that use "Publish without Acknowledge" transaction the subset of metrics valid for C to E directions is (L_{EC}, L_{TC}, L_{ET}), while the subset of metrics valid for C to E direction is (L_{CE}, L_{TE}, L_{CT}). L_{CC} and L_{TT} do not apply.

B. Synchronization

The proposed setup of Fig. 2 is a distributed measurement system. The measurement of traverse latency is affected by the drift and the offset between the time references of the devices that take the source and destination timestamps. In this work, as proposed in [21], all the devices must be time synchronized to compensate for the effect of drift and offset. For the synchronization a specialized time transfer protocols called Network Time Protocol (NTP) is used. The NTP synchronization protocol is based on exchanging packets between clients and servers, through the determination of: i) offset of the client's local clock respect to the server's clock, ii) latency of the network connection. Observing the clock offset, the client can correct its local clock to match the server's time.

Still referring to [21] it's possible to evaluate the synchronization standard uncertainty as follows in (1):

$$u_{sm} = \sqrt{\mu_{sm}^2 + \sigma_{sm}^2} \tag{1}$$

Where σ_{sm} represents the standard deviation of the device *m* that takes the timestamp, and μ_{sm} is the systematic error that it is necessary to consider because in the experimental setup no calibration is performed.

The standard uncertainty u_{mn} of any latencies evaluated between two devices (*m* and *n*), introduced in Section IV. , is calculated as in (2):

$$u_{\mathrm{mn}} = \sqrt{u_{\mathrm{sm}}^2 + u_{\mathrm{sn}}^2} \tag{2}$$

When the evaluated latency is calculated between the same device (2) becomes equal to $u_{mm} = \sqrt{2u_{sm}^2}$.

V. The Use Case

The goal of the use case is to demonstrate the effectiveness of the proposed methodology. Currently, most virtual PLC solutions available on the market are built upon open-source Soft PLC IEC61131-3 compliant platforms and are executed on vendor-dependent Automation Platforms and/or Hardware. For example: i) PLCNext by Phoenix Contact, featuring the PLCNext Control PLC based on the Linux kernel [23]; ii) ctrlX by Bosch Rexroth, offering a PLC App that supports target platforms based on ARM64 or x64, and Linux Ubuntu Core with real-time extension (called ctrlX OS [24]); iii) Software-defined Automation solutions, previously introduced in [14].

The aforementioned solutions do not allow for a direct comparison between virtual PLCs and their real counterparts; thus, they are not the best choice for evaluating the proposed methodology.

A. The industrial system used in the use case

In this use case, the Siemens virtual PLC CPU1582v is specifically addressed as an "equivalent" to real PLCs of the S7-1500 product family. This virtual PLC runs within a Docker container on the Siemens Industrial Edge (SIE) Platform. The main components of the SIE platform are as follows: i) the Industrial Edge Hub (IEH), located in the Cloud, which serves as a repository for documentation and containerized applications available on the marketplace; ii) the Industrial Edge Management (IEM), which runs locally and oversees the configuration and setup of edge devices and applications; iii) The Industrial Edge Devices (IEDs), which refer to the actual machines running the containerized applications.

B. Experimental setup for the use case

In this use case, the PLC models S7-1512C-1 PN (v2.6) and CPU1582v (v0.30) assume the role of C and E in to Fig. 2. Three types of devices have been used/implemented:

- PLC S7-1512C-1 PN (called R) is a compact and powerful PLC from Siemens SIMATIC S7-1500 series. It features a fast CPU, expandable I/O modules, 250 Kbyte program memory, 1 Mbyte data memory and support for real-time protocols (e.g. Profibus, Profinet, and Ethernet/IP).
- CPU1582V hosted by Simatic IPC227E (called V_{IPC}, shown in Fig. 3 (b)), a compact industrial PC boasting a quadcore Intel Celeron N2930 processor running at 1.83 GHz (burst frequency 2.16 GHz), 8GB of main memory and a 240 GB SATA SSD. It executes the "IED OS" (version ied-os-1.9.0-27-amd64), which includes the Mentor Industrial OS (based on Debian Linux real-time) and the additional middleware for containerization.
- CPU1582V hosted by desktop PC (called V_{PC}, shown in Fig. 3 (a)), boasting a CPU Intel Core i7-7700b running at 3.60GHz, 16GB of main memory and 500 GB SATA SSD and Windows 10 pro as Host OS. VMware Workstation pro 17.1 hypervisor is installed on the SSD it executes the "IED OS" (version 1.9.0-5-a-rc2). At the virtual machine are assigned 4GB of memory, 2 processors and two network bridge adapters (one for the connection with supervision IEM and the other for connection to the field level).

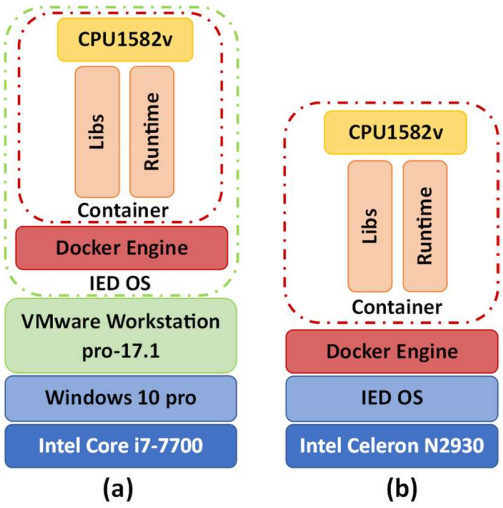

Fig. 3. Architectures of different virtualization solutions for CPU1582v : (a) commercial PC, (b) IPC227E.

For the use case experiment, two devices of each type are used, to evaluate all the combinations in TABLE I. At each experiment is assigned a reference number # used in the following.

TABLE I
EXPERIMENT MATRIX FOR THE COMPARISON (R: REAL PLC, V_{IPC}: VPLC ON IPC, V_{PC}: VPLC ON DESKTOP PC)

Ref. #	*Client (C)*	*Edge (E)*
1	R	R
2	R	V_{IPC}
3	R	V_{PC}
4	V_{IPC}	R
5	V_{IPC}	V_{IPC}
6	V_{IPC}	V_{PC}
7	V_{PC}	R
8	V_{PC}	V_{IPC}
9	V_{PC}	V_{PC}

Referring again to Fig. 2 the network of the use case is Ethernet, and the network access is obtained using a Ethernet Tap (Profitap C1AP-100). The T duplicates all Ethernet traffic on the link and forwards it to an embedded system (Siemens IOT2050), which assigns timestamps and store each packets.

C. The M2M protocol used in the use case (S7comm)

To enable seamless communication between PLCs and supervisory systems, specific communication protocols are often employed. Siemens developed S7comm, that is the primary communication protocol for M2M, C2C, and SCADA. It is used by Siemens S7-300, S7-400, S7-1200 and S7-1500 families and external devices [25]. The protocol runs on ISO transport services on top of the TCP (TPKT) and all the communications occur on the port 102. S7comm data are encapsulated in a COTP (Connection oriented Transport Protocol) packets. The protocol incorporates security mechanisms such as authentication, integrity checks, and confidentiality using encryption algorithms. However, S7comm has faced vulnerabilities and attacks ([26][27][28]).

There are 3 steps to establish a S7 connection with the PLCs [29]: i) establish a COTP connection by sending a request and receiving the corresponding ACK, ii) S7 communication setup and iii) exchange of S7 function code related to the transaction. In Fig. 4 it's shown an example of GET_DB instruction, it gets data from the desired data area of the server PLC and assigns them to the data area in the client PLC.

The S7 request is handled by the PLC operating system and it allows access to the specified memory area without disturbing the normal behavior of the PLC program.

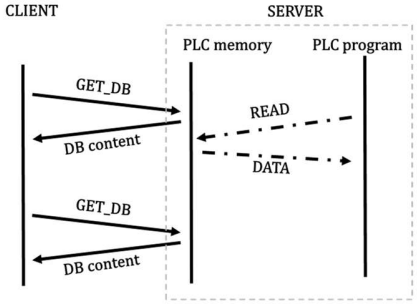

Fig. 4. Sequence diagram of S7comm GET Request.

D. Synchronization for the use case

In the proposed measurement setup, the connection with NTP Server (ntp1.inrim.it) is implemented as follows: i) R via IOT2040; ii) V_{IPC} and V_{PC} via settings in SIE device management; iii) Tap analyzer via IOT2050.

The performance statistics are taken through the NTP daemon, as shown in TABLE II:

TABLE II
NTP STATISTICS

	Poll time (s)	$Average\ Offset\ (ms)$	$Std.dev.\ Offset\ (ms)$
R	32	0.001	0.02
V_{IPC}	256	-0.258	0.52
V_{PC}	1024	-0.051	0.14
Tap	32	-0.016	0.56

From the values reported in TABLE II, the synchronization standard uncertainty is evaluated following Section IV.

The resulting standard uncertainty and the corresponding expanded uncertainty, U, for this use case, calculated with a coverage factor of $k = 2$ are reported in TABLE IV. The obtained resolution ranges from 0.1 ms to 1.7 ms, depending on the experiment. These values align with round-trip time latencies for industrial applications, as reported in [30].

In this paper, the timestamping uncertainty, u_{tm}, is not taken into account because, with the considered hardware, it has a negligible impact (Note that in previous work [21] the timestamping uncertainty has been estimated on the order of 0.01 ms).

E. Data validation and pre-processing

In this use case, the data exchange is implemented using S7

communication between C and E. The transaction type is "Request and Response"; specifically, C generates a GET Request to E for reading a variable (T3) inside one of its internal DataBlock (DB). The transaction is repeated every 10 seconds plus a random time between 0 and 999 ms.

In TABLE III statistics relative to S7 protocol are listed for all evaluated experiments. In details, TABLE III shows that in the different experiments the rates of S7 packets are similar, while the number of the total packets may change between the experiments.

TABLE III
S7 PROTOCOL PACKETS ANALYSIS

#	*Total*	*S7*	*S7 (%)*	*S7 Rate (packets/s)*
1	3222	504	15.6	0.2
2	17592	1376	7.8	0.2
3	11037	1480	13.4	0.2
4	116687	980	0.2	0.2
5	30444	708	2.3	0.2
6	15542	1723	11.1	0.2
7	11457	1387	12.1	0.3
8	18797	1959	10.4	0.2
9	19628	1128	5.7	0.2

During the data analysis, it is observed that:

- when the two partner C and E are both a real PLC, the transaction is very fast and the variability is extremely low. Such a situation suggests that the real PLC handles the IP stack (on which the S7 protocol lays) on interrupt.
- when one of the communication partners is a virtual PLC (either V_{IPC} or V_{PC}), a large variability of the latency metrics is found. Such a situation triggered a deeper analysis about the reason of this behavior.

The analysis of the sequences of network packets, captured by the Ethernet Tap T, revealed that the implementation of the containerized CPU1582V (v0.30) seems to have an IP stack with a polling cycle equal to 100 ms ($L_{\text{stack\_e}}$ *in* Fig. 5 (a)).

For sake of clarity, Fig. 5 illustrates a temporal diagram that represents an example of data exchange in experiment #2 (TABLE I). During normal operation (see Fig. 5 (a)), C generates the GET Request at time T1. The request received from E is then processed at the end of the cycle introducing, at every request, a delay equal to the time required to reach the end of the IP stack cycle.

While investigating, another, less frequent, behavior of the IP stack has been observed. As shown in Fig. 5 (b), C creates as usual the S7 GET Request at time T1 but, every 60 s, this request is queued because the IP stack of C is busy sending an ARP request and then waiting until the ARP response arrives. This situation, in addition to the normal data exchange, results in a further delay on the order of hundreds of milliseconds.

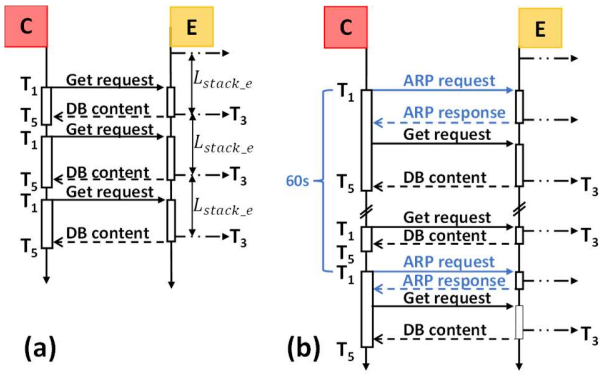

Fig. 5. Example of exhange in experiment #2: (a) normal exchange (b) with ARP activity.

In this paper, for the metrics evaluation, only data pertaining to the normal behavior are considered and the data exchanges involving ARP activity are filtered out with pre-processing. This approach is justified by the sporadic nature of this behavior, and by the possibility to increase (or even eliminate) the ARP request configuring the IP parameters of clients. For sake of completeness not filtered metrics are reported in TABLE VII and TABLE VIII. In Fig. 6 is shown the comparison between L_{CC} distribution with ARP activity Fig. 6 (a) and filtered L_{CC} distribution Fig. 6 (b).

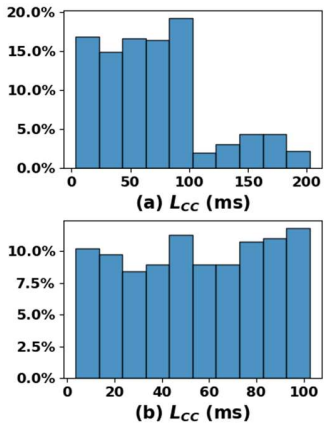

Fig. 6. L_{CC} distribution for experiment #2 : (a) with ARP activity (b) excluding ARP activity.

VI. THE EXPERIMENTAL RESULTS FOR THE USE CASE

The primary goal of the proposed methodology is to provide useful insights of the system under test. For this use case, TABLE V and TABLE VI report the results of the experiments. It is possible to observe the described behaviour in Section V.E (previous subsection): when one of the communication partners is a virtual PLC, the round-trip time L_{cc} is on the order of 100 ms. Two examples of distributions are shown in Fig. 7 (a) and (b).

In Fig. 7 (a) is possible to observe the evaluated latency for the experiment #1 (i.e. only real PLCs) that works as reference. The round-trip time (L_{CC}) has an average value of 3.3 ms.

In Fig. 7 (b) is shown the results of experiment #5. Two

cycles of 100 ms exist in the communication, one for C and one for E. The cycle on E is described by the L_{ET} latency, in particular is characterized by: average value of 51 ms, maximum value of 100 ms and a minimum value lower than U (see TABLE IV). Similarly, the cycle on C is described by L_{CT} latency that is characterized by: an average value of 65 ms, maximum value of 74 ms and a minimum value of 4 ms. These values are confirmed analyzing L_{CC} (which is the sum of L_{CT} and L_{ET}) having an average value of 118 ms. In the latency L_{ET}, thanks to the random time added to the GET Request, all the possible values of the delays are tested and, as shown in Fig. 7 (b), a uniform distribution is obtained. On the contrary, in L_{CT} the situation is different because there is not the possibility to add random time to the response in the S7 protocol. Thus, the response arrives more frequently in the same point of the cycle of C, leading to a distribution where most of the samples are at 70 ms.

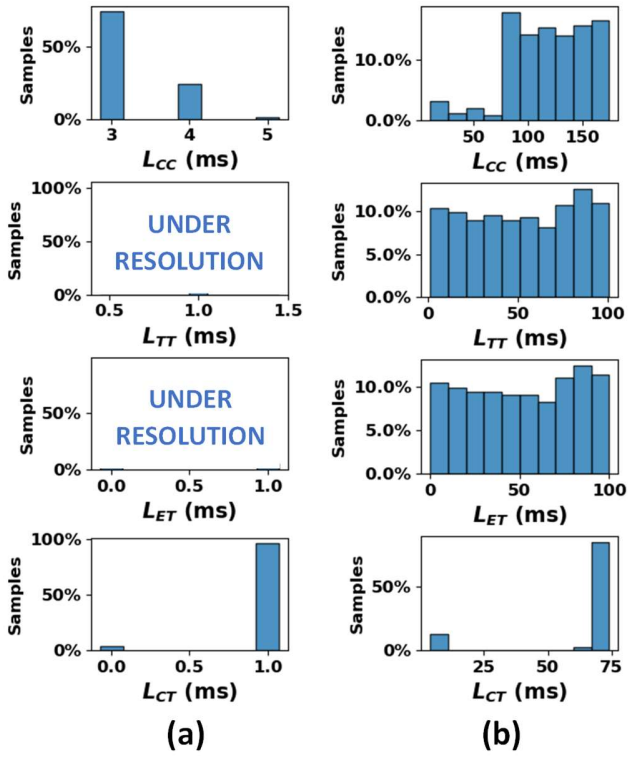

Fig. 7. Distributions of the evaluated latencies: (a) experiment #1, (b) experiment #5.

These cycles are well described in TABLE VI. When C is a V_x (experiments #4,5,6,7,8,9), L_{CT} latency has average values that ranged from 37 to 102 ms; these values may change depending on the synchronization status of C and E when the experiment is started. When E is based on V_x (experiments #2,3,5,6,8,9), L_{ET} latency has average values that range from 47 to 54 ms and are characterized by a uniform distribution.

Again, the validity of the setup is confirmed by observing in TABLE V that the average value of L_{CC} is the sum of the average value of L_{EC} and L_{CE} (considering the expanded uncertainty).

In conclusion, comparing the average performance in this use case, the real PLC appears to be faster than the virtual PLC in completing the S7 transaction. However, if the minimum values of the latencies are compared, it is clear that the virtual PLC is as fast as the real PLC; the additional delay depends only on the implementation of the IP stack. Moreover, no noticeable differences are visible between the two types of virtual environments used for the virtual PLC. Hence, for this specific use case, the main suggestion for boosting performance is to ask developers to focus in enhancing the IP stack implementation of the virtual PLC.

A. The derived analytical model of the use case

The proposed methodology allows for creating analytical model of the systems under test. For instance, considering the experimental results and referring to Fig. 5 , an analytical model for the evaluation of the round-trip time L_{cc} can be derived:

$$L_{cc,max} = L_{elab\_c,\max} + L_{elab\_e,\max} + L_{cable\_req,\max} + L_{\text{cable\_res,max}} + \max(L_{\text{stack\_c}}) + \max(L_{\text{stack\_e}}) \quad (3)$$

where L_{elab_c} and L_{elab_e} represent the elaboration time of the C and E, L_{cable_req} and L_{cable_res} the request/response transmission time on the cable, L_{stack_c} and L_{stack_e} the time taken by the IP stack of PLCs to read the request/response. The (3) is the worst case scenario.

Considering PLCs elaboration time equal for C and E, we can accumulate it into a single variable L_{elab}. The same can be done for the transmission time on cable L_{cable}.

Since communication in real PLCs works with interrupts, for the real PLCs is considered L_{stack}= 0 ms. Due to these considerations it's possible to simplify (3) in (4):

$$L_{cc,max} = 2\,L_{\text{elab,max}} + 2\,T_{\text{cable,max}} + j\,max(T_{\text{stack}}) \quad (4)$$

where j is equal to the number of virtual PLCs involved in the scenario to be modeled.

A model of the average latency can be also obtained. As previously described if the GET Request is not correlated to the response of the previous transaction, a uniform distribution is obtained for the latency. As it is possible to see in TABLE VI, the average time spent in the stack is equal the half of L_{stack} . Thus the model can be written as:

$$L_{cc,avg} = 2\,L_{\text{elab,avg}} + 2\,L_{\text{cable,avg}} + j\,\frac{L_{\text{stack}}}{2} \quad (5)$$

From (5) it is clear that with two real PLCs involved in the communication ($j = 0$), the round-trip time is due only to the elaboration time of PLCs and the propagation time on the cables. On the other hand, with two virtual PLCs ($j = 2$), L_{cc} is mainly given by the latency introduced by the IP stack.

TABLE IV

STANDARD UNCERTAINTY AND EXPANDED UNCERTAINTY FOR EVERY METRIC AND EXPERIMENT.

	L_{CC} (ms)		L_{TT} (ms)		L_{TC} (ms)		L_{ET} (ms)		L_{TE} (ms)		L_{CT} (ms)		L_{EC} (ms)		L_{CE} (ms)	
#	u_c	U	u_c	U	u_c	U	u_c	U	u_c	U	u_c	U	u_c	U	u_c	U
1	0.03	0.1	0.79	1.6	0.64	1.3	0.64	1.3	0.64	1.3	0.64	1.3	0.03	0.1	0.03	0.1
2	0.03	0.1	0.79	1.6	0.64	1.3	0.87	1.7	0.87	1.7	0.64	1.3	0.59	1.2	0.59	1.2
3	0.03	0.1	0.79	1.6	0.64	1.3	0.66	1.3	0.66	1.3	0.64	1.3	0.14	0.3	0.14	0.3
4	0.74	1.5	0.79	1.6	0.87	1.7	0.64	1.3	0.64	1.3	0.87	1.7	0.59	1.2	0.59	1.2
5	0.74	1.5	0.79	1.6	0.87	1.7	0.87	1.7	0.87	1.7	0.87	1.7	0.74	1.5	0.74	1.5
6	0.74	1.5	0.79	1.6	0.87	1.7	0.66	1.3	0.66	1.3	0.87	1.7	0.60	1.2	0.60	1.2
7	0.20	0.4	0.79	1.6	0.66	1.3	0.64	1.3	0.64	1.3	0.66	1.3	0.14	0.3	0.14	0.3
8	0.20	0.4	0.79	1.6	0.66	1.3	0.87	1.7	0.87	1.7	0.66	1.3	0.60	1.2	0.60	1.2
9	0.20	0.4	0.79	1.6	0.66	1.3	0.66	1.3	0.66	1.3	0.66	1.3	0.20	0.4	0.20	0.4

TABLE V

ROUND-TRIP LATENCIES, FORWARD AND BACKWARD LATENCIES ('---' MEANS UNDER RESOLUTION).

	L_{CC}(ms)				L_{TT}(ms)				L_{EC}(ms)				$L_{CE}(ms)$			
#	*Avg*	*max*	*min*	*Std*	*Avg*	*max*	*min*	*Std*	*Avg*	*max*	*min*	*Std*	*Avg*	*max*	*min*	*Std*
1	3.3	4.6	2.6	0.3	---	---	---	---	0.9	2.2	0.2	0.4	2.41	3.4	1.3	0.42
2	54.6	102.8	3.3	28.8	52	100	---	29	51	100	---	29	2	4	---	---
3	50.8	103.6	3.4	28.7	48	101	---	29	48	101	1	29	1	3	---	0.5
4	45	102	3	32	---	2	---	---	---	3	---	---	38	97	---	30
5	118	175	10	37	52	101	1	30	51	100	---	30	61	71	---	22
6	134	188	82	29	52	106	1	29	52	107	---	29	82	83	81	29
7	57	114	4	29	---	2	---	---	1	2	---	---	60	116	7	29
8	151	201	95	30	53	101	---	29	54	104	---	30	96	98	95	---
9	137	191	86	29	48	100	---	29	50	106	---	29	95	98	93	---

TABLE VI

TRANSMISSION LATENCIES BETWEEN DEVICES ('---' MEANS UNDER RESOLUTION).

	L_{TC}(ms)				$L_{ET}(ms)$				$L_{TE}(ms)$				$L_{CT}(ms)$			
#	*Avg*	*max*	*min*	*Std*	*Avg*	*max*	*min*	*Std*	*Avg*	*max*	*min*	*Std*	*Avg*	*max*	*min*	*Std*
1	---	2	---	---	---	---	---	---	---	2	---	---	---	---	---	---
2	2	2	1	---	50	98	---	29	3	4	---	---	---	2	---	---
3	2	2	---	---	46	99	---	29	---	2	---	---	---	2	---	---
4	---	3	---	---	---	2	---	---	---	3	---	---	37	97	---	29
5	2	3	---	---	51	100	---	30	---	2	---	---	65	74	4	22
6	---	---	---	---	54	109	2	29	---	2	---	---	79	79	77	---
7	---	---	---	---	---	2	---	---	2	3	---	---	59	116	6	29
8	---	---	---	---	47	99	---	30	---	2	---	---	102	103	101	---
9	---	2	---	---	52	100	6	23	3	4	---	---	91	94	89	---

TABLE VII

ROUND-TRIP LATENCIES, FORWARD AND BACKWARD LATENCIES WITH ARP ACTIVITY ('---' MEANS UNDER RESOLUTION).

	L_{CC}(ms)				L_{TT}(ms)				L_{EC}(ms)				$L_{CE}(ms)$			
#	*Avg*	*max*	*min*	*Std*	*Avg*	*max*	*min*	*Std*	*Avg*	*max*	*min*	*Std*	*Avg*	*max*	*min*	*Std*
1	3.3	4.6	2.6	0.3	---	---	---	---	0.9	2.2	0.2	0.4	2.41	3.4	1.3	0.42
2	70.5	202.4	3.5	46.6	60	100	---	32	68	200	1	47	2	4	---	---
3	67.8	203.4	3.4	47.8	57	101	---	33	64.8	200.8	0.5	47.8	1.5	2.8	0	0.5
4	58	201	3	47	8	101	---	21	---	3	---	---	59	207	0	47
5	118	175	10	37	52	101	1	30	51	100	---	30	61	71	---	22
6	134	188	82	29	52	106	1	29	52	107	---	29	82	83	81	29
7	71.6	208.9	4.0	45.8	9	107	---	21	1	2	0	---	68	207	6	46
8	151	201	95	30	53	101	---	29	54	104	---	30	96	98	95	---
9	137	191	86	29	48	100	---	29	50	106	---	29	95	98	93	---

TABLE VIII
TRANSMISSION LATENCIES BETWEEN DEVICES WITH ARP ACTIVITY ('---' MEANS UNDER RESOLUTION).

	L_{TC}(ms)				$L_{ET}(ms)$				$L_{TE}(ms)$				$L_{CT}(ms)$			
#	*Avg*	*max*	*min*	*Std*	*Avg*	*max*	*min*	*Std*	*Avg*	*max*	*min*	*Std*	*Avg*	*max*	*min*	*Std*
1	---	2	---	---	---	---	---	---	---	2	---	---	---	---	---	---
2	10	101	---	22	57	100	---	32	3	4	---	---	---	2	---	---
3	10	103	---	22	54	99	---	32	---	2	---	---	---	---	---	---
4	---	3	---	---	1	3	---	---	8	102	---	21	43	100	---	33
5	2	3	---	---	51	100	---	30	---	2	---	---	65	74	4	22
6	---	---	---	---	54	109	2	29	---	2	---	---	79	79	77	---
7	---	---	---	---	---	---	---	---	10	108	---	2	66	114	5	32
8	---	---	---	---	47	99	---	30	---	2	---	---	102	103	101	---
9	---	2	---	---	52	100	6	23	3	4	---	---	91	94	89	---

VII. CONCLUSION

The new virtual PLCs can be executed on any platform that support containers, assuring independency from both the hardware and the operating system. They are maintainable, scalable, traceable and open to new concepts of micro service architectures in industrial automation. But what is the performance of virtual PLCs compared to (proven in use) real PLCs?

This paper, the first of a multistage research work, provides a methodology for the evaluation of the communication performance of virtual PLCs when exchanging data for supervision, coordination and control with other machines. For comparing the performance, a set of metrics has been defined corresponding to the round trip time of the transaction, and to the transmission latencies between the devices. The methodology is completed by the proposal of a general experimental setup for measuring the relevant metrics across a distributed measurement environment. The synchronization of the devices under test (and of ancillary devices) is discussed, and their expanded uncertainty is taken into account.

The effectiveness of the proposed has been demonstrated by a use case, where real and virtual PLCs are compared. In details, Siemens virtual PLC (CPU1582V) and real PLCs of the S7-1500 family are used in a scenario where they exchange supervision data by means of S7comm protocol. The virtual PLC is hosted by two different virtual environments, allowing for a comparative assessment of the type of virtualization.

Utilizing the recommended approach, the analysis of the use case results points out that: a virtual PLC could work as fast as a real PLC with average data exchange latencies L_{CC} in the order of 3 ms (and almost identical time distributions), but the IP stack implementation introduces a higher delay up to maximum of 100 ms and an average of 50 ms. This insight information is useful for developer of virtual PLC that can work on reducing such delay.

Last, the full access to the network traffic given by the proposed setup, combined with the fully synchronized timestamping, allows for the creation of analytical model of latencies of the use case under test. The analytical model can be used for simulators or worst case analysis.

In conclusion in this paper a general methodology for the evaluation of the performance of virtual PLC is provided , when exchanging data with other machines and SCADA in a M2M scenario.

The second stage of the ongoing project will involve evaluating the real-time communication performance between the PLC and sensors/actuators. This evaluation will be conducted using advanced time measuring devices, including direct 1-PPS synchronization signal and GPS receivers, with a resolution of around 50 microseconds. However, it is important to note that this may require additional hardware and incur extra costs, resulting in a more expensive setup for real-time measurement.

Future evolution of virtual components (like virtual PLC, but not limited to) can greatly benefit from measurement and test procedures (like the ones described in this paper) of their performance. As a matter of fact, virtual PLC implementations (being a full software approach) can be quickly improved by means of combined cycle of design, test and redesign.

ACKNOWLEDGMENT

The authors would thank Siemens Italy for the support during the implementation of the use case.

REFERENCES

[1] W. Shi, J. Cao, Q. Zhang, Y. Li and L. Xu, "Edge Computing: Vision and Challenges," in IEEE Internet of Things Journal, vol. 3, no. 5, pp. 637-646, Oct. 2016

[2] T. Goldschmidt, S. Hauck-Stattelmann, S. Malakuti, S. Grüner, "Container-based architecture for flexible industrial control applications," Journal of Systems Architecture, Volume 84, 2018, pp. 28-36

[3] D. Taibi, V. Lenarduzzi and C. Pahl, "Processes, Motivations, and Issues for Migrating to Microservices Architectures: An Empirical Investigation," in *IEEE Cloud Computing*, vol. 4, no. 5, pp. 22-32, Sept./Oct. 2017

[4] S. Sarkar, G. Vashi and P. P. Abdulla, "Towards Transforming an Industrial Automation System from Monolithic to Microservices," 2018 IEEE 23rd International Conference on Emerging Technologies and Factory Automation (ETFA), Turin, Italy, 2018, pp. 1256-1259

[5] P. Ferrari et al.,"On the Use of LoRaWAN and Cloud Platforms for Diversification of Mobility-as-a-Service Infrastructure in Smart City Scenarios ", IEEE Trans, Instrumentation and Measurements, 2022

[6] Rui Queiroz, Tiago Cruz, Jérôme Mendes, Pedro Sousa, and Paulo Simões. 2023. Container-based Virtualization for Real-time Industrial Systems—A Systematic Review. ACM Comput. Surv. 56, 3, Article 59 (March 2024), 38 pages

[7] B. Scholten, The Road to Integration: A Guide to Applying the Isa-95 Standard in Manufacturing. ISA, 2007

[8] D. Pliatsios, P. Sarigiannidis, T. Lagkas, A.G. Sarigiannidis, "A Survey on SCADA Systems: Secure Protocols, Incidents, Threats and Tactics", IEEE Communications Surveys and Tutorials, vol. 22, n. 3, art. no. 9066892, 2020, pp. 1942 – 1976

[9] Ayiad M., Maggioli E., Leite H., Martins H., "Communication Requirements for a Hybrid VSC Based HVDC/AC Transmission Networks State Estimation," Energies 2021, 14, 1087

[10] A. Eckhardt, S. Müller, “Analysis of the Round Trip Time of OPC UA and TSN based Peer-to-Peer Communication”, IEEE Int. Conf on Emerging Technologies and Factory Automation, ETFA, Sept., art. no. 8869060, 2019, pp. 161 – 167

[11] T. Borangiu et al. "Digital transformation of manufacturing through cloud services and resource virtualization.", Computers in Industry 108 (2019): 150-162

[12] N. Carvalho et al. "Manufacturing in the fourth industrial revolution: A positive prospect in sustainable manufacturing." Procedia Manufacturing 21 (2018): 671-678

[13] E. Sisinni et al., "Assessment of Time Performance of Lightweight Virtualization for Edge Computing Applications," 2023 IEEE 19th International Conference on Factory Communication Systems (WFCS), Pavia, Italy, 2023, pp. 1-4

[14] D. Javier Perez, J. Waltl, L. Prenzel and S. Steinhorst, "How Real (Time) Are Virtual PLCs?," 2022 IEEE 27th International Conference on Emerging Technologies and Factory Automation (ETFA), Stuttgart, Germany, 2022, pp. 1-8

[15] J. Mellado, F. Núñez, "Design of an IoT-PLC: A containerized programmable logical controller for the industry 4.0,"Journal of Industrial Information Integration,Volume 25, 2022, n. 100250

[16] T. Cruz, P. Simões and E. Monteiro, "Virtualizing Programmable Logic Controllers: Toward a Convergent Approach," in IEEE Embedded Systems Letters, vol. 8, no. 4, pp. 69-72, Dec. 2016

[17] W. Dai, Y. Zhang, L. Kong, J. H. Christensen and D. Huang, "Design of Industrial Edge Applications Based on IEC 61499 Microservices and Containers," in IEEE Transactions on Industrial Informatics, vol. 19, no. 7, pp. 7925-7935, July 2023

[18] C. Pallasch et al., "Edge Powered Industrial Control: Concept for Combining Cloud and Automation Technologies," 2018 IEEE International Conference on Edge Computing (EDGE), San Francisco, CA, 2018, pp. 130-134,

[19] M. Sollfrank, F. Loch, S. Denteneer and B. Vogel-Heuser, "Evaluating Docker for Lightweight Virtualization of Distributed and Time-Sensitive Applications in Industrial Automation," in IEEE Transactions on Industrial Informatics, vol. 17, no. 5, pp. 3566-3576, May 2021

[20] L. Catuogno, C. Galdi and N. Pasquino, "An Effective Methodology for Measuring Software Resource Usage," in IEEE Trans. on Instrumentation and Measurement, vol. 67, no. 10, pp. 2487-2494, Oct. 2018.

[21] Lorenzo Rosa, Andrea Garbugli, Lorenzo Patera, and Luca Foschini. 2023. Supporting vPLC Networking over TSN with Kubernetes in Industry 4.0. In Proceedings of the 1st Workshop on Enhanced Network Techniques and Technologies for the Industrial IoT to Cloud Continuum (IIoT-NETs '23). Association for Computing Machinery, New York, NY, USA, 15–21

[22] P. Ferrari, A. Flammini, E. Sisinni, S. Rinaldi, D. Brandão and M. S. Rocha, "Delay Estimation of Industrial IoT Applications Based on Messaging Protocols," in IEEE Transactions on Instrumentation and Measurement, vol. 67, no. 9, pp. 2188-2199, Sept. 2018

[23] https://www.plcnext-community.net/infocenter/home/

[24] https://docs.automation.boschrexroth.com/doc/1368933374/ctrlx-core-runtime-application-manual/latest/en/

[25] S7comm (wireshark.org) (accessed on 01 July 2023) .

[26] L. Martín-Liras, M. A. Prada, J. J. Fuertes, A. Morán, S. Alonso, M. Domínguez, "Comparative analysis of the security of configuration protocols for industrial control devices, International Journal of Critical Infrastructure Protection, Volume 19, 2017, Pages 4-15

[27] Beresford, Dillon. "Exploiting Siemens Simatic s7 plcs." Black Hat USA 16.2, 2011, pp. 723-733

[28] R. Spenneberg, M. Brüggemann, and H. Schwartke. "Plc-blaster: A worm living solely in the plc.", Black Hat Asia 16, 2016, pp. 1-16

[29] Xiao, Feng, Enhong Chen, and Qiang Xu. "S7commtrace: A high interactive honeypot for industrial control system based on s7 protocol." Information and Communications Security: 19th International Conference, ICICS 2017, Beijing, China, December 6-8, 2017, Proceedings 19. Springer International Publishing, 2018

[30] Figueroa-Lorenzo S., Añorga J., Arrizabalaga S."A role-based access control model in modbus SCADA systems. A centralized model approach," Sensors (Switzerland), vol. 19, art. no. 4455, 2019