

UNIVERSIDAD DE BURGOS

DOCTORAL THESIS

A Framework for Dataset Generation, Anomaly Detection and Cybersecurity in Industrial Cyber-Physical Systems

Juan Vicente Martín Fraile

Doctoral Programme in Industrial, Information, and Civil Engineering and Technologies

Supervisors

Jesús Enrique Sierra García

Nuño Basurto Hornillos

July 2026

A mis padres.

Acknowledgments

I would like to express my sincere gratitude to my thesis supervisors, Jesús Enrique Sierra and Nuño Basurto, for their guidance, motivation, and continuous support throughout this doctoral journey.

I am grateful to the University of Burgos and to all colleagues who have contributed, directly or indirectly, to the development of this research. I would also like to express my appreciation to the head of the department, Álvaro Herrero, for his support, advice, and valuable contributions throughout this work.

Finally, I would like to thank my family and friends for their trust, understanding, and unconditional support throughout these years.

Abstract

Industrial Cyber-Physical Systems (CPS) increasingly rely on data-driven techniques for anomaly detection, predictive maintenance, and cybersecurity monitoring. However, the availability of realistic and reproducible datasets remains a significant challenge, particularly in industrial environments where fault conditions and cyberattacks are difficult to observe and label. This thesis addresses this challenge through the development of a comprehensive framework for the generation of structured datasets in industrial CPS.

The proposed framework is based on a hybrid and flexible experimental testbed that integrates industrial hardware, digital twin technologies, event-based data acquisition, and controlled fault injection mechanisms. A reproducible methodology is introduced in which anomalies are generated through predefined experimental configurations and ground truth is established using a hybrid strategy that combines externally defined fault information with runtime operational context. The resulting event-based representation incorporates process states, GEMMA operational modes, and temporal information through inter-event time intervals (Δt).

Using this framework, the DAYPSCI dataset was generated and organized into controlled experimental batches covering both normal operation and fault-injection scenarios. Preliminary machine learning studies demonstrated the suitability of the generated data for anomaly detection and classification tasks involving sensor and actuator faults. In addition, cybersecurity experiments based on Man-in-the-Middle (MitM) attacks against PROFINET I/O communications showed that process semantics can be manipulated without introducing measurable deviations in conventional communication-level indicators, highlighting the limitations of network-based monitoring approaches.

The results demonstrate that the combination of hybrid CPS testbeds, event-based data acquisition, controlled fault injection, and hybrid ground-truth generation provides a robust and reproducible foundation for anomaly detection and cybersecurity research. The proposed methodology contributes a realistic experimental platform, a systematic dataset-generation process, and a reusable dataset that supports the development and evaluation of data-driven techniques for industrial CPS.

Resumen

Los Sistemas Ciberfísicos Industriales (CPS) dependen cada vez más de técnicas basadas en datos para la detección de anomalías, el mantenimiento predictivo y la monitorización de la ciberseguridad. Sin embargo, la disponibilidad de conjuntos de datos realistas y reproducibles continúa siendo un desafío significativo, especialmente en entornos industriales donde las condiciones de fallo y los ciberataques son difíciles de observar y etiquetar. Esta tesis aborda este problema mediante el desarrollo de un marco integral para la generación de conjuntos de datos estructurados en CPS industriales.

El marco propuesto se basa en un banco de pruebas experimental híbrido y flexible que integra hardware industrial, tecnologías de gemelo digital, adquisición de datos basada en eventos y mecanismos de inyección controlada de fallos. Se propone una metodología reproducible en la que las anomalías se generan mediante configuraciones experimentales predefinidas y el ground truth se establece mediante una estrategia híbrida que combina información de fallos definida externamente con el contexto operativo obtenido durante la ejecución del sistema. La representación resultante, basada en eventos, incorpora estados del proceso, modos operativos GEMMA e información temporal mediante intervalos de tiempo entre eventos (Δt).

Utilizando este marco, se generó el conjunto de datos DAYPSCI y se organizó en lotes experimentales controlados que incluyen tanto escenarios de funcionamiento normal como escenarios de inyección de fallos. Estudios preliminares de aprendizaje automático demostraron la idoneidad de los datos generados para tareas de detección de anomalías y clasificación de fallos en sensores y actuadores. Asimismo, experimentos de ciberseguridad basados en ataques Man-in-the-Middle (MitM) sobre comunicaciones PROFINET I/O demostraron que la semántica del proceso puede manipularse sin introducir desviaciones medibles en indicadores convencionales de comunicación, poniendo de manifiesto las limitaciones de los enfoques de monitorización basados exclusivamente en la red.

Los resultados obtenidos demuestran que la combinación de bancos de pruebas CPS híbridos, adquisición de datos basada en eventos, inyección controlada de fallos y generación híbrida de ground truth proporciona una base sólida y reproducible para la investigación en detección de anomalías y ciberseguridad.

La metodología propuesta aporta una plataforma experimental realista, un proceso sistemático de generación de conjuntos de datos y un conjunto de datos reutilizable que facilita el desarrollo y la evaluación de técnicas basadas en datos para sistemas ciberfísicos industriales.

Table of Contents

Acknowledgments ..... **iii**

Table of Contents ..... **ix**

List of Figures ..... **xiii**

List of Tables ..... **xv**

List of Acronyms ..... **xvii**

1. Introduction **1**

1.1. Context and Motivation 1

1.2. State of the Art 2

1.3. Objectives 9

1.4. Methodology 9

1.5. **Contributions** 16

1.6. Thesis Structure 17

2. Fundamentals of Industrial Cyber-Physical Systems **19**

2.1. Industrial CPS 19

2.2. Discrete Event Systems (DES) 20

2.3. GEMMA Model 23

2.4. PLC Programming 25

2.5. Digital Twins 27

2.6. Event-Based Data and Anomaly Detection 29

2.7. Chapter Summary 31

3. Development of the Experimental Framework **33**

3.1. CPS Testbed Overview 33

3.2. Development of the Fault Injection Tool 35

3.3. Event-Based Dataset Generation 38

3.4. Chapter Summary 39

4. Testbed-Based Validation for Anomaly Detection 41
4.1. Introduction 41
4.2. Use Case Definition 42
4.3. Dataset Description for Validation 43
4.4. Methodology for Anomaly Detection 45
4.5. Results and Discussion 46
4.6. Chapter Summary 50
5. Cybersecurity-Oriented Extension of the CPS Framework 53
5.1. Introduction 53
5.2. Motivation: Cybersecurity in Industrial CPS 54
5.3. Logical MitM Model and Threat Definition 55
5.4. Industrial Use Case Under Attack 57
5.5. Definition of Attack Scenarios 58
5.6. Communication-Level Validation 64
5.7. Results and Discussion 67
5.8. Chapter Summary 68
6. Industrial CPS Dataset Generation 71
6.1. Introduction 71
6.2. Dataset Overview 72
6.3. Data Acquisition 73
6.4. Dataset Structure 74
6.5. Anomaly Scenarios 76
6.6. Ground Truth Labelling 77
6.7. Dataset Organization 79
6.8. Chapter Summary 81
7. Conclusions and Future Work 83
7.1. Conclusions 83
7.2. Future Work 85
References 86

Appendix A. Publications Derived from This Thesis **89**

A.1. Main Publications 89

A.2. Other Journal Publications 91

A.3. Conference Publications 92

A.4. Published Datasets 92

Appendix B. Full Publications **93**

B.1. Paper I (DYNA Journal). 93

Testbed for Evaluation of Machine Learning Techniques in Anomaly Detection in Industrial Cyber Physical Systems 93

B.2. Paper II (Applied Sciences Journal). 105

A Hybrid Digital-Twin-Based Testbed for Real-Time Manipulation of PROFINET I/O: A Practical Man-in-the-Middle Attack Implementation ... 105

B.3. Paper III (CEA Conference). 139

Banco de pruebas para evaluación de técnicas IA en la detección de anomalías en sistemas ciberfísicos industriales..... 139

List of Figures

Figure 1.1. GEMMA-based CPS framework for dataset generation 11
Figure 2.1. Controlled system and DES-based control model interaction. 21
Figure 2.2. Evolution of system states X(k) driven by discrete events E(k). 22
Figure 2.3. Standard GEMMA diagram illustrating operational modes and transitions. 24
Figure 2.4. Event-driven transitions between GEMMA operational modes. 25
Figure 2.5. DES-based control logic implemented in Siemens TIA Portal using S7-Graph. 27
Figure 3.1. Hybrid CPS testbed architecture for dataset generation and cybersecurity experimentation. 33
Figure 3.2. Reference CPS use case: industrial part-marking system with actuators, sensors, presence detectors (B1, B2), and GEMMA-based HMI..... 34
Figure 3.3. Dataset builder interface for batch configuration and fault injection candidate selection. 36
Figure 3.4. Dataset Builder interface for parametric fault definition and batch execution. 37
Figure 3.5. Event-based dataset stored in the SQL database with Δt and ground truth. 38
Figure 4.1. Radar plot comparison of KNN, Random Forest, and SVM performance..... 47
Figure 4.2. Multiclass ROC curves for KNN, Random Forest, and SVM..... 48
Figure 5.1. Communication flow between PCSimu, PLC2, and PLC1 under normal and MitM attack. 55
Figure 5.2. Industrial process workflow highlighting attack entry points at sensor (DI) and actuator (DO) levels. 59
Figure 5.3. DI-bit manipulation (B1 and B2) in PLC 2. 60
Figure 5.4. Correlation of SQL Server, GEMMA and PROFINET data in scenario A1 61
Figure 5.5. Wireshark capture of PROFINET RTC1 traffic in scenario A1 62
Figure 5.6. Correlation of SQL Server, GEMMA and PROFINET data in scenario A2. 63
Figure 5.7. ICMP RTT measurements for the four experimental conditions..... 65
Figure 5.8. PROFINET RTC1 cycle-time distribution for PLC 1-PLC 2 traffic. .. 67

Figure 6.1. Event-based data acquisition pipeline with SQL-based Δt computation..... 74
Figure 6.2. Fault injection and hybrid ground-truth generation process based on GEMMA. 78
Figure 6.3. Directory structure and organization of the DAYPSCI_v1 dataset.. 80

List of Tables

Table 3.1. Hardware components used in the testbed..... 35
Table 3.2. Software components used in the testbed. 35
Table 4.1. Summary of classification results 47
Table 5.1. PROFINET throughput statistics under four conditions. 66
Table 6.1. Summary of dataset variables and their descriptions 76

List of Acronyms

AI	Artificial Intelligence
AUC	Area Under the Curve
CPS	Cyber-Physical System
CSV	Comma-Separated Values
DES	Discrete Event System
DI	Digital Input
DO	Digital Output
DT	Digital Twin
GEMMA	Guide d'Étude des Modes de Marches et d'Arrêts
HMI	Human–Machine Interface
ICMP	Internet Control Message Protocol
ICS	Industrial Control System
IDS	Intrusion Detection System
IEC	International Electrotechnical Commission
IIoT	Industrial Internet of Things
IT	Information Technology
KNN	k-Nearest Neighbours
LD	Ladder Diagram
ML	Machine Learning
MitM	Man-in-the-Middle
MTBF	Mean Time Between Failures
OT	Operational Technology
PCAPNG	Packet Capture Next Generation
PLC	Programmable Logic Controller
PPS	Packets Per Second
RBF	Radial Basis Function
RF	Random Forest
ROC	Receiver Operating Characteristic
RTT	Round-Trip Time
SCL	Structured Control Language
SFC	Sequential Function Chart
SQL	Structured Query Language
SDU	Service Data Unit
SVM	Support Vector Machine
Δt	Time Difference Between Consecutive Events

Chapter 1

Introduction

This chapter introduces the research context and motivation of the thesis, focusing on the challenges associated with dataset generation, anomaly detection, and cybersecurity in industrial Cyber-Physical Systems (CPS). It also presents the state of the art, defines the research objectives, describes the adopted methodology, highlights the main contributions, and outlines the structure of the thesis.

1.1. Context and Motivation

Industrial Cyber-Physical Systems have become a fundamental pillar of modern manufacturing and a key enabler of Industry 4.0 paradigms, such as smart factories, by enabling the seamless integration of physical processes, computational intelligence, and communication networks [1-3]. In this context, CPS are often complemented by Digital Twins, which provide virtual representations of physical systems and support real-time monitoring, simulation, and data-driven analysis of industrial processes [4, 5]. This convergence has led to significant improvements in productivity, flexibility, and automation capabilities in industrial environments [1, 5].

However, this increased level of integration also introduces additional system complexity, making industrial environments more susceptible to faults, anomalous system behaviour, and cybersecurity threats. In particular, the strong coupling between cyber and physical components implies that disruptions in one domain may propagate to the other, potentially affecting the overall system performance and safety [6-8]. In such tightly interconnected systems, anomalies or cyberattacks targeting sensors, actuators, or communication networks can lead to production inefficiencies, degradation of product quality, or even unsafe operating conditions [9, 10]. Furthermore, the growing interconnection of industrial devices significantly expands the attack surface of CPS, reinforcing the need for robust monitoring and advanced anomaly detection mechanisms [6, 11].

In this context, data-driven approaches based on machine learning have emerged as promising solutions for anomaly detection in industrial control systems [12-14]. These techniques enable the identification of complex and previously unseen anomalies by capturing temporal and multivariate relationships present in industrial process data [12, 13]. However, the effectiveness of such approaches strongly depends on the availability of high-quality datasets that accurately represent both normal and anomalous system behaviours [8, 13].

Despite their importance, acquiring such datasets in real industrial environments remains particularly challenging. The limited availability of labelled fault data, the risks associated with injecting faults into industrial operational systems, and the difficulty of obtaining reliable ground truth annotations significantly hinder dataset generation and validation processes [8, 11, 13]. In many cases, ground truth labelling is tightly coupled with the operational process itself, which may introduce inconsistencies and reduce the reliability of the resulting datasets [8, 11]. These limitations ultimately constrain the development, evaluation, and reproducibility of data-driven anomaly detection models in industrial CPS.

Moreover, although artificial intelligence enhances detection capabilities, it also introduces new security concerns, as machine learning models can themselves become targets of adversarial attacks, further increasing the attack surface of these systems [14, 15].

To address these challenges, controlled experimental environments, particularly industrial testbeds, have been proposed as a viable solution for reliable dataset generation [10, 16, 17]. These platforms enable the systematic and reproducible acquisition of data under both nominal and faulty operating conditions. Importantly, they facilitate a clear separation between the operational CPS and the experimental layer used for ground truth generation and labelling. This separation preserves system integrity while ensuring the consistency, reliability, and traceability of the generated datasets, thereby supporting the development of robust and reproducible data-driven approaches for anomaly detection in industrial environments.

1.2. State of the Art

This section reviews the current state of the art related to Cyber-Physical Systems in industrial environments, with a particular focus on dataset generation,

anomaly detection, and cybersecurity. The objective is to identify existing approaches, their limitations, and the open challenges associated with the development of data-driven techniques in industrial CPS.

To this end, the analysis is structured into four main areas. First, the fundamental concepts and architectures of CPS in industrial environments are introduced. Next, the challenges associated with dataset generation are discussed. Then, anomaly detection techniques are reviewed. Finally, cybersecurity aspects are analysed, highlighting the increasing importance of secure and reliable industrial operation.

1.2.1. Cyber-Physical Systems in Industrial Environments

The emergence of Industry 4.0 has driven the integration of physical processes with digital technologies, leading to the development of Cyber-Physical Systems, which integrate computational and physical components to monitor and control real-world processes, constituting a disruptive paradigm across multiple industrial domains [1, 3, 18]. These systems combine sensing, computation, communication, and actuation to enable real-time monitoring and control of industrial processes [6].

In industrial environments, CPS architectures are typically composed of interconnected components such as sensors, actuators, programmable logic controllers (PLCs), industrial communication networks, and supervisory systems [19]. Among these components, PLCs play a central role, as they execute control logic and coordinate the operation of the physical process [19].

The tight integration between the cyber and physical domains improves efficiency, flexibility, and adaptability in industrial systems [1]. However, it also significantly increases system complexity, as CPS must handle heterogeneous data sources, strict real-time constraints, and the interaction between discrete and continuous dynamics [6].

Furthermore, the increasing connectivity introduced by the Industrial Internet of Things (IIoT) has expanded the capabilities of CPS while simultaneously increasing their exposure to faults and cyber threats [11]. Systems that were traditionally isolated are now interconnected with enterprise and external networks, which introduces new vulnerabilities and potential attack vectors [6].

As a result, modern CPS must not only ensure operational efficiency but also support advanced functionalities such as anomaly detection, predictive maintenance, and cybersecurity analysis [11]. These requirements highlight the need for controlled experimental environments capable of capturing system behaviour under both normal and abnormal conditions, motivating the research presented in this thesis.

In addition to data-driven approaches, some works have explored the use of structured industrial methodologies to formally represent system behaviour. In particular, methodologies integrating GEMMA guidelines with the GRAFCET formal modelling language have been proposed to enable digital twin architectures through real-time coupling between PLCs and digital models [20]. This approach provides a structured and unambiguous representation of operational modes, states, and transitions, ensuring a direct correspondence between control logic and its formal description.

While such approaches primarily focus on control architecture and system formalization, they highlight the importance of structured process representation as a foundation for reliable data acquisition and further integration with data-driven techniques. However, the adoption of these methodologies remains limited in the context of anomaly detection and cybersecurity in industrial CPS.

1.2.2. Dataset Generation in Industrial CPS

The development of data-driven techniques in industrial CPS strongly depends on the availability of representative and high-quality datasets. However, dataset generation in industrial environments remains a major challenge.

In contrast to many research domains where large-scale public datasets are widely available, industrial datasets are typically scarce, proprietary, and difficult to access due to confidentiality constraints and safety concerns [8, 21].

Additionally, real industrial systems are not designed for experimentation. Any attempt to generate specific operating conditions or faults may interfere with production, compromise safety, or reduce system availability [8]. Consequently, most available datasets primarily represent normal operating conditions, while anomalous data is limited.

Another critical limitation is the lack of labelled datasets. In industrial environments, anomaly labelling often requires expert knowledge and manual analysis, which is time-consuming and subjective [8, 11]. This issue becomes more severe when dealing with rare or complex anomalies.

Moreover, existing datasets frequently focus on specific types of scenarios, such as isolated faults or network-based attacks, without providing an integrated view of system behaviour. In particular, datasets combining normal operation, induced faults, and cybersecurity-related anomalies within a single framework are extremely limited [21, 22].

To overcome these limitations, several approaches have been proposed. In addition to simulation-based approaches, descriptor-based methodologies have also been proposed to generate synthetic datasets by replicating statistical, structural, and temporal properties extracted from real-world data. A representative example is GeMSyD [23], which models system behaviour through the extraction of descriptors that capture data distributions, patterns, and noise characteristics.

These approaches enable the generation of synthetic datasets that preserve key properties of real systems, making them suitable for data-driven applications. However, they are generally not designed for industrial CPS environments and do not explicitly address the integration with operational control systems or the definition of ground truth through controlled experimental mechanisms.

Simulation-based methods and digital twins allow generating synthetic data under controlled conditions [4, 5]. However, purely simulated data often lacks realism and may not accurately represent the complexity of real industrial processes.

More recently, hybrid approaches combining real industrial hardware with simulation environments have gained attention. These approaches enable controlled experimentation while preserving a high degree of realism [10, 16].

In this context, industrial testbeds have emerged as a key solution to overcome dataset scarcity and enable controlled experimentation. Testbeds provide environments for generating realistic datasets, simulating anomalies and cyberattacks, and validating detection techniques under reproducible conditions.

Recent surveys highlight the growing importance of testbeds in industrial cybersecurity, emphasizing their role in data generation, vulnerability

assessment, and anomaly detection evaluation. These works also classify testbeds according to their implementation, including physical, simulated, virtual, and hybrid approaches, each offering different levels of fidelity and flexibility.

However, despite their relevance, existing testbeds often present important limitations, such as reduced fidelity, limited integration of industrial protocols, and insufficient support for comprehensive dataset generation in complex CPS environments [17].

Despite these advances, challenges remain regarding dataset reliability, reproducibility, and ground truth definition. These limitations highlight the need for systematic methodologies for dataset generation in CPS, which is one of the main contributions of this thesis.

1.2.3. Anomaly Detection in Industrial Systems

Anomaly detection is a key component in industrial CPS, enabling the identification of abnormal behaviour, faults, and performance degradation at early stages.

Traditional approaches rely on rule-based methods and threshold monitoring, where deviations from predefined limits trigger alarms. Although simple, these methods lack flexibility and cannot capture complex relationships between system variables [6, 11].

In recent years, machine learning approaches have become increasingly popular for anomaly detection in industrial systems. Supervised methods, such as Support Vector Machines (SVM), k-Nearest Neighbours (KNN), and Random Forests, have been widely used when labelled data is available [13]. However, their effectiveness strongly depends on the quality and availability of ground truth.

Unsupervised approaches, including clustering techniques and autoencoders, aim to detect anomalies as deviations from normal patterns without requiring labelled anomalous data [13]. These methods are particularly suitable for industrial environments but may struggle to distinguish between different types of anomalies.

A critical challenge in CPS is the temporal nature of industrial processes. System behaviour depends on sequences of events and timing relationships, rather than

on individual observations. As a result, models capable of capturing temporal dependencies, such as recurrent neural networks and sequence-based approaches, have gained increasing attention [12].

Furthermore, industrial CPS involve heterogeneous data sources, including discrete signals, continuous variables, and event-based information. This heterogeneity increases modelling complexity and requires integrated approaches [6].

Despite these advancements, major limitations remain, including dataset scarcity, class imbalance, and the lack of reliable ground truth. These challenges motivate the need for structured approaches to data generation and validation, as proposed in this thesis.

1.2.4. Cybersecurity in Industrial CPS

The increasing connectivity of industrial CPS has significantly expanded their attack surface. Systems that were once isolated are now exposed to external networks, making them increasingly vulnerable to cyberattacks [6]. This expansion of the attack surface is further intensified by the tight coupling between cyber and physical layers in CPS, where attacks targeting communication or control mechanisms can propagate to the physical domain, potentially causing safety-critical failures and large-scale disruptions [24].

Programmable Logic Controllers (PLCs), as central components of industrial control systems, have become primary targets. As core components responsible for controlling physical processes, they present vulnerabilities at both the component level, including control logic, firmware, and memory, and the system level, involving communication protocols and supervisory software. In particular, industrial communication protocols such as PROFINET have been identified as critical attack surfaces due to limited built-in security mechanisms and increasing connectivity in OT networks. Such vulnerabilities are exacerbated by the increasing integration of IT and OT systems, further expanding the attack surface and enabling more sophisticated cyber threats [19, 25-27].

Traditional cybersecurity approaches include network-based intrusion detection systems (IDS) and traffic analysis. However, these methods are often insufficient to detect attacks targeting internal logic or physical processes [11, 25]. In

particular, advanced attacks may mimic normal behaviour or remain dormant for long periods before activation, making them difficult to detect [25].

Another major limitation is the lack of comprehensive datasets that integrate both cyber and physical aspects of industrial systems. Most existing datasets focus either on network traffic or isolated components, without capturing the full CPS behaviour [21].

Recent works have explored the use of testbeds and digital twins to simulate cyberattacks in controlled environments [14, 16, 28]. These approaches enable safe experimentation but still lack standardized methodologies for integrating anomaly detection and cybersecurity analysis.

These limitations highlight the need for unified and systematic frameworks that combine dataset generation, anomaly detection, and cybersecurity within the same CPS environment, enabling comprehensive and reproducible evaluation of detection approaches, as addressed in this thesis.

1.2.5. Summary and Identified Limitations

This section has reviewed CPS architectures, dataset generation challenges, anomaly detection approaches, and cybersecurity issues in industrial environments.

The analysis reveals several key limitations:

- Limited availability of industrial datasets
- Difficulty in obtaining reliable ground truth
- Absence of controlled and reproducible experimental environments
- Lack of integrated datasets combining physical and cyber anomalies

These gaps motivate the research presented in this thesis, which proposes a unified and systematic framework for dataset generation, anomaly detection, and cybersecurity analysis in industrial CPS, enabling the generation of structured and reliable data for the evaluation of machine learning techniques applied to anomaly detection and cybersecurity in industrial Cyber-Physical systems.

1.3. Objectives

The main objective of this thesis is to develop and validate a unified and systematic framework for ground truth-based dataset generation, anomaly detection, and cybersecurity analysis in industrial Cyber-Physical Systems. This framework aims to address the limitations identified in the literature regarding data availability, the lack of reliable ground truth, and the difficulty of evaluating machine learning techniques in realistic industrial environments.

To achieve this objective, the following specific objectives are defined:

- To develop a structured methodology for dataset generation in industrial CPS environments, based on controlled anomaly injection and the explicit separation between the CPS process and the ground truth generation layer, thereby ensuring data reliability, traceability, and validity.
- To design and implement a hybrid CPS testbed integrating real industrial hardware and digital twins, capable of reproducing realistic and controlled industrial scenarios, and enabling the practical implementation of the proposed methodology for both research and educational purposes.
- To generate industrial-grade datasets that integrate normal operation, induced anomalies in process data from sensors and actuators, as well as cybersecurity-related anomalies, within a unified, reproducible, and event-based environment.
- To validate the proposed framework by evaluating machine learning techniques for anomaly detection and cybersecurity applications, demonstrating the effectiveness of the generated datasets.

1.4. Methodology

This section presents the methodology proposed in this thesis for the systematic generation of structured datasets in industrial Cyber-Physical Systems. The approach defines a conceptual framework that enables the generation of both normal and anomalous behaviour while ensuring the definition of reliable and unambiguous ground truth.

The methodology is designed to address the limitations identified in the state of the art, particularly those related to dataset scarcity, lack of reproducibility, and ambiguity in anomaly labelling. To this end, the proposed approach integrates a

controlled CPS environment, a structured anomaly generation strategy, and a clear separation between the industrial process and the ground truth layer.

While existing methodologies, such as GeMSyD [23], focus on descriptor-based synthetic data generation by extracting statistical, structural, and temporal characteristics from real-world data, they do not explicitly consider the integration of these descriptors within an operational CPS environment, nor the systematic definition of ground truth through controlled experimental mechanisms.

In contrast, the methodology proposed in this thesis extends this paradigm by embedding data generation within an operational CPS environment and introducing an external experimental layer for controlled anomaly and attack injection together with explicit ground-truth generation. This approach enables the systematic generation of realistic, labelled, and reproducible event-based datasets tailored for anomaly detection and cybersecurity evaluation.

To structure the proposed methodology, a GEMMA-based CPS dataset generation framework is introduced. This framework follows a sequential and iterative process that integrates data acquisition, controlled anomaly and attack injection, ground truth definition, and dataset construction.

Figure 1.1 illustrates the proposed framework, organized around a central CPS environment that enables controlled execution of the industrial process and supports reproducible experimentation.

The workflow includes data acquisition, anomaly and attack injection, hybrid ground truth generation, and dataset construction based on an event-driven representation. The framework highlights the interaction between system operation and experimental control, enabling systematic dataset generation under diverse conditions.

A key aspect of the methodology is the explicit separation between the CPS process and the ground truth generation mechanism. This decoupling ensures consistent labelling, eliminates ambiguity, and enables precise traceability between injected anomalies and their observable effects.

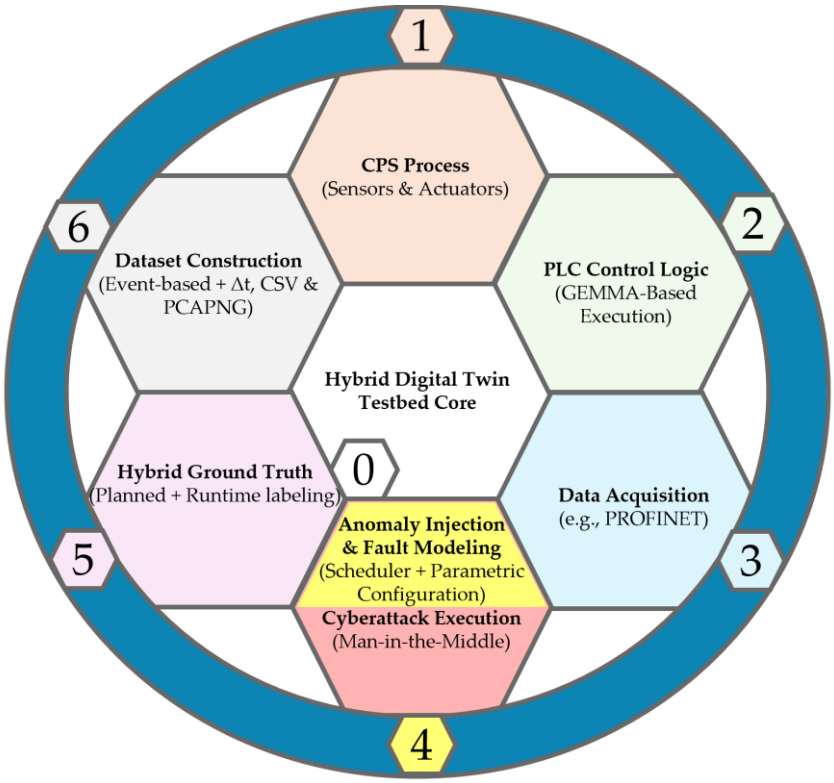


Figure 1.1. GEMMA-based CPS framework for dataset generation

The dataset construction stage enables the generation of structured datasets suitable for different types of analysis, including machine learning and network-level evaluation.

Within this process, descriptor extraction is implicitly performed by modelling temporal behaviour (Δt), operational sequences such as GEMMA states, and anomaly patterns. This ensures consistency with descriptor-based dataset generation principles while preserving the dynamics of the CPS.

The proposed approach emphasizes the generation of reliable datasets through the incorporation of event-based and temporal information, as well as the explicit separation between the CPS process and the ground truth layer.

Overall, this methodology, referred to as the GEMMA-based CPS dataset generation framework, provides a structured, reproducible, and scalable approach for dataset generation and evaluation in industrial CPS environments.

Detailed descriptions of the CPS implementation, data acquisition mechanisms, anomaly injection strategies, and dataset structure are provided in Chapters 3 and 6.

1.4.1. Overall Framework

The proposed framework is based on the interaction of three main conceptual components: the CPS process, the data acquisition layer, and the ground truth generation mechanism. These components are integrated within a unified experimental CPS environment, represented as the central element of the framework in Figure 1.1.

This layered architecture enables a clear decoupling between system behaviour and its interpretation, which is essential for the systematic generation of reliable datasets and consistent anomaly labelling.

The CPS process represents the operational execution of the industrial system, where interactions between system variables and control logic define its operational behaviour. On top of this process, the data acquisition layer captures system variables, enabling a structured representation of system dynamics.

A key aspect of the proposed framework is the explicit separation between the CPS process and the ground truth generation mechanism. This separation allows anomalies to be defined independently of observable system signals, ensuring consistent labelling and enabling accurate evaluation of anomaly detection methods.

Overall, the proposed methodology provides a structured and modular foundation for dataset generation, supporting reproducible experimentation and reliable validation of data-driven approaches in anomaly detection and cybersecurity for CPS.

The implementation of the CPS environment and its components is detailed in Chapter 3, while the dataset construction and labelling process are described in Chapter 6.

1.4.2. CPS Process and Control

This subsection presents the CPS process and control stages illustrated in Figure 1.1, where these stages correspond to elements (1) and (2) of the framework.

Within the proposed methodology, the CPS process represents the operational layer in which the industrial system is executed under controlled conditions. This process defines the dynamic behaviour of the system through the interaction

between system variables and control logic, enabling the reproduction of different operational scenarios.

The CPS layer provides the basis for controlled experimentation, allowing the system to operate under both normal and anomalous conditions. In particular, it supports the representation of different operational modes and state transitions, enabling the systematic exploration of system behaviour.

This controlled execution environment enables the generation of representative system dynamics, which are essential for dataset construction and subsequent analysis. By operating under reproducible conditions, the CPS process ensures consistency across experimental runs and supports reliable evaluation of anomaly detection and cybersecurity approaches.

The implementation of the CPS environment and its integration within the experimental testbed are described in detail in Chapter 3.

1.4.3. Data Acquisition

The data acquisition stage, represented by element (3) in Figure 1.1, involves recording process data and operational states during industrial system execution. The acquisition follows an event-driven approach, whereby data are captured whenever changes occur in system variables rather than at fixed sampling intervals. This strategy provides an efficient and compact representation of system dynamics while preserving the sequence of events and their temporal relationships.

Special attention is given to the integration of temporal information, enabling the characterization of inter-event times (Δt) and the reconstruction of system evolution over time. This representation captures both the logical execution of the control process and its temporal behaviour, which are essential for modelling industrial CPS dynamics.

The resulting data provide a consistent and structured description of system operation, enabling the extraction of temporal dependencies and state transitions required for the development and evaluation of anomaly detection and cybersecurity methods.

The specific implementation of the data acquisition mechanisms is described in Chapter 6.

1.4.4. Anomaly Injection and Attack Execution

This subsection focuses on the anomaly injection and attack execution stage of the proposed framework, as illustrated in Figure 1.1, where it corresponds to element (4).

Within the proposed methodology, anomalies are introduced through controlled mechanisms that enable the systematic generation of abnormal system behaviour under reproducible conditions. These mechanisms allow the definition of anomalous scenarios affecting different components of the CPS, ensuring diversity and realism in the generated data.

Anomaly generation is performed in a structured and traceable manner, where each injected condition is associated with a predefined identifier. This approach enables consistent integration with the ground truth generation process and supports precise traceability between injected anomalies and their observable effects.

In addition to fault-based anomalies, the experimental framework incorporates controlled execution of cybersecurity scenarios, enabling the generation of data that reflect both operational faults and adversarial conditions. The separation between fault injection and cyberattack execution ensures that different types of anomalies are generated through independent and well-defined mechanisms.

Overall, this design enables fine-grained control over anomaly generation, ensures reproducibility across experimental scenarios, and provides the traceability required for reliable dataset construction and evaluation in anomaly detection and cybersecurity.

The specific implementation of anomaly injection mechanisms and attack scenarios is detailed in Chapters 3, 4 and 5.

1.4.5. Ground Truth Generation

This subsection introduces the ground truth generation stage of the proposed framework, as illustrated in Figure 1.1, where it corresponds to element (5).

A central contribution of this methodology is the explicit and controlled definition of ground truth as a component decoupled from direct CPS observations. This decoupling allows anomalies to be defined independently of

observable system signals, eliminating ambiguity and ensuring consistent labelling across the dataset.

The ground truth is generated as a distinct and independent module that captures the actual system condition over time, including normal operation, fault scenarios, and cybersecurity events. It is constructed by combining externally defined anomaly information with process-aware context derived from the control system, enabling a consistent interpretation of system behaviour.

This hybrid approach integrates predefined anomaly definitions with contextual system information, allowing each data instance to be accurately labelled while preserving its operational context. As a result, the framework supports a precise and consistent characterization of both normal and anomalous system behaviour.

Furthermore, this design provides traceability between injected anomalies and their observable effects, enabling the association of each anomaly with its underlying cause. This capability is essential for the reliable evaluation and validation of anomaly detection and cybersecurity approaches in CPS.

The implementation of the ground truth generation mechanism and its integration with the dataset are described in detail in Chapter 6.

1.4.6. Dataset Construction

This subsection outlines the final dataset construction stage of the proposed framework, as illustrated in Figure 1.1, where it corresponds to element (6).

The dataset is generated by integrating the data acquired from the CPS with the corresponding ground truth information, resulting in a structured and fully annotated representation of system behaviour.

Special attention is given to the incorporation of event-based and temporal information, allowing the dataset to capture the dynamic evolution of the system. This design preserves both state transitions and temporal dependencies, which are essential for accurately modelling industrial CPS behaviour.

The resulting datasets support different types of analysis, including process-level modelling and network-level evaluation, enabling the study of both operational behaviour and cybersecurity scenarios within a unified framework.

Overall, this structured dataset construction approach ensures reproducibility, preserves the temporal and logical properties of the CPS, and enables consistent

benchmarking and fair evaluation of data-driven methods for anomaly detection and cybersecurity.

The detailed dataset structure, formats, and labelling processes are described in Chapter 6.

1.5. Contributions

This thesis presents several contributions to the field of anomaly detection and cybersecurity in industrial Cyber-Physical Systems. These contributions are structured into three categories: technical, methodological, and experimental, and are supported and validated through the scientific publications included in the Papers section.

1.5.1. Technical Contributions

The technical contributions focus on the development of a comprehensive experimental infrastructure for CPS research:

- Development of a hybrid CPS testbed integrating real industrial hardware (Siemens S7-1200 and S7-1500 PLCs), digital twin environments (based on PCSimu, with compatibility for alternatives such as RobotStudio and SIMIT), and industrial communication networks.
- Design of a scalable and flexible architecture capable of reproducing multiple industrial scenarios, supporting different operational modes (e.g., GEMMA states), and enabling both normal and faulty system conditions.
- Implementation of a data acquisition framework supporting event-based data capture, integration of process signals and GEMMA state information, and structured storage tailored for machine learning applications.
- Integration of fault injection and cybersecurity experimentation within a unified and controllable platform.

1.5.2. Methodological Contributions

The methodological contributions constitute the main scientific contribution of this thesis, particularly in the generation, labelling, and validation of datasets:

- Proposal of a novel ground truth generation paradigm based on the explicit decoupling between system observations and anomaly definition, ensuring

consistent and unambiguous labelling, extending and generalizing the concepts introduced in GeMSyD [23].
- Definition of a structured dataset generation methodology, including controlled execution of industrial processes, systematic data acquisition, and reproducible fault injection strategies.
- Design of a hybrid ground truth mechanism that combines process-aware information derived from the control system (e.g., GEMMA states) with externally defined annotations driven by controlled anomaly scheduling.
- Incorporation of event-based data and temporal dynamics (Δt between events), enabling accurate modelling of system behaviour and enhancing the suitability of datasets for machine learning techniques.

1.5.3. Experimental Contributions

The experimental contributions validate the proposed framework and demonstrate its applicability in realistic CPS environments:

- Generation of industrial datasets with controlled and fully traceable ground truth, encompassing normal operation, induced faults, and cybersecurity scenarios, suitable for training, validation, and testing of anomaly detection models.
- Experimental evaluation of machine learning techniques (e.g., KNN, SVM, Random Forest) under controlled and reproducible conditions, demonstrating the effectiveness of the proposed framework in enabling reliable anomaly detection and realistic industrial scenario modelling.

The experimental validation of the proposed contributions is supported by the scientific publications included in this thesis.

1.6. Thesis Structure

This thesis is structured as a unified work that integrates conceptual, methodological, and experimental components.

Chapter 1 introduces the context and motivation of the research, presents the state of the art, and defines the objectives, methodology, and main contributions of the thesis.

Chapter 2 provides the theoretical foundations required for the development of this work, including industrial Cyber-Physical Systems, discrete-event modelling, control systems, digital twins, and event-based data representation.

Chapter 3 describes the development of the experimental framework, describing the overall CPS-based system that supports the proposed methodology, including the testbed architecture, the configuration and planning of anomaly scenarios, and the integration of the dataset generation process.

Chapter 4 introduces the initial validation of the proposed approach through the design and application of a CPS testbed for anomaly detection using machine learning techniques.

Chapter 5 extends the experimental framework by incorporating advanced cybersecurity scenarios, including Man-in-the-Middle (MitM) attacks.

Chapter 6 describes the dataset derived from the proposed framework, detailing its structure, acquisition process, labelling strategy, and organisation, enabling reproducible evaluation of anomaly detection and cybersecurity approaches.

Finally, Chapter 7 presents the main conclusions of the thesis and outlines directions for future research.

Chapter 2

Fundamentals of Industrial Cyber-Physical Systems

This chapter introduces the fundamental concepts underlying industrial Cyber-Physical Systems, which constitute the theoretical and technological basis of the proposed approach.

2.1. Industrial CPS

Cyber-Physical Systems represent a paradigm that integrates physical processes with computation, control, and communication capabilities [1, 2]. In industrial environments, CPS enable the seamless interaction between mechanical components, sensing devices, control units, and digital infrastructures, forming interconnected systems capable of monitoring and controlling processes in real time.

In the context of industrial automation, CPS are typically composed of several interconnected layers [1]. The physical layer includes sensors and actuators that interact directly with the production process. The control layer is typically implemented using programmable logic controllers, which execute control logic to govern system behaviour. The communication layer enables data exchange between devices through industrial protocols, while the cyber layer encompasses data processing, storage, and higher-level applications such as monitoring, optimization, and decision-making.

This layered architecture allows industrial CPS to support advanced functionalities associated with Industry 4.0 [1], including increased flexibility, real-time supervision, and integration with digital services. In particular, the availability of continuous data streams from sensors and controllers provides the foundation for data-driven techniques, such as anomaly detection, predictive maintenance, and system optimization [21].

Despite these advantages, the modelling of industrial CPS remains a challenging task due to the hybrid nature of these systems, where physical dynamics coexist with discrete control logic [29, 30]. However, in many industrial applications, system behaviour is predominantly governed by discrete events, such as sensor activations, execution of control commands, or transitions between operational states. Therefore, discrete event abstractions provide a suitable framework for representing the evolution of industrial processes in a structured and interpretable manner [29].

This observation motivates the use of Discrete Event Systems (DES) as the primary modelling paradigm in this work [29]. By focusing on event-driven dynamics, DES enable a formal representation of system behaviour that is consistent with the logic implemented in PLC-based controllers. This event-driven perspective establishes the foundation for the subsequent sections, where industrial processes are first modelled using DES and then implemented through PLC programming and evaluated within a digital twin environment.

2.2. Discrete Event Systems (DES)

Discrete Event Systems provide a formal and well-established framework for modelling, analysing, and controlling systems whose dynamics are driven by asynchronous events [29, 30]. In this paradigm, the state of the system changes only at discrete instants defined by the occurrence of events, such as sensor activations, command executions, or state transitions. This characteristic makes DES particularly suitable for representing industrial processes governed by programmable logic controllers, where control logic is inherently event-driven.

In industrial Cyber-Physical Systems, DES enable the abstraction of complex interactions between physical components and control software into a structured representation based on states and transitions. Typically, DES models are formalized using finite automata (state machines), Petri nets, and GRAFCET (Sequential Function Charts), as well as other state-based representations, which capture both the logical sequencing of operations and the conditions under which transitions occur [30].

Figure 2.1 illustrates the conceptual interaction between the controlled system and its DES-based control model within an industrial CPS framework. The physical process evolves through its interaction with the control logic, which is

implemented using a state-based representation. Inputs from the system are processed by the control model, while outputs generated by the model influence the system behaviour, establishing a closed-loop interaction. This abstraction provides a clear separation between the physical system and its logical representation, which is essential for both analysis and implementation.

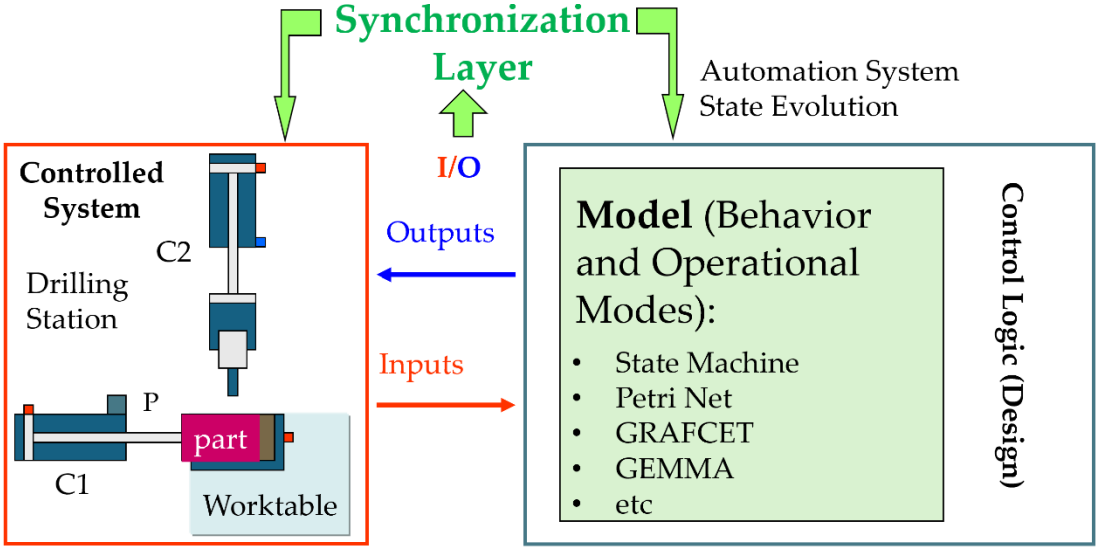


Figure 2.1. Controlled system and DES-based control model interaction.

These models facilitate system verification, validation, and analysis, allowing engineers to reason about system behaviour under both nominal and faulty conditions.

A key advantage of DES in industrial environments is their compatibility with the execution semantics of PLC-based control systems. Since PLC programs operate by scanning inputs, executing logic, and updating outputs in cyclic or event-driven manners, DES models naturally align with this operational structure. As a result, DES-based representations can be directly mapped to PLC implementations, reducing the gap between system design and deployment.

Additionally, DES provide a suitable foundation for the systematic generation of datasets for data-driven approaches. By representing system behaviour as sequences of discrete events, it becomes possible to capture temporal dependencies, state transitions, and fault-related behaviours in a structured manner.

This event-based representation is particularly suitable for anomaly detection, as deviations from expected event sequences can be systematically identified.

Figure 2.2 depicts the evolution of the system state X(k) as a function of discrete events E(k). The system transitions between discrete states only upon the occurrence of events, resulting in a piecewise-constant state trajectory over discrete time steps. Each state represents a specific operational condition of the system, while events correspond to observable changes such as sensor activations or command executions. This event-driven representation forms the basis for constructing structured, time-indexed datasets, where sequences of events and corresponding system states can be systematically captured and analysed for anomaly detection tasks.

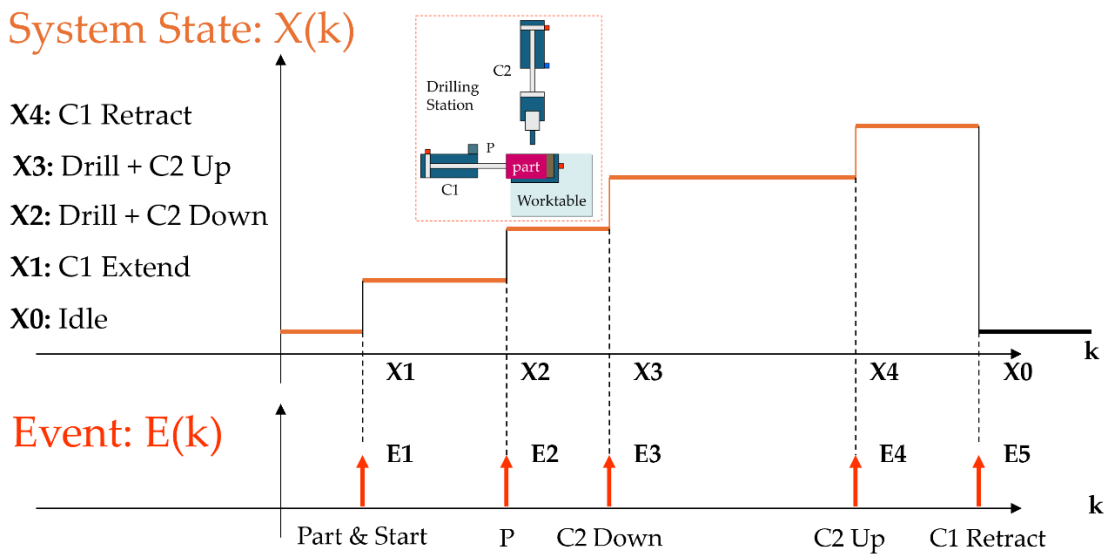


Figure 2.2. Evolution of system states X(k) driven by discrete events E(k).

This representation is particularly relevant for anomaly detection tasks, where deviations from expected event sequences can indicate abnormal or faulty behaviour.

In this work, DES are employed to model industrial processes following a structured operational logic, enabling the identification of state transitions and event patterns. These models are subsequently used to generate event-based datasets with well-defined temporal resolution (Δt), which are later leveraged for the training and evaluation of machine learning models. Furthermore, the DES framework facilitates the controlled injection of faults and anomalies, ensuring

the availability of labelled data while maintaining a clear separation between the physical process and the experimental layer used for ground truth generation.

This integration of DES modelling, PLC implementation, and dataset generation constitutes a coherent pipeline that supports both system analysis and the development of robust anomaly detection methods within industrial CPS.

2.3. GEMMA Model

The GEMMA (Guide d'Étude des Modes de Marches et d'Arrêts) framework provides a standardized methodology for defining and organizing the operational modes of industrial automated systems. It is widely used in industrial automation to describe the different states of operation of a system, including normal functioning, fault conditions, and recovery procedures. By structuring system behaviour into well-defined modes, GEMMA facilitates both the design and analysis of control logic in complex industrial environments.

In the context of Discrete Event Systems, GEMMA can be interpreted as a higher-level abstraction that organizes system behaviour into operational modes, while DES provide the formal representation of state transitions driven by events. Each GEMMA mode can be associated with a subset of system states and transitions, enabling a hierarchical representation where operational logic and state evolution are coherently integrated.

A key advantage of GEMMA is its ability to explicitly represent transitions between different operational conditions, such as automatic operation, manual control, fault detection, and recovery phases.

Figure 2.3 presents the standard GEMMA diagram, which organizes the operational modes of an industrial system into three main categories: operating procedures, stopping procedures, and fault-handling procedures. The diagram highlights the possible transitions between these modes, providing a comprehensive view of system behaviour under both nominal and abnormal conditions.

From a discrete-event system perspective, each GEMMA mode can be associated with a subset of system states and event-driven transitions. This enables a direct mapping between high-level operational modes and the underlying state–event representation introduced in Figure 2.2 [20, 30].

These transitions are typically triggered by discrete events, including operator commands, sensor signals, or internal system conditions. As a result, GEMMA naturally complements discrete-event system modelling by providing a structured framework for organizing event-driven behaviour in industrial systems [20].

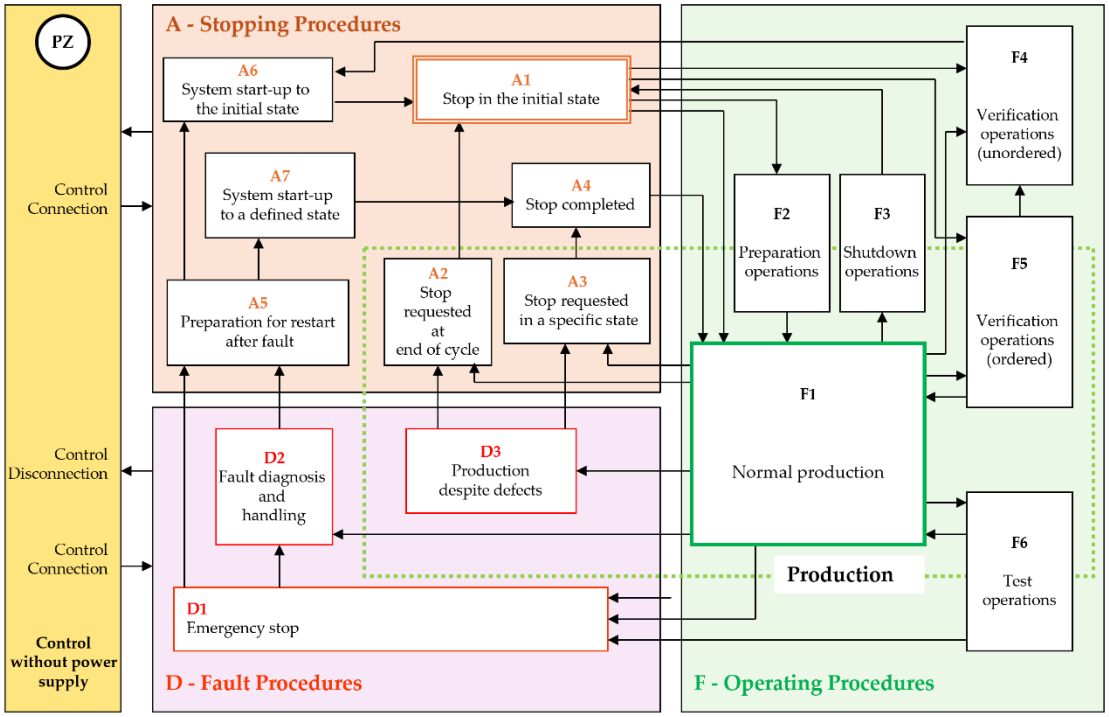


Figure 2.3. Standard GEMMA diagram illustrating operational modes and transitions.

Building upon the general principles of GEMMA described above, this work adopts a structured operational framework to guide system evolution under both normal and fault conditions.

In particular, the considered operational sequence follows a cyclic pattern that includes normal operation, fault occurrence, recovery, and reinitialization phases. This sequence is consistent with typical industrial behaviour and supports the controlled generation of datasets including both normal and anomalous conditions. By leveraging GEMMA, controlled perturbations can be introduced in specific modes, enabling systematic fault injection while preserving full traceability of the system state. This approach, aligned with recent methodologies that integrate GEMMA and model-based control frameworks for Digital Twin environments [20], enables the generation of event-based datasets

that reflect realistic industrial scenarios, where transitions between operational modes are explicitly captured. Furthermore, the clear separation between the physical process and the experimental definition of operational modes ensures that the ground truth labelling is consistent, reproducible, and independent of the underlying CPS execution.

An example of GEMMA operational modes and transitions adapted to the considered framework is shown in Figure 2.4.

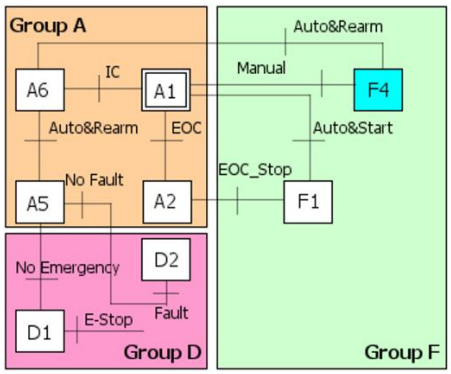


Figure 2.4. Event-driven transitions between GEMMA operational modes.

This representation highlights how GEMMA can be used to structure and analyse the operational behaviour of industrial systems. In order to realize this integration, the control logic must be implemented using appropriate programming techniques, which are discussed in the following section.

2.4. PLC Programming

The implementation of Discrete Event Systems in industrial environments is commonly achieved through programmable logic controllers, which provide a reliable and standardized platform for executing control logic. PLCs are widely used in industrial automation due to their robustness, real-time capabilities, and direct integration with sensors, actuators, and Human–Machine Interfaces (HMIs).

In this context, the DES models described in previous sections are translated into executable control logic within the PLC. This translation involves mapping states, events, and transitions into structured programming constructs that define system behaviour. The cyclic execution model of PLCs, based on continuous

scanning of inputs, logic execution, and output updates, naturally supports the implementation of event-driven systems.

The control logic is typically developed using the programming languages defined in the IEC 61131-3 standard [31], such as Ladder Diagram (LD), Structured Text (ST), and Sequential Function Charts (SFC). In the proposed system, multiple languages are combined to exploit their respective advantages. Ladder Diagram (LD) is used for low-level control and signal handling, Structured Text, implemented as Structured Control Language (SCL) in Siemens environments, is used for more complex logic and data processing, and Sequential Function Charts (SFC) are used to represent the sequential behaviour of the system.

In particular, the GRAFCET-based control sequences are implemented using Siemens S7-Graph, which provides a graphical environment for defining steps and transitions. This approach enables a direct correspondence between the conceptual models introduced in Sections 2.2 and 2.3 and their practical implementation at the PLC level.

In the proposed system, the DES-based control logic is implemented in a PLC that manages the execution of the physical process within the digital twin environment. Each system state is encoded through a set of internal variables, while transitions between states are triggered by events derived from sensor signals, operator inputs, or internal conditions. This structure ensures that the execution of the control program faithfully reproduces the state–event dynamics defined in the DES model.

Furthermore, the integration of GEMMA operational modes into the PLC logic allows the system to dynamically switch between different modes of operation, such as manual control, automatic production, fault handling, and recovery procedures. These modes are implemented through supervisory logic that coordinates the activation of sequences and state transitions, ensuring that the PLC executes the appropriate control routines and process sequences according to the active GEMMA operational mode.

This implementation also enables the systematic acquisition of process data, where relevant variables such as sensor states, actuator signals, and operational modes are continuously monitored and recorded. By associating these variables with discrete time steps defined by the system sampling interval Δt, it becomes

possible to construct structured, time-indexed datasets that capture the evolution of system behaviour. These time-indexed observations are inherently event-driven, reflecting changes in system states and transitions over time.

As a result, the PLC not only serves as the execution platform for the DES and GEMMA-based control logic, but also as a key component for data generation. The data collected during system operation provides the foundation for subsequent analysis, including anomaly detection and machine learning applications, thereby closing the loop between system design, implementation, and data-driven evaluation.

Figure 2.5 shows a representative excerpt of the DES-based control logic implemented in the Siemens TIA Portal environment. The sequential behaviour of the system is defined using S7-Graph, where control logic is structured in terms of steps and transitions.

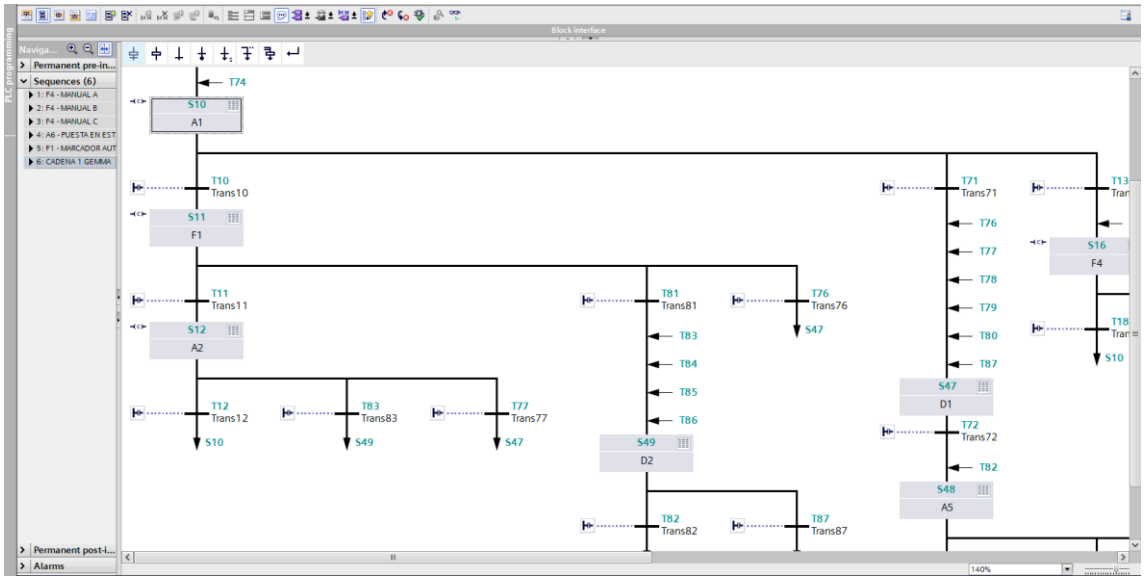


Figure 2.5. DES-based control logic implemented in Siemens TIA Portal using S7-Graph.

2.5. Digital Twins

Digital Twins have emerged as a key technology in the development of modern Industrial Cyber-Physical Systems, enabling the creation of virtual representations of physical processes that evolve in parallel with the real system [4, 5, 32]. A Digital Twin can be defined as a dynamic digital counterpart of a

physical asset, continuously updated through data exchange, allowing monitoring, simulation, and analysis of system behaviour in real time [4, 32].

In industrial CPS, the Digital Twin establishes a synchronized interaction with the physical process, where relevant variables such as sensor signals, actuator states, and control commands are exchanged between both domains [4, 5, 32]. This synchronization ensures that the virtual model accurately reflects the current state of the system, while also allowing the evaluation of alternative scenarios without interfering with the physical operation [3, 32]. As a result, the Digital Twin provides a consistent framework for analysing system behaviour under different operating conditions [5, 32].

Although Digital Twins are often associated with continuous synchronization between physical and virtual domains, their ability to reproduce the operational dynamics of a process also makes them valuable experimentation environments, even when real-time synchronization is not required [32]. In such cases, the Digital Twin can be used to execute controlled scenarios, analyse system behaviour, and generate representative datasets while preserving the operational characteristics of the industrial process.

In the context of Discrete Event Systems, the Digital Twin enables the representation and tracking of system evolution in terms of discrete states and events, preserving the coherence of the process model and its event-driven representation [4, 30]. Each state transition can be reproduced within the Digital Twin, maintaining the event-driven dynamics introduced in Section 2.2. This capability is particularly relevant when considering the temporal discretization of the system through a sampling interval Δt, ensuring consistency between process execution and data representation.

In this work, the Digital Twin is integrated into the developed testbed as a key component that links the PLC-based control system and the data acquisition infrastructure. The Digital Twin serves as the experimentation environment within the testbed, enabling the execution of controlled experiments and the systematic observation of system behaviour. This configuration allows the reproduction of different operational conditions, including both nominal operation and fault scenarios, under controlled and reproducible settings.

A major advantage of this approach is the ability to perform controlled fault injection and repeatable experimentation. By using the Digital Twin, it is possible

to introduce anomalies in a structured manner while maintaining full traceability of system states and transitions. This capability is particularly important for the generation of datasets, as it ensures that both normal and anomalous behaviours are captured under well-defined and reproducible conditions.

Furthermore, the integration of the Digital Twin enables a clear separation between the physical process and the experimental layer used for data generation and labelling. This separation allows the definition of ground truth independently from the industrial process, avoiding interference with the control system while ensuring accurate and consistent annotation of system behaviour.

As a result, the Digital Twin provides the foundation for generating structured datasets that capture the evolution of the system in terms of discrete events and state transitions. These datasets are particularly suitable for anomaly detection tasks, where identifying deviations from expected behaviour is essential for ensuring system reliability and security [4].

As illustrated in Figure 2.4, the industrial process is represented through a Digital Twin that emulates the behaviour of the real physical process, including sensors and actuators, while being directly integrated with the control logic and HMI interface. This configuration provides a consistent CPS framework for controlled experimentation and data generation.

2.6. Event-Based Data and Anomaly Detection

In industrial Cyber-Physical Systems, system behaviour can be represented using different data abstraction approaches, commonly categorized as time-based or event-based representations [4, 30]. Time-based data consists of signals sampled at fixed intervals, providing a continuous view of system variables over time. In contrast, event-based data captures discrete changes in system behaviour, where information is recorded only when relevant events occur, such as changes in sensor and actuator states, control commands, or state transitions.

In this work, an event-driven representation is adopted, in line with the Discrete Event Systems framework introduced in Section 2.2. Under this paradigm, system behaviour is described as sequences of events E(k) and associated states X(k), where changes occur at discrete time instants [30]. This representation enables a more compact and structured description of industrial processes, focusing only on meaningful changes in system dynamics. Additionally, the use

of a sampling interval Δt ensures temporal consistency and allows the alignment of event sequences with time-based measurements when required.

The generation of datasets in industrial environments presents several challenges. One of the main difficulties is the inherent imbalance between normal and faulty operating conditions, as industrial systems typically operate under nominal conditions for long periods, while faults occur rarely [15]. This imbalance can significantly impact the performance of data-driven models, particularly in anomaly detection tasks. Furthermore, in many real-world scenarios, the absence of reliable ground truth complicates the development and evaluation of machine learning models, as system states and fault conditions are not explicitly labelled [4, 5].

The proposed approach addresses these challenges by leveraging the capabilities of the Digital Twin and the structured modelling of system behaviour through DES and GEMMA. By organizing system operation through the operational modes defined by GEMMA and controlled process sequences, it becomes possible to systematically generate datasets that include both normal and anomalous conditions. In particular, fault scenarios can be introduced in a controlled manner through experimental design, enabling the capture of representative event sequences associated with fault conditions and abnormal behaviour.

A key aspect of this approach is the explicit separation between the physical (or simulated) process and the experimental layer used for data labelling. Ground truth is defined externally to the industrial process, based on the known conditions under which the system is operated and the controlled injection of faults. This strategy ensures accurate and consistent labelling, avoiding ambiguities that commonly arise in real industrial datasets.

As a result, the generated datasets consist of structured sequences of events and state transitions, which can be directly used for the development and evaluation of anomaly detection algorithms [4]. In this context, anomalies are interpreted as deviations from the expected behaviour of the system, regardless of whether they affect event sequences, state transitions, temporal dynamics, or process signals. These event sequences are inherently linked to the operational modes defined by GEMMA, allowing anomaly detection to be performed not only at the signal level but also at the level of system operation modes.

The combination of event-based data representation, controlled dataset generation, and reliable ground truth provides a robust foundation for data-driven analysis. This framework enables the application of machine learning techniques for anomaly detection, while ensuring that the underlying data accurately reflects the operational behaviour of the system.

2.7. Chapter Summary

This chapter presented the theoretical foundations of the proposed research. It introduced the fundamental concepts of Industrial Cyber-Physical Systems, Discrete Event Systems, GEMMA operational modes, PLC-based control, Digital Twins, and event-based data representation, establishing the conceptual framework that supports the development of the proposed methodology. Discrete Event Systems were identified as the primary modelling paradigm, while GEMMA provided a structured framework for organizing operational modes under normal and fault conditions. The chapter also discussed the implementation of these concepts through PLC-based control systems and highlighted the role of Digital Twins in enabling controlled experimentation, fault injection, and dataset generation.

Finally, the advantages of event-based data representation for anomaly detection were examined, emphasizing its ability to capture system behaviour through structured sequences of events and state transitions. Together, these concepts provide the theoretical basis for the development of reliable datasets and data-driven anomaly detection approaches in industrial CPS. Building upon these foundations, the next chapter presents the experimental framework developed for this research, including the CPS testbed architecture, fault-injection mechanisms, and dataset generation process used throughout this thesis.

Chapter 3

Development of the Experimental Framework

This chapter presents the experimental framework and system architecture developed in this research to implement the CPS-based dataset generation approach introduced in Chapter 1. The proposed framework is based on a hybrid CPS testbed that integrates real industrial hardware with simulation environments, enabling controlled experimentation under realistic operating conditions.

3.1. CPS Testbed Overview

The experimental framework is structured around a hybrid Cyber-Physical system architecture that integrates process simulation, PLC-based control, supervisory systems, and data acquisition through PROFINET into a unified platform. This architecture is illustrated in Figure 3.1.

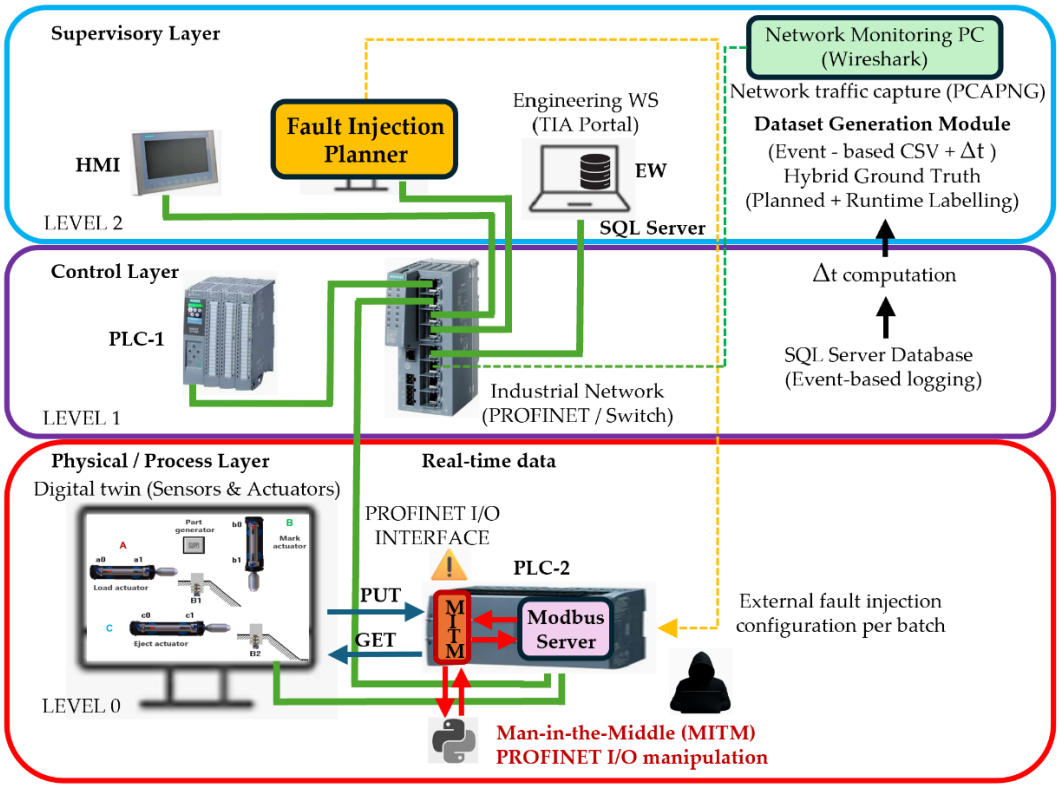


Figure 3.1. Hybrid CPS testbed architecture for dataset generation and cybersecurity experimentation.

As shown in Figure 3.1, the proposed CPS testbed follows a hierarchical multi-layer structure composed of the physical/process layer, control layer, and supervisory layer.

At the physical layer, the industrial process is represented by a digital twin interacting with sensors and actuators through a PROFINET I/O interface, enabling real-time operation. The control layer is responsible for process execution through PLC-based logic, ensuring deterministic system behaviour. At the supervisory level, higher-level functionalities such as fault planning, data storage, and dataset generation are integrated.

In particular, the architecture enables a clear data flow from the physical process to the data acquisition modules, supporting the transformation of real-time signals into structured, event-based datasets.

This layered architecture provides a suitable framework for the controlled injection of faults and the systematic acquisition of process data, enabling the generation of structured datasets under different operating conditions.

The industrial process corresponding to the physical layer, implemented through a digital twin, is depicted in Figure 3.2. The system represents a part-marking process composed of three main stations (loading, marking, and ejection), each associated with an actuator (A, B, and C), together with the corresponding position sensors and presence detectors. The figure also includes the HMI used for supervision, providing a direct visualization of the interaction between the physical process and the control logic.

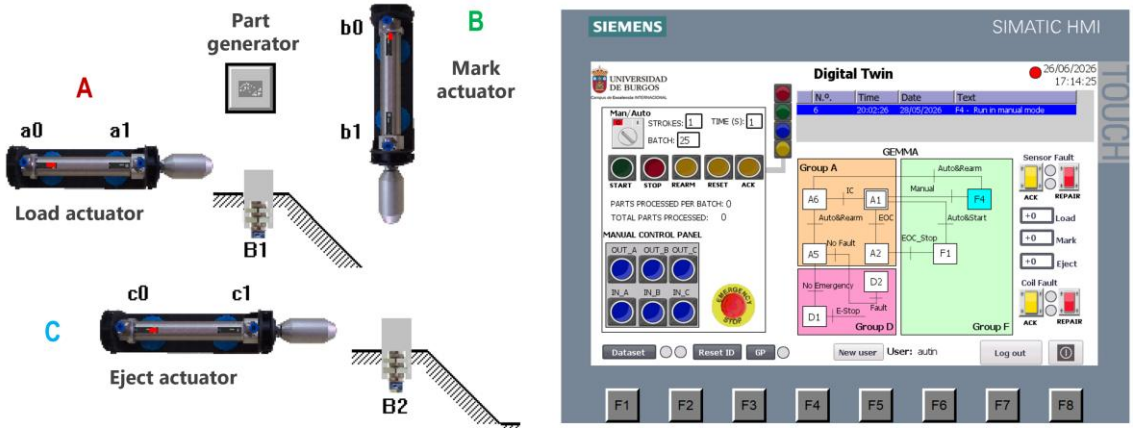


Figure 3.2. Reference CPS use case: industrial part-marking system with actuators, sensors, presence detectors (B1, B2), and GEMMA-based HMI.

To complement the architectural description and provide a detailed view of the experimental setup, Tables 3.1 and 3.2 summarize the main hardware and software components of the proposed CPS testbed.

Table 3.1. Hardware components used in the testbed.

Component	Model	Manufacturer	Role in the testbed
PLC 1	6ES7512-1CK01-0AB0	Siemens	Main Controller
PLC 2	6ES7215-1BG40-0XB0	Siemens	PROFINET I-Device
HMI Panel	6AV2 123-2JB03-0AX0	Siemens	HMI panel
Switch	6GK5208-0BA00-2AC2	Siemens	Industrial Ethernet
Engineering PC	PC (Intel i7, 32 GB RAM)	-	TIA Portal, MSSQL
Digital Twin Host	PC (Intel i5, 8 GB RAM)	-	Runs PCSIMU
Attacker PC	PC (intel i7, 16 GB RAM)	-	Python MitM, Wireshark

Table 3.2. Software components used in the testbed.

Component	Version	Platform / OS	Role in the testbed
TIA Portal	V18 Update 5	Windows 11	PLC engineering
PCSIMU	V3.0	Windows 10	Digital twin
Python	3.13.5	Windows 10	MitM module
Pymodbus	3.11.4	Python	Modbus client
MB_SERVER block	Standard lib.	PLC S7 1200	Modbus server
MSSQL Server Express	2019	Windows 11	Event logging
SSMS	20.2.30.0	Windows 11	SQL administration
Wireshark	4.6.3	Windows 10	Network capture
Kali Linux tools	Rolling	Linux	Reconnaissance

In addition to its role in the MitM module reported in Table 3.2, Python-based tools are also used for fault planning within the dataset generation process.

These components collectively define the experimental infrastructure used throughout the thesis.

3.2. Development of the Fault Injection Tool

To enable the controlled generation of anomalous conditions, the proposed CPS testbed incorporates a fault injection tool developed for this research and based on predefined fault planning strategy at the batch level. This tool is specifically designed for the generation of labelled datasets, allowing the systematic introduction of faults affecting sensors and actuators.

While the testbed also supports cybersecurity-oriented experiments, including Man-in-the-Middle (MitM) attacks, these represent a distinct mechanism from the fault injection process. Fault injection is specifically used for the controlled generation of labelled datasets, particularly for sensor and actuator anomalies. Cybersecurity scenarios are addressed as an extension of the proposed CPS framework in Chapter 5, and may also complement the experimental evaluation by introducing network-level anomalies.

The fault injection mechanism is supported by a dedicated planning interface that allows the configuration of fault scenarios, as illustrated in Figures 3.3 and 3.4, which correspond to the batch configuration and selection stage and the subsequent planning and execution stages.

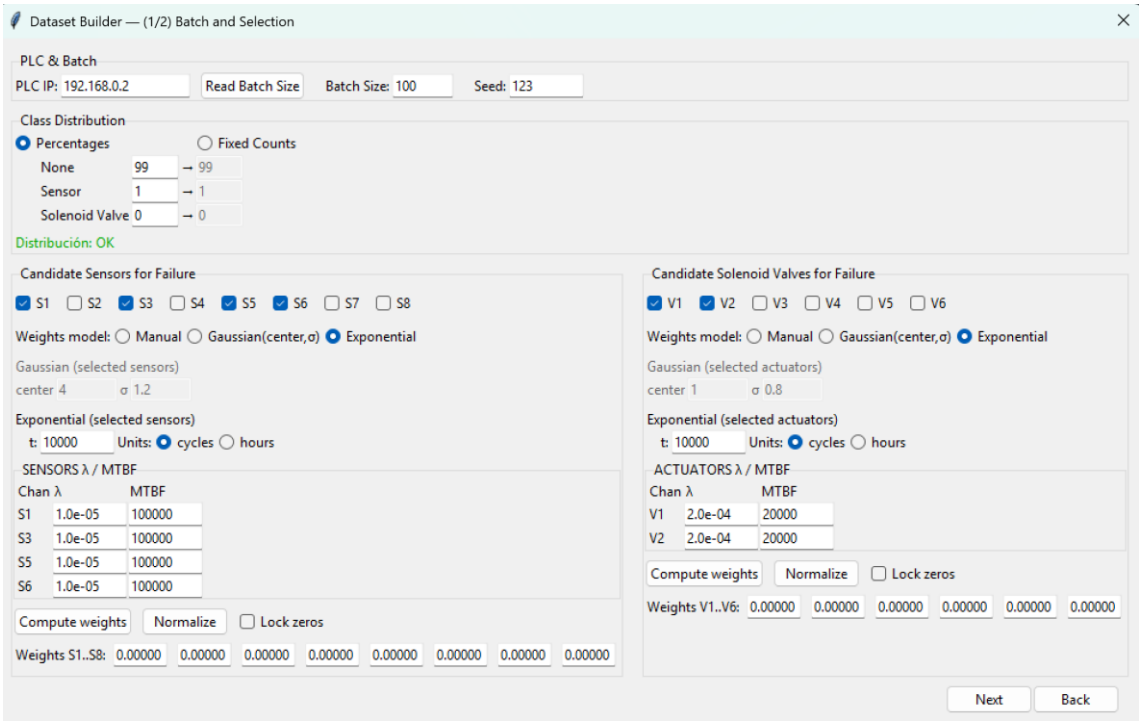


Figure 3.3. Dataset builder interface for batch configuration and fault injection candidate selection.

This figure shows the first stage of the dataset builder interface, which is used to define the experimental configuration for fault injection. In this stage, batch parameters, including the batch size and seed value, are specified to ensure reproducibility of the generated scenarios.

The interface enables the definition of class distribution, either through percentages or fixed counts, allowing a controlled proportion of normal and anomalous instances within each batch. This includes the specification of

anomalies affecting sensors and actuators, which can be independently configured.

To configure the occurrence of faults, the system provides different mechanisms, including manual weighting as well as parameterized probabilistic options such as Gaussian and exponential distributions. These approaches are parameterized through variables such as λ and MTBF, enabling flexible control over the frequency and distribution of anomalies.

Importantly, these mechanisms are not intended to model failure behavior itself, but to provide a configurable and reproducible way of generating diverse experimental conditions.

Overall, this stage defines the experimental space in which fault scenarios are generated, establishing the conditions under which anomalies may occur before specifying their detailed behavior.

Following the definition of the experimental configuration, the next stage focuses on the detailed specification of fault behavior, as illustrated in next Figure 3.4.

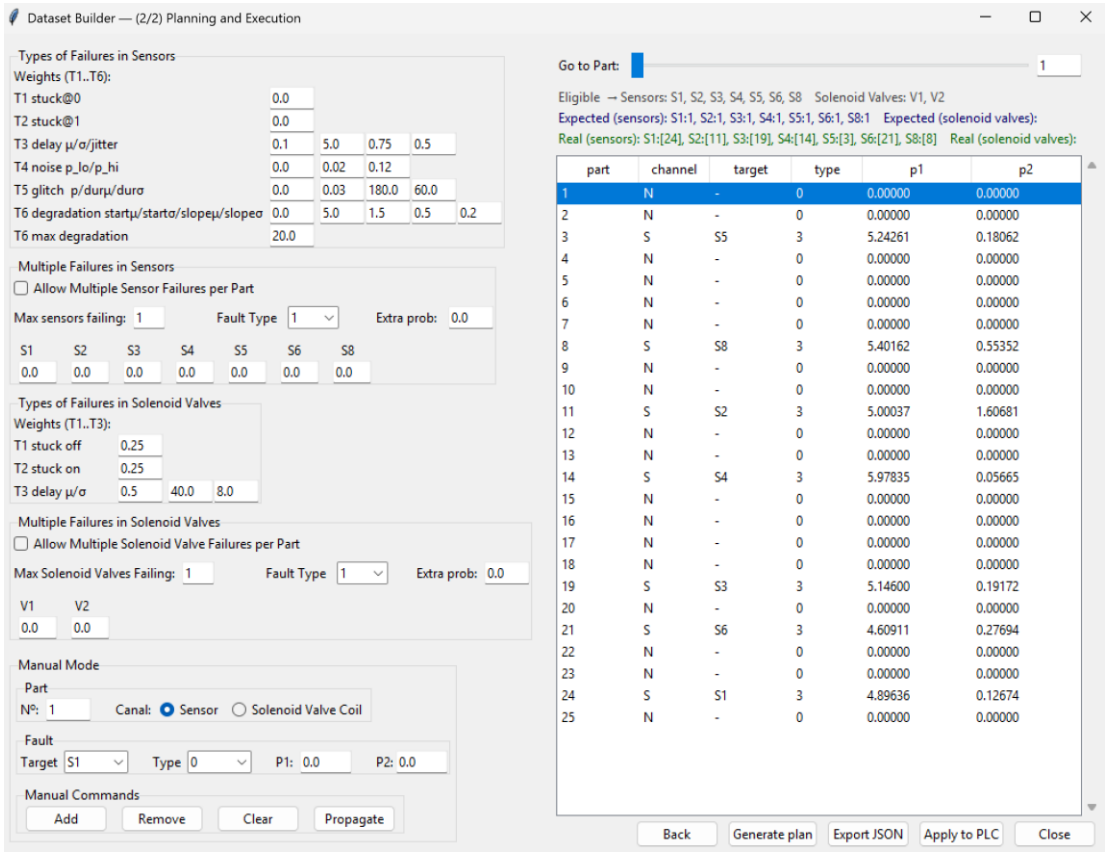


Figure 3.4. Dataset Builder interface for parametric fault definition and batch execution.

In this stage, fault scenarios are instantiated and scheduled through the planning interface, enabling the configuration of different failure types affecting sensors and actuators. Each fault is described by a set of parameters, allowing the representation of diverse behaviours such as delays, degradations, or transient perturbations.

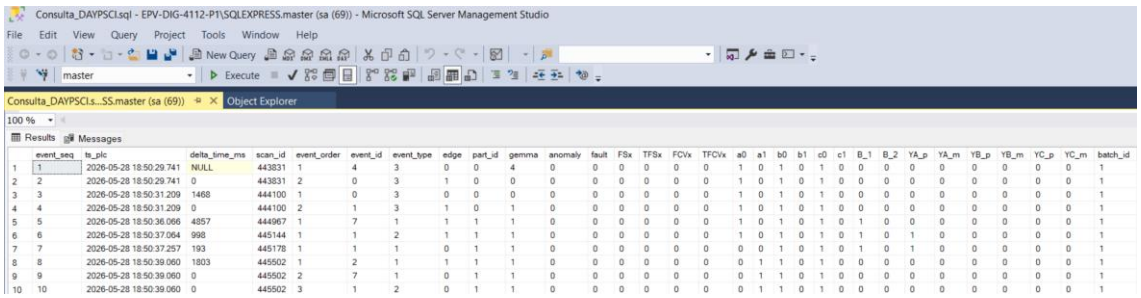
Faults are scheduled at the batch level, where each part in the production sequence is associated with a predefined fault configuration. This results in an explicit fault plan in which anomalies are assigned to specific components and time steps within the batch.

This approach ensures the systematic and reproducible generation of anomalous conditions across different executions.

The generated fault plan is subsequently applied during system operation, establishing a direct correspondence between planned scenarios and observed system behaviour. This ensures full traceability and supports the consistent construction of ground truth for dataset generation.

3.3. Event-Based Dataset Generation

The dataset generation process is based on integrating the data acquired during system operation with the corresponding ground truth information derived from the fault injection mechanism, combining planned configurations with runtime labelling within a hybrid ground truth framework, as depicted in Figure 3.1.



	event_seq	ts_plc	delta_time_ms	scan_id	event_order	event_id	event_type	edge	part_id	gamma	anomaly	fault	FSx	TFSx	FCVx	TFCVx	a0	a1	b0	b1	c0	c1	B_1	B_2	YA_p	YA_m	YB_p	YB_m	YC_p	YC_m	batch_id
1	1	2026-05-28 18:50:29.741	NULL	443831	1	4	3	0	0	4	0	0	0	0	0	0	1	0	1	0	1	0	0	0	0	0	0	0	0	0	1
2	2	2026-05-28 18:50:29.741	0	443831	2	0	3	1	0	0	0	0	0	0	0	0	1	0	1	0	1	0	0	0	0	0	0	0	0	0	1
3	3	2026-05-28 18:50:31.209	1468	444100	1	0	3	0	0	0	0	0	0	0	0	0	1	0	1	0	1	0	0	0	0	0	0	0	0	0	1
4	4	2026-05-28 18:50:31.209	0	444100	2	1	3	1	0	1	0	0	0	0	0	0	1	0	1	0	1	0	0	0	0	0	0	0	0	0	1
5	5	2026-05-28 18:50:36.066	4857	444967	1	7	1	1	1	1	0	0	0	0	0	0	1	0	1	0	1	0	1	0	0	0	0	0	0	0	1
6	6	2026-05-28 18:50:37.064	998	445144	1	1	2	1	1	1	0	0	0	0	0	0	1	0	1	0	1	0	1	0	1	0	0	0	0	0	1
7	7	2026-05-28 18:50:37.257	193	445178	1	1	1	0	1	1	0	0	0	0	0	0	0	0	1	0	1	0	1	0	1	0	0	0	0	0	1
8	8	2026-05-28 18:50:39.060	1803	445502	1	2	1	1	1	1	0	0	0	0	0	0	0	1	1	0	1	0	0	0	0	0	0	0	0	0	1
9	9	2026-05-28 18:50:39.060	0	445502	2	7	1	0	1	1	0	0	0	0	0	0	0	1	1	0	1	0	0	0	0	0	0	0	0	0	1
10	10	2026-05-28 18:50:39.060	0	445502	3	1	2	0	1	1	0	0	0	0	0	0	0	1	1	0	1	0	0	0	0	0	0	0	0	0	1

Figure 3.5. Event-based dataset stored in the SQL database with Δt and ground truth.

An example of the generated dataset is presented in Figure 3.5, illustrating the event-based representation where each record corresponds to an event-triggered state transition in the system.

For each event, both the system state and its associated timestamp are stored, enabling the computation of the time difference (Δt) between consecutive events.

The dataset includes information from sensors and actuators, along with process state indicators such as the GEMMA operational modes. In addition, ground truth labels are incorporated based on the predefined fault injection plan, enabling the identification of anomalous conditions at the event level.

This representation preserves the temporal dynamics of the process while avoiding redundancy associated with fixed-rate sampling. Furthermore, the direct correspondence between planned faults and recorded data ensures full traceability and consistency in dataset annotation.

A detailed description of the dataset structure, labelling strategy, and acquisition pipeline is provided in Chapter 6.

3.4. Chapter Summary

This chapter presented the experimental framework developed to support dataset generation and anomaly detection research within an industrial Cyber-Physical System environment. The proposed framework is based on a hybrid CPS testbed that integrates industrial hardware, PLC-based control, a digital twin, supervisory systems, and data acquisition mechanisms into a unified and flexible architecture. This infrastructure enables controlled experimentation under realistic operating conditions while ensuring repeatability and traceability throughout the data generation process.

The chapter also described the fault injection methodology adopted to generate labelled anomaly scenarios affecting sensors and actuators. By planning and scheduling faults at the batch level, the proposed approach enables the systematic and reproducible creation of experimental scenarios while maintaining a clear correspondence between injected anomalies and observed system behaviour. Furthermore, the dataset generation process was introduced, highlighting the integration of event-based data acquisition, temporal information (Δt), and ground-truth labelling to produce structured datasets suitable for data-driven analysis.

Together, these elements establish the experimental foundation of this research, providing a flexible and reproducible environment for the generation of event-

based industrial datasets. Building upon this foundation, the next chapter presents the first validation study of the proposed CPS testbed, assessing its capability to support anomaly detection through the generation of representative datasets and the evaluation of machine learning models.

Chapter 4

Testbed-Based Validation for Anomaly Detection

4.1. Introduction

The increasing integration of Cyber-Physical Systems within industrial environments has significantly enhanced the automation and efficiency of production processes. However, this integration has also introduced new challenges, particularly in terms of system reliability, fault detection, and cybersecurity. In this context, anomaly detection has emerged as a critical capability for ensuring the safe and efficient operation of Industrial Control Systems, enabling the early identification of abnormal behaviours associated with faults or potential cyber incidents[18, 33].

In previous chapters, a hybrid CPS testbed has been presented as a flexible and scalable experimental framework for research in anomaly detection and industrial cybersecurity. Specifically, Chapter 3 has detailed the testbed architecture, the integration of physical and virtual components, the implementation of fault injection mechanisms based on the GEMMA methodology, and the event-based dataset generation process supported by a SQL Server database. This experimental framework enables the controlled simulation of both normal and anomalous operating conditions, as well as the systematic acquisition of high-resolution process data.

Building upon this foundation, the objective of this chapter is to validate the proposed testbed as a reliable environment for the evaluation of machine learning techniques applied to anomaly detection in industrial CPS. To this end, a representative use case is defined, and a set of supervised learning models is evaluated using datasets generated within the testbed under controlled conditions. This validation focuses not only on the performance of the learning algorithms, but also on the suitability of the generated datasets for supporting anomaly detection tasks.

In contrast to the previous chapter, which focused on the design and implementation of the experimental framework, this chapter emphasizes its practical application and analytical capabilities. The testbed is therefore assessed from a functional perspective, demonstrating its ability to support the complete experimental workflow: from controlled data acquisition to model training, validation, and performance evaluation.

Furthermore, this chapter contributes to the objectives defined in this thesis. In particular, it provides experimental validation of the proposed framework by evaluating machine learning techniques for anomaly detection, while also demonstrating the capability of the testbed to generate valid datasets and support different industrial use cases.

The remainder of this chapter is structured as follows. Section 4.2 presents the definition of the selected use case. Section 4.3 describes the dataset employed for validation purposes. Section 4.4 details the methodology for anomaly detection, including preprocessing and model configuration. Section 4.5 presents the experimental results and provides a detailed interpretation of the findings. Finally, Section 4.6 summarizes the main conclusions of the chapter.

4.2. Use Case Definition

To validate the proposed experimental framework, a representative industrial use case has been defined within the CPS testbed. The selected process corresponds to a part marking system, chosen for its simplicity, reproducibility, and suitability for controlled anomaly generation. Despite its reduced complexity, it captures key characteristics of industrial automation systems, including sequential operation, interaction between sensors and actuators, and dependence on PLC-based control logic.

The system automates the marking of a batch of parts, including feeding, marking, and ejection stages. This sequence is coordinated through sensors and actuators, whose states evolve according to the control program and operating conditions. The process behaviour is governed by discrete state transitions, which are recorded in the dataset as event-based data.

The use case is implemented through a digital twin integrated within the testbed, enabling controlled simulation of an industrial marking process. The system includes three actuators (A, B, and C), position sensors, and presence detectors

for part tracking, corresponding to the industrial process represented in Chapter 3 (see Figure 3.2). This configuration defines a multi-stage industrial process in which correct operation depends on both sensor feedback and actuator performance.

A key feature of this use case is the controlled injection of anomalies at the level of sensors and actuators. In the experimental setup considered in this study, both sensor and actuator anomalies are implemented as stuck-at-0 faults, leading to persistent incorrect readings in sensors and inactive states in the electromechanical valves. This controlled approach enables precise and reliable labelling, ensuring a clear distinction between normal and faulty operation.

Based on this configuration, the anomaly detection task is formulated as a supervised classification problem, where each part is assigned to one of three classes: normal operation (Class 0), sensor anomalies (Class 1), and actuator anomalies (Class 2).

Each part is represented by a sequence of time-ordered events describing its full processing cycle, including sensor states, actuator commands, and system modes. This representation provides a detailed view of process behaviour.

Overall, this use case enables the evaluation of machine learning models under controlled but realistic industrial conditions, supporting the analysis of anomaly detection performance using event-based data.

4.3. Dataset Description for Validation

The validation of the proposed anomaly detection approach is carried out using datasets generated within the CPS testbed described in Chapter 3. As previously detailed, the dataset generation process enables the controlled acquisition of process data through an event-based logging mechanism supported by a SQL Server database, where all relevant changes in system variables are recorded as events.

Rather than focusing on the dataset generation procedure itself, this section characterizes the datasets employed for validation purposes and their suitability for machine learning-based anomaly detection.

The dataset is composed of discrete events associated with the operation of the industrial process. Each event corresponds to a change in the system state and

includes information such as timestamp, duration, sensor states, actuator commands, system operating mode (GEMMA), anomaly type, and part identifier. These variables provide a comprehensive description of both the logical and temporal dynamics of the process.

A key feature of the dataset is its event-based nature, which differs from conventional time-sampled industrial datasets. Instead of recording data at fixed time intervals, the dataset registers only relevant state transitions. This results in a compact yet information-rich representation, where the temporal evolution of the process is encoded through the sequence of events and the associated duration between them (Δt). This representation is particularly well aligned with industrial systems governed by discrete control logic.

For validation, a mixed dataset has been constructed by combining parts generated under normal operating conditions, sensor fault scenarios, and actuator fault scenarios. Accordingly, each part is associated with one of the anomaly classes defined in Section 4.2: normal operation (Class 0), sensor anomalies (Class 1), and actuator anomalies (Class 2). These datasets were generated under controlled experimental conditions, ensuring consistency, reproducibility, and precise labelling.

Importantly, the ground truth is generated outside the operational industrial process, enabling full control over anomaly definition and ensuring a clear and unambiguous distinction between normal and anomalous behaviours.

Each part is represented as a sequence of time-ordered events corresponding to its full processing cycle. Due to the variable duration and number of events associated with each part, the dataset consists of sequences with different lengths, reflecting the inherent dynamics of the process, particularly under anomalous conditions.

To enable the application of supervised learning techniques, the dataset is divided into three subsets following a standard protocol: training set (70%), validation set (15%), and test set (15%), ensuring a robust and unbiased evaluation of model performance.

From a data quality perspective, the dataset offers several advantages: controlled anomaly injection enables precise labelling, the inclusion of both normal and faulty behaviours supports realistic evaluation scenarios, and the sequential structure preserves temporal dependencies critical for anomaly detection.

Overall, the dataset provides a consistent, well-structured, and representative basis for evaluating machine learning techniques in industrial CPS, supporting the validation of the proposed framework.

4.4. Methodology for Anomaly Detection

This section describes the methodology used to identify anomalies in the industrial process defined in Section 4.2, based on the datasets introduced in Section 4.3. The objective is to evaluate the performance of supervised machine learning techniques in classifying industrial parts according to their operational behaviour.

The anomaly detection problem is formulated as a multi-class classification task, where each part is assigned to one of three classes: normal operation (Class 0), sensor anomalies (Class 1), and actuator anomalies (Class 2). This formulation enables the evaluation of the ability of machine learning models to detect and distinguish different types of faults.

The methodology follows a structured pipeline consisting of data preprocessing, model training, and performance evaluation.

4.4.1. Data Preprocessing and Representation

The dataset consists of event-based sequences representing the evolution of the industrial process. Each part is described by a sequence of time-ordered events including sensor states, actuator commands, and operating modes (GEMMA).

Due to the variable length of these sequences, a transformation into a uniform representation is required. The preprocessing stage includes:

- ***Grouping by part:*** events corresponding to each part are aggregated to represent its full processing cycle.
- ***Sequence padding***: the maximum sequence length is determined from the dataset, and shorter sequences are padded with zeros to obtain a uniform representation.
- ***Feature scaling***: variables are normalized to ensure balanced contribution during model training.

This transformation enables the application of conventional supervised learning models by converting variable-length event sequences into a uniform feature

representation while preserving information about the temporal evolution of the process.

4.4.2. Model Selection and Configuration

Three supervised learning techniques are evaluated:

- *K-Nearest Neighbours (KNN)*
- *Random Forest (RF)*
- *Support Vector Machine (SVM)*

These models provide a representative comparison between instance-based, ensemble, and kernel-based approaches.

Model parameters are initially selected based on common configurations and subsequently refined using grid search with cross-validation. The final configurations correspond to those achieving the best balance between accuracy and generalization.

4.4.3. Evaluation Metrics

Model performance is assessed using standard classification metrics:

- ***Accuracy***: proportion of correctly classified instances.
- ***Precision***: proportion of correct positive predictions.
- ***Recall***: proportion of correctly identified positive instances.
- ***F1-score***: balance between precision and recall.
- ***AUC-ROC****:* ability to discriminate between classes.

These metrics provide a comprehensive evaluation of model performance, including robustness in detecting different types of anomalies and handling potential class imbalance.

4.5. Results and Discussion

This section presents the experimental results and provides a detailed interpretation of the performance of the evaluated machine learning models in the context of industrial anomaly detection. The performance of the three selected classifiers, K-Nearest Neighbours (KNN), Random Forest (RF), and Support

Vector Machine (SVM), was assessed using the test dataset following the methodology previously described.

4.5.1. Comparative Performance of Models

The overall performance of the evaluated models is summarized in Figure 4.1, which presents radar plots comparing different configurations of KNN, Random Forest, and SVM across the selected evaluation metrics.

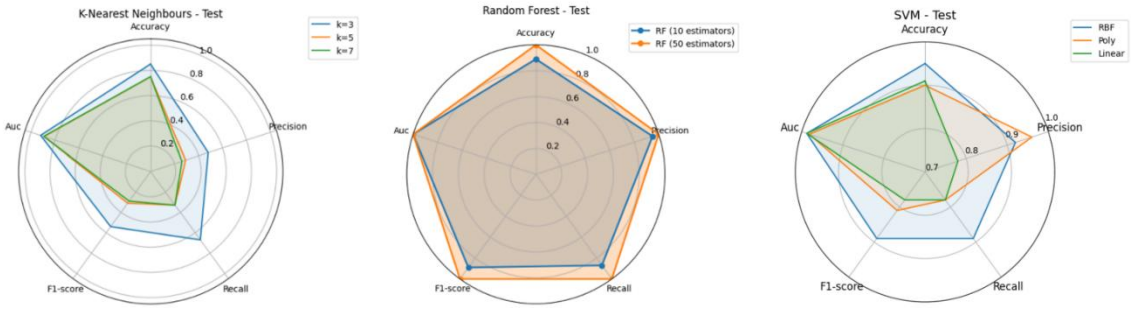


Figure 4.1. Radar plot comparison of KNN, Random Forest, and SVM performance.

The results show that KNN achieves its best performance with k = 3, while higher values (k = 5 and k = 7) lead to a slight degradation in performance. Random Forest demonstrates consistently strong performance across all configurations, with improvements in recall and F1-score as the number of estimators increases. SVM, particularly with the RBF kernel, achieves competitive results, outperforming other kernel configurations in most metrics.

To provide a clearer quantitative comparison, Table 4.1 summarizes the best-performing configuration of each model along with its corresponding test-set results. Random Forest achieves perfect performance across all evaluation metrics, while SVM shows strong performance with slightly lower recall and F1-score. In contrast, KNN exhibits significantly lower performance, particularly in precision and F1-score.

Table 4.1. Summary of classification results

Technique	Accuracy	Precision	Recall	F1-score	AUC-ROC
KNN (k=3)	0.85	0.48	0.67	0.54	0.92
RF (n=50)	1.0	1.0	1.0	1.0	1.0
SVM (RBF)	0.95	0.92	0.89	0.89	0.99

These results indicate that ensemble-based methods such as Random Forest are particularly well suited for the classification of event-based industrial data, while instance-based methods such as KNN are more sensitive to variability and high-dimensional representations.

4.5.2. Class Separability Analysis (ROC)

To further analyse the discriminative capability of the models, multiclass ROC curves were computed using a one-vs-rest approach, as shown in Figure 4.2.

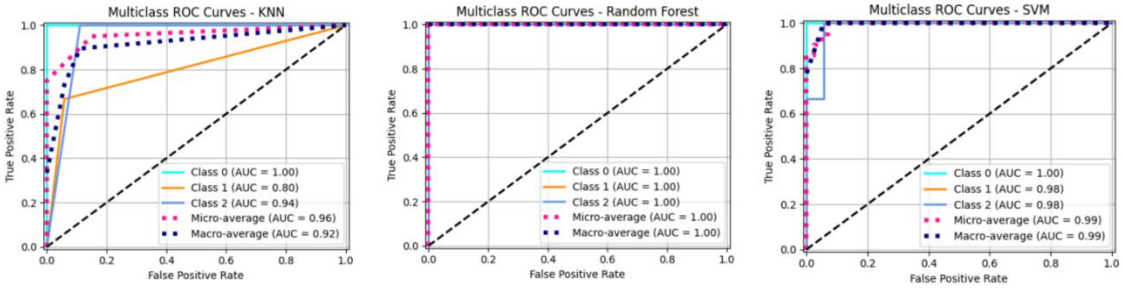


Figure 4.2. Multiclass ROC curves for KNN, Random Forest, and SVM.

The ROC analysis confirms that Random Forest achieves perfect class separation, with an AUC-ROC value of 1.0 across all classes. SVM provides near-perfect discrimination, with very high AUC values, confirming its robustness. KNN, although competitive in terms of AUC-ROC, exhibits less consistent separation, particularly for minority classes.

These results reinforce the robustness of Random Forest not only in terms of accuracy but also in its ability to maintain strong discrimination across all anomaly classes, making it particularly reliable for industrial anomaly detection tasks.

4.5.3. Analysis of Model Behaviour

The superior performance of the Random Forest model can be explained by both the characteristics of the algorithm and the structure of the dataset. Random Forest is particularly effective at modelling non-linear relationships and decision boundaries derived from combinations of discrete process variables. In this

study, the dataset encodes structured sequences of discrete events representing meaningful transitions in the industrial process, which provide discriminative patterns for distinguishing between normal operation, sensor anomalies, and actuator anomalies.

In contrast, KNN shows limitations due to its reliance on distance-based similarity measures, which may not adequately capture the complex event patterns associated with different fault conditions. SVM demonstrates strong overall performance but presents slightly lower recall and F1-score compared to Random Forest, indicating minor limitations in detecting all anomalous cases.

Overall, these findings suggest that Random Forest is the most suitable approach among the evaluated models for the classification of event-based industrial data generated in the proposed use case, as it effectively captures complex event patterns while maintaining excellent classification performance across all classes.

4.5.4. Impact of Dataset Design (event-based, Δt)

A key factor contributing to the performance of the evaluated models is the event-based nature of the dataset. Unlike conventional time-sampled datasets, the proposed representation captures only relevant state transitions, each associated with a duration (Δt).

This results in a compact and information-rich representation that preserves both the logical and temporal behaviour of the industrial process. By directly reflecting the execution of the control logic, the dataset enables models to learn meaningful patterns with reduced redundancy.

The results obtained, particularly with Random Forest and SVM, suggest that this representation effectively captures the essential dynamics of the system and supports reliable anomaly detection.

4.5.5. Role of GEMMA and process dynamics

The integration of GEMMA operational modes into the dataset plays an important role in enhancing model performance. By incorporating information about system states, such as normal operation, fault conditions, and recovery phases, the dataset provides contextual information that complements sensor and actuator data.

This allows machine learning models to distinguish between normal transitions associated with operational modes and transitions caused by faults. As a result, the dataset represents a structured sequence of process states and operational modes, improving both pattern recognition and generalization capability.

4.5.6. Validation of the CPS Framework

The results obtained in this chapter provide convincing evidence supporting the validity of the proposed CPS-based framework for anomaly detection.

From a system perspective, the hybrid testbed enables the reproducible configuration of industrial scenarios under both normal and faulty conditions. The integration of real industrial control systems with a digital twin allows realistic behaviour to be reproduced while maintaining full experimental control.

From a data perspective, the framework supports the generation of structured and representative datasets, including both normal operation and controlled anomaly scenarios. The use of event-based acquisition and hybrid ground-truth generation ensures data reliability, traceability, and clear interpretation.

From a machine learning perspective, the high performance achieved by Random Forest indicates that the dataset contains sufficiently discriminative features to distinguish between normal and anomalous behaviours.

Overall, these findings validate the proposed framework as a robust and effective approach for dataset generation and anomaly detection evaluation in industrial CPS, providing a solid foundation for its extension to more complex and cybersecurity-oriented scenarios in subsequent chapters.

4.6. Chapter Summary

This chapter has presented the validation of the proposed CPS-based testbed as a practical and effective framework for anomaly detection in industrial environments. A representative use case based on a part marking system has been defined, and a structured methodology has been applied to evaluate the performance of supervised machine learning techniques using datasets generated under controlled conditions.

The results demonstrate that the proposed framework enables the generation of structured, reliable, and well-labelled datasets, supporting a robust evaluation of

anomaly detection models. Among the evaluated approaches, Random Forest achieved superior performance across all metrics, highlighting its suitability for classification tasks involving event-based industrial data. Support Vector Machine also showed strong performance, while KNN achieved comparatively lower results, indicating a reduced ability to capture the complex patterns present in the dataset.

The integration of event-based dataset design and the incorporation of temporal information through Δt were shown to be key factors in capturing the dynamics of the industrial process. Additionally, the inclusion of GEMMA operational modes provided contextual information that improved the ability of machine learning models to distinguish between normal operation and anomalous behaviour.

Overall, the findings confirm that the proposed framework effectively supports the complete workflow of anomaly detection in industrial CPS, from controlled data acquisition to model evaluation. This chapter provides a solid experimental validation of the framework, establishing a foundation for its extension to more complex scenarios, including cybersecurity-oriented applications, which are addressed in the following chapter.

Chapter 5

Cybersecurity-Oriented Extension of the CPS Framework

5.1. Introduction

In Chapter 4, the proposed CPS-based framework was validated for the detection of process anomalies using machine learning techniques, demonstrating its capability to generate structured datasets and support data-driven analysis.

However, modern industrial environments are increasingly exposed to cybersecurity threats that go beyond traditional process faults. The integration of industrial networks, standardized communication protocols, and interconnected control systems has significantly expanded the attack surface of Cyber-Physical Systems [6, 24].

To address this challenge, this chapter extends the previously validated experimental framework by incorporating a cybersecurity dimension into the methodological pipeline. In particular, an additional stage corresponding to cyberattack execution is introduced, enabling controlled adversarial interaction with the CPS.

The focus is placed on the analysis of cybersecurity threats affecting industrial CPS, where communication protocols such as PROFINET are used to exchange information between sensors, controllers, and actuators. To support this analysis, attack mechanisms based on a logical Man-in-the-Middle (MitM) model are incorporated, allowing the injection of malicious modifications without disrupting network-level communication.

This enables the study of semantic manipulation, where the process behaviour is altered while maintaining communication integrity, impacting both data acquisition and control execution within the CPS loop.

Through this approach, the chapter extends the original framework into a unified experimental environment that integrates anomaly detection and cybersecurity

analysis, supporting multi-layer evaluation across process, control, and communication levels.

5.2. Motivation: Cybersecurity in Industrial CPS

The increasing digitalization of industrial environments has significantly transformed the landscape of Cyber-Physical Systems, particularly within Industrial Control Systems. The adoption of standardized communication protocols, the integration of distributed control architectures, and the convergence of operational technology (OT) with information technology (IT) networks have introduced new capabilities, while simultaneously raising critical cybersecurity challenges [6, 15].

In this context, industrial communication protocols such as PROFINET have become widely adopted due to their real-time capabilities and deterministic behaviour. However, many of these protocols were originally designed with a primary focus on performance and interoperability rather than security. As a result, they often lack native mechanisms for authentication, encryption, or integrity verification, making them vulnerable to malicious manipulation [7, 9, 26].

This vulnerability is further exacerbated by the increasing exposure of industrial networks driven by remote access requirements, system integration, and Industry 4.0 initiatives [5, 15]. Consequently, the attack surface of CPS environments has expanded considerably, enabling adversaries to exploit both network-level and process-level weaknesses [6, 7].

From a regulatory and standardization perspective, frameworks such as the NIS2 Directive and the IEC 62443 series emphasize the need to incorporate cybersecurity measures into industrial systems. These initiatives highlight the importance of monitoring, detection, and response mechanisms capable of addressing advanced threats in CPS environments [34].

Despite these efforts, a key challenge remains: many existing cybersecurity approaches focus primarily on network traffic analysis, potentially overlooking attacks that do not disrupt communication patterns but instead alter the semantics of the industrial process [6, 7].

This limitation motivates the need for experimental frameworks capable of analysing cybersecurity threats from a multi-layer perspective, combining network-level observation with process-level understanding. In this regard, the CPS-based testbed presented in this thesis enables the systematic study of such threats, particularly those targeting the logical behaviour of industrial systems.

5.3. Logical MitM Model and Threat Definition

Building upon these limitations, a logical Man-in-the-Middle (MitM) model is defined to enable the controlled study of cybersecurity threats that cannot be detected through conventional network-level analysis.

Rather than modifying the CPS architecture introduced in Chapter 3, the proposed approach exploits the existing control and communication structure to manipulate process variables at the logical level.

This manipulation targets the semantic interpretation of process data, as it alters the meaning of system states without modifying the underlying communication behaviour.

Figure 5.1 illustrates the proposed logical Man-in-the-Middle (MitM) model, comparing the normal operation of the CPS with the adversarial scenario.

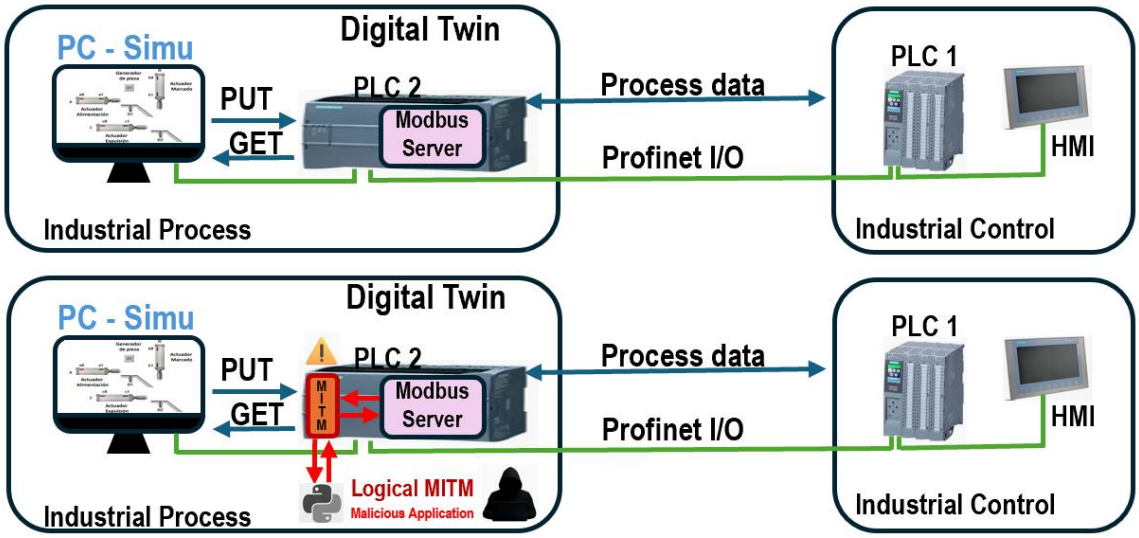


Figure 5.1. Communication flow between PCSimu, PLC2, and PLC1 under normal and MitM attack.

In this model, the attacker is not positioned within the physical communication infrastructure but instead acts through a controlled manipulation layer

implemented in the secondary controller (PLC2). This design allows the interception, modification, and injection of process variables without altering the underlying network communication.

From the perspective of the system, all communications over the PROFINET network remain valid and consistent with expected timing and protocol behaviour. Therefore, traditional network-based detection mechanisms are unable to identify the presence of the attack.

The attacker is assumed to have logical access to selected process variables, enabling the modification of sensor inputs (digital inputs, DI) and actuator commands (digital outputs, DO). This access is implemented through industrial communication interfaces such as Modbus, which provide a controlled channel for variable manipulation within the experimental environment.

Based on this model, the attack strategy focuses on altering the semantic meaning of the exchanged data rather than disrupting communication itself. This class of attacks, referred to as semantic manipulation, targets the interpretation of system states by the control logic, leading to incorrect process behaviour while maintaining normal network operation.

From a threat modelling perspective, the proposed approach assumes: (i) the attacker has access to the logical layer of the system but does not compromise the physical network infrastructure, (ii) communication protocols operate normally without packet loss, delay anomalies, or malformed traffic, and (iii) the control program of the primary controller (PLC1) is not modified.

Under these conditions, the attack remains stealthy at the network level while producing observable effects in the physical process. This distinction between communication integrity and process integrity is a key aspect of the proposed model and serves as the foundation for the analysis presented in the following sections.

This model enables reproducible and controlled experimentation, facilitating the evaluation of detection strategies that go beyond conventional network-based approaches.

The following section builds upon this model by applying it to the industrial use case, enabling the analysis of its impact under adversarial conditions.

5.4. Industrial Use Case Under Attack

The cybersecurity analysis presented in this chapter is conducted using the same industrial use case introduced in Chapter 4, namely a part-marking system defined within the CPS-based experimental framework. While Chapter 4 focused on anomaly detection and the validation of machine learning models, this chapter analyses the same industrial process under adversarial conditions, ensuring methodological consistency and enabling direct comparison between normal operation and attack scenarios.

The process involves a sequence of operations controlled by a programmable logic controller (PLC), integrating sensor input information with actuator output commands. Digital inputs (DI) are used to capture the status of key elements such as part presence and positioning, while digital outputs (DO) control actuators responsible for mechanical movements and marking operations.

The system behaviour is structured according to the GEMMA model, where different operational states define the progression of the industrial process, including automatic operation, fault conditions, and recovery phases. This logical structure provides a well-defined framework for analysing deviations induced by external manipulations.

This structured representation of system states is particularly relevant for identifying inconsistencies between expected and observed behaviour under attack conditions.

From a cybersecurity perspective, the relevance of this use case lies in the tight coupling between sensor inputs, control logic, and actuator responses. This strong interdependency makes the system particularly susceptible to semantic manipulation attacks, where the meaning of process data is altered without modifying the underlying communication.

As a result, any alteration in the value of these process variables can lead to incorrect process execution, even if communication between system components remains unchanged.

In the context of the proposed logical MitM model, the attack surface is defined by the set of process variables accessible through the manipulation layer. In particular, sensor signals (DI) and actuator commands (DO) are identified as

critical points where adversarial modifications can produce observable effects in the system behaviour.

By maintaining the original control logic and physical configuration unchanged, this use case provides a realistic and controlled environment to evaluate the impact of semantic manipulation attacks on industrial CPS.

The following section builds upon this scenario to define specific attack cases targeting sensors and actuators.

5.5. Definition of Attack Scenarios

Based on the logical MitM model defined in Section 5.3, a set of controlled attack scenarios is established to evaluate the impact of adversarial manipulation on the industrial process.

The objective of these scenarios is to systematically analyse how controlled modifications at the level of process variables affect system behaviour, while maintaining normal communication conditions. To this end, the attack design focuses on the manipulation of selected signals within the control loop, specifically targeting sensor inputs and actuator commands.

Two primary categories of attack scenarios are defined. The first category involves the manipulation of digital input signals (DI), where the attacker alters the perceived state of the system by injecting incorrect sensor values. The second category focuses on digital output signals (DO), where the attacker interferes with actuator commands, leading to unintended or unsafe system actions.

Each scenario is designed to introduce controlled deviations from the expected process sequence defined by the GEMMA model. These deviations may result in incorrect state transitions, improper actuation, or inconsistencies between the physical process and the control logic.

Importantly, all attack scenarios are executed within the constraints of the proposed threat model: the communication network remains unaffected, the control program of the primary PLC is not modified, and all manipulations are performed through the logical interaction layer provided by the secondary controller.

This approach enables reproducible and isolated experimentation, where the effects of each attack type can be analysed independently. The following

subsections describe in detail the implementation of sensor-level and actuator-level attack scenarios.

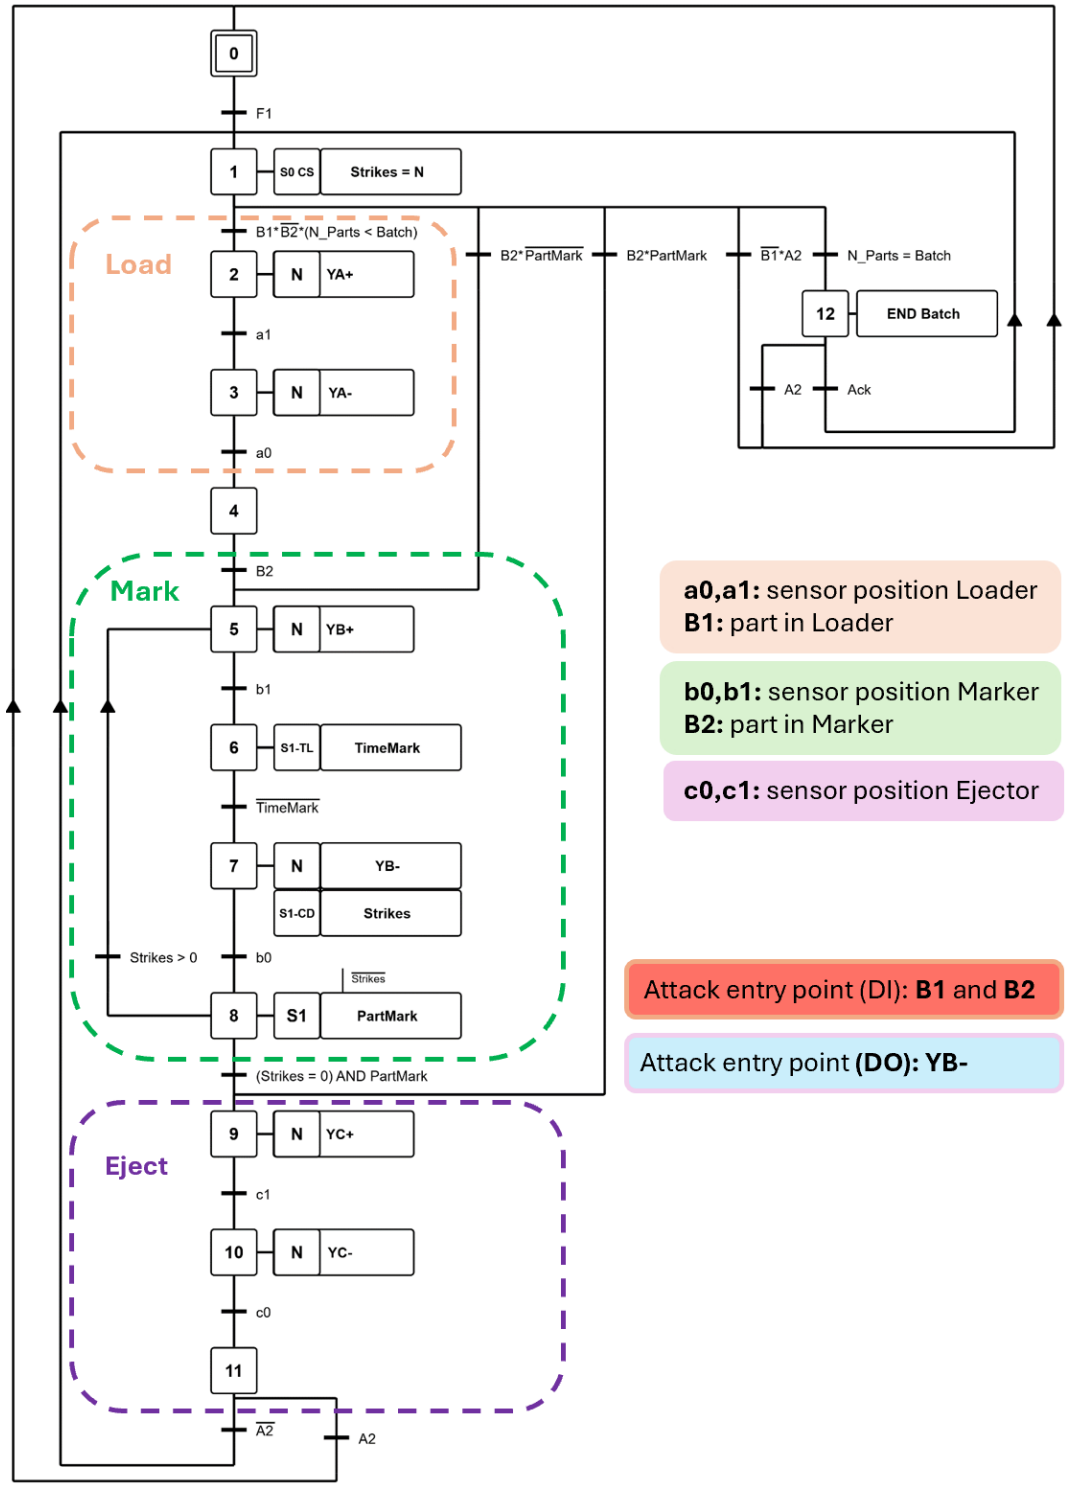


Figure 5.2. Industrial process workflow highlighting attack entry points at sensor (DI) and actuator (DO) levels.

Figure 5.2 presents the industrial process workflow, including the main control states, sensor inputs, and actuator outputs involved in the system. The process is

structured into three main stages: loading, marking, and ejection, each governed by a sequence of logical transitions based on the GEMMA model, with critical attack entry points explicitly identified within the control loop.

From a cybersecurity perspective, this representation enables the identification of key variables susceptible to manipulation. In particular, sensor signals related to part detection (e.g., B1 and B2) and actuator commands associated with the marking operation (e.g., YB−) are highlighted as primary attack targets.

These entry points establish a direct mapping between the logical MitM model and its observable impact on the industrial process, supporting the systematic definition of the attack scenarios considered in this work.

5.5.1. Scenario A1 — Sensor Manipulation (DI)

The first attack scenario focuses on the manipulation of digital input signals (DI), aiming to alter the perceived state of the industrial process without modifying its physical configuration or communication infrastructure.

In this scenario, the attacker targets specific sensor signals that provide critical information to the control logic regarding the state of the system. These signals typically indicate conditions such as part presence, position verification, or operation completion, and are essential for ensuring the correct progression of the process sequence.

Through the logical MitM model, the attacker modifies the value of selected DI signals before they are processed by the primary controller (PLC1), as illustrated in Figure 5.3. As a result, the control program receives incorrect information about the system state, leading to decisions based on false perceptions rather than actual physical conditions.

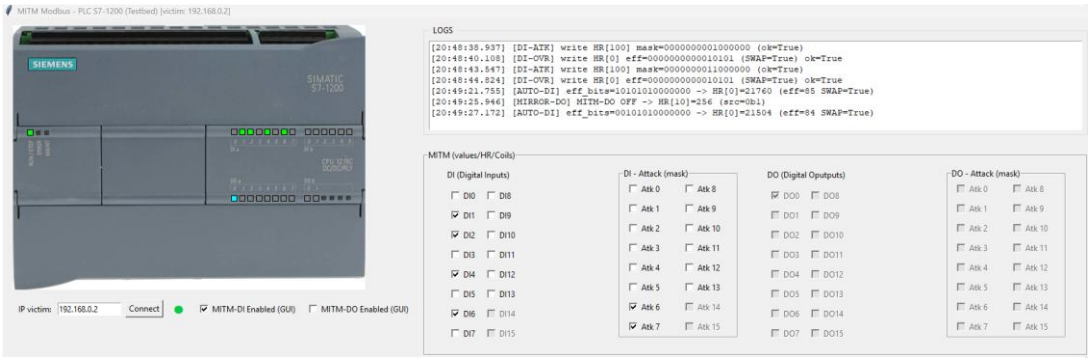


Figure 5.3. DI-bit manipulation (B1 and B2) in PLC 2.

This type of manipulation directly affects the semantic interpretation of the process. For example, a sensor indicating the presence of a part may be forced to an active state even when no part is available, or conversely, a valid signal may be suppressed. These alterations lead to inconsistencies in the execution of the GEMMA-based sequence, potentially triggering premature transitions, skipped operations, or unintended states. In the experimental implementation, this scenario is exemplified by the manipulation of part-detection signals (e.g., B1 and B2), enabling the controller to execute the production cycle despite the absence of a physical workpiece.

Importantly, the physical behaviour of the system remains unchanged at the sensor level, meaning that the actual environment does not correspond to the values processed by the controller. However, since the communication over the PROFINET network is not altered, the transmitted data appears valid from a protocol perspective.

From a cybersecurity standpoint, this scenario exemplifies a semantic manipulation attack, where the integrity of the communication is preserved while the meaning of the transmitted information is compromised. As a consequence, conventional network-based monitoring tools are unable to detect the attack, whereas its effects become evident only through deviations in process behaviour.

This controlled manipulation of sensor inputs provides a reproducible mechanism to analyse how incorrect state perception impacts system operation, serving as a foundation for evaluating detection strategies based on process-aware or data-driven approaches.

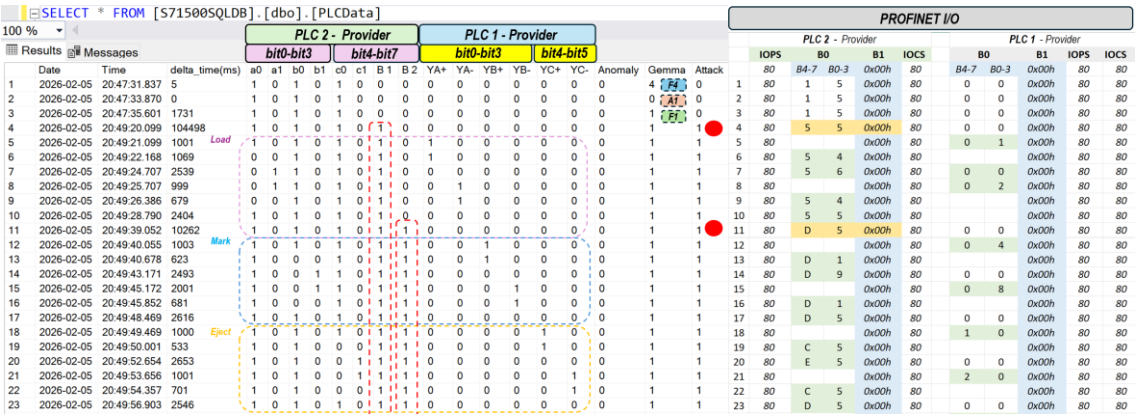


Figure 5.4. Correlation of SQL Server, GEMMA and PROFINET data in scenario A1

The effects of this manipulation can be observed in the dataset recorded at the process level in SQL Server, together with the corresponding communication exchange. Figure 5.4 shows the evolution of key process variables, including sensor inputs, control states according to the GEMMA model, and PROFINET I/O data, highlighting the inconsistencies introduced by the attack.

While the effects of the attack are clearly observable at the process and data level, the underlying network communication remains unaffected.

Figure 5.5 shows the PROFINET communication during the attack, where all frames maintain a valid structure and remain consistent with expected protocol operation, without any indication of anomalous behaviour at the network level.

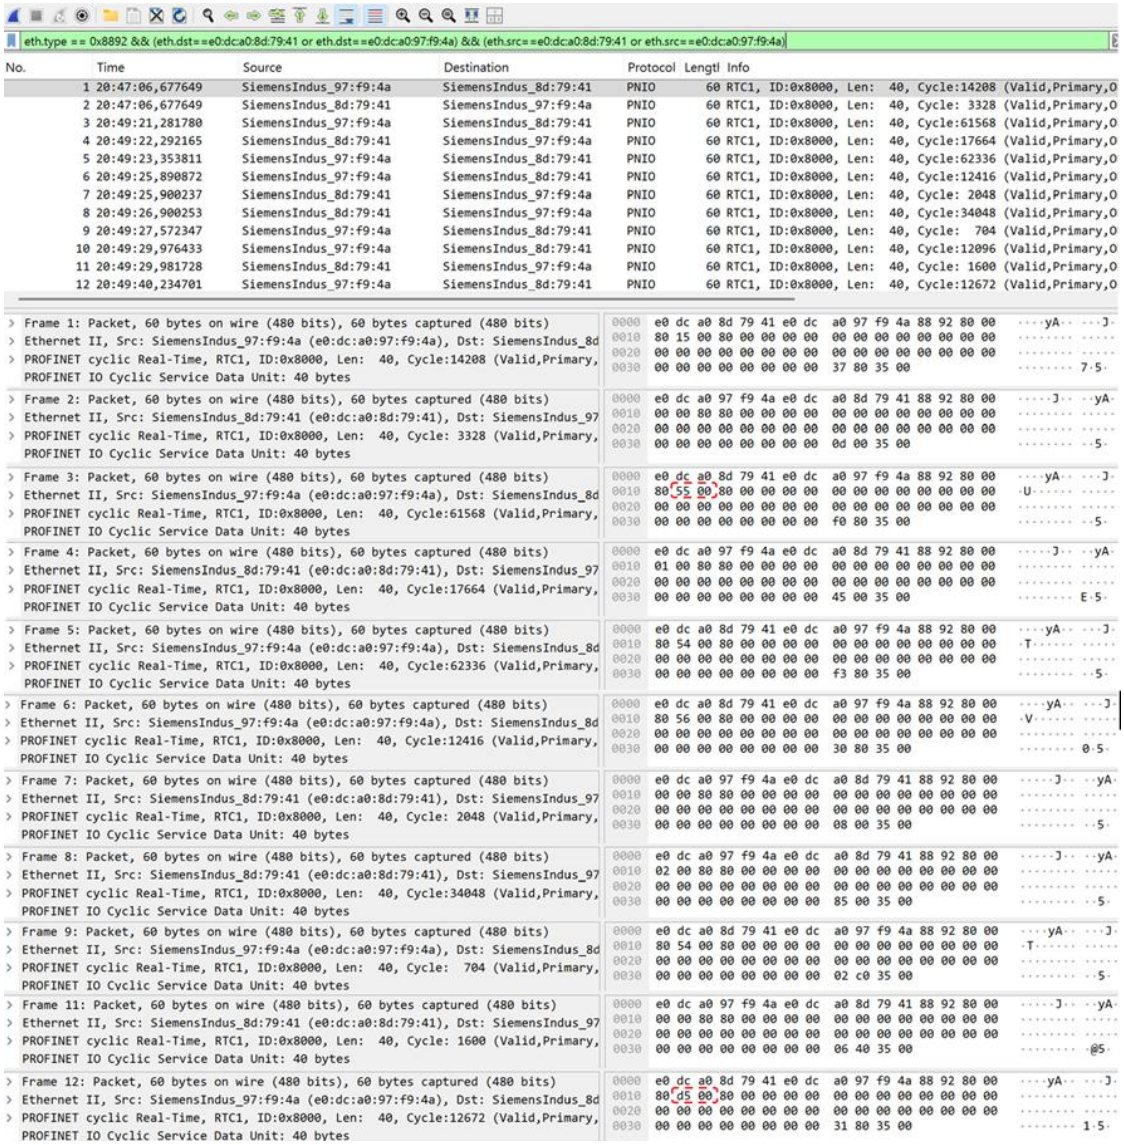


Figure 5.5. Wireshark capture of PROFINET RTC1 traffic in scenario A1

5.5.2. Scenario A2 — Actuator Manipulation (DO)

The second attack scenario focuses on the manipulation of digital output signals (DO), aiming to directly alter the behaviour of the industrial process by interfering with actuator commands.

In this scenario, the attacker targets control signals responsible for driving physical actions within the system, such as the activation or deactivation of actuators involved in the marking operation. In particular, commands associated with the marking stage (e.g., YB-) are identified as critical variables through which the process behaviour can be modified.

As shown in the previously presented attack interface (Figure 5.3), the manipulation capabilities extend to both digital inputs and outputs, enabling direct intervention in actuator-level commands.

Unlike the sensor manipulation scenario, where the attack affects the perceived state of the system, this approach directly impacts the execution of physical actions. Through the logical MitM model, the attacker modifies actuator commands before they are executed or selectively overrides them, leading to discrepancies between the intended control logic and the actual system behaviour.

As a result, the system may perform unintended actions, such as incomplete marking operations, incorrect actuator movements, or deviations in the expected process sequence defined by the GEMMA model (see Figure 5.6).

SQLQuery1.sql - EPV...SS.master (sa (66))* | Object Explorer

```
SELECT * FROM [S71500SQLDB].[dbo].[PLCDataset_v1]
```

100 %

Results | Messages

				PLC 2 - Provider								PLC 1 - Provider								
				Bit0-bit3				Bit4-bit7				Bit0-bit3				Bit4-bit5				
	date	time	delta_time(ms)	a0	a1	b0	b1	c0	c1	B 1	B 2	YA+	YA-	YB+	YB-	YC+	YC-	Anomaly	Gemma	Attack
1	2026-03-15	13:42:45.168	5	1	0	1	0	1	0	0	0	0	0	0	0	0	0	0	4 F4	0
2	2026-03-15	13:42:47.795	0	1	0	1	0	1	0	0	0	0	0	0	0	0	0	0	0 A1	0
3	2026-03-15	13:42:52.370	4575	1	0	1	0	1	0	0	0	0	0	0	0	0	0	0	1 F1	0
4	2026-03-15	13:43:19.087	26718	1	0	1	0	1	0	1	0	0	0	0	0	0	0	0	1	0
5	2026-03-15	13:43:20.089	1001	1	0	1	0	1	0	1	0	1	0	0	0	0	0	0	1	0
6	2026-03-15	13:43:20.482	393	0	0	1	0	1	0	1	0	1	0	0	0	0	0	0	1	0
7	2026-03-15	13:43:21.890	1407	0	1	1	0	1	0	0	0	0	0	0	0	0	0	0	1	0
8	2026-03-15	13:43:22.888	999	0	1	1	0	1	0	0	0	0	1	0	0	0	0	0	1	0
9	2026-03-15	13:43:23.282	394	0	0	1	0	1	0	0	0	0	1	0	0	0	0	0	1	0
10	2026-03-15	13:43:24.649	1368	1	0	1	0	1	0	0	0	0	0	0	0	0	0	0	1	0
11	2026-03-15	13:43:35.392	10741	1	0	1	0	1	0	0	1	0	0	0	0	0	0	0	1	2
12	2026-03-15	13:43:36.392	1002	1	0	1	0	1	0	0	1	0	0	1	0	0	0	0	1	2
13	2026-03-15	13:43:49.398	13005	1	0	1	0	1	0	0	1	0	0	1	0	0	0	0	7 D2	2
14	2026-03-15	13:43:51.577	2179	1	0	1	0	1	0	1	1	0	0	0	0	0	0	0	7	2
15	2026-03-15	13:43:56.025	2179	1	0	1	0	1	0	1	1	0	0	0	0	0	0	0	7	2
16	2026-03-15	13:43:56.233	4454	1	0	1	0	1	0	1	1	0	0	0	0	0	0	0	5 A5	2

Figure 5.6. Correlation of SQL Server, GEMMA and PROFINET data in scenario A2.

In this case, the control logic in the primary controller (PLC1) continues to operate based on valid internal states, but the physical response of the system does not reflect those intended actions.

From a cybersecurity perspective, this scenario represents a direct manipulation of the actuation layer, where the integrity of control actions is compromised while communication behaviour remains unchanged. Similar to the previous case, the PROFINET network remains unaffected, and all exchanged data conforms to the expected protocol behaviour.

This type of attack reinforces the concept of semantic manipulation, as the relationship between control decisions and physical execution is disrupted. Consequently, the attack remains undetectable at the network level and can only be identified through inconsistencies in process behaviour or deviations from expected operational patterns.

The controlled manipulation of actuator outputs complements the sensor-level scenario, providing a comprehensive basis for evaluating the impact of adversarial actions on both perception and execution layers within the CPS framework.

Together, both scenarios illustrate the dual nature of semantic attacks in CPS environments, affecting both system perception and physical actuation.

The monitoring infrastructure described in this chapter enables the generation of synchronized multi-layer datasets, combining process-level variables and communication data. These datasets extend the experimental basis established in Chapter 4 for the analysis and detection of anomalies in industrial CPS environments.

5.6. Communication-Level Validation

This section evaluates whether the execution of the proposed attack scenarios affects the stability of the communication channel. As the attack focuses on manipulating process data rather than altering the communication infrastructure itself, it is important to verify that network-level indicators remain unaffected.

To this end, four conditions were analysed:

- **Baseline**: Normal system operation without attack.
- **Attack-Idle**: MitM mechanism active without data manipulation.

- **Attack-DI**: Sensor-level manipulation corresponding to Scenario A1.
- **Attack-DO**: Actuator-level manipulation corresponding to Scenario A2.

5.6.1. Round-Trip Time (RTT)

RTT measurements were obtained from ICMP ping tests performed under baseline and attack conditions. The results indicate that communication latency remained stable across all evaluated scenarios.

Figure 5.7 presents the mean RTT values and their variability across baseline and attack scenarios, showing that communication latency remained stable under all tested conditions.

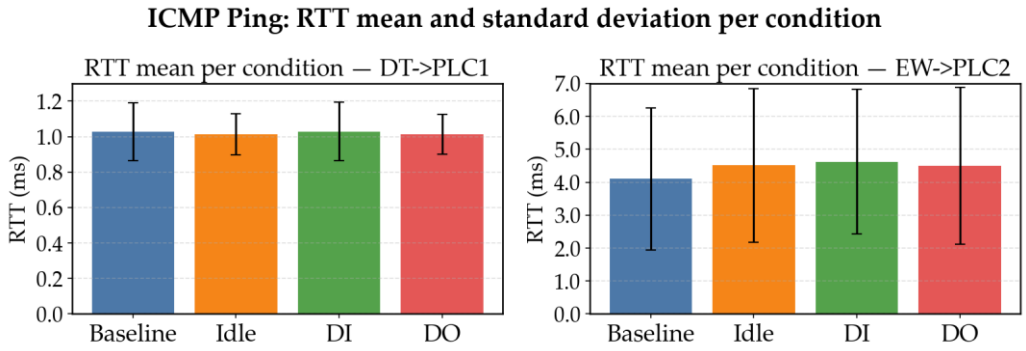


Figure 5.7. ICMP RTT measurements for the four experimental conditions.

This behaviour confirms that the attack does not introduce delays or disruptions in the communication cycle between the PLC and field devices. From a network perspective, the system operates within its expected temporal constraints. Overall, the ICMP service remained stable under all scenarios, confirming that the proposed attack did not introduce observable effects on network-layer connectivity, communication latency, or RTT variability.

5.6.2. Throughput Analysis

The throughput of the PROFINET network was evaluated to assess whether the attack affects the volume or frequency of transmitted data under the different experimental conditions.

Table 5.1 summarizes the throughput statistics for the evaluated scenarios, including packet rate (PPS) and data volume (kB/s).

Table 5.1. PROFINET throughput statistics under four conditions.

Scenario	PPS mean	PPS SD	kB/s mean	kB/s SD	ΔPPS(%)
Baseline	4073.01	307.47	248.30	18.74	--
Attack-Idle	4068.86	302.12	248.10	18.42	-0.10%
Attack-DI	4082.76	166.56	248.93	10.15	0.24%
Attack-DO	4072.78	282.94	248.34	17.25	-0.01%

The results show that throughput metrics remained consistent across all evaluated conditions. Variations in packet rate (ΔPPS) remained within ±0.3%, which is within the expected operating range of a lightly loaded PROFINET RT network.

These variations were minor compared with the inherent variability of the communication system and could not be attributed to the attack mechanism.

Overall, the throughput analysis confirms that the proposed attack did not introduce observable effects on network load or traffic volume. The monitored communication-level indicators remained within their expected operational bounds across all evaluated scenarios.

5.6.3. PROFINET Cycle Timing

The communication cycle timing, a critical parameter in real-time industrial networks, was analysed to determine whether the implemented attack scenarios introduce timing disturbances in the cyclic I/O exchange.

Cycle time was estimated from the inter-arrival times of PROFINET RTC1 frames captured during network monitoring across all experimental conditions.

The results show that the mean cycle time remains effectively constant across baseline and attack scenarios, with no systematic deviations observed. Figure 5.8 illustrates both the distribution of cycle times and their corresponding mean values, confirming the stability of the communication cycle under all evaluated conditions.

The distribution of cycle times remained consistent across all scenarios, with similar variability observed in both baseline and attack conditions. Occasional deviations in inter-arrival times appeared in all cases and were attributable to measurement-side jitter rather than to the attack mechanism.

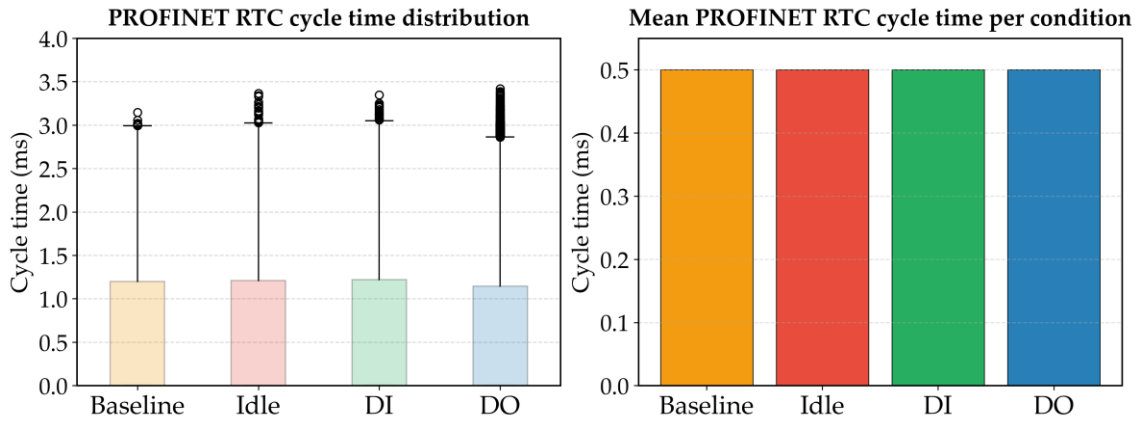


Figure 5.8. PROFINET RTC1 cycle-time distribution for PLC 1-PLC 2 traffic.

The cyclic exchange of input/output data between the PLC and distributed devices therefore remained stable and deterministic under all evaluated conditions. These results confirm that the proposed attack did not affect the timing behaviour of the PROFINET communication cycle or its real-time scheduling behaviour. The communication cycle remained within its expected operational parameters across all scenarios, indicating that the attack did not introduce observable effects on the temporal characteristics of the communication process.

5.6.4. Summary of RTT, Throughput and Cycle Timing

Taken together, the RTT, throughput, and PROFINET RTC1 cycle-timing analyses indicate that the proposed attack did not produce observable deviations in the monitored communication-level indicators, despite causing significant changes in process behaviour.

5.7. Results and Discussion

The results obtained from the experimental evaluation demonstrate that the proposed attack scenarios effectively compromise the consistency between the real physical state of the process and the information used by the control system, as well as between the control actions commanded by the controller and those actually executed in the process. In both cases, the system operates on logically valid data that do not correspond to the actual physical state of the process,

resulting in incorrect behaviour without triggering conventional fault mechanisms.

At the same time, the communication-level validation showed that key network indicators, including RTT, throughput, and PROFINET RTC1 cycle timing, remained stable across all experimental conditions. No relevant deviations were observed between baseline and attack scenarios, confirming that the communication layer operated within its nominal performance bounds.

This discrepancy between process-level impact and network-level stability highlights a fundamental limitation of conventional monitoring approaches in industrial control systems. While communication integrity is preserved, the semantic consistency of process data can be compromised without being reflected in the evaluated communication-level indicators.

These results demonstrate the need for multi-layer monitoring strategies that incorporate process-aware information, providing the necessary context to identify semantic anomalies that cannot be detected through communication-level analysis alone.

This confirms that the proposed attack paradigm exploits a structural gap between communication integrity and process semantics, highlighting the limitations of relying exclusively on network-level monitoring.

5.8. Chapter Summary

This chapter has presented the design and experimental evaluation of attack scenarios targeting both the perception and execution layers of an industrial cyber-physical system. Through Scenarios A1 and A2, it has been shown that sensor and actuator manipulation can introduce significant deviations in process behaviour while maintaining valid communication exchanges.

The communication-level validation demonstrated that these attacks did not produce observable deviations in standard network indicators, including latency, throughput, and PROFINET RTC1 cycle timing.

These findings highlight the importance of integrating process-level and communication-level information for effective anomaly detection and cybersecurity analysis in CPS environments. The results demonstrate that communication data alone are insufficient to identify semantic attacks affecting

industrial CPS. Consequently, datasets integrating communication and process information become essential for identifying process anomalies and abnormal behaviours that remain invisible to conventional network-level monitoring. The generation of such datasets is addressed in Chapter 6.

Chapter 6

Industrial CPS Dataset Generation

6.1. Introduction

This chapter implements the methodology introduced in Section 1.4, providing a detailed description of its main components, including data acquisition, anomaly scenario definition, ground-truth generation, and dataset structure.

The dataset is designed to capture system behaviour in terms of event sequences and operational modes, following the Discrete Event Systems and GEMMA-based modelling framework introduced in Chapter 2. This event-based representation enables the characterization of system dynamics through state transitions and the events associated with them.

Data are obtained using an event-driven acquisition strategy, enabling the construction of structured datasets suitable for anomaly detection. In this context, deviations from expected behaviour can be identified through the analysis of event sequences and operational modes, which correspond directly to the stages defined in the proposed framework.

Overall, this chapter presents the DAYPSCI dataset as a structured and reproducible resource for the evaluation of anomaly detection approaches in industrial Cyber-Physical Systems. Although the validation study presented in Chapter 4 was conducted using datasets generated during the progressive development of the experimental framework, those datasets represent earlier stages of the same methodology. The consolidated version is formalized in this chapter through an event-based acquisition strategy, the incorporation of temporal information (Δt), GEMMA operational modes, planned anomaly injection, and hybrid ground-truth generation.

6.2. Dataset Overview

The dataset developed in this work, referred to as DAYPSCI, has been generated using the CPS testbed described in Chapter 3. It captures the behaviour of a PLC-controlled industrial process through an event-based data acquisition strategy, enabling a detailed representation of system dynamics in terms of discrete events and operational state changes.

The dataset includes both normal operation and controlled anomaly scenarios generated through systematic experimentation using the Digital Twin and the PLC-based control system. This configuration enables the reproduction of different operating conditions in a consistent and repeatable manner, providing a reliable basis for data-driven analysis.

Unlike conventional time-sampled datasets, DAYPSCI follows an event-driven representation in which data are recorded only when changes occur in relevant system variables. Each event is associated with its corresponding inter-event time (Δt), explicitly capturing the temporal dynamics of the system. In addition, scan-level information is incorporated to represent the logical execution of the PLC program, including the ordering of events within each control cycle.

The recorded data include digital sensor signals, actuator control variables, and process states derived from the GEMMA model, enabling the representation of system behaviour at both the signal and operational levels. Furthermore, the dataset incorporates ground-truth labels that distinguish between normal operation and different anomaly classes, including sensor faults and actuator faults.

The dataset is organized into independent experimental batches, each corresponding to a specific operating condition or anomaly configuration. Each batch contains structured data files together with metadata supporting reproducibility and traceability of the experimental process. This organization facilitates the comparison of different scenarios and the evaluation of anomaly detection methods.

Overall, DAYPSCI provides a comprehensive and structured representation of an industrial CPS operating under controlled conditions. Its event-based nature, combined with temporal information (Δt), operational modes, and ground-truth labels, makes it particularly suitable for the development and evaluation of anomaly detection and cybersecurity approaches. Although the primary focus of

the dataset is the representation of process-level behaviour and anomalies, the proposed CPS framework simultaneously captures synchronized network traces, including PROFINET traffic, thereby enabling future cybersecurity-oriented studies and multi-layer analyses that combine process and communication data.

6.3. Data Acquisition

Data acquisition in the proposed CPS testbed is performed at the PLC level using an event-based logging strategy, enabling the direct capture of system behaviour through variable changes rather than continuous signal sampling. This approach represents industrial processes as sequences of discrete events, consistent with the Discrete Event Systems framework described in Section 2.2.

Unlike conventional time-based acquisition methods, where system variables are sampled at fixed intervals, the proposed approach records data only when relevant signals change. These signals include digital sensor states, actuator control commands, and GEMMA-based operational state variables used to represent the current operating mode of the system. Consequently, the acquired data reflect the logical evolution of the process while avoiding the redundancy commonly associated with fixed-rate sampling.

Each event is generated within the PLC program through edge-detection mechanisms, capturing both rising and falling transitions of binary variables. Every detected change is associated with a timestamp obtained from the PLC system clock, ensuring precise temporal tracking of system operation. Events are also sequentially indexed, enabling the reconstruction of ordered event sequences corresponding to the execution of the industrial process.

A key feature of the proposed acquisition strategy is the explicit representation of temporal information through the inter-event time (Δt), defined as the elapsed time between consecutive events. This feature provides a direct characterization of process dynamics by linking timing information to actual process events, resulting in a compact and information-rich description of system behaviour, particularly in industrial environments where events occur irregularly.

Figure 6.1 illustrates the event-based data acquisition pipeline, showing how PLC binary signal changes are transformed into structured event records and how temporal features such as the inter-event time (Δt) are computed.

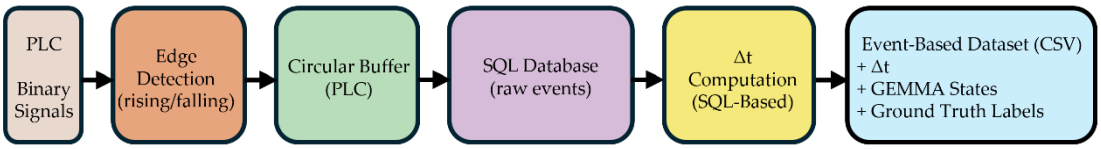


Figure 6.1. Event-based data acquisition pipeline with SQL-based Δt computation.

This pipeline highlights the transformation of raw PLC binary signals into a structured dataset in which both temporal and logical aspects of system operation are explicitly represented. It also illustrates how the inter-event time (Δt) is obtained as an intrinsic feature of the process rather than being imposed through fixed sampling intervals.

In addition to temporal information, the logging system captures the execution semantics of the PLC program through scan-level identifiers. Each event is associated with a scan cycle identifier (*scan_id)*, representing the PLC cycle in which the event occurred, and an execution order index (*event_order*), indicating the relative order of events within the same scan cycle. This enables reconstruction of the logical execution sequence of the control program, even when multiple events occur during a single scan cycle.

The generated events are temporarily stored within the PLC using an internal buffering mechanism, ensuring that no data are lost during high-frequency event generation. The data are subsequently transmitted to a SQL database, where raw event records are stored in a structured format. Additional temporal features, including the computation of Δt, are derived during this stage without modifying the original event records, thereby preserving data integrity.

Overall, this acquisition strategy provides a compact, structured, and information-rich representation of industrial process behaviour, capturing both temporal dynamics and control logic execution. As a result, the acquired data are particularly suitable for sequence-based analysis, anomaly detection, and cybersecurity-oriented studies.

6.4. Dataset Structure

The DAYPSCI dataset is structured as a collection of event-based sequences, where each record corresponds to a discrete change in the state of the system. This representation allows modelling the industrial process as an ordered sequence of events, reflecting both the temporal and logical evolution of the CPS.

Each event in the dataset is defined by a set of attributes that describe its occurrence, context, and associated system conditions. These attributes include temporal information, process variables, control states, and labelling fields, enabling a comprehensive characterization of system behaviour.

Temporal information is captured through multiple fields, including the event timestamp, a global event sequence index, and the inter-event time Δt, which represents the elapsed time between consecutive events and explicitly captures system temporal dynamics. This representation preserves the evolution of the system while avoiding the redundancy associated with fixed-rate sampling.

In addition to temporal features, the dataset explicitly represents the execution semantics of the PLC program through scan-level variables. Each event is associated with a scan cycle identifier (*scan_id)* and an execution order index (*event_order*), allowing the reconstruction of the logical sequence of operations performed by the control system and preserving the relative ordering of events within each scan cycle.

The dataset includes process-level variables derived from digital sensor signals and actuator control commands, representing the physical and control states of the system. Additionally, system operation is described through states based on the GEMMA model, enabling the interpretation of system behaviour at the level of operational modes.

To support supervised analysis, the dataset incorporates ground-truth labels that distinguish between normal operation and different types of anomalous conditions, including sensor faults, actuator faults, and combined scenarios. These labels are complemented by contextual information derived from the execution of the PLC program, the event-based representation of the process, and the GEMMA operational modes. This structured labelling enables both binary and multi-class classification tasks, providing flexibility for different analytical approaches.

At a higher level, the dataset is organized into sequences associated with individual processed parts, identified by a unique *part_id.* Each sequence represents the evolution of the system during the processing of a single unit, forming the basis for sequence-level analysis.

The dataset is further grouped into experimental batches, where each batch corresponds to a specific operating condition or fault scenario. This organization

facilitates reproducibility, comparison across scenarios, and systematic evaluation of anomaly detection methods.

Table 6.1 summarizes the main variables included in the dataset, along with their data types and corresponding descriptions.

Table 6.1. Summary of dataset variables and their descriptions

Label	Type	Description
event_seq	Integer	Event sequence index
ts_plc	DateTime	PLC Timestamp
delta_time_ms	Integer	Time between consecutive events (Δt) in milliseconds
scan_id	Integer	Identifier of the PLC scan cycle
event_order	Integer	Execution order within the scan cycle
event_id	Integer	Unique event identifier
event_type	Binary	Edge type: rising (1) or falling (0)
part_id	Integer	Part identifier
gemma	Integer	GEMMA state (encoded)
anomaly	Integer	Anomaly label
fault	Integer	Fault presence indicator
FSx	Integer	Sensor fault indicator (per sensor channel)
TFSx	Integer	Sensor fault type descriptor
FCVx	Integer	Actuator fault indicator (per actuator channel)
TFCVx	Integer	Actuator fault type descriptor
sensor_signals	Binary	Binary states of sensors (a0–c1, B1, B2)
actuator_signals	Binary	Binary states of actuators (YA_p, YA_m, YB_p, YB_m, YC_p, YC_m)
batch_id	Integer	Identifier of the experimental batch (fault configuration)
experiment_id	Integer	Identifier of the experiment instance

6.5. Anomaly Scenarios

The experimental scenarios used to generate the dataset were designed to capture both normal operation and operation under controlled fault injection of the industrial CPS.

The data are organized into a series of independent batches, each corresponding to a specific operating condition or anomaly configuration. These scenarios include normal system operation, as well as controlled anomaly scenarios affecting different components of the process, enabling the representation of diverse system behaviours while ensuring consistency across experiments.

The anomaly scenarios are designed to reflect representative industrial fault scenarios, including sensor faults, actuator malfunctions, and combined scenarios affecting multiple components simultaneously (see Figure 3.2). This structured definition ensures that the dataset captures realistic and heterogeneous anomaly patterns.

The anomalies are introduced through external configuration mechanisms managed by the fault scheduler presented in Chapter 3 (see Figure 3.3), ensuring that the control logic of the system remains unchanged across all scenarios. This design guarantees that variations in system behaviour are exclusively due to the defined anomaly conditions, preserving consistency across the dataset.

During the execution of each scenario, the system evolves according to the operational modes defined by the GEMMA model, including normal operation, fault occurrence, and recovery phases. These transitions are represented in the dataset through event sequences and associated state information, enabling a structured characterization of system dynamics.

Each scenario is associated with a predefined configuration describing the type of anomaly, the affected components, and the experimental conditions. This structured description enables traceability between anomaly definitions and observed system behaviour and supports the subsequent ground-truth generation and labelling process.

Overall, the structured definition of anomaly scenarios ensures that the dataset systematically captures both normal and abnormal system behaviours in a reproducible and coherent manner.

6.6. Ground Truth Labelling

The definition of ground truth is a critical component of the proposed approach, as it directly determines the reliability and validity of subsequent data-driven analysis. In industrial environments, ground truth is often unavailable or unreliable, since faults are not explicitly recorded and system states are not systematically labelled during operation.

In this work, ground truth is defined through a structured and controlled experimental methodology, ensuring clear and consistent labelling of system behaviour. The labelling process follows a hybrid strategy that combines

externally defined experimental configurations with contextual information derived from the execution of the control system.

A key principle of the proposed approach is the explicit separation between the CPS execution layer and the labelling process. Fault conditions are defined externally through the experimental setup, rather than being inferred from system behaviour. This ensures that labels are independent of the control logic and prevents ambiguities that may arise when attempting to derive anomaly information directly from process signals.

During each experimental batch, the configuration defines the type of fault introduced, the affected components, and the timing of fault activation. This information is systematically associated with the recorded data, enabling each event to be labelled according to its corresponding operating condition. In this way, the dataset maintains full traceability between the injected cause (fault configuration) and the observed effect (system behaviour).

Figure 6.2 illustrates the fault injection and ground-truth generation process, showing the relationship between externally defined fault conditions and the observable system behaviour captured through event sequences and operational states.

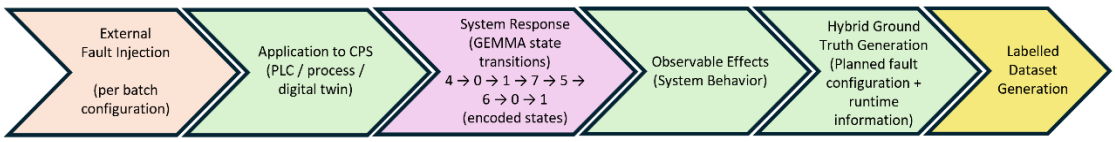


Figure 6.2. Fault injection and hybrid ground-truth generation process based on GEMMA.

This figure highlights how externally defined fault configurations are integrated with runtime system behaviour to generate consistent and interpretable labels. It also illustrates the role of GEMMA states in contextualizing anomalies within the operational modes of the system, enabling a structured interpretation of system behaviour under both normal and faulty conditions.

As illustrated in Figure 6.2, the ground-truth generation process is initiated outside the operational process and subsequently enriched with runtime information associated with the corresponding event sequences and system states.

In addition to externally defined labels, runtime information from the control system is incorporated to enrich the labelling process. In particular, system states derived from the GEMMA model provide contextual information about operational modes, allowing the differentiation between normal operation, fault occurrence, and recovery phases. This dual representation enables a more comprehensive interpretation of anomalies.

An important aspect of this approach is the distinction between fault activation and observable anomaly manifestation. While a fault may be introduced at a specific point in time, its effects on system behaviour may appear later due to process dynamics. The labelling strategy explicitly accounts for this temporal decoupling, enabling a more accurate representation of cause–effect relationships within the dataset.

The dataset incorporates multiple labelling fields, including anomaly indicators and fault descriptors, enabling both binary and multi-class classification tasks. These labels provide information about the presence of anomalies, the affected components, and the type of fault scenario.

Overall, by combining externally configured fault definitions with runtime information from the control system, the proposed strategy ensures accurate, consistent, and reproducible labelling. This approach reduces uncertainty in anomaly interpretation and provides a reliable foundation for the development and evaluation of anomaly detection methods in industrial CPS.

6.7. Dataset Organization

At the storage and distribution level, the dataset is organized into a collection of independent experimental batches, each corresponding to a specific operating condition or fault configuration. This structure enables the systematic representation of both normal operation and operation under controlled fault injection, facilitating reproducibility and controlled comparison across different experimental conditions.

Each batch represents a complete and self-contained experiment, capturing the full sequence of events associated with the execution of the industrial process under a given configuration. This includes both normal operation and fault scenarios affecting sensors, actuators, or both.

For each experimental batch, the dataset includes structured data files containing event-based records, where each row corresponds to a discrete event as defined in previous sections. These files store temporal information, GEMMA-based control states, labelling fields, as well as *batch_id* and *experiment_id* identifiers defined in Table 6.1, enabling detailed analysis of system behaviour.

In addition to the event-based process data, the dataset incorporates synchronized network traffic captures obtained from the industrial communication infrastructure through a switch port mirroring configuration. These captures include PROFINET process data exchanged between the CPS components and are provided in PCAPNG format. The availability of both process-level and network-level data enables joint analysis of control execution, process dynamics, and communication behaviour within the CPS environment.

To support reproducibility and traceability, supplementary information is also provided. This includes metadata describing the experimental setup, fault injection configuration, affected components, and the resulting system behaviour. This information ensures that each dataset instance can be unambiguously interpreted and reproduced.

The dataset is publicly available through an institutional repository [35], ensuring accessibility and supporting reproducible research. The organization into well-defined batches, combined with the availability of both process-level and network-level data, provides a consistent and comprehensive framework for the evaluation of anomaly detection and industrial cybersecurity methods, as illustrated in Figure 6.3.

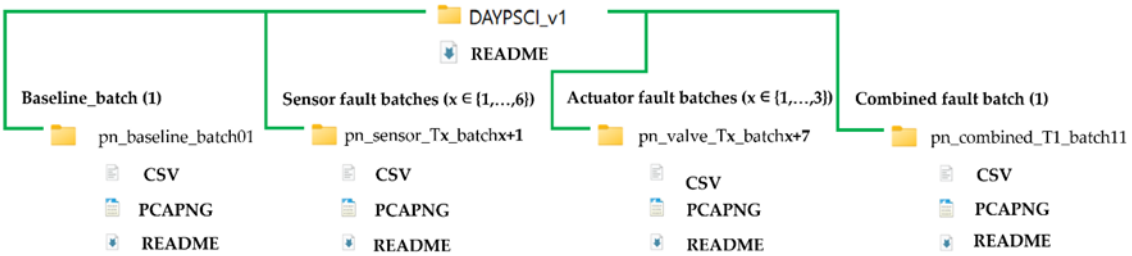


Figure 6.3. Directory structure and organization of the DAYPSCI_v1 dataset.

Overall, the structured organization, detailed documentation, and public availability of the dataset enable its reuse as a benchmark for data-driven analysis in industrial CPS. This organization not only ensures reproducibility but also

enables the systematic evaluation of anomaly detection techniques under controlled and well-defined industrial scenarios.

6.8. Chapter Summary

This chapter has presented the DAYPSCI dataset generated using the proposed CPS-based experimental framework, describing its structure, data acquisition process, experimental scenarios, and ground-truth generation methodology.

The dataset is built upon an event-based representation, capturing system behaviour through sequences of discrete events and associated temporal information (Δt), thereby enabling an accurate characterization of system dynamics. Data are acquired directly at the PLC level and enriched with GEMMA operational states, providing both temporal and contextual information about the execution of the industrial process.

A key contribution of the proposed approach is the generation of ground truth through a hybrid strategy that combines controlled fault-injection configurations with contextual information obtained during system execution. This strategy ensures full traceability between fault definitions, system behaviour, and dataset labels, while reducing ambiguity in anomaly interpretation. The chapter has also described the organization of the dataset into reproducible experimental batches, enabling systematic evaluation under controlled operating conditions.

Overall, DAYPSCI constitutes a structured, reproducible, and well-documented resource for the analysis of industrial CPS behaviour and for the development and evaluation of anomaly detection and cybersecurity approaches.

Chapter 7

Conclusions and Future Work

7.1. Conclusions

This thesis has presented a comprehensive methodology for the design, implementation, and validation of a hybrid Cyber–Physical system testbed aimed at the generation of structured datasets for anomaly detection and cybersecurity research in industrial environments.

The proposed framework integrates industrial hardware, digital twin technologies, event-based data acquisition, and fault injection mechanisms within a unified and flexible experimental platform. This integration enables the systematic reproduction of industrial operating conditions under controlled and repeatable scenarios, providing a reliable environment for the study of industrial CPS behaviour.

A major contribution of this work is the definition of a reproducible methodology for dataset generation based on controlled experimental scenarios. Unlike conventional approaches that rely on passive monitoring of industrial systems, the proposed methodology allows anomalies and faults to be intentionally introduced according to predefined experimental configurations. This approach ensures full traceability between the injected fault conditions and the resulting system behaviour.

Another important contribution is the explicit separation between CPS operation and ground-truth generation. Rather than being inferred solely from process data, ground truth is generated through a hybrid strategy that combines externally defined fault configurations with runtime information, including operational states and execution context. This approach guarantees consistent labelling, reduces ambiguity in anomaly interpretation, and improves dataset reproducibility.

The work has also demonstrated the suitability of event-based representations for modelling industrial CPS behaviour. By combining event sequences, PLC

execution semantics, GEMMA operational modes, and temporal features such as the inter-event time (Δt), the proposed approach captures both the logical and temporal characteristics of industrial processes. This representation provides a compact and information-rich description of system behaviour suitable for data-driven analysis.

The resulting dataset, DAYPSCI, constitutes a structured and reproducible resource that combines process-level information, operational context, hybrid ground-truth labels, and network-level observations. The organization of the dataset into well-defined experimental batches facilitates systematic comparison across scenarios and supports the evaluation of anomaly detection and cybersecurity techniques under controlled conditions.

Furthermore, the practical applicability of the proposed framework has been demonstrated through preliminary machine learning studies using datasets generated by the testbed. The obtained results showed that the generated datasets are suitable for anomaly detection and classification tasks, achieving high classification performance in scenarios involving sensor and actuator faults. These findings provide an additional level of validation for both the generated datasets and the proposed methodology.

Beyond anomaly-detection applications, the experimental results obtained from the cybersecurity analysis demonstrate that cyberattacks may significantly alter the behaviour of an industrial CPS while remaining difficult to identify using conventional network-level monitoring approaches alone. In particular, the evaluated Man-in-the-Middle (MitM) attacks showed that process semantics can be manipulated without introducing measurable deviations in communication-level indicators such as PROFINET cycle timing, throughput, or network latency. While communication properties remained largely unchanged, observable deviations were produced in the physical process and control behaviour. These findings highlight the importance of incorporating process-aware information into cybersecurity analysis and reinforce the need for cyber–physical approaches that jointly consider communication behaviour, control execution, and process dynamics.

Overall, this thesis demonstrates that the combination of hybrid CPS testbeds, controlled fault injection, event-based data acquisition, and hybrid ground-truth generation provides a solid and reproducible foundation for anomaly detection

and cybersecurity research in industrial environments. The proposed framework contributes not only a practical experimental platform but also a systematic methodology for generating realistic industrial datasets and evaluating data-driven approaches in industrial Cyber–Physical systems.

The main contribution of this thesis is the definition of a reproducible CPS-oriented methodology for generating labelled industrial datasets through the explicit separation between process execution and hybrid ground-truth generation. By combining controlled anomaly injection, event-based data acquisition, GEMMA operational modes, and temporal information (Δt), the proposed framework provides a systematic foundation for the development and evaluation of anomaly detection and cybersecurity approaches in industrial CPS.

7.2. Future Work

The research presented in this thesis opens several promising directions for future work in the areas of industrial anomaly detection, cybersecurity, and Cyber–Physical systems.

A first research direction involves extending the CPS testbed by incorporating additional industrial assets and production processes. In particular, the integration of robotic cells and more complex manufacturing systems would increase the realism and diversity of the generated datasets, enabling the evaluation of anomaly detection techniques in more representative industrial environments.

A second direction focuses on the exploitation of the generated datasets for advanced machine learning applications. Future studies may investigate deep learning architectures specifically designed to process temporal and event-based information, including recurrent neural networks, temporal convolutional networks, transformers, and graph neural networks. These approaches may enable the identification of complex temporal and structural patterns that are difficult to capture using conventional machine learning techniques.

A third direction concerns the development of integrated cyber–physical anomaly detection strategies that combine process-level and communication-level information. The availability of both event-based process data and synchronized PROFINET traffic captures provides an opportunity to develop

multi-source detection frameworks capable of correlating network activity with process behaviour.

Future work may also explore additional cybersecurity scenarios involving different attack vectors, communication protocols, and industrial architectures. Extending the fault injection framework to incorporate a wider variety of cyberattacks would contribute to a more comprehensive evaluation of cybersecurity monitoring and detection mechanisms.

Another relevant direction involves the deployment and validation of the proposed methodology in operational industrial environments. While the current work demonstrates the effectiveness of the approach within a controlled experimental testbed, its application in real production systems would provide valuable insights into scalability, robustness, and adaptability under practical constraints.

Finally, future research may further refine the event-based representation and hybrid ground-truth generation methodology. The incorporation of additional contextual information, richer temporal descriptors, and more sophisticated hybrid labelling strategies could enhance dataset quality and support the development of increasingly accurate anomaly detection models.

Overall, these research directions highlight the potential of the proposed framework as a foundation for future advances in industrial CPS, enabling the generation of realistic datasets, the development of intelligent monitoring systems, and the study of cybersecurity challenges through reproducible and controlled experimentation.

References

[1] S. J. Oks *et al.*, "Cyber-Physical Systems in the Context of Industry 4.0: A Review, Categorization and Outlook," *Information Systems Frontiers,* vol. 26, no. 5, pp. 1731-1772, 2022, doi: 10.1007/s10796-022-10252-x.

[2] N. Jazdi, "Cyber Physical Systems in the Context of Industry 4.0," presented at the 2014 IEEE INTERNATIONAL CONFERENCE ON AUTOMATION, QUALITY AND TESTING, ROBOTICS, 2014.

[3] F. Pires *et al.*, "Positioning Cyber-Physical Systems and Digital Twins in Industry 4.0," presented at the 2024 IEEE 7TH INTERNATIONAL CONFERENCE ON INDUSTRIAL CYBER-PHYSICAL SYSTEMS, ICPS 2024, 2024.

[4] X. Liu *et al.*, "A systematic review of digital twin about physical entities, virtual models, twin data, and applications," *ADVANCED ENGINEERING INFORMATICS,* vol. 55, JAN 2023, Art no. 101876, doi: 10.1016/j.aei.2023.101876.

[5] E. Zancul, A. Oliveira, Y. B. Wang, K. Schützer, and B. Schleich, "A digital twin design and implementation approach for industrial application leveraging programmable logic controllers," *INTERNATIONAL JOURNAL OF INTERACTIVE DESIGN AND MANUFACTURING - IJIDEM,* vol. 19, no. 5, pp. 3479-3487, MAY 2025, doi: 10.1007/s12008-024-01969-x.

[6] H. Harkat, L. M. Camarinha-Matos, J. Goes, and H. F. T. Ahmed, "Cyber-physical systems security: A systematic review," *COMPUTERS & INDUSTRIAL ENGINEERING,* vol. 188, FEB 2024, Art no. 109891, doi: 10.1016/j.cie.2024.109891.

[7] W. Alsabbagh and P. Langendörfer, "Security of Programmable Logic Controllers and Related Systems: Today and Tomorrow," *IEEE OPEN JOURNAL OF THE INDUSTRIAL ELECTRONICS SOCIETY,* vol. 4, pp. 659-693, 2023, doi: 10.1109/OJIES.2023.3335976.

[8] M. Cook, A. Marnerides, C. Johnson, and D. Pezaros, "A Survey on Industrial Control System Digital Forensics: Challenges, Advances and Future Directions," *IEEE COMMUNICATIONS SURVEYS AND TUTORIALS,* vol. 25, no. 3, pp. 1705-1747, 2023, doi: 10.1109/COMST.2023.3264680.

[9] W. Alsabbagh and P. Langendörfer, "No need to be online to attack-exploiting s7-1500 plcs by time-of-day block," 2022: IEEE, pp. 1-8.

[10] R. Ramirez, C. K. Chang, and S. H. Liang, "PLC Cybersecurity Test Platform Establishment and Cyberattack Practice," *ELECTRONICS,* vol. 12, no. 5, MAR 2023, Art no. 1195, doi: 10.3390/electronics12051195.

[11] Q. Gulzar and K. Mustafa, "An analytical survey of cyber-physical systems in water treatment and distribution: Security challenges, intrusion detection, and future directions," *SECURITY AND PRIVACY,* vol. 7, no. 6, NOV 2024, doi: 10.1002/spy2.440.

[12] Z. Ahmad, A. Petrovski, and Ieee, "Securing Cyber-Physical Systems with Two-level Anomaly Detection Strategy," presented at the 2024 IEEE 7TH INTERNATIONAL CONFERENCE ON INDUSTRIAL CYBER-PHYSICAL SYSTEMS, ICPS 2024, 2024.

[13] E. A. Boateng, J. W. Bruce, and D. A. Talbert, "Anomaly Detection for a Water Treatment System Based on One-Class Neural Network," *IEEE ACCESS,* vol. 10, pp. 115179-115191, 2022, doi: 10.1109/ACCESS.2022.3218624.

[14] S. Krishnaveni, S. Sivamohan, B. Jothi, T. M. Chen, and M. Sathiyanarayanan, "TwinSec-IDS: An Enhanced Intrusion Detection System in SDN-Digital-Twin-Based Industrial Cyber-Physical Systems," *CONCURRENCY AND COMPUTATION-PRACTICE & EXPERIENCE,* vol. 37, no. 3, DEC 10 2024, Art no. e8334, doi: 10.1002/cpe.8334.

[15] G. S. Gaba, A. Sari, I. Butun, P. Singh, A. Gurtov, and M. Liyanage, "A survey on security and privacy of Industry 4.0 and beyond: Technical aspects, use cases, challenges, and research directions," *IEEE Open Journal of the Communications Society,* 2025.

[16] M. Akeel, S. Hosseinzadeh, M. Zeeshan, H. Homatash, N. Owoh, and M. Ashawa, "Virtual Testbed for Cyber-Physical System Security Research and Education: Design, Evaluation, and Impact," *ELECTRONICS,* vol. 15, no. 3, JAN 29 2026, Art no. 582, doi: 10.3390/electronics15030582.

[17] O. Pospisil, P. Blazek, K. Kuchar, R. Fujdiak, and J. Misurec, "Application Perspective on Cybersecurity Testbed for Industrial Control Systems," *SENSORS,* vol. 21, no. 23, DEC 2021, Art no. 8119, doi: 10.3390/s21238119.

[18] D. Serpanos, "The Cyber-Physical Systems Revolution," *COMPUTER,* vol. 51, no. 3, pp. 70-73, MAR 2018, doi: 10.1109/MC.2018.1731058.

[19] W. Alsabbagh, C. Kim, and P. Langendörfer, "Investigating the Security of OpenPLC: Vulnerabilities, Attacks, and Mitigation Solutions," *IEEE ACCESS,* vol. 12, pp. 11561-11583, 2024, doi: 10.1109/ACCESS.2024.3356051.

[20] G. Barbieri and D. A. Gutierrez, "A GEMMA-GRAFCET Methodology to enable Digital Twin based on Real-Time Coupling," presented at the PROCEEDINGS OF THE 2ND INTERNATIONAL CONFERENCE ON INDUSTRY 4.0 AND SMART MANUFACTURING (ISM 2020), 2021.

[21] M. Conti, D. Donadel, and F. Turrin, "A Survey on Industrial Control System Testbeds and Datasets for Security Research," *IEEE COMMUNICATIONS SURVEYS AND TUTORIALS,* vol. 23, no. 4, pp. 2248-2294, 2021, doi: 10.1109/COMST.2021.3094360.

[22] K. Cabaj *et al.*, "Cyber4OT dataset: Network traces for cyber-security vulnerability evaluation in industrial control systems," *SOFTWAREX,* vol. 31, SEP 2025, Art no. 102196, doi: 10.1016/j.softx.2025.102196.

[23] R. Tolas, R. Portase, and R. Potolea, "GeMSyD: Generic Framework for Synthetic Data Generation," *DATA,* vol. 9, no. 1, JAN 2024, Art no. 14, doi: 10.3390/data9010014.

[24] K. Sun *et al.*, "A Survey of Attack Detection and Security Control for Advanced Cyber-Physical Systems and Their Applications," *IEEE TRANSACTIONS ON SYSTEMS MAN CYBERNETICS-SYSTEMS,* Article 2026 2026, doi: 10.1109/TSMC.2026.3690568.

[25] W. Alsabbagh and P. Langendöerfer, "A new injection threat on s7-1500 PLCs-disrupting the physical process offline," *IEEE Open Journal of the Industrial Electronics Society,* vol. 3, pp. 146-162, 2022.

[26] Z. B. Wang, Y. F. Zhang, Y. L. Chen, H. R. Liu, B. L. Wang, and C. H. Wang, "A Survey on Programmable Logic Controller Vulnerabilities, Attacks, Detections, and Forensics," *PROCESSES,* vol. 11, no. 3, MAR 2023, Art no. 918, doi: 10.3390/pr11030918.

[27] A. Walz *et al.*, "PROFINET Security: A Look on Selected Concepts for Secure Communication in the Automation Domain," presented at the 2023 IEEE 21ST INTERNATIONAL CONFERENCE ON INDUSTRIAL INFORMATICS, INDIN, 2023.

[28] T. T. Nguyen, R. Kadavil, and H. Hooshyar, "A Real-Time Cyber-Physical Simulation Testbed for Cybersecurity Assessment of Large-Scale Power Systems," *IEEE TRANSACTIONS ON INDUSTRY APPLICATIONS,* vol. 60, no. 6, pp. 8329-8340, NOV 2024, doi: 10.1109/TIA.2024.3457877.

[29] B. Brandin, R. Su, and L. Lin, "Supervisory control of time-interval discrete event systems," *IEEE Transactions on Automatic Control,* vol. 69, no. 5, pp. 3080-3095, 2023.

[30] B. Hrúz and M. Zhou, *Modeling and Control of Discrete-Event Dynamic Systems*. Springer, 2007.

[31] International Electrotechnical Commission, "IEC 61131-3: Programmable controllers – Part 3: Programming languages," Geneva, Switzerland, 2013.

[32] W. Kritzinger, M. Karner, G. Traar, J. Henjes, and W. Sihn, "Digital Twin in manufacturing: A categorical literature review and classification," *Ifac-PapersOnline,* vol. 51, no. 11, pp. 1016-1022, 2018.

[33] A. Dehlaghi-Ghadim, M. H. Moghadam, A. Balador, and H. Hansson, "Anomaly Detection Dataset for Industrial Control Systems," *IEEE ACCESS,* vol. 11, pp. 107982-107996, 2023, doi: 10.1109/ACCESS.2023.3320928.

[34] K.-H. Niemann, "PROFINET–Future OT security requirements: What do NIS2, CER, CRA and IEC 62443 require?," 2025.

[35] J. V. Martín-Fraile, N. Basurto Hornillos, J. E. Sierra García, and Á. Herrero Cosío. *DAYPSCI v1: An Event-Based Dataset of Fault Injection Scenarios in PLC-Controlled Industrial Cyber-Physical Systems*, doi: 10.71486/2drf-1g53.

Appendix A. Publications Derived from This Thesis

A.1. Main Publications

This subsection presents the main scientific publications derived from the research carried out during this doctoral thesis. These publications are closely related to the contributions presented in the manuscript and provide additional validation of the proposed methodologies and experimental results.

Paper I (DYNA)

A Testbed for Evaluation of Machine Learning Techniques in Anomaly Detection in Industrial Cyber Physical Systems

Authors: J.V. Martín-Fraile, N. Basurto, J.E. Sierra-García, A. Herrero

Published in DYNA (2025), JCR Impact Factor 0.8 (Q3), DOI: 10.52152/D11516.

Summary: This paper presents a hybrid industrial cyber-physical system (CPS) testbed for evaluating artificial intelligence techniques in anomaly detection. The proposed approach integrates real industrial hardware, including Siemens S7-1200 and S7-1500 PLCs, with a digital twin to reproduce realistic industrial processes. A use case demonstrates that supervised machine learning algorithms (KNN, Random Forest, and SVM) achieve high accuracy in detecting normal behaviour and actuator faults, with acceptable performance for sensor anomalies.

Contribution: The author led the design and implementation of the hybrid CPS testbed, including the integration of industrial PLCs and digital twin environments. The author was also responsible for data acquisition design, experimental evaluation, and the analysis of machine learning techniques for anomaly detection.

Paper II (Applied Sciences)

A Hybrid Digital-Twin-Based Testbed for Real-Time Manipulation of PROFINET I/O: A Practical Man-in-the-Middle Attack Implementation

Authors: J.V. Martín-Fraile, N. Basurto, J.E. Sierra-García, A. Herrero

Published in Applied Sciences (2026), JCR Impact Factor 2.5 (Q2), DOI: 10.3390/app1607353.

Summary: This paper presents the implementation of a Man-in-the-Middle (MitM) attack on PROFINET I/O industrial control systems using a hybrid digital-twin-based testbed for real-time manipulation. The proposed approach enables real-time manipulation of SDU-level data within RTC1 cyclic communication while preserving protocol integrity and timing, and experimental results demonstrate that both input-level (DI) and output-level (DO) manipulations produce deterministic and reproducible effects on PLC behaviour without affecting PROFINET communication indicators, remaining undetectable at the protocol level.

Contribution: The author led the conceptualization, methodology, and implementation of the proposed Man-in-the-Middle attack framework using a hybrid digital-twin-based testbed. The author was also responsible for software development, experimental design and execution, data analysis, and dataset preparation, and contributed to visualization and manuscript writing.

Paper III (Data in Brief)

DAYPSCI: Event-Based Dataset for Anomaly Detection through Fault Injection in PLC-Controlled Industrial Cyber-Physical Systems

Authors: J.V. Martín-Fraile, N. Basurto, J.E. Sierra-García, A. Herrero

Under review in Data in Brief (2026). JCR Impact Factor 1.9 (Q2). DOI not yet available.

Summary: This paper presents DAYPSCI, an event-based dataset generated from a hybrid industrial cyber-physical system (CPS) testbed based on PLC-controlled processes. The dataset integrates process data from a digital twin with process states defined according to the GEMMA methodology. Data acquisition is performed using an event-driven logging approach, where each event is timestamped and only changes in system variables are recorded together with their associated inter-event duration (Δt).

Contribution: The author led the conceptualization, methodology, and implementation of the proposed dataset generation approach. The author was responsible for software development, experimental design and execution, data acquisition and structuring, and contributed to validation, visualization, and manuscript writing.

A.2. Other Journal Publications

This subsection includes additional journal articles co-authored by the candidate during the doctoral period, including submitted manuscripts related to the research area of this thesis.

Detecting Performance Anomalies in a Robot by Applying Missing Values Imputation and Balancing Methods

Authors: Nuno Basurto, Juan V. Martín-Fraile, Angel Arroyo, J. Enrique Sierra-Garcia, Alvaro Herrero

Under review in Intelligent & Robotic Systems (2026). JCR Impact Factor 2.8 (Q3). DOI not yet available.

A.3. Conference Publications

This subsection includes conference publications co-authored by the candidate that are related to the research area of this thesis. These contributions present intermediate results and developments associated with the research conducted during the doctoral period.

Banco de pruebas para evaluación de técnicas IA en la detección de anomalías en sistemas ciberfísicos industriales

Authors: J.V. Martín-Fraile, J.E. Sierra-García, N. Basurto,

Presented at the XX Simposio CEA de Control Inteligente (2025). DOI not available.

A.4. Published Datasets

This subsection presents published datasets co-authored by the candidate that are related to the research presented in this thesis, providing reusable resources for anomaly detection and cybersecurity studies in industrial CPS.

DAYPSCI v1: An Event-Based Dataset of Fault Injection Scenarios in PLC-Controlled Industrial Cyber-Physical Systems

Authors: J.V. Martín-Fraile, J.E. Sierra-García, N. Basurto, Á. Herrero

Available in the Institutional Repository of the University of Burgos (Riubu) (2026). DOI: 10.71486/2drf-1g53.

Description: Public dataset comprising event-based industrial CPS data, including CSV files, PROFINET network captures (PCAPNG), and experimental metadata for anomaly detection and cybersecurity research.

Appendix B. Full Publications

B.1. Paper I (DYNA Journal).

Testbed for Evaluation of Machine Learning Techniques in Anomaly Detection in Industrial Cyber Physical Systems

A TESTBED FOR EVALUATION OF MACHINE LEARNING TECHNIQUES IN ANOMALY DETECTION IN INDUSTRIAL CYBER PHYSICAL SYSTEMS

Juan-Vicente Martín-Fraile1, Nuño Basurto2, Jesús-Enrique Sierra-García1 y Álvaro Herrero2

Universidad de Burgos. Escuela Politécnica Superior. 1 Grupo de Investigación en Automatización, Robótica, Control y Optimización (ARCO) y 2 Grupo de Inteligencia Computacional Aplicada (GICAP). Av. Cantabria, s/n - 09006, Burgos (España).

Received: 06/aug/25 – Reviewed: 09/sep/25 -- Accepted: 09/sep/26 - DOI: https://doi.org/10.52152/D11516

To cite this article: MARTÍN-FRAILE, Juan-Vicente; BASURTO, Nuño; SIERRA-GARCÍA, Jesús-Enrique; HERRERO, Álvaro. A TESTBED FOR EVALUATION OF MACHINE LEARNING TECHNIQUES IN ANOMALY DETECTION IN INDUSTRIAL CYBER PHYSICAL SYSTEMS. DYNA March April 2026. Vol. 101, n.2, pp. xxx | DOI: https://doi.org/10.52152/D11516

ABSTRACT:

The increasing digitalization of industrial environments, driven by Industry 4.0, has fostered the integration of Cyber-Physical Systems (CPS), which combine Information Technology (IT) and Operational Technology (OT) to enhance efficiency, automation, and real-time decision-making. However, this convergence has also increased the exposure of Industrial Control Systems (ICS) to cyber threats, highlighting the need for tools and frameworks capable of detecting anomalies and cyberattacks. In this context, the present work proposes a testbed for evaluating Artificial Intelligence techniques aimed at the automatic detection of anomalies in industrial processes. The testbed enables the simulation of both normal and anomalous conditions following the GEMMA methodology. A use case is presented in which a supervised machine learning approach—specifically, the K-Nearest Neighbours (KNN), Random Forest (RF), and Support Vector Machine (SVM) algorithms—is applied to data generated in the Testbed for Anomaly Detection and Protection of Industrial Cyber-Physical Systems (DAyPSCI). The results demonstrate the feasibility of the proposed approach for classifying instances based on their operational behaviour, achieving high accuracy in identifying normal patterns and actuator failures, and acceptable performance in detecting sensor anomalies. These findings validate the testbed as a realistic and effective environment for the development and evaluation of intelligent anomaly-detection systems in industrial CPS. Moreover, the proposed testbed can be considered a solid foundation for future research in predictive maintenance, process optimization, and industrial cybersecurity.

Keywords: *Testbed, Anomaly detection, Cybersecurity, Digital Twin, Cyber-physical system, Programmable Logic Controller (PLC), Machine Learning, Industrial control.*

1. - INTRODUCTION

Cyber-Physical Systems (CPS) integrate computational and physical processes to automate complex real-world tasks [1][2]. These systems leverage technological advances in sensors, communication networks, and embedded computing, particularly through the use of Programmable Logic Controllers (PLCs). They combine software and hardware components to interact with the physical environment in an automatic and efficient manner. Key functionalities include real-time monitoring and control of physical processes, such as energy management in smart grids or the optimization and automation of production in smart factories.

However, this ongoing technological evolution is not without risks. In 2024, the number of cyberattacks targeting cyber-physical systems increased significantly compared to the previous year [3][4]. This rise highlights their growing vulnerability and exposure, reinforcing the urgent need to implement more robust cybersecurity measures to protect Industrial Control Systems (ICS) and mitigate the impact of cyber threats [5][6].

According to the report by the Computer Emergency Response Team of the National Cryptologic Center (CCN-CERT) [7], as well as studies by Pospisil et al. [8], Domínguez et al. [9], and Benka et al. [10], the growing integration of Information Technology (IT) and Operational Technology (OT) with systems such as SCADA (Supervisory Control and Data Acquisition), HMI (Human-Machine Interface), PLCs (Programmable Logic Controllers), and DCS (Distributed Control Systems) has significantly increased the exposure of Industrial Control Systems (ICS) to cyber threats.

The implementation of Industry 4.0 has accelerated the IT/OT convergence, enabling companies to enhance operational efficiency, make real-time data-driven decisions, and reduce costs through advanced automation and predictive maintenance. This integration provides a holistic view of both production processes and business operations, resulting in greater adaptability, competitiveness, and resource optimization.

ISSN: 0012-7361 eISSN: 1989-1490 / DYNA Vol.XX nºX DOI: https://doi.org/10.6036/XXXX

Ingeniería e Industria	A Testbed for Evaluation of Machine Learning Techniques in Anomaly Detection in Industrial Cyber Physical Systems	COMPUTER SCIENCES
RESEARCH ARTICLE	Juan Vicente Martín Fraile, Nuño Basurto, Jesús Enrique Sierra García y Álvaro Herrero	Artificial Intelligence and Simulation

This article presents preliminary results obtained by applying Artificial Intelligence (AI) techniques, Machine Learning (ML) more specifically, to a set of datasets generated using the Testbed for Anomaly Detection and Protection of Industrial Cyber-Physical Systems (DAyPSCI) [11], a platform designed for research and training in cybersecurity applied to Industrial Control Systems (ICS). Specifically, a supervised machine learning approach is employed, widely validated in the literature for its effectiveness in fault detection tasks in industrial environments [12][13]. DAyPSCI serves a dual purpose: on one hand, it facilitates the development and validation of new mechanisms for anomaly detection and protection against cyber threats; on the other hand, it provides an accessible experimental environment for teaching anomaly detection and industrial cybersecurity.

To guide this study, the following **research questions** were formulated:

- **RQ1**: Does the proposed testbed possess sufficient flexibility to configure anomaly detection use cases within Industrial Cyber-Physical Systems (ICPS)?
- **RQ2**: Is the proposed testbed capable of generating valid datasets for the evaluation of machine learning techniques in anomaly detection?
- **RQ3**: How can the testbed be leveraged for educational purposes?

These contributions aim to bridge the gap between academic research and practical training in industrial cybersecurity and anomaly detection.

The remainder of this article is structured as follows. Section 2 presents the testbed developed for the generation of industrial CPS datasets, describing its architecture and components. Section 3 introduces a use case, including the digital twin, the human-machine interface, the dataset generation and labelling process and the methodology applied for anomaly detection. Section 4 presents the results. Section 5 provides a discussion of these results, highlighting key insights, limitations, and educational implications. Finally, Section 6 concludes the article with final remarks and outlines future work.

2. - TESTBED FOR THE GENERATION OF INDUSTRIAL CPS DATASETS

This section presents the testbed developed for generating datasets resulting from the production of a batch of 100 parts, parameterized by the operator, as well as the methodology followed for their analysis.

In general, industrial CPS testbeds are classified according to the components they integrate into the following types [2]:

- **Virtual:** These are solutions based exclusively on software, which provides them with great flexibility and scalability.
- **Physical:** These include real components such as sensors, actuators, network devices, controllers, etc., meaning they reflect an actual ICS.
- **Hybrid:** These combine virtual and physical components. They are suitable for testing integrated industrial systems without the high cost of physical test equipment.

The testbed presented in this article falls within the hybrid category, as it combines both physical and virtual components.

2.1.- ARCHITECTURE

The testbed presented in this article features an open and flexible architecture based on the Purdue Enterprise Reference Architecture (PERA) model, enabling the recreation of conditions typical of real industrial environments. This is made possible by the fact that most of its components are equivalent to those found in real automated systems, with the exception of the physical industrial process (sensors, actuators, parts, etc.), which is implemented through a digital twin. A digital twin is a virtual representation of objects, processes, or systems that exist and operate in real time [14].

To simulate more realistic industrial communication conditions, this digital twin uses a Siemens S7-1215C PLC (PLC-2) as a Profinet interface with the control system (Siemens S7-1512C-1PN PLC), which is in turn connected to an industrial switch. In this way, the Level 1 PLC (PLC-1) can send commands and receive input signals from the process through the Level 0 PLC (PLC-2), as if it were controlling a physical process via Profinet I/O. Additionally, the system includes an HMI interface that allows the operator to interact with the virtual process.

This initial version covers only the levels corresponding to the process control zone, as shown in Figure 1. However, its modular and scalable architecture allows for the integration of new components, facilitating expansion to other zones and levels of the PERA model.

ISSN: 0012-7361 eISSN: 1989-1490 / DYNA Vol.XX nºX DOI: https://doi.org/10.6036/XXXX

DYNA Ingeniería e Industria	A Testbed for Evaluation of Machine Learning Techniques in Anomaly Detection in Industrial Cyber Physical Systems	COMPUTER SCIENCES
RESEARCH ARTICLE	Juan Vicente Martín Fraile, Nuño Basurto, Jesús Enrique Sierra García y Álvaro Herrero	Artificial Intelligence and Simulation



Fig. 1. Network Architecture of the Testbed.

With the aim of using the testbed to evaluate AI techniques for anomaly detection and the protection of cyber-physical systems—thereby contributing to their improved resilience—a database managed by Microsoft SQL Server 2022 has been integrated. This database enables the generation and storage of datasets for subsequent analysis. Although it is hosted by default on the engineering workstation, its location is not mandatory, as it can be deployed in other compatible environments.

The testbed also allows the connection of a third monitoring PC for capturing communication frames and packets using tools such as Wireshark or similar. Additionally, it supports the connection of a device running Kali Linux, enabling the execution and evaluation of various attack techniques for testing and validating defence mechanisms.

2.2.- COMPONENTS

The following Table 1 shows the components that make up the current version of the Testbed:

Nº	Device	Model	Description
1	PS	6EP1333-4BA00	Power supply
2	PLC-1	6ES7512-1CK01-0AB0	Test bench control PLC
3	PC-1	S.O. Windows	PC with Digital Twin
4	PLC-2	215-1BG40-0XB0	PLC as Profinet I/O Interface
5	Switch	208-0BA00-2AC2	Industrial switch
6	HMI	6AV2128-3GB06-0AX1	HMI panel
7	PC-2	S.O. Windows	Engineering station with TIAP
8	PC-3	S.O. Windows	Equipment with Kali Linux, Wireshark, etc.

Table 1. Testbed Components.

3.- USE CASE

The following presents a use case of the testbed to demonstrate its applicability. For this purpose, a very simple industrial process has been selected: a part marking system. The objective of this system is to feed parts into the marker, perform the marking operation, and then eject the processed parts from the system.

3.1.- DIGITAL TWIN AND HUMAN-MACHINE INTERFACE

In order to have information from sensors, parts, pre-actuators, and actuators—that is, process data comparable to those of a real process—along with their corresponding animated visualization, the PCSimu application [15] has been used. This free tool allows design creation by providing a 2D visualization. For this reason, it was selected and used to generate the digital twin of the marking system.

As shown in Figure 2, the system is composed of three actuators (A, B, and C), each with their respective magnetic position sensors (a0, a1, b0, b1, c0, and c1), a parts generator, and two inductive presence detectors (B1 and B2), located in the feeding and marking areas, respectively. The actuators are configured as double-acting cylinders, controlled by bistable 5/2 solenoid valves.

The PCSimu application also allows adjustment of the actuators' forward and backward speeds, contributing to a more realistic simulation of the process's dynamic behaviour. Additionally, as shown in Figure 1, the digital twin is connected to PLC-2, which acts as a Profinet I/O (PNIO) interface, via its connection to port 1 of the industrial switch. This configuration enables information exchange between the digital twin and PLC-1 through Profinet I/O communication, thus replicating a decentralized peripheral architecture typical of modern industrial environments.

ISSN: 0012-7361 eISSN: 1989-1490 / DYNA Vol.XX nºX DOI: https://doi.org/10.6036/XXXX

Ingeniería e Industria	A Testbed for Evaluation of Machine Learning Techniques in Anomaly Detection in Industrial Cyber Physical Systems	COMPUTER SCIENCES
RESEARCH ARTICLE	Juan Vicente Martín Fraile, Nuño Basurto, Jesús Enrique Sierra García y Álvaro Herrero	Artificial Intelligence and Simulation

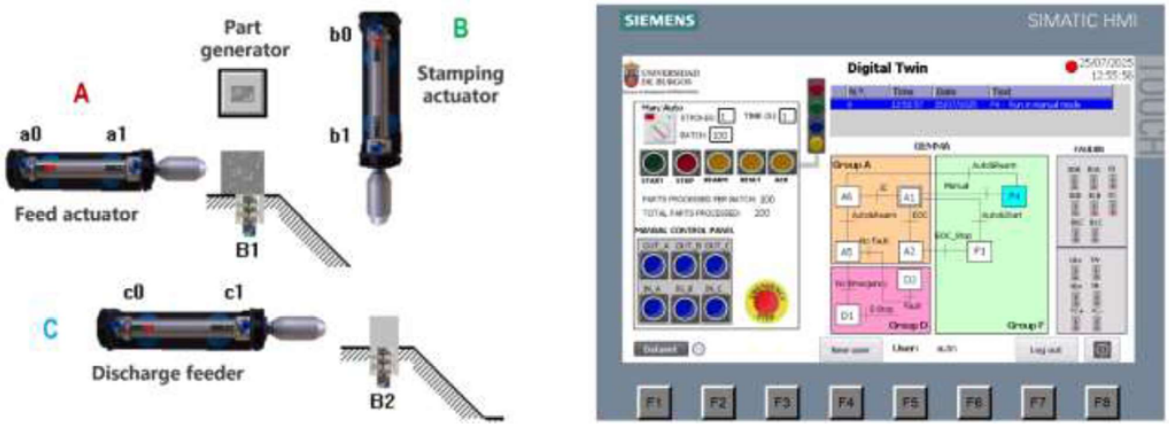


Fig. 2. Digital Twin and Operator Panel HMI Interface.

To interact with the ICS system, an interface has been implemented on the operator panel as shown in the previous Figure 2. Through this interface, the operator can adjust parameters such as the actuation time in seconds of the die on the part, the number of marking strikes per part, and the quantity of parts to be marked in each batch. Additionally, the **MAN/AUT** selector allows choosing between manual and automatic operation modes of the system. In manual mode, the **OUT_x** and **IN_x** pushbuttons (where **x** = **A**, **B**, or **C**), with Normally Open (NO) contacts, enable the operator to send advance and retract commands to the actuators when necessary. The interface also includes other pushbuttons such as **start**, **stop**, **reset**, **rearm**, and **ack**, which serve to issue the following commands: start automatic mode, stop at the end of the cycle, rearm the system in case of a fault or defect, reset the completed batch of parts, and acknowledge events requested by the control system, respectively.

Furthermore, the system includes an **emergency stop** pushbutton with a Normally Closed (NC) contact. For signalling purposes, it features a light beacon composed of four coloured indicators—red, green, blue, and yellow—that provide visual information about the system's operational status.

As previously mentioned, the equipment is also designed for educational purposes and therefore incorporates the **GEMMA** methodology (Study Guide for the Modes of Start and Stop of an Automated System). **Group A** includes stop procedures: **A1** – stop in the initial state, **A2** – stop requested at the end of the cycle, **A5** – preparation for startup after a fault, and **A6** – return of the system to the initial state. **Group F**, which covers the process in operation, includes **F1** – normal production and **F4** – verification runs without order or manual mode. Finally, **Group D**, process in fault, includes **D1** – emergency stop and **D2** – diagnosis and/or treatment of faults.

Emergency stop management and reset are considered the main part of a PLC program [16]. To facilitate reset operations for the operator, the system is equipped, among other features, with a reset pushbutton and a notification system that guides and informs the operator about the system's status, making the reset operation easier.

Another objective is anomaly detection within the system. For this purpose, it includes switches **B0x**, **B1x**, **Yx+**, and **Yx**—(where **x** = **A**, **B**, and **C**) that simulate faults in sensors (B) and pre-actuators—that is, simulate faults in the solenoid valves (Y) controlling the actuators. This allows consideration of faults or anomalies as they occur in real systems and adds this information to the dataset.

When a fault is simulated—such as a sensor failure—the control system detects it at the appropriate point within its normal operation sequence, triggering a change from the current operation mode to mode D2. In this mode, the operator is notified via the HMI of the fault or defect and is prompted to confirm that it has been resolved. Once the operator's acknowledgment (ACK) is received, the system switches to mode F5, waiting either for a request to change to manual mode F4 or for the reset pushbutton to be activated. This allows the system to transition to mode A6 and return to the initial state A1, ready to receive a new start command and continue with the programmed batch. All this information is recorded by the system and incorporated into the dataset for further processing and analysis.

3.2.- DATASET GENERATION AND LABELLING

As previously mentioned, the test bench is equipped with a database that records all programmed events occurring in the system, enabling the generation of datasets for further analysis. The collected variables are described in Table 2, indicating the data type and a brief description of the information they record.

ISSN: 0012-7361 eISSN: 1989-1490 / DYNA Vol.XX nºX DOI: https://doi.org/10.6036/XXXX

DYNA Ingeniería e Industria

A Testbed for Evaluation of Machine Learning Techniques in Anomaly Detection in Industrial Cyber Physical Systems

COMPUTER SCIENCES

RESEARCH ARTICLE

Juan Vicente Martín Fraile, Nuño Basurto, Jesús Enrique Sierra García y Álvaro Herrero

Artificial Intelligence and Simulation

Variable	Type	Description
Date	Date	Event registration date (yy-mm-dd)
Time	Time	Event registration time (hh:mm:ss:ms).
Duration	Int	Units in milliseconds.
Sensors	Boolean	B_x0, B_x1, where x = a, b, and c; B1 and B2
Actuators	Boolean	YA+, YA-, YB+, YB-, YC+ and YC-.
Anomalies	Int	0 – No anomaly, 1 – Sensors, 2 – Actuators.
Gemma	Int	Modes: 0-A1, 1-F1, 2-A2, 3-D1, 4-F4, 5-A5, 6-A6 and 7-D2.
part_ID	Int	Identification code of the part being manufactured

Table 2. Dataset Variables.

To create the dataset, the first step is to start PC-1 and launch the digital twin, establishing communication between PLC-1 and the plant under control via Profinet I/O. Next, the operator must configure the batch parameters to be processed by the system through the human-machine interface: number of parts, marking actuation time, and number of marking strikes per part. Once the batch is configured, the operator should activate the event logging system by using the dataset switch on the HMI, which begins saving the events occurring in the system to the SQL Server database.

Figure 3 below shows the legend and the first 7 rows obtained from an Excel file connected to the Test Bench database. Observing the data, you can see how data recording starts with the system in the initial position (Gemma = 0, row 2), then the system switches to automatic mode (Gemma = 1, row 3), followed by the registration of a part arriving at the system (B1 set to true, row 4). In the next row, row 5, the system issues the part load command (YA+ set to true), the loading actuator advances (B_a0 set to false in row 6 and B_a1 set to true in row 7), and after loading is completed, the retract command for the loader is given (YA- set to true, row 8).

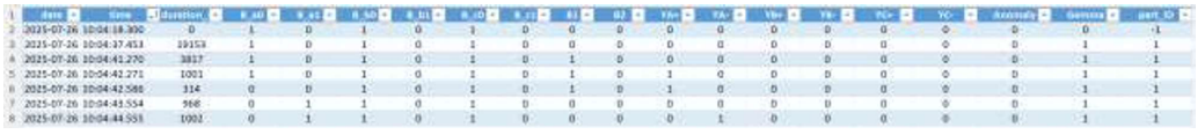


Fig. 3. Sample of Variable Logging in Dataset Creation and Labelling.

The labelling of the dataset aims to provide suitable data for evaluation using AI techniques focused on anomaly detection and the protection of cyber-physical systems. Accordingly, datasets have been generated based on the information collected in Table 2: one corresponding to a process without anomalies, another with controlled anomalies introduced in the sensors, another with anomalies in the actuators, and finally a mixed dataset for evaluation.

- **Without anomalies:** This dataset contains information from the test bench configured with a batch of 100 parts, applying one strike per part and a tightening time of one second per part (**Anomaly = 0**).
- **With type 1 anomalies:** The test bench configuration was maintained, but for a batch of 15 parts in which sensor faults were deliberately introduced in a controlled manner (**Anomaly = 1**).
- **With type 2 anomalies:** Same test bench configuration as in type 1, but faults were introduced in the solenoid valves that control the actuators (**Anomaly = 2**).

3.3.- METHODOLOGY FOR ANOMALY DETECTION

This section describes the proposed methodology for identifying anomalies that occur during the manufacturing process of the defined batch, using automatic part classification through supervised machine learning.

The objective of this work is to evaluate the performance of three supervised classification techniques: K-Nearest Neighbours (**KNN**), Random Forest (**RF**), and Support Vector Machine (**SVM**), using datasets obtained from the testbed, in which anomalies have been deliberately induced in a controlled manner. These techniques have been employed in various studies for classification tasks in different application domains [17] [18]. The ultimate goal is to develop an automatic system capable of classifying the manufactured industrial parts into three categories or classes: **Class 0** for good (normal) parts**, Class 1** for parts with type 1 anomalies (sensor faults) and **Class 2** for parts with type 2 anomalies (actuator faults).

ISSN: 0012-7361 eISSN: 1989-1490 / DYNA Vol.XX nºX DOI: https://doi.org/10.6036/XXXX

Ingeniería e Industria	A Testbed for Evaluation of Machine Learning Techniques in Anomaly Detection in Industrial Cyber Physical Systems	COMPUTER SCIENCES
RESEARCH ARTICLE	Juan Vicente Martín Fraile, Nuño Basurto, Jesús Enrique Sierra García y Álvaro Herrero	Artificial Intelligence and Simulation

The procedure followed is structured in the following stages:

Step 1. Data sources:
A mixed dataset was used, representing different types of parts: without anomalies, with type 1 anomalies, and with type 2 anomalies.
Step 2. Preprocessing:
Since each part consists of multiple time-stamped rows in the dataset, the rows were grouped by part, determining the maximum length among all parts. Parts with fewer rows were padded with zeros to reach this maximum length, and the data was normalized.
Step 3. Dataset Splitting:
The complete dataset is divided into three subsets to ensure a robust evaluation of the model:

- **Training (70%):** used to fit the model parameters.
- **Validation (15%):** used to tune hyperparameters and prevent overfitting.
- **Testing (15%):** reserved for evaluating the model's final performance on unseen data.

Step 4. Performance Evaluation:
The model is evaluated using the mixed test set with standard classification metrics:

- **Accuracy:** Proportion of correct predictions over the total number of samples.
- **Precision:** Proportion of true positives among the predicted positives.
- **Recall:** Proportion of true positives among the actual positives.
- **F1-score**: Harmonic mean of precision and recall, useful when there is class imbalance.
- **AUC-ROC**(**Area Under the Receiver Operating Characteristic Curve**): This metric measures the model's ability to distinguish between classes, providing an aggregate performance indicator across all classification thresholds. A higher AUC value indicates a better capacity to separate normal parts from those with anomalies.

3.3.1- Hyperparameter tuning

The models were evaluated using the metrics described in step 4 of this methodology, selected for their ability to assess performance under class imbalance, consistent with the structure of the DAyPSCI dataset.

Hyperparameter tuning was conducted in two phases: an initial exploratory phase using traditional search based on configurations commonly found in the literature, followed by a systematic optimization using grid search with cross-validation, implemented through **GridSearchCV**, following a methodology similar to that proposed by [18].

In the first phase, for **KNN**, values of k = 3, 5, and 7 were evaluated, observing that k = 3 provided the best result. For **RF**, configurations with 10, 50, and 100 estimators were tested, confirming that performance stabilizes from 50 estimators, without being significantly affected by the tree depth. For **SVM**, linear, polynomial, and RBF kernels were selected, along with values of C = 0.1, 1.0, and 10 combined with gamma = 'scale', 'auto', 0.01, 0.1, and 1, experimentally obtaining the best result with RBF, C = 1, and gamma = 'scale'.

The second phase confirmed the adequacy of the initial parameter ranges and selections, with the exception of SVM, where an optimal value of C = 10 was suggested. However, this configuration was discarded as it did not yield improvements over the previously selected value.

4. - RESULTS

The performance of the three classification models—K-Nearest Neighbours (KNN), Random Forest (RF), and Support Vector Machines (SVM)—was evaluated on the test set using five widely adopted metrics: accuracy, precision, recall, F1-score, and AUC-ROC. These metrics provide a comprehensive view of each model's ability to correctly classify instances, balance false positives and false negatives, and maintain discriminative power across classes. The results obtained are shown below, visualizing the models' performance on the test set.

ISSN: 0012-7361 eISSN: 1989-1490 / DYNA Vol.XX nºX DOI: https://doi.org/10.6036/XXXX

DYNA Ingeniería e Industria	A Testbed for Evaluation of Machine Learning Techniques in Anomaly Detection in Industrial Cyber Physical Systems	COMPUTER SCIENCES
RESEARCH ARTICLE	Juan Vicente Martín Fraile, Nuño Basurto, Jesús Enrique Sierra García y Álvaro Herrero	Artificial Intelligence and Simulation

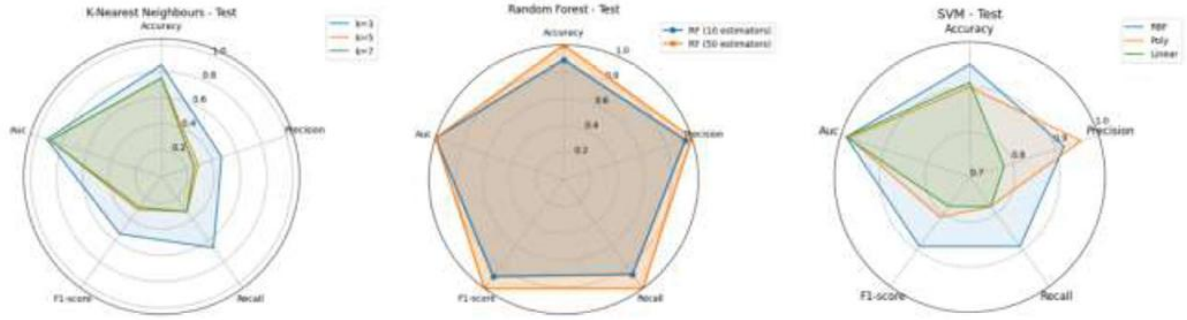


Fig. 4. Radar plots comparing the test set performance of KNN (k = 3, 5, 7), Random Forest (10 and 50 estimators), and SVM (RBF, polynomial, and linear kernels) across five metrics: accuracy, precision, recall, F1-score and Auc.

As illustrated in **Figure 4**, **KNN** achieved its best results with **k = 3**, while configurations with k =5 and k=7 produced slightly lower scores. **RF** exhibited strong and consistent performance across all configurations, with **50 estimators** yielding notable improvements in recall and F1-score. For **SVM**, the **RBF** kernel delivered the highest overall performance, except for precision, where the polynomial kernel performed slightly better.

Beyond the radar plots, ROC curves shown in **Figure 5** provide additional evidence of class separation and model robustness. In this figure, RF achieved perfect separation across all classes, while SVM showed near-perfect discrimination. KNN, although competitive in AUC-ROC, exhibited less consistent separation, particularly for minority classes.

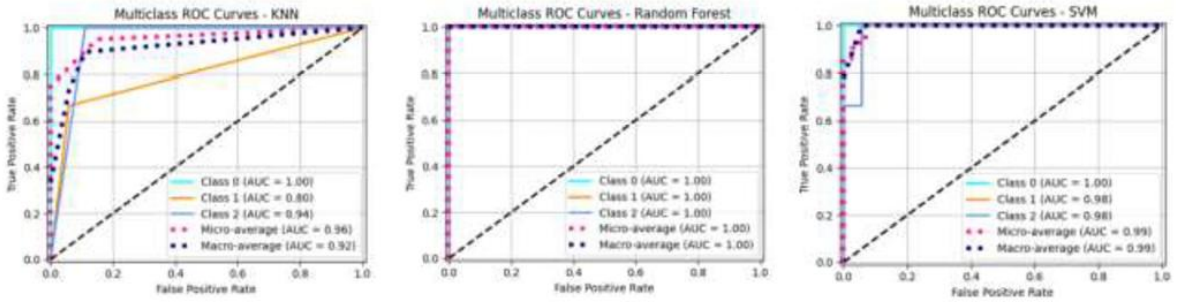


Fig. 5. Multiclass ROC curves for KNN (k = 3), Random Forest (50 estimators), and SVM (RBF kernel) using a one-vs-rest approach. Each plot displays class-wise, micro-average, and macro-average ROC curves.

The following table summarizes the best results obtained for each model on the test set, providing a clear comparison of their performance across all metrics.

Technique	Accuracy	Precision	Recall	F1-score	AUC-ROC	Support
KNN	0.85	0.48	0.67	0.54	0.92	20
RF	1.0	1.0	1.0	1.0	1.0	20
SVM	0.95	0.92	0.89	0.89	0.99	21

Table 3. Summary of classification results on the test set for KNN (k = 3), Random Forest (50 estimators), and SVM (RBF kernel). Random Forest achieved perfect performance, SVM showed slightly lower but still competitive results, while KNN performed significantly worse, particularly in precision and F1-score.

As presented in **Table 3**, RF achieved perfect scores across all metrics (accuracy, precision, recall, F1-score, and AUC-ROC), confirming its robustness even with a relatively small number of estimators. SVM demonstrated competitive performance, particularly in F1-score (0.89) and AUC-ROC (0.99), making it a viable alternative when computational efficiency or interpretability is prioritized. KNN, while simpler, showed limitations in precision (0.48), recall (0.67), and F1-score (0.54), which may restrict its applicability in scenarios requiring high reliability.

ISSN: 0012-7361 eISSN: 1989-1490 / DYNA Vol.XX nºX DOI: https://doi.org/10.6036/XXXX

DYNA Ingeniería e Industria	A Testbed for Evaluation of Machine Learning Techniques in Anomaly Detection in Industrial Cyber Physical Systems	COMPUTER SCIENCES
RESEARCH ARTICLE	Juan Vicente Martín Fraile, Nuño Basurto, Jesús Enrique Sierra García y Álvaro Herrero	Artificial Intelligence and Simulation

As a result, it can be concluded that RF is the most recommended approach for anomaly detection in the part marking system, implemented through a digital twin within the testbed. This validation is supported by the radar plots and ROC curves presented in Figures 4 and 5.

5. – DISCUSSION

The results obtained in this study are analysed below in relation to the research questions posed. Specifically, the flexibility of the testbed to configure different anomaly detection scenarios is discussed (RQ1), as well as its ability to generate valid datasets for evaluating machine learning techniques (RQ2). Finally, the potential of the testbed as an educational tool in the field of anomaly detection and industrial cybersecurity is addressed (RQ3).

RQ1: Does the proposed testbed possess sufficient flexibility to configure anomaly detection use cases within Industrial Cyber-Physical Systems (ICPS)?

The developed testbed demonstrates notable flexibility in configuring various use cases related to anomaly detection in ICPS. Its hybrid architecture integrates real industrial hardware components (Siemens S7-1200 and S7-1500 PLCs, etc.) with simulation software such as PCSimu—used in this study—and is compatible with other tools like RobotStudio and SIMIT. This setup allows for the simulation of a wide range of operational scenarios, including normal operation, sensor failures and actuator faults. This configuration capability extends to both automatic and manual modes, enabling the simulation of operational conditions typical of industrial processes, such as those described in the GEMMA guide.

Moreover, in addition to enabling dynamic anomaly generation, it performs real-time acquisition of all system events and stores them in a Microsoft SQL Server database.

RQ2: Is the proposed testbed capable of generating valid datasets for the evaluation of machine learning techniques in anomaly detection?

The flexibility provided by the testbed has proven capable of generating valid and high-quality datasets, suitable for evaluating machine learning techniques in anomaly detection. From a data quality perspective, the datasets generated with DAyPSCI ensure comprehensive coverage of relevant cases by intentionally including the production of both normal parts and parts with sensor and actuator faults. The class distribution can be adjusted according to the study objectives, allowing for both balanced scenarios and realistic representations of fault frequency.

The dataset comprises discrete and sequential data derived from industrial control events, characterized by high granularity and temporal consistency. These events are sequential, delineating transitions in the machine's states, thereby offering significant potential for analytical applications, including temporal analysis, deviation detection, and process cycle time optimization.

The clear structure of the dataset—with precise delimitation of each labelled work cycle—facilitates its use in classification and anomaly detection tasks. Reproducibility is guaranteed through the ability to repeat the generation conditions, and temporal representativeness is addressed by including variability in cycle times and operating modes. Finally, the generation process is precisely documented, enabling clear data interpretation and supporting reuse in comparative studies.

RQ3: How can the testbed be leveraged for educational purposes?

The results highlight the practical relevance of the testbed for both research and education. By replicating real industrial conditions, it helps bridge the gap between academic research and applied training. In educational contexts, the testbed can be integrated into courses related to industrial automation, industrial processes control, and industrial cybersecurity, among others, supporting the development of competencies in undergraduate, master's, and doctoral programs. Likewise, in vocational training, it can be applied in advanced modules on PLC programming and industrial cybersecurity, contributing to hands-on learning and skill acquisition in real-world scenarios.

6. – CONCLUSIONS AND FUTURE WORK

The results presented in this study validate the proposed testbed as an effective and realistic environment for the application of machine learning techniques in anomaly detection and the protection of industrial cyber-physical systems. The high classification accuracy achieved across different models—particularly in scenarios involving induced anomalies—demonstrates the system's ability to replicate relevant operational behaviour and failure conditions. This supports the use of the testbed as a valuable tool for research and development in industrial AI, including fault detection, predictive maintenance and cybersecurity.

Future work will focus on expanding the system's applicability to more complex and dynamic production scenarios. One key objective is to extend the dataset to include all the modes contemplated in the GEMMA guide, as well as other non-standard operating conditions.

ISSN: 0012-7361 eISSN: 1989-1490 / DYNA Vol.XX nºX DOI: https://doi.org/10.6036/XXXX

This will enable the models to learn from a broader range of behaviours and improve their robustness for deployment in real-world environments.

Another important line of development is the real-time monitoring and logging of manufacturing process efficiency, along with automatic notifications to the quality control team in the event of anomalies during the production of parts, thus enabling immediate inspection and traceability

From a methodological perspective, future research will explore more advanced ML techniques, such as Long Short-Term Memory (LSTM) networks and Graph Neural Networks (GNNs). These models are particularly well-suited for capturing temporal dependencies and relational patterns in sequential sensor-actuator data, and may offer improved performance in anomaly detection and predictive maintenance tasks.

REFERENCES

[1] Serpanos D, "The Cyber-Physical Systems Revolution", *IEEE Computer*. March 2018. Vol. 51-3, p.70–73. DOI: https://doi.org/10.1109/MC.2018.1731058

[2] Hosseinzadeh S, Voutos D, Barrie D, et al. "Design and development considerations of a cyber physical testbed for operational technology research and education". *Sensors*. June 2024. Vol. 24-12. art.3923. DOI: https://doi.org/10.3390/s24123923

[3] Fortinet, 2024 State of Operational Technology and Cybersecurity Report. Fortinet. June 2024. Available at: https://www.fortinet.com/content/dam/fortinet/assets/reports/report-state-ot-cybersecurity.pdf

[4] Faelli M, Informe global de ciberataques del primer trimestre de 2025 de Check Point Software. Check Point Research. April 2025. Available at: https://www.itsitio.com/seguridad/informe-global-de-ciberataques-del-primer-trimestre-de-2025-de-check-point-software

[5] Han S, Lee K, Cho S, et al. "Anomaly detection based on temporal behavior monitoring in programmable logic controllers", *Electronics*. May 2021. Vol. 10-10, art. 1218. DOI: https://doi.org/10.3390/electronics10101218

[6] Nankya M, Chataut R and Akl R, "Securing industrial control systems: Components, cyber threats, and machine learning-driven defense strategies", *Sensors*. October 2023. Vol. 23-21. art.8840. DOI: https://doi.org/10.3390/s23218840

[7] CCN-CERT. CCN-CERT IA-04/24: Ciberamenazas y tendencias. Edición 2024. Centro Criptológico Nacional. October 2024. Available at: https://www.ccn-cert.cni.es/es/informes/informes-ccn-cert-publicos/7274-ccn-cert-ia-04-24-ciberamenazas-y-tendencias-edicion-2024/file.html

[8] Pospisil O, Blazek P, Kuchar K, et al. "Application perspective on cybersecurity testbed for industrial control systems". *Sensors*. December 2021. Vol. 21-23. p.8119. DOI: https://doi.org/10.3390/s21238119

[9] Domínguez M, Fuertes J, Prada J, et al. "Design of platforms for experimentation in industrial cybersecurity". *Applied Sciences*. June 2022. Vol. 12-13. p.6520. DOI: https://doi.org/10.3390/app12136520

[10] Benka D, Horvath D, Spendla L, et al. "Machine learning-based detection of anomalies, intrusions, and threats in industrial control systems". *IEEE Access*. January 2025. Vol. 13. p.12502–12514. DOI: https://doi.org/10.1109/ACCESS.2025.3530902

[11] Martín-Fraile JV, Sierra-García JE, Basurto N. "Banco de pruebas para evaluación de técnicas IA en la detección de anomalías en sistemas ciberfísicos industriales". *XX Simposio CEA de Control Inteligente*. June 2025. Vol. 1-2. DOI: https://doi.org/10.64117/simposioscea.v1i2.79

[12] Fernandes M, Corchado JM, Marreiros G. "Machine learning techniques applied to mechanical fault diagnosis and fault prognosis in the context of real industrial manufacturing use-cases: a systematic literature review", *Applied Intelligence*. March 2022. Vol. 52. p.12345–12367. DOI: https://doi.org/10.1007/s10489-022-03344-3

[13] Chia YK, Saptoro A, Putra ZA, et al. "Supervised deep learning algorithms for process fault detection and diagnosis under different temporal subsequence length of process data", *Applied Intelligence*. July 2025. Vol. 55, art. 883. DOI: https://link.springer.com/article/10.1007/s10489-025-06711-y

[14] Javid M, Haleem A, Suman R. "Digital Twin applications toward Industry 4.0: A review". *Cognitive Robotics*. April 2023. Vol. 3. p.71–92. DOI: https://doi.org/10.1016/j.cogr.2023.04.003

[15] PCSimu. PCSimu: Industrial Automation Simulator (Software Version 3.0). 2023. Available at: https://cade-simu.com/pc-simu

ISSN: 0012-7361 eISSN: 1989-1490 / DYNA Vol.XX nºX DOI: https://doi.org/10.6036/XXXX

[16] Piedrafita R. Ingeniería de la Automatización Industrial. Madrid: Ra-Ma. 2004.

[17] Kim C, Park JH, Lee JY. "AI-based betting anomaly detection system to ensure fairness in sports and prevent illegal gambling". *Scientific Reports*. April 2024. Vol. 14. art.6470. DOI: https://doi.org/10.1038/s41598-024-57195-8

[18] Khan FA, Jamil A, Khan SA, et al. "Enhancing robotic manipulator fault detection with advanced machine learning techniques". *Engineering Research Express*. April 2024. Vol. 6-2. art. 025204. DOI: https://doi.org/10.1088/2631-8695/ad3dae

ACKNOWLEDGEMENTS

This publication is part of the AI4SECIoT project ("Artificial Intelligence for Securing IoT Devices" - C032.23), funded by the National Cybersecurity Institute (INCIBE), derived from a collaboration agreement signed between the National Institute of Cybersecurity (INCIBE) and the University of Burgos. This initiative is carried out within the framework of the Recovery, Transformation and Resilience Plan funds, financed by the European Union (Next Generation), the project of the Government of Spain that outlines the roadmap for the modernization of the Spanish economy, the recovery of economic growth and job creation, for solid, inclusive and resilient economic reconstruction after the COVID19 crisis, and to respond to the challenges of the next decade.

B.2. Paper II (Applied Sciences Journal).

A Hybrid Digital-Twin-Based Testbed for Real-Time Manipulation of PROFINET I/O: A Practical Man-in-the-Middle Attack Implementation

applied sciences

MDPI

Article

A Hybrid Digital-Twin-Based Testbed for Real-Time Manipulation of PROFINET I/O: A Practical Man-in-the-Middle Attack Implementation

Juan V. Martín-Fraile [1,*], Jesús E. Sierra García [1], Nuño Basurto [2] and Álvaro Herrero [2]

[1] Grupo de Investigación en Automatización, Robótica, Control y Optimización (ARCO), Departamento de Digitalización, Escuela Politécnica Superior, Universidad de Burgos, Av. Cantabria s/n, 09006 Burgos, Spain; jesierra@ubu.es

[2] Grupo de Inteligencia Computacional Aplicada (GICAP), Departamento de Digitalización, Escuela Politécnica Superior, Universidad de Burgos, Av. Cantabria s/n, 09006 Burgos, Spain; nbasurto@ubu.es (N.B.); ahcosio@ubu.es (Á.H.)

* Correspondence: jvmartin@ubu.es

Abstract

This study presents a practical methodology for executing Man-in-the-Middle (MitM) attacks on industrial control systems that utilize PROFINET I/O—a communication layer that remains largely underexplored in ICS cybersecurity research. A hybrid digital-twin-based testbed is developed by integrating Siemens S7-1500 and S7-1200 PLCs with a process replica implemented in PCSimu, together with a malicious application that modifies specific process data before it is delivered through the PROFINET I/O channel, enabling controlled falsification of process information in real time. The attacker operates through a Modbus TCP control channel while injecting the manipulated values into the 40-byte Real-Time Class 1 (RTC1) cyclic process-data payload while preserving frame integrity and protocol-level validity indicators. Experimental results show that SDU-level modifications on the 2-ms RTC1 cycle produced deterministic and fully reproducible effects on PLC-level behavior, including forced actuator confirmations and falsified process states, demonstrating the feasibility of both DI- and DO-level manipulation scenarios. Network captures and MSSQL-based event logs provide bit-level correlation between the injected SDU modifications and their impact on the automation sequence, confirming the reliability of the proposed manipulation mechanism. The testbed also supports the systematic generation of labeled datasets for training and evaluating machine-learning-based intrusion and anomaly-detection methods, and offers direct applicability to research, education, and operator-training activities in industrial cybersecurity. Overall, the proposed platform offers a secure, reproducible, and practically applicable environment for vulnerability assessment, attack simulation, and the development of detection techniques in industrial PROFINET networks.

Keywords: PROFINET I/O; RTC1 cyclic real-time communication; Man-in-the-Middle (MitM) attack; SDU manipulation; digital twin testbed; PLC security; Industrial cybersecurity; dataset generation

Check for updates

Academic Editor: Dimitri Konstantas

Received: 4 March 2026
Revised: 27 March 2026
Accepted: 1 April 2026
Published: 3 April 2026

Copyright: © 2026 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license.

1. Introduction

Industrial Control Systems (ICS) constitute the operational backbone of essential sectors such as manufacturing, energy, and transportation, as well as critical infrastructures

including electric grids, gas pipelines, and water distribution systems. These systems enable the automation of physical processes through the interaction of controllers, sensors, and actuators. The progressive digitalization driven by Industry 4.0 has intensified the convergence between Information Technologies (IT) and Operational Technologies (OT), giving rise to Industrial Cyber–Physical Systems (ICPS) characterized by the tight integration of computational and physical components in real time [1,2].

While these transformations have improved efficiency, monitoring, and remote operation, they have also expanded the attack surface, increasing the exposure of industrial infrastructures to cyber threats. Several studies have shown that traditional SCADA environments exhibit structural vulnerabilities that can be exploited [3–5]. Reference [3] reports significant weaknesses in a SCADA testbed, demonstrating that replay, denial-of-service, and traffic-manipulation attacks can easily compromise Modbus-based systems. Further work [4] confirms that these vulnerabilities persist due to the lack of native authentication, integrity, and encryption mechanisms. More recent analyses [5] show that these deficiencies are still present even in critical infrastructures such as electrical substations, where attackers may capture unencrypted credentials, intercept SCADA traffic, or inject malicious commands via Man-in-the-Middle (MitM) attacks.

Among these threats, MitM attacks are particularly critical because they enable the interception and manipulation of data without clear indicators of intrusion. The absence of built-in security features in commonly used industrial protocols facilitates the injection of false commands or the covert alteration of process signals. Certain MitM attacks can alter operational data without generating anomalies detectable solely through network-traffic analysis, underscoring the need to incorporate physical-process information into detection mechanisms [6,7].

Progress in analysis and detection techniques is further constrained by the limited availability of realistic datasets for ICS environments. Although some works [8] provide large datasets captured from physical testbeds, they primarily focus on DoS, ARP spoofing, or MitM attacks on generic protocols and do not include real-time industrial communication such as PROFINET I/O or the direct manipulation of process signals. This limitation hinders the evaluation and comparison of detection methods under realistic OT conditions.

Moreover, the industrial sector is currently facing concrete protection challenges related to PROFINET-based communication. Industrial Ethernet protocols—including PROFINET—have historically lacked integrated security mechanisms, making them susceptible to data manipulation, spoofing, and Man-in-the-Middle attacks at the protocol level [9]. In parallel, manufacturers, integrators, and operators must now comply with increasingly stringent OT-security regulations, including NIS2, the CER Directive, the Cyber Resilience Act, and the IEC 62443 series, which explicitly mandate authenticity, integrity, and secure configuration of industrial communication systems [10]. These regulatory and technical pressures have motivated the ongoing development of PROFINET Security Classes 2 and 3, which introduce certificate-based mutual authentication, EAP-TLS key establishment, and cryptographic protection of cyclic and acyclic traffic. However, recent studies reveal that secure PROFINET Application Relations (ARs) introduce non-negligible timing overheads—ranging from tens to hundreds of milliseconds depending on the cryptographic configuration—posing challenges for time-critical industrial processes and highlighting the need for empirical evaluation under realistic conditions [11]. Additionally, attempts to extend PROFINET communication beyond the factory network, such as through VPN or cloud-based controllers, expose further limitations and security risks due to PROFINET's reliance on OSI Layer-2 frames and its lack of native end-to-end protection mechanisms [12]. Together, these industrial pain points underscore the relevance of testbeds capable of reproducing PROFINET communication, supporting controlled

https://doi.org/10.3390/app16073533

signal manipulation, and enabling the assessment of security-related behaviors under operationally realistic scenarios.

Direct experimentation with MitM attacks in live ICS remains limited due to operational risks, high costs, and constraints on process interruption. This has motivated the development of cyber–physical testbeds. However, many existing platforms focus exclusively on network traffic or reproduce simplified scenarios, lacking the capability to actively manipulate input and output signals. Recent contributions highlight the need for platforms that integrate industrial communications with the physical behavior of the process [13–15].

The combination of real industrial hardware and digital twins offers a balanced and safe approach for studying cyber threats. Prior analyses demonstrate the value of such environments for analyzing attacks and detection mechanisms without risking production systems [16,17]. Experimental studies further show that MitM attacks remain feasible even in hardened environments, allowing attackers to modify commands and process data without detection [18].

Within this context, the present work introduces a hybrid testbed based on a digital twin capable of executing practical Man-in-the-Middle (MitM) manipulation techniques in industrial environments. The platform integrates a physical PLC (PLC 1), a second PLC configured as a PROFINET I/O iDevice (PLC 2), and a dedicated manipulation module that injects falsified values into the RTC1 process-data payload. Rather than operating as a physical network intermediary, the attack module cooperates with PLC 2 to alter the process truth before it is delivered to PLC 1, enabling real-time modification of sensor readings or actuator states without placing the physical infrastructure at risk.

Unlike contributions focused solely on detection techniques, this work centers on an experimental infrastructure that enables direct manipulation of process behavior through selective modification of I/O signals, with full control over the duration, persistence, and scope of the injected values. This capability makes the platform a versatile tool for applied research, dataset generation, and industrial cybersecurity training.

The main technical contributions of this work are as follows:

(CT1) A hybrid industrial testbed combining real hardware (Siemens PLCs (Siemens AG, Munich, Germany) and a PROFINET iDevice) with a digital twin, enabling safe and realistic experimentation on PROFINET I/O communications.

(CT2) An active MitM attack module capable of manipulating industrial signals in real time through direct modification of the RTC1 Service Data Unit (SDU)—the PROFINET I/O payload containing process-data bytes together with the standard IOPS/IOCS status indicators—without altering network traffic.

(CT3) A set of representative attack scenarios that demonstrate the scope, flexibility, and practical relevance of the proposed platform.

(CT4) Integrated monitoring through port mirroring, enabling the capture and analysis of PROFINET I/O traffic during attacks using tools such as Wireshark, and facilitating the study of how signal manipulation affects PROFINET cycle timing—a topic rarely addressed in the current literature.

1.1. Scientific Contribution

This work provides a scientific contribution at three complementary levels. First, it empirically demonstrates that semantic manipulation of PROFINET I/O at the SDU layer can alter the behavior of an industrial process without affecting RTC1 cycle timing, protocol-level validity indicators (IOPS/IOCS), or availability metrics. This reveals a critical blind spot in current ICS integrity validation mechanisms, which ensure structural and temporal correctness but do not verify the semantic coherence of process data values.

https://doi.org/10.3390/app16073533

Second, the study introduces a generalizable experimental methodology that integrates physical process modeling, real hardware execution, SDU-level payload manipulation, and multi-layer validation (RTT, throughput, cycle timing, and GEMMA-based behavioral analysis). This methodology is not tied to the specific PLC–digital-twin pair used here and can be adapted to other industrial communication protocols (e.g., PROFIBUS, INTERBUS, DeviceNet, etc.), as long as an intermediate control node provides access to process data manipulation.

Third, the proposed platform offers a reproducible research framework capable of generating synchronized physical, logical, and network-level datasets for evaluating anomaly-detection approaches, semantic-validation techniques, and intelligent intrusion-detection mechanisms. Taken together, these capabilities position the testbed as a scientifically grounded instrument that enables the development of intelligent and machine-learning-based protection strategies for industrial cyber-physical systems.

1.2. Research Questions

Building on the gaps identified above, this work addresses the following research questions:

- RQ1. Does the proposed hybrid testbed provide sufficient flexibility to reproduce realistic PROFINET I/O communication and enable controlled Man-in-the-Middle (MitM) manipulation of industrial signals without risking physical infrastructure?
- RQ2. Does the execution of MitM attacks in the proposed testbed introduce measurable deviations in standard communication metrics—such as ICMP round-trip times, throughput, or PROFINET cycle timing—compared to baseline operation?
- RQ3. Is the testbed capable of generating reproducible datasets suitable for supporting anomaly detection, attack analysis, and training activities in industrial cybersecurity?

The remainder of this article is structured as follows: Section 2 reviews related work on ICS cybersecurity and MitM attack modeling. Section 3 describes the architecture of the testbed. Section 4 presents the implementation of the MitM attack, including the attack logic, the preceding reconnaissance phase, and the experimental scenarios implemented in the testbed. Section 5 analyzes the experimental results. Section 6 discusses the findings and their implications. Finally, Section 7 summarizes the conclusions and outlines future research directions.

2. Related Work

The increasing digitalization of Industrial Control Systems (ICS) and SCADA environments has substantially expanded their exposure to cyberattacks, particularly those targeting or intercepting communication between critical devices. For more than a decade, studies have shown that SCADA systems exhibit structural vulnerabilities in both protocols and architectures, mainly due to the absence of authentication, integrity, and encryption mechanisms in widely used protocols such as Modbus TCP. Early analyses demonstrated multiple exploitable weaknesses in SCADA testbeds, showing that denial-of-service, replay, and traffic-manipulation attacks can easily compromise Modbus-based systems [3]. Subsequent work confirmed that these vulnerabilities persist because the protocol lacks essential security primitives [4].

Recent literature has examined these deficiencies more comprehensively. A systematic assessment of Modbus TCP from an ICT perspective highlights that the absence of cryptographic protections enables impersonation, frame manipulation, and tampering with industrial traffic [4]. These findings reinforce the notion that Modbus TCP continues to represent one of the most vulnerable components in modern ICS architectures.

In parallel, the scientific community has progressed toward developing testbeds capable of reproducing realistic attack scenarios, since experimentation on operational

https://doi.org/10.3390/app16073533

industrial systems carries high risk, high cost, and operational constraints. Recent work demonstrates that such environments are essential for assessing attack impact, generating reliable datasets, and validating detection strategies [13–15]. One line of research presents cyber–physical testbeds that integrate both physical-process behavior and industrial communication in real time, while other contributions combine electrical simulators with Modbus-based communication layers to study operational impacts under representative conditions.

Efforts toward high-fidelity experimental platforms continue to advance the field. A recent physical testbed built on Mitsubishi R04ENCPU PLCs and SCALANCE industrial switches executed seven categories of attacks—including DoS, ARP spoofing, replay, scanning, and ARP-based MitM—to produce a realistic dataset exceeding 2.6 million labeled records [8]. The study showed that virtual datasets fail to capture the temporal, electrical, and communication dynamics of operational ICS, reinforcing the need for physical testbeds when evaluating advanced cybersecurity techniques.

While public datasets such as ICSCASD MPLC [8] provide valuable flow-based network traffic captured from real industrial testbeds—including TCP/UDP communications, ARP-based attacks, and multiple intrusion scenarios—they operate at the packet and flow-statistics level and do not expose real-time industrial I/O dynamics. In contrast, the dataset generated in this work targets a different analytical layer: cyclic PROFINET I/O RTC1 communication, SDU-level payloads, and the direct manipulation of DI/DO signals synchronized with the evolution of the physical process. This cyber–physical alignment, supported by a hybrid digital-twin plus real-PLC environment, provides process-level features not present in existing public datasets and enables the study of fine-grained, real-time manipulation attacks on PROFINET-based industrial I/O.

Among the threats affecting ICS, Man-in-the-Middle (MitM) attacks remain especially critical due to their ability to alter operational data flowing between PLCs, HMIs, and field devices while preserving protocol-level validity indicators, often without generating clear signs of compromise. Existing studies report that MitM and replay attacks pose major challenges for detection methods relying solely on network analysis [6], and later work confirms that MitM attacks can modify process values without causing detectable anomalies at the network level. These findings underscore the need for detection approaches that combine network-traffic information with process-level data [7].

MitM attacks have also been validated experimentally in real industrial environments. Experiments conducted on hardened systems have also shown that Man-in-the-Middle techniques remain feasible even in environments with encrypted control traffic, demonstrating operational impact and system response [18]. This provides strong evidence that MitM attacks remain feasible even in systems equipped with additional protection mechanisms.

Other contributions have examined DoS and basic MitM-related attacks on real PLCs in energy settings. Analyses performed on Siemens S7-1200 and Schneider M241 devices showed that both controllers become inaccessible under port-saturation attacks and that the lack of encryption enables the extraction of sensitive information using tools such as Wireshark [19]. However, these studies focus on low-level network effects and do not address selective manipulation of process data or alterations in cyclic, real-time industrial protocols such as PROFINET I/O.

In distributed energy systems, the impact of MitM attacks has also been examined in DERMS platforms, where manipulation of critical data was shown to cause operational instability in modern infrastructures [5]. Although focused on a different domain, these findings reinforce the relevance of studying integrity-focused attacks in industrial environments.

Beyond attack characterization, several studies propose detection mechanisms. Random-Forest-based approaches have been used to identify MitM attacks in advanced

https://doi.org/10.3390/app16073533

metering infrastructures [20], while LSTM autoencoders have demonstrated strong performance in detecting anomalies even when manipulated values appear legitimate at the network level [7]. Additional work reports superior performance from tree-based models when analyzing real SCADA datasets [21]. However, these techniques are generally evaluated on protocols such as Modbus/TCP and do not address semantic manipulation of cyclic, real-time industrial traffic such as PROFINET I/O.

Overall, the reviewed literature highlights three main trends:

1. Widely deployed industrial protocols—particularly Modbus/TCP—continue to exhibit significant vulnerabilities to manipulation and lack built-in integrity mechanisms.
2. Hybrid and cyber–physical testbeds have become essential for safe experimentation and dataset generation.
3. Research is increasingly combining network-level information with process-level behavior and machine-learning techniques to enhance detection capabilities.

However, none of these contributions address PROFINET I/O in its real-time RTC1 modality, and no prior work evaluates Man-in-the-Middle attacks that selectively manipulate the SDU before cyclic transmission. This gap motivates the hybrid platform proposed in this study.

To address this limitation, the proposed platform enables semantic manipulation of PROFINET I/O at the SDU level using real hardware and cyclic I/O exchange. Unlike prior approaches, the testbed integrates physical-process modeling, controlled signal manipulation, and synchronized physical, logical, and network-level monitoring, providing a unique foundation for evaluating the impact of semantic perturbations within real-time industrial communication. Furthermore, while existing detection-oriented studies rely mostly on high-level flow statistics or packet traces, the proposed platform generates aligned datasets that capture process semantics, controller behavior, and PROFINET RTC1 traffic simultaneously. These characteristics differentiate the proposed approach from State-of-the-Art systems and contextualize its contribution within current ICS cybersecurity research.

3. Testbed Architecture

The architecture of the testbed follows a hybrid, open, and flexible approach that integrates real industrial hardware, a digital-twin representation of the process implemented in PCSimu [22], and a logical manipulation module that operates on the signals exchanged between the twin and PLC 2. The main objective is to provide a safe and reproducible environment for evaluating Man-in-the-Middle attacks that target the integrity of process data, as well as for dataset generation and training activities. The system is organized into three functional layers—virtualized process, distributed control, and supervision—whose interaction enables the recreation of realistic operating conditions without jeopardizing any physical infrastructure. Figure 1 shows the overall architecture of the testbed.

3.1. System Overview

The testbed is organized into three functional levels that reproduce the interaction between the virtualized process, the distributed control layer, and the system-level supervision.

Level 0 consists of the digital twin and PLC 2, configured as a PROFINET I/O iDevice, which together form the virtual representation of the plant. The twin reproduces the dynamic behavior of the process, and exchanges simulated I/O values with PLC 2 through PUT/GET (S7 communication). Using these signals, PLC 2 generates the process data later delivered to PLC 1 over the PROFINET I/O channel, emulating a real decentralized device. In addition, PLC 2 runs the standard MB_SERVER block, exposing Modbus registers that the Python application can read and write. These registers support the logical MitM mechanism: in each PLC cycle, internal logic merges the real signals from PCSimu with the

https://doi.org/10.3390/app16073533

manipulated values supplied by Python, selecting which bits are forwarded depending on the attack configuration.

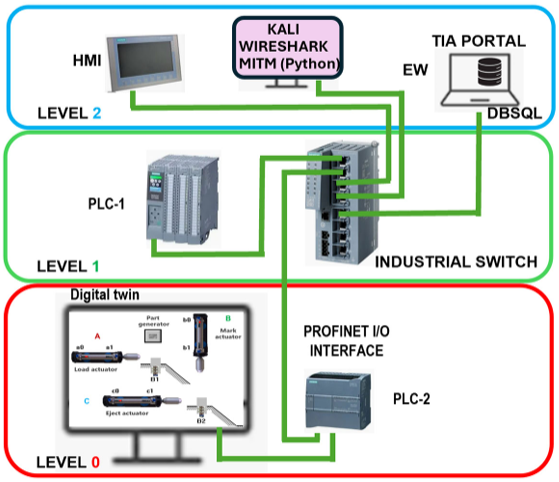


Figure 1. General Testbed Architecture.

Level 1 consists solely of PLC 1, which acts as the main controller of the process. It executes the automation logic using the signals received from PLC 2 over the PROFINET I/O channel, without performing additional integrity checks. As a result, any logical alteration introduced at PLC 2 is directly reflected in PLC 1's control cycle.

Level 2 integrates the supervision and monitoring components: the HMI for visualization, the engineering workstation with TIA Portal [23], the optional MSSQL server for event logging, and analysis tools such as Kali Linux and Wireshark used during the discovery and evaluation phases. This layer also hosts the Python application that implements the logical MitM module, which communicates with PLC 2 through Modbus TCP. By reading and writing the registers exposed by MB_SERVER, the application cooperates with the PLC logic to determine which manipulated or real values are sent to PLC 1—and thus influence its perception of the process. This layer also contains the industrial switch that interconnects all devices in the testbed.

This layered design reproduces a realistic ICS environment by combining industrial hardware, an operational digital twin, and a controlled mechanism for logical signal manipulation. Together, these elements ensure safe experimentation and enable precise analysis of the effects of SDU-level attacks.

Real-Time Performance Parameters

To ensure authenticity and reproducibility, the testbed operates under a set of real-time execution parameters required to characterize its behavior. The PROFINET RTC1 communication between PLC 1 and PLC 2 uses a 2 ms update cycle, corresponding to the default value automatically assigned when the S7-1200 is configured as an I-Device. During operation, PLC 1 (S7-1500) executes the OB1 cyclic program with a typical scan time of 5.4 ms (range: 3.6–5.4 ms, with occasional peaks caused by communication servicing), while PLC 2 (S7-1200) exhibits a typical scan time of 4.0 ms (range: 2–5 ms).

The digital-twin module (PCSIMU) exchanges process values with PLC 2 through PUT/GET communication services, which only requires enabling the option "Permit access with PUT/GET communication from remote partner" in the PLC settings. PCSIMU does not expose a deterministic internal execution cycle. Therefore, the exact timing between

https://doi.org/10.3390/app16073533

updates cannot be retrieved. In practice, PCSIMU simulates the process logic: when PLC 1 sends commands to the plant (e.g., YA+) by energizing the solenoid valve that controls an actuator, the virtual actuator in PCSIMU performs the corresponding movement according to its configured valve type and speed, and once the end position is reached, the associated digital-input signal (a1) is activated. The time required for each transition is not deterministic and may vary between cycles—just as in real industrial installations, where pneumatic or mechanical dynamics introduce natural fluctuations. What remains relevant is the sequence integrity, while the PROFINET RTC1 channel enforces the deterministic 2 ms timing under which the MitM scenarios are executed.

3.2. Hardware Components

The testbed integrates the following physical elements:

- PLC 1 (main controller): Acts as the central controller and executes the automation logic based on the signals provided by PLC 2 through the PROFINET I/O channel. Since it performs no additional integrity checks, any modification introduced upstream is directly reflected in its control cycle.
- PLC 2 (PROFINET iDevice): Configured as a PROFINET I/O device, it cyclically exchanges simulated sensor and actuator states with the digital twin via PUT/GET (S7 communication), emulating the behavior of a decentralized peripheral. PLC 2 also runs the MB_SERVER block, exposing Modbus TCP registers to the Python application. By reading and writing these registers, the Python module cooperates with the PLC logic to determine which real or manipulated values are inserted into the process-data payload ultimately delivered to PLC 1.
- Industrial switch: Provides the Ethernet communication backbone interconnecting PLC 1, PLC 2, the HMI, the engineering workstation, and the Python module. It is not used as an attack-insertion point, but simply as the physical networking component required for coordinated system operation.
- HMI: Enables process supervision, status visualization, and operator interaction. Its display reflects either the real or manipulated behavior of the system, depending on the values propagated through the controllers during each scenario.

Although these hardware components form the physical foundation of the testbed, their behavior depends on a set of software modules that virtualize the process, coordinate communication between controllers, and implement the logical MitM attack. These modules are described in the following section.

In addition to these industrial devices, the testbed also includes three host PCs (Engineering Workstation, Digital-Twin Host, and Attacker PC), which support the execution of TIA Portal, PCSIMU, event-logging services, and the Python-based MitM module.

To ensure full reproducibility, Table 1 lists the complete hardware configuration of the testbed, including the exact models, manufacturer information, and roles of all physical devices described in Section 3.2. This information provides the necessary level of detail for replicating the hardware setup used in the experiments.

Table 1. Hardware components used in the testbed.

Component	Model	Manufacturer	Role in the Testbed
PLC 1	6ES7512-1CK01-0AB0	Siemens AG, Munich, Germany	Main Controller
PLC 2	6ES7215-1BG40-0XB0	Siemens AG, Munich, Germany	PROFINET I-Device

https://doi.org/10.3390/app16073533

Table 1. *Cont.*

Component	Model	Manufacturer	Role in the Testbed
HMI Panel	6AV2 123-2JB03-0AX0	Siemens AG, Munich, Germany	HMI panel
Switch	6GK5208-0BA00-2AC2	Siemens AG, Munich, Germany	Industrial Ethernet
Engineering PC	PC (Intel i7, 32 GB RAM)	-	TIA Portal, MSSQL
Digital Twin Host	PC (Intel i5, 8 GB RAM)	-	Runs PCSIMU
Attacker PC	PC (intel i7, 16 GB RAM)	-	Python MitM, Wireshark

3.3. Software Components

The system is complemented by the following logical components, distributed across three independent host machines as depicted in Figure 1:

- Engineering Workstation: Host platform used for TIA Portal programming, network diagnostics, PROFINET monitoring, and execution of tools such as MSSQL Server Express for event-log storage, data analysis, and dataset generation. Although it runs on a standard PC rather than industrial hardware, it provides the supervisory and analysis environment required for experimentation.
- Digital twin: Executed on a dedicated PC using the industrial simulation tool PCSimu, it reproduces the dynamic behavior of the process and updates the process image of PLC 2 through PUT/GET (S7 communication). Through this mechanism, PCSimu provides virtual sensors and actuators whose signals are transferred to PLC 2 as if they originated from a real peripheral device. These signals—either genuine or manipulated—are then forwarded to PLC 1 via PROFINET I/O.
- Man-in-the-Middle (MitM) module–Python application (Attacker PC): Run on a separate host with network access to PLC 2, the Python application implements a logical MitM mechanism without acting as a network proxy between PLC 1 and PLC 2. It communicates with PLC 2 over Modbus TCP through the MB_SERVER block, reading and writing the registers that contain the packaged process signals. In collaboration with internal PLC 2 logic, the module determines whether the values received from PCSimu are propagated unchanged or replaced by manipulated values (e.g., bit masks or forced states). This modification alters the "process truth" that PLC 2 subsequently forwards to PLC 1 over PROFINET I/O.
- Monitoring and Reconnaissance Tools (Attacker PC): Tools such as Wireshark, Nmap/Kali Linux, and diagnostic utilities are executed on the attacker PC. They support reconnaissance, validation of the manipulation mechanism, and packet-level analysis of the PROFINET I/O exchange through mirrored-port captures.

These software components provide the supervisory, simulation, manipulation, and monitoring capabilities required for executing and analyzing the attack scenarios.

Table 2 summarizes the software components and version information required to reproduce the experimental environment, complementing the detailed description provided in Section 3.3. These specifications ensure that the testbed configuration and the MitM experiments can be replicated accurately.

https://doi.org/10.3390/app16073533

Table 2. Software components used in the testbed.

Component	Version	Platform/OS	Role in the Testbed
TIA Portal	V18 Update 5	Windows 11	PLC engineering
PCSIMU	V3.0	Windows 10	Digital twin
Python	3.13.5	Windows 10	MitM module
Pymodbus	3.11.4	Python	Modbus client
MB_SERVER block	Standard lib.	PLC S7 1200	Modbus server
MSSQL Server Express	2019	Windows 11	Event logging
SSMS	20.2.30.0	Windows 11	SQL administration
Wireshark	4.6.3	Windows 10	Network capture
Kali Linux tools	Rolling	Linux	Reconnaissance

3.4. Interfaces and Protocols

The testbed integrates several industrial communication channels that accurately reproduce the behavior of a real ICS environment. Each channel fulfills a specific role within the data flow between the digital twin, the controllers, and the Python-based manipulation module.

- PROFINET I/O (PLC 1 ↔ PLC 2): The communication between PLC 1 and PLC 2 via PROFINET I/O constitutes the cyclic I/O channel on which the execution of the control logic relies. In this configuration, PLC 1 acts as the IO-Controller and PLC 2 as the iDevice, meaning that PLC 1 receives the signals generated or manipulated at PLC 2 as if they originated from a real decentralized peripheral.
- RTC1 cyclic communication and the Provider–Consumer model: The PROFINET I/O exchange is based on 40-byte RTC1 telegrams whose fields follow the Provider–Consumer model. Each submodule transmits process data together with two status indicators: the IO-Provider State (IOPS) and the IO-Consumer State (IOCS). IOPS indicates the validity of the data provided by the sender, whereas IOCS specifies whether the receiver has processed it correctly. As a direct consequence of this model, the SDU layout depends on the direction of the flow: PLC 2 → PLC 1: IOPS–Process Data–IOCS. PLC 1 → PLC 2: Process Data–IOPS–IOCS. This organization matches the RTC1 captures obtained in the testbed (Figure 2) and corresponds to the cyclic-update mechanism described in the official PROFINET IO-Device stack specification [24]. Since IOPS/IOCS are automatically managed by the controllers, the Python module can modify only the process-data bytes without affecting protocol-level validity.
- Modbus TCP (Python module ↔ PLC 2): The Python application communicates with PLC 2 through the MB_SERVER block, which exposes a set of Holding Registers containing the packaged process signals. By reading and writing these registers, the application injects manipulated values before PLC 2 assembles the RTC1 payload and forwards it to PLC 1, forming the insertion point of the logical MitM mechanism.
- S7 PUT/GET Communication (PCSimu ↔ PLC 2): PCSimu exchanges virtualized sensor and actuator states with PLC 2 using PUT/GET operations (S7 Comm), reading and updating the PLC's process image through standard I/O addresses (e.g., I0.0, Q0.0). This mechanism allows the digital twin to emulate process behavior without requiring additional physical equipment.

The digital twin is not synchronized with a physical plant but instead operates in a simulated environment, allowing controlled evaluation of SDU-level manipulations and anomalous behaviors in a safe and reproducible setting.

https://doi.org/10.3390/app16073533

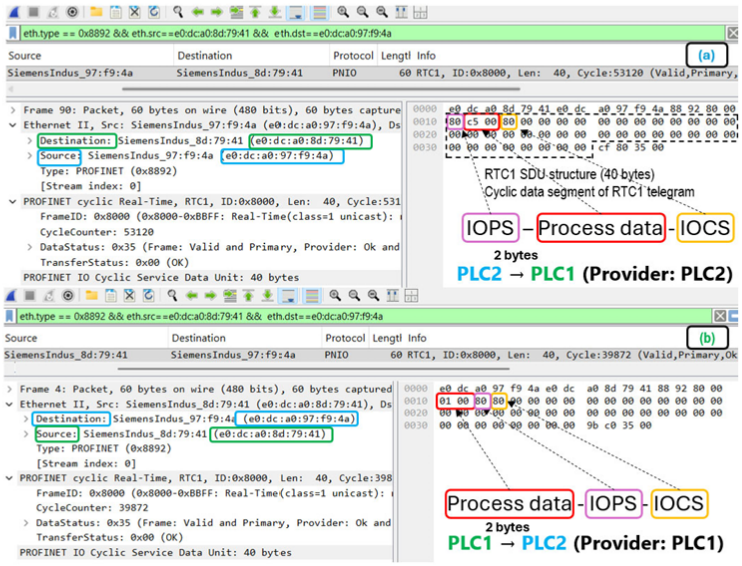


Figure 2. RTC1 cyclic telegrams in both directions: PLC2→PLC1 (**a**) and PLC1→PLC2 (**b**).

3.5. Communication Flow and MitM Insertion Point

The operation of the testbed during a MitM experiment is structured around a communication flow in which the manipulation does not occur on the network between the controllers, but rather at the point where PLC 2 merges the signals originating from the digital twin with the values supplied by the Python application. This results in a logical MitM, where the alteration is applied to the "process truth" before it is transmitted to PLC 1 over PROFINET I/O.

The operational sequence proceeds as follows. First, PCSimu updates the process-image values in PLC 2 through PUT/GET (S7 communication), writing virtualized sensor and actuator states directly into its memory. Next, the Python application accesses the Holding Registers exposed via MB_SERVER to read and modify the data words where the process signals are packaged. PLC 2 then executes an internal logic block that combines, cycle by cycle, the values received from PCSimu with the manipulated values coming from Python. Using bit masks, the PLC determines which bits are preserved and which are replaced.

The resulting signals—whether genuine or altered—are then sent to PLC 1 via PROFINET I/O, where they are interpreted as if they originated from a real decentralized peripheral. Since PLC 1 does not apply additional integrity checks, it incorporates these values directly into its control logic with no ability to detect prior manipulation.

This design, based on coordinated cooperation between the Python module and PLC 2, enables a MitM attack without modifying the network topology or intercepting traffic between PLCs. By acting directly on the process data exchanged between the digital twin and PLC 2, the system reproduces the effect of an integrity-manipulation attack in a precise and controlled manner, while preserving the determinism and stability of the experimental environment.

Figure 3 graphically summarizes this flow, contrasting normal operation with the scenario in which the logical MitM mechanism is active.

https://doi.org/10.3390/app16073533

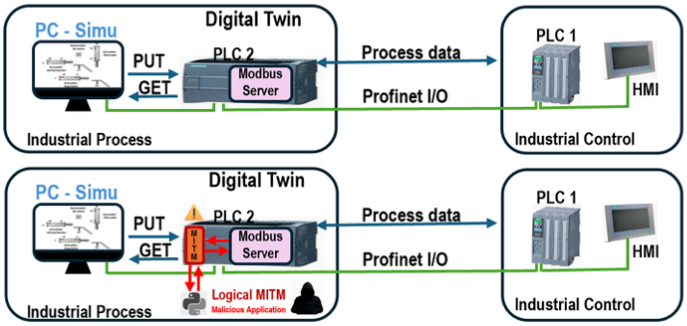


Figure 3. Communication flow between PCSimu, PLC 2, and PLC 1 during normal operation (**top**) and under a logical MitM attack (**bottom**).

3.6. Generalizable Framework for Semantic Manipulation and Multi-Layer Validation

Building on the architectural elements described above, the proposed experimental setup can be abstracted into a generalizable framework that extends beyond the specific PLC–digital-twin configuration used in the testbed. This framework is organized into four methodological pillars that can be replicated across a wide range of industrial control environments.

- Physical-Process Modeling or Acquisition: The framework begins with a representation of the physical process, which may be implemented using a digital twin, a hardware-in-the-loop simulator, or a real industrial asset. Although the present implementation focuses on digital input/output signals, the methodology is not restricted to this subset; it can be extended to any process variable, whether digital, analog, or structured, whose semantics influence the control sequence.
- Intermediate Control Node for Data Manipulation: A key element of the framework is the presence of an intermediate controller capable of merging genuine process signals with externally supplied manipulated values. In the proposed testbed, this role is fulfilled by PLC 2, but equivalent functionality can be achieved with any device that offers cyclic I/O exchange and provides programmatic or register-level access to its process image. This abstraction allows the methodology to be adapted to other industrial communication protocols (e.g., PROFIBUS, INTERBUS, DeviceNet), as long as the intermediate node can intercept and selectively rewrite I/O data before it reaches the primary controller. As the framework is not tied to a particular vendor or control architecture, the primary controller may be implemented using PLCs or PACs from manufacturers such as Allen-Bradley, Siemens, Omron, and others.
- Semantic Manipulation Layer: The logical MitM mechanism operates at the semantic level, modifying I/O data without altering the underlying communication frames. This layer can be implemented through Modbus/TCP, vendor-specific APIs, or any interface that enables controlled injection of process-data values. The abstraction relies on the ability to manipulate the process "truth" upstream of the consumer controller, independent of the communication stack used for cyclic exchange.
- Multi-Layer Monitoring and Validation: To ensure reproducibility and scientific rigor, the framework integrates physical-state logging, controller-behavior monitoring, and network-traffic acquisition. This multi-layer monitoring and validation approach provides synchronized evidence of both the semantic manipulation and its consequences on the industrial process. In addition, structured event logging through a database system (e.g., MSSQL, MySQL, PostgreSQL, or an equivalent solution) enables precise correlation between GEMMA transitions, manipulated values, and PROFINET RTC1 traffic. The design is independent of the specific controller or industrial communi-

https://doi.org/10.3390/app16073533

cation protocol employed, enabling the framework to support comparative studies across heterogeneous ICS environments.

Overall, this generalizable framework formalizes the methodology introduced in this work and demonstrates how the proposed architecture can be transferred to different control platforms, communication buses, and cyber-physical system configurations, both industrial and non-industrial. In this way, it provides a reusable foundation for the systematic study of semantic manipulation, anomaly detection, and intelligent protection mechanisms across a wide range of environments.

4. Man-in-the-Middle Attack Implementation

This section describes how the Man-in-the-Middle (MitM) attack is implemented within the hybrid digital-twin-based testbed developed for the industrial part-marking system. The attack does not intervene in the network between PLC 1 and PLC 2; instead, it operates as a logical MitM applied at the point where PLC 2 combines the virtualized signals from the digital twin with the manipulated values provided by the Python application. In this way, the alteration affects the "process truth" received by PLC 1 over PROFINET I/O while leaving the underlying communication layers intact.

In a real system, an attacker would achieve this level of access only after an initial compromise of the OT network or an engineering workstation. Once inside the subnet, PLC 2 becomes discoverable through standard reconnaissance techniques such as ARP scanning, TCP SYN probing, or PROFINET DCP identification. If a maintenance service such as Modbus TCP (port 502) remains unintentionally exposed and configured to allow write operations, the attacker can legitimately interact with the PLC's memory areas. This represents a realistic post-compromise scenario frequently observed in industrial environments and provides the required level of access for the logical MitM manipulation demonstrated in our testbed.

Building on the architecture detailed in Section 3, the description below covers the operation of the process, the digital signals involved, the target of the manipulation, and the mechanism through which the attack logic selectively modifies input and output values. The activation logic, signal-mask configuration, and coordination between the Python module and PLC 2 are also detailed to illustrate how controlled falsification of operational data is achieved.

4.1. Industrial Process Description

The process corresponds to a part-marking system represented through a digital twin developed in PCSimu. The digital twin provides a functional simulation of the process, generating virtual sensor and actuator signals with timing and sequencing behavior comparable to an industrial cell, but without reproducing the full dynamics of a real physical system.

As illustrated in Figure 4, the process includes three actuators (A, B, and C), each equipped with magnetic position sensors (a0, a1, b0, b1, c0, c1), together with two inductive detectors (B1 at the loading station and B2 at the marking station) and a workpiece generator that feeds the production cycle. The digital twin exchanges virtualized sensor and actuator states with PLC 2 via PUT/GET (S7 communication), with PLC 2 configured as a PROFINET I/O iDevice that forwards these values to PLC 1, which acts as the main controller. This configuration reproduces a decentralized control architecture suitable for controlled experimentation in a safe environment.

Interaction with the system is performed through an HMI that allows adjustment of operational parameters—such as marking time, number of impacts per workpiece, and batch size—selection between manual and automatic modes, issuing extend/retract

https://doi.org/10.3390/app16073533

commands to the actuators (OUT_x/IN_x, with x = A, B, C), and controlling start, stop, reset, rearm, and acknowledge operations. The system also incorporates safety elements, including an emergency stop and a four-color signal beacon. From a control-system perspective, the process follows the GEMMA methodology, with stop (A), operation (F), and fault (D) modes.

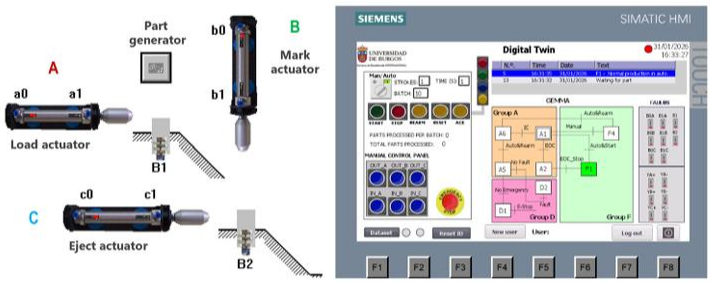


Figure 4. PCSimu digital twin (**left**) and corresponding HMI panel (**right**).

To facilitate a better understanding of the process, Figure 5 presents the partial GRAFCET corresponding to the automatic normal production operating mode, in accordance with the GEMMA guide.

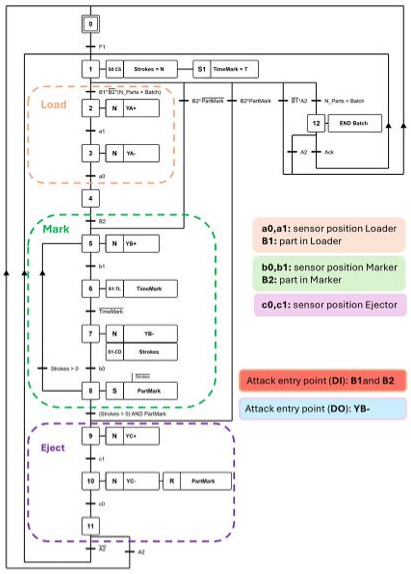


Figure 5. Partial GRAFCET representation of the automatic production cycle (F1).

In this GRAFCET, it can be observed that the initial step remains waiting for mode F1. Once mode F1 is activated, step 1 proceeds to load the parameters set by the operator for the part (strokes and time mark).

Next, the system evaluates the state of the marking unit to determine whether it is necessary to load a part into the marking station (detector B2 inactive) and, in that case, when a part is available at the loading station (detector B1 active), it issues the YA+ command to the solenoid valve so that actuator A advances and performs the part loading.

Once confirmation of the forward position is received (sensor a1 active), the controller withdraws YA+ and commands the return to the initial position via YA−.

https://doi.org/10.3390/app16073533

Subsequently, when the part is detected by B2 at the marking station, the marking phase begins by activating YB+ until its magnetic sensor b1 is triggered. At that moment, YB+ is withdrawn, and the actuator remains in that position until the configured marking time elapses; afterwards, YB− is commanded to return to the initial position (b0 active), either to execute a new impact or to proceed to the part evacuation phase by means of YC+ and YC−.

Once the part has been evacuated, the controller evaluates whether the cycle should continue until the batch is completed or whether there is a stop request at the end of the cycle from the operator (GEMMA state A2).

As can be observed, based on the program logic and the information received from sensors and detectors (inputs), the controller issues commands (outputs) to the solenoid valves in order to actuate the actuators. These I/O signals constitute the target of the MitM attack described in the following sections, whereas HMI-level commands remain outside the scope of manipulation. Nevertheless, Figure 5 highlights the input and output signals that are targeted by the attack.

4.2. Attack Motivation and Target I/O Signals

The logical MitM attack is oriented toward the selective manipulation of the digital input (DI) and digital output (DO) signals exposed by PLC 2 through its Modbus TCP server (MB_SERVER). The DI values correspond to virtual sensor states generated by the digital twin, while the DO values originate from PLC 1 as actuator-command signals. Both sets of signals are packaged in PLC 2 and can be modified by the Python application before PLC 2 generates the final values to be transmitted through the PROFINET I/O channel—either incorporating DI values received from the digital twin or forwarding DO commands originating from PLC 1, depending on the direction of the flow. The specific DI and DO signals that can be targeted by the attack are summarized in Table 3.

Table 3. Digital signals of the process vulnerable to MitM manipulation.

Process Data	Symbol	Direction	Type	Comment
Digital Input (DI)	a0	I 0.0	Bool	Loader backward (actuator A)
Digital Input (DI)	a1	I 0.1	Bool	Loader forward
Digital Input (DI)	b0	I 0.2	Bool	Marker up (actuator B)
Digital Input (DI)	b1	I 0.3	Bool	Marker down
Digital Input (DI)	c0	I 0.4	Bool	Ejector backward (actuator C)
Digital Input (DI)	c1	I 0.5	Bool	Ejector forward
Digital Input (DI)	B1	I 0.6	Bool	Part detector at the loader
Digital Input (DI)	B2	I 0.7	Bool	Part detector at the marker
Digital Output (DO)	YA+	Q 0.0	Bool	Loader forward
Digital Output (DO)	YA−	Q 0.1	Bool	Loader backward
Digital Output (DO)	YB+	Q 0.2	Bool	Marker down
Digital Output (DO)	YB−	Q 0.3	Bool	Marker up
Digital Output (DO)	YC+	Q 0.4	Bool	Ejector forward
Digital Output (DO)	YC−	Q 0.5	Bool	Ejector backward

The digital inputs from magnetic position sensors (a0, a1, b0, b1, c0, and c1; see Figure 4), together with inductive detectors B1 and B2, allow PLC 1 to determine in real time the position of the actuators and the presence or absence of workpieces in the loading and marking zones. The digital outputs (Yx+ and Yx−, with x = A, B, and C), associated with the extension and retraction commands of the actuators, directly determine the evolution of the automated sequence and contribute to the operational safety of the system.

The manipulation of these signals is particularly critical for three main reasons:

https://doi.org/10.3390/app16073533

1. Immediate impact on the control logic: Altering a single bit can cause PLC 1 to interpret a nonexistent state, leading to unintended movements, incorrect transitions between stages, or even interrupting the work cycle.
2. Low local detection capability: In a decentralized architecture, PLC 1 fully trusts the data provided by PLC 2. Since Modbus/TCP lacks native mechanisms for authentication, integrity verification, or cross-checking, a logical MitM can remain unnoticed as long as the communication format remains valid and the cyclic exchange remains stable.
3. High potential for operational damage: Manipulating sensors or actuators can cause anything from automatic-sequence blockages to mechanical collisions or damage to the workpiece or the machine itself, depending on the experimental scenario.

Taken together, these digital I/O signals represent a high-value target for an attacker aiming to compromise the integrity, coherence, and reliability of the automated process. Their controlled manipulation within the testbed enables the safe evaluation of the effects of a MitM attack in a decentralized industrial environment.

4.3. MitM Attack Logic

The attack is not physically inserted into the network between PLC 1 and PLC 2; instead, it is implemented as a logical MitM on the PCSimu ↔ PLC 2 signal path. This mechanism is achieved through the interaction of two coordinated components:

- A Python application that reads and writes values via Modbus/TCP to the MB_SERVER block in PLC 2.
- An internal logic block in PLC 2 that determines, cycle by cycle, whether the signals exposed to the PROFINET I/O exchange with PLC 1 correspond to the values received from the digital twin or to the values manipulated by Python.

The manipulation is restricted to a subset of digital inputs (DI) and digital outputs (DO) relevant to the marking cycle. For each scenario, the Python application defines a mask specifying which bits are affected and the forced values (0/1). Although the architecture is prepared to support additional parameters—such as persistence, pulse patterns, duty cycle, or timing offsets—these capabilities were not required in the experiments conducted in this work.

The logic embedded in PLC 2 synchronizes the write operations issued by Python, applies the selected mask, and ensures that, when no attack is active, the real signals are forwarded without modification. This mechanism operates symmetrically for DI and DO: DI values originating from the digital twin and DO commands originating from PLC 1 can both be intercepted and modified in PLC 2 before being incorporated into the corresponding data flow—either toward PLC 1 via PROFINET I/O or toward the digital twin via S7 PUT/GET.

This design enables reproducing the effect of a Man-in-the-Middle attack on process data while preserving system stability and ensuring full reproducibility of the experimental scenarios. Figure 6 illustrates the internal operation for the DI-manipulation case, showing how PLC 2 applies the bit-mask logic to merge genuine inputs from the digital twin with the manipulated values supplied by the Python application before exposing the resulting data to PLC 1 via PROFINET I/O. This mechanism establishes the basis for the attack scenarios presented next, starting with the reconnaissance phase described in Section 4.4.



Figure 6. PLC 2-bit mask logic used for DI manipulation in the logical MitM attack.

https://doi.org/10.3390/app16073533

4.4. Pre-Attack Reconnaissance for the MitM Attack

Before activating each manipulation scenario, the initial reconnaissance phase that an attacker would typically perform after gaining access to the testbed's OT network is reproduced. This phase enables the identification of the devices present in the industrial segment, their IP addresses, and the services they expose—without requiring intrusive techniques. The testbed faithfully replicates the information an adversary could obtain through passive traffic inspection or superficial device probing, providing a realistic starting point for the subsequent manipulation of PROFINET I/O signals.

4.4.1. Identification of Devices in the Industrial Network

Figure 7 shows the ARP/Ping scan executed from Kali Linux (nmap-sn), which identifies active devices on the 192.168.0.0/24 subnet. The analysis reveals several industrial components, including the Siemens S7-1200 PLC (192.168.0.2), which exposes the process-related digital signals via PROFINET I/O. This information allows selecting the target on which the RTC1-manipulation scenarios will be carried out in the following sections.

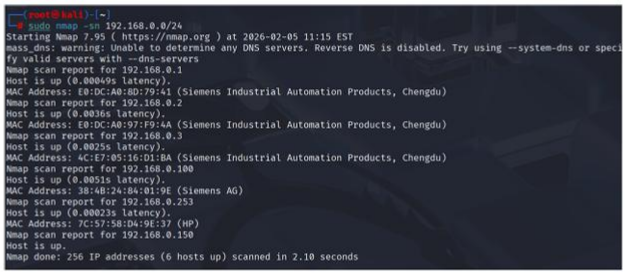


Figure 7. ARP/Ping scan showing active devices in the 192.168.0.0/24 subnet.

4.4.2. Discovery of Open Ports

After identifying the target device's IP address during the initial reconnaissance phase, the next step is to analyze the industrial services exposed by that device. For this purpose, a port scan is performed using Nmap from Kali Linux, focusing on TCP ports 102 and 502, associated with the S7comm (Siemens) and Modbus/TCP protocols, respectively. Figure 8 shows the results of the analysis, confirming that both services are accessible on the Siemens S7-1200 PLC (192.168.0.2). This validates its role as an IO Device and helps determine which communication channels may be manipulated in the subsequent attack scenarios.

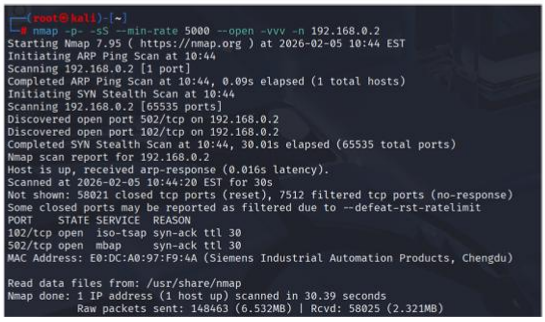


Figure 8. Port scan showing open services on the device.

https://doi.org/10.3390/app16073533

4.4.3. Fingerprinting of Exposed Industrial Services

After identifying the target host, a preliminary inspection of the accessible services is performed to determine the type of device and the industrial protocols enabled. The Siemens S7-1200 PLC located at 192.168.0.2 exposes the following relevant services:

- 102/TCP—S7comm (ISO TSAP): The service response confirms that the device is a genuine Siemens PLC and reveals hardware and firmware information (e.g., Module 6ES7 215-1BG40-0XB0, Firmware 4.3.1). This fingerprinting step confirms the device's identity and its role as a PROFINET I/O Device within the testbed.
- 502/TCP—Modbus TCP: Port associated with the MB_SERVER block running on PLC 2, which exposes the internal registers used by the Python application for DI/DO manipulation. Its presence reflects the hybrid nature of the testbed, where PROFINET and Modbus/TCP coexist for experimental purposes.

The detected MAC address—assigned to Siemens Industrial Automation Products (Chengdu)—further confirms the correspondence between the real PLC hardware and the information obtained during fingerprinting. Figure 9 shows the results of a version-detection scan (nmap-sCV), which identifies the active industrial protocols and provides detailed module and firmware information, relevant for characterizing the device's role within the testbed.

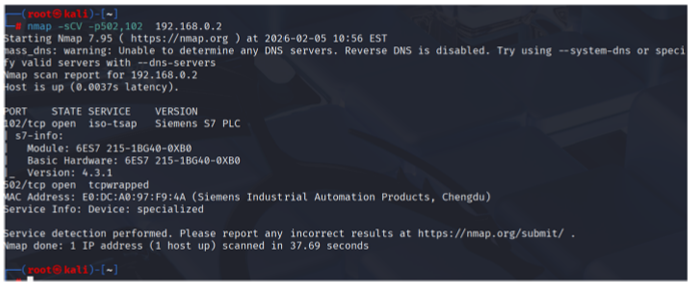


Figure 9. Service fingerprinting scan showing open ports on the device.

4.4.4. Relation to the Manipulation Scenarios

The information collected during the reconnaissance phase directly supports the MitM manipulations developed in the following sections. In particular:

- Identification of the PLC exposing key industrial services: The attacker determines which device publishes ports associated with process control (e.g., PROFINET, S7comm, or Modbus/TCP), allowing the Siemens S7-1200 PLC to be profiled as the primary target for signal manipulation.
- Detection of the Modbus/TCP service: The presence of the Modbus server (MB_SERVER on PLC 2 enables the modification of the DI/DO values packaged in its internal registers, complementing the manipulation performed over the PROFINET I/O channel and providing the insertion point for the logical MitM mechanism.
- Confirmation of the S7comm service and the PLC's role: The exposure of port 102/TCP (S7comm) confirms the device's identity, its communication role within the PROFINET cycle, and the coherence between the fingerprinted services and the RTC1 traffic analyzed during the MitM scenarios.

Based on this information, scenarios A1–A2 are designed to demonstrate how selective manipulation of DI/DO signals transmitted through PROFINET I/O can induce incorrect behaviors in PLC 1 without producing visible inconsistencies at the physical layer of the process.

https://doi.org/10.3390/app16073533

4.5. Experimental MitM Attack Scenarios

This section describes the attack scenarios implemented in the testbed. Each scenario specifies which signals are manipulated, under which conditions the attack is triggered, and the intended impact on PLC 1. The manipulation is carried out by the Python application, which alters a subset of the DI/DO signals exposed by PLC 2 through the Modbus/TCP interface before they are incorporated into the PROFINET I/O cyclic exchange.

4.5.1. Scenario A1—Manipulation of Sensors B1 and B2

This scenario represents an attacker who forces the DI bits corresponding to sensors B1 and B2 while still observing their real values. The manipulation alters the GEMMA-based logic in PLC 1 without generating PROFINET alarms, thereby implementing a realistic DI-level logical MitM attack.

- Objective: Force PLC 1 to advance in the automatic sequence despite the absence of a physical workpiece. The attacker simulates the sequential activation of B1 (workpiece at loading) and B2 (workpiece in the marking area), causing the process to run through loading, marking, and ejection without any real physical intervention.
- Initial State and Affected Signals: The system is in automatic mode (GEMMA F1), waiting for a workpiece at the loading position (see Figure 4). DI bits B1 and B2 are selected for manipulation.
- Mechanism of the MitM Attack: The Python application modifies the DI values of B1 and B2 in the MB_SERVER registers of PLC 2. This PLC then applies its internal bit-mask logic to replace the real DI values with the forged ones before assembling the RTC1 process-data payload that is sent to PLC 1 over PROFINET I/O.

Figure 10 illustrates an example of DI-bit manipulation affecting signals B1 and B2 in PLC 2 during the logical MitM attack.

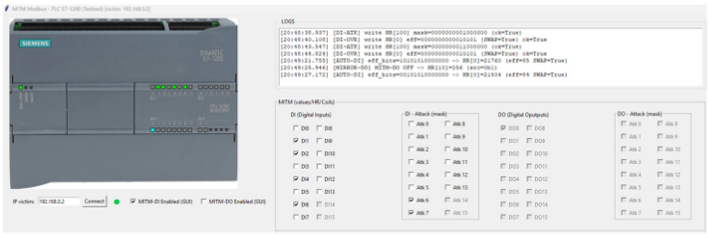


Figure 10. Example of DI-bit manipulation (B1 and B2) in PLC 2.

The effect of the manipulation appears immediately in the internal behavior of the system. Figure 11 shows the online view of PLC 2, where the real process value, the active attack mask, and the overridden value propagated to PLC 1 are displayed simultaneously. This confirms that the altered bits are generated inside PLC 2 and do not originate from the physical process.

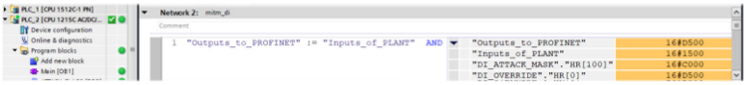


Figure 11. Online view of PLC 2 showing real and manipulated values.

- Effect: PLC 1 remains in GEMMA macro-state F1 (automatic mode) but proceeds incorrectly through the operational sequence (loading → marking → ejection) due to the forced DI bits (Figure 12). PLC 1 interprets the manipulated signals as genuine and executes the full cycle even though no physical workpiece exists.

https://doi.org/10.3390/app16073533

- Evidence in PROFINET I/O traffic: Port-mirrored captures (Figure 13) show that the RTC1 frames contain the manipulated B1/B2 bits in the SDU, even though the digital twin did not generate these events. This confirms that the logical MitM attack modifies the SDU-level content while preserving protocol-level validity.

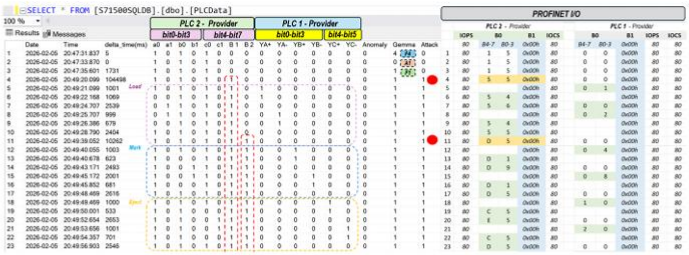


Figure 12. Correlation of SQL Server, GEMMA, and PROFINET data in scenario A1.

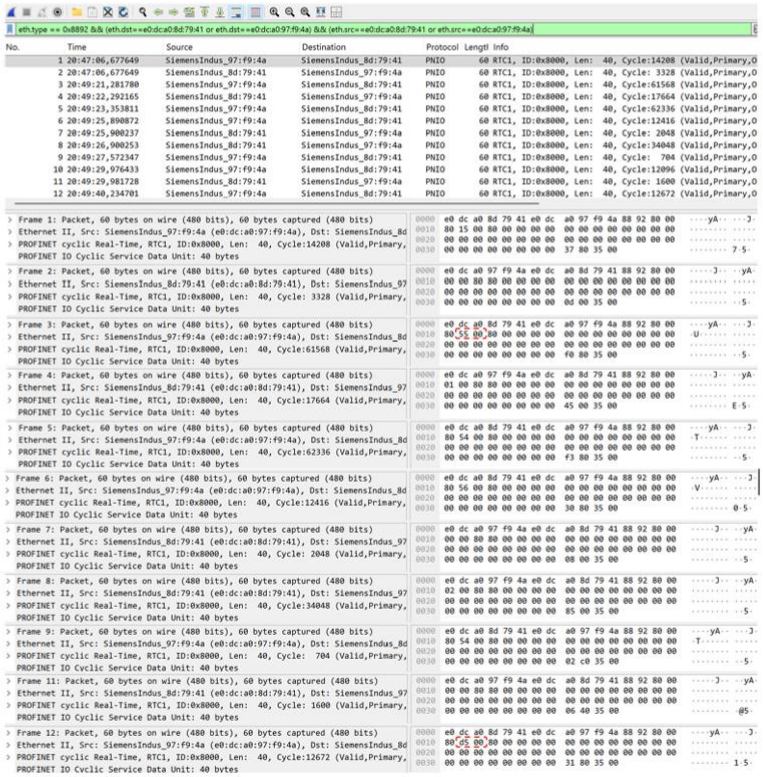


Figure 13. Wireshark capture of PROFINET RTC1 traffic in scenario A1.

Once the affected byte is identified, Figure 13 shows the corresponding forced bits in the RTC1 frames captured in Wireshark, offering clear packet-level evidence of the manipulation propagated to PLC 1.

4.5.2. Scenario A2—Manipulation of Actuator B (Marking Device)

In addition to scenario A1, which focuses on the manipulation of digital input (DI) signals by forcing selected bits, this scenario represents an attacker manipulating one of

https://doi.org/10.3390/app16073533

the digital output bits responsible for actuating the bistable solenoid valve that governs the marking actuator. The attack is carried out by a malicious application through the Modbus/TCP interface, forcing a premature activation of the marking up command while PLC 1 continues issuing its legitimate command. By injecting a false value into the digital output prior to the actual process transition, the attack alters the mechanical behavior of the actuator without triggering PROFINET alarms, thereby implementing a realistic MitM attack at the digital output (DO) level.

- Objective: To prevent the marking operation from being correctly completed during the automatic sequence. By prematurely issuing a false command to the solenoid valve that triggers marking up, the attacker causes the valve to become blocked when the process later sends the legitimate marking down command, due to the reception of conflicting instructions.
- Initial State and Affected Signals: The system is operating in automatic mode (GEMMA F1), and the process is producing parts under normal conditions. Digital output DO 3, responsible for commanding the marking actuator to the up position, is selected for manipulation.
- MitM Attack Mechanism: While the process is running, the Python-based application writes a false value to the Modbus register corresponding to DO 3 = true (YB−) in PLC 2. This artificial activation is injected milliseconds before the legitimate command generated by PLC 1 during the transition to the marking stage (DO 2 = true (YB+). The premature activation interferes with the actual command, preventing the actuator from executing the marking operation and causing PLC 1 to generate a process fault due to an unexpected actuator response.

Figure 14 below shows the attacker interface injecting the malicious activation command through the aforementioned application.

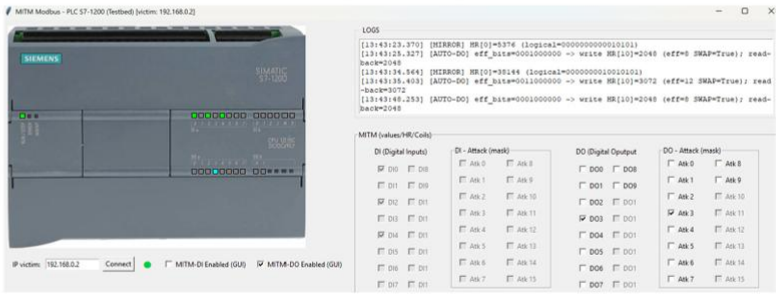


Figure 14. Manipulation of the DO signal YB− in PLC 2.

The outcome of the attack is recorded in the dataset stored in SQL Server, as well as in the PROFINET IO packet captures obtained using Wireshark, which are presented in Figures 15 and 16 below.

These figures, together with the previous one, demonstrate how the testbed captures all process events in real time. It simultaneously records the anomalous conditions generated in the process, the exact correspondence between each applied modification and its physical effect on the control sequence, the GEMMA operating mode in which each deviation occurs, and the type of detected attack (1 and 2 for DI and DO manipulation, respectively).

In this context, Figure 15 illustrates how these recorded variables evolve during the attack through the corresponding GEMMA sequence: F4 (system in manual mode) → A1 (system stopped in the initial state) → F1 (automatic normal production) → D2 (fault) → A5 (preparation for restart after a fault). To facilitate the interpretation of the attack,

https://doi.org/10.3390/app16073533

particular attention is drawn to the state of sensor b0 (indicating the marker in the up position, with no change despite the issued command), the field Attack = 2 (indicating an output-level attack), and the marking command (YB+ = true) issued by PLC 1 to the solenoid valve controlling the marking unit during its automatic production sequence.

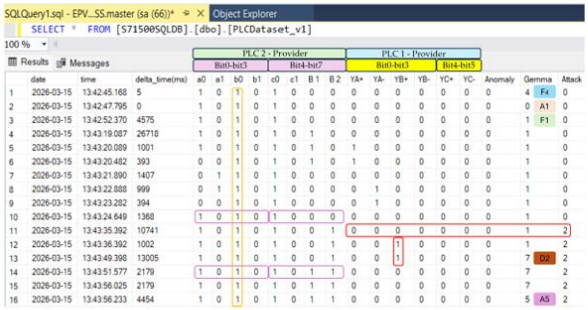


Figure 15. Correlation of SQL Server, GEMMA, and PROFINET data in scenario A2.

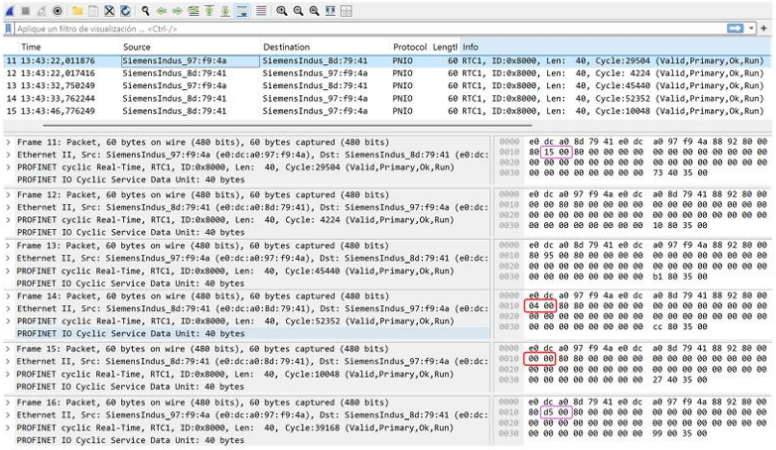


Figure 16. Wireshark capture of PROFINET RTC1 frames during scenario A2.

The fact that PLC 1 does not receive positive feedback within a specified time after issuing the command to the plant (marker down) causes the controller program itself to trigger a transition to state D2, thereby indicating to the operator a fault in actuator B (marker). However, no actual fault exists; rather, the solenoid valve is blocked as a consequence of the attack. This condition is not known to PLC 1.

Figure 16 shows a Wireshark capture of PROFINET IO RTC1 cyclic frames, which is particularly useful for identifying SDU-level manipulation consistent with a MitM attack on digital outputs. It illustrates both the command issued by PLC 1 to initiate the downward movement of the marking actuator and the corresponding state of the plant sensors before and after the attack.

By analyzing the sensor data transmitted by the plant, it can be observed that sensor b0 of actuator B does not change state, despite PLC 1 having issued the command YB+ = true (04–00 hex = 00000100; see Figure 15).

This behavior is a direct consequence of the attack, in which a forged digital output is injected into the plant, resulting in a malfunction of actuator B (marker).

https://doi.org/10.3390/app16073533

Likewise, the testbed supports combined DI/DO manipulation, in which both masks are applied simultaneously to the SDU and the plant. While this variant represents a more aggressive attack configuration, it does not introduce a fundamentally different mechanism from those described in Scenarios A1 and A2, as it corresponds to a combination of input- and output-level manipulations.

4.6. Validation of Communication Integrity (Ping and PROFINET Timing)

This section evaluates whether the execution of the MitM attacks affects the stability of the communication channel. Since the proposed mechanism modifies only the RTC1 SDU content, it is essential to confirm that the underlying communication layers remain unaffected. Four conditions were analyzed:

- Baseline: No MitM module active.
- Attack-Idle: MitM module active but without manipulating SDU content.
- Attack-DI: Manipulation of the digital input (DI) byte (scenario A1).
- Attack-DO: Manipulation of the digital output (DO) byte (scenario A2).

The assessment comprises three complementary measurements: ICMP round-trip time (RTT), PROFINET RTC1 cycle timing derived from Wireshark captures, and throughput evaluation. Together, these measurements address RQ2 by determining whether the attack introduces measurable deviations compared with baseline behavior.

4.6.1. ICMP Round-Trip Time Analysis

ICMP echo requests were used to assess whether enabling the MitM module—or performing SDU-level manipulation—induces measurable changes in L3 connectivity. Two communication paths were evaluated: EW → PLC 2 and DT → PLC 1, using 250 ICMP probes per condition. As summarized in Table 4, both paths exhibited zero packet loss across all four scenarios (baseline, idle, DI, and DO).

Table 4. ICMP RTT statistics for the four experimental conditions.

Condition	Direction	Mean (ms)	SD (ms)	Min (ms)	Max (ms)	Loss (%)
Baseline	DT→PLC1	1.027	0.162	1	2	0.0
Baseline	EW→PLC2	4.100	2.163	1	12	0.0
Attack-Idle	DT→PLC1	1.014	0.116	1	2	0.0
Attack-Idle	EW→PLC2	4.508	2.33	1	13	0.0
Attack-DI	DT→PLC1	1.028	0.165	1	2	0.0
Attack-DI	EW→PLC2	4.624	2.196	1	12	0.0
Attack-DO	DT→PLC1	1.013	0.112	1	2	0.0
Attack-DO	EW→PLC2	4.496	2.374	1	12	0.0

RTT measurements for the DT → PLC 1 direction remained tightly concentrated around their nominal values, with only minor variations attributable to normal switched-Ethernet jitter. These results indicate that neither enabling the MitM module nor performing SDU manipulation introduces additional L3 latency on this path.

In contrast, the EW → PLC 2 direction presented higher dispersion in RTT values, a behavior consistently observed across all conditions and reflecting the inherent variability of this route rather than any attack-induced effects. This wider spread is also consistent with host-side effects on the Engineering Workstation, where TIA Portal and the MSSQL server operate concurrently, introducing CPU load and kernel-level scheduling delays that affect ICMP timestamping. Such effects are well-known to increase RTT variability without indicating any degradation in network-layer performance or any influence from the MitM mechanism. Importantly, the differences between baseline, idle, and DI/DO

https://doi.org/10.3390/app16073533

Appl. Sci. **2026**, *16*, 3533

scenarios remained within normal operational fluctuations, and no systematic latency increase was observed.

Figure 17 consolidates these findings by displaying the RTT mean and standard deviation for both communication paths. The figure highlights the stability of the DT → PLC 1 path and the modest, non-systematic variability of the EW → PLC 2 path. Overall, the ICMP service remains stable under all scenarios, confirming that the MitM module operates exclusively at the SDU level without impacting network-layer connectivity, jitter, or packet delivery.

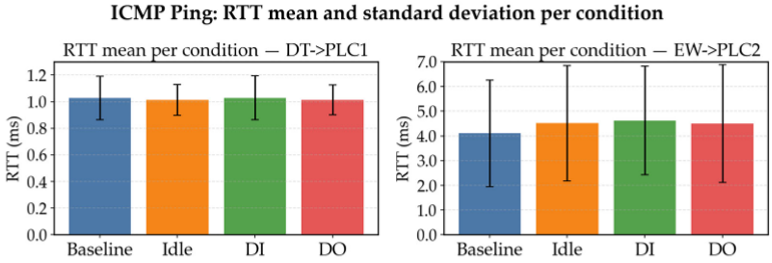


Figure 17. ICMP RTT measurements for the four experimental conditions.

4.6.2. PROFINET I/O Cycle Timing Analysis

RTC1 PROFINET traffic was analyzed to determine whether enabling the MitM module—or performing SDU-level manipulation—introduces timing disturbances in the cyclic I/O exchange. For each experimental condition, all RTC1 frames captured through port mirroring were extracted, and their inter-arrival times were used to estimate the effective cycle time.

As summarized in Table 5, the mean cycle time remained effectively constant across all scenarios, with no systematic deviation from the baseline. Jitter values were likewise comparable between conditions and consistent with the timestamp variability expected when captures are performed on a non-real-time operating system. Minimum inter-arrival times were bounded by timestamp resolution. In contrast, occasional maximum spikes appeared in all scenarios—including the baseline—indicating that such variations originate from measurement-side jitter rather than from the MitM activity.

Table 5. PROFINET RTC1 cycle-time statistics across the four experimental conditions.

Condition	Frames	Mean (ms)	SD (ms)	Min (ms)	Max (ms)
Baseline	351,972	0.5	0.728	0.0	3.147
Attack-Idle	387,239	0.5	0.782	0.0	3.363
Attack-DI	488,238	0.5	0.782	0.0	3.346
Attack-DO	419,406	0.5	0.834	0.0	3.417

Overall, the results show that neither enabling the MitM module nor performing DI-based or DO-based SDU manipulation affects PROFINET RTC1 cycle timing or its real-time scheduling behavior. The cyclic I/O exchange remains stable and deterministic under all tested conditions.

The inter-arrival times obtained from the pooled PLC 1 ↔ PLC 2 capture are also consistent with the packet-rate measurements observed in Wireshark. During the mirrored-port acquisition, the traffic reached approximately 2 kpps, reflecting the combined RTC1 exchanges of both controllers (≈1000 packets/s per PLC). Because the inter-arrival time corresponds to the inverse of the packet rate, a throughput of 2000 packets/s yields an

https://doi.org/10.3390/app16073533

expected mean interval of 1/2000 s (≈0.5 ms), in agreement with the distribution shown in Figure 18. Occasional jitter produces isolated inter-arrival values exceeding 1 ms, explaining the extended whiskers in the boxplot while having no effect on the mean value.

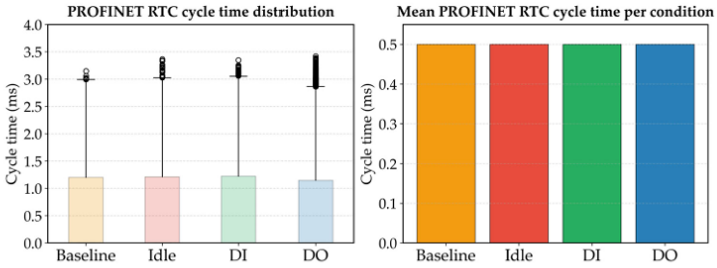


Figure 18. PROFINET RTC1 cycle-time distribution for PLC 1 ↔ PLC 2 traffic.

4.6.3. Throughput Evaluation

Throughput was evaluated by analyzing the RTC1 packet rate (PPS) under the four experimental conditions. As summarized in Table 6, the baseline exhibits a stable packet-exchange pattern with an average throughput of approximately 4 kpps. The attack-idle condition presents the largest deviation relative to the baseline, whereas the DI- and DO-based manipulation scenarios introduce only marginal differences. Across all attack modes, throughput variations remain within ±0.3% for both PPS and Bytes/s, and the increased variability observed in the idle and DO scenarios reflects normal second-level binning effects and the micro-timing introduced by the MITM relay.

Table 6. PROFINET throughput statistics for the four experimental conditions.

Scenario	PPS Mean	PPS SD	kB/s Mean	kB/s SD	ΔPPS (%)
Baseline	4073.01	307.47	248.30	18.74	0.00%
Attack-Idle	4068.86	302.12	248.10	18.42	−0.10%
Attack-DI	4082.76	166.56	248.93	10.15	0.24%
Attack-DO	4072.78	282.94	248.34	17.25	−0.01%

These deviations are modest and compatible with the expected behavior of a lightly loaded PROFINET RT network. Importantly, this throughput metric aggregates all RTC1 frames counted per second on the mirrored port, in contrast to the pooled RTC1 inter-arrival analysis in Section 4.6.2, which reflects an effective rate of ≈2 kpps. The observed differences, therefore, represent measurement-scope effects rather than protocol-level disturbances.

Numerically, the Idle and DO scenarios exhibit PPS values marginally below the baseline (ΔPPS = −0.10% and −0.01%, respectively), whereas the DI scenario is slightly above it (+0.24%). These shifts are two orders of magnitude smaller than the within-scenario dispersion (SD ≈ 283–307 PPS) and fall within the normal jitter of lightly loaded PROFINET RT networks. Scenarios with reduced process activity—such as Idle and DO—are inherently more sensitive to timestamp granularity, mirrored-port timing, and second-level binning effects, which can slightly lower the PPS count in individual one-second measurement windows. Conversely, the higher DI activity increases the probability of capturing additional RTC1 frames within the same interval. As similar fluctuations also appear in the baseline condition, these differences cannot be attributed to the logical MitM manipulation, but rather to measurement-side variability.

Overall, the measurements confirm that neither enabling the MITM module nor manipulating DI or DO values compromises RTC1 timing, protocol validity, or network

https://doi.org/10.3390/app16073533

load. Although throughput analysis is primarily employed here to confirm the execution of the attack and verify the absence of traffic-level side effects, the testbed also supports future research on automated anomaly and attack detection in PROFINET I/O systems using AI-based methods—an area that remains largely unexplored.

5. Results

5.1. Scenario A1—Manipulation of DI Signals

Scenario A1 validates the capability of the MitM module to selectively manipulate the PROFINET I/O exchange between PLC 2 and PLC 1, altering the SDU content delivered to the controller without affecting the timing or validity of the cyclic RTC1 frames. The RTC1 traffic captured through port mirroring confirms that the attacker successfully modifies the SDU bytes corresponding to signals B1 and B2, which appear as active despite the absence of any physical events in the digital twin. These injected values are embedded in otherwise valid RTC1 frames—preserving structural integrity, temporal coherence, and correct IOPS/IOCS indicators—thereby preventing detection at the communication layer. Figure 13 illustrates the modified SDU and its seamless integration within the cyclic exchange.

The dataset recorded in MSSQL (Figure 12) provides a synchronized view of the physical process state, the PROFINET-delivered values, the associated GEMMA transitions, and the timing of the attack activation. Throughout the manipulation interval, the forged bits remain continuously active, causing PLC 1 to infer the presence of a workpiece even though no physical sensor asserts that condition. This confirms the consistency between the modified network traffic and the values logged at the application layer, demonstrating that the attack yields a logically valid yet physically incorrect process state.

Inspection of the control logic in TIA Portal (Figure 11) shows that PLC 1 integrates the received PROFINET digital inputs directly into its execution cycle without applying additional validation mechanisms. As a result, the forged input values propagate through the GEMMA state machine and the control chain without triggering alarms or fault conditions. From the controller's perspective, the manipulated SDU represents a legitimate and temporally coherent state transition.

5.2. Scenario A2—Manipulation of DO Signals

Scenario A2 confirms that the attacker can manipulate the digital-output layer of the control chain without altering the PROFINET SDU exchanged between PLC 1 and PLC 2. The RTC1 cyclic communication remains valid and unmodified, as the attacker neither intercepts nor rewrites the DO bytes transmitted from PLC 1. Instead, the forged YB− command is injected directly into the process image of PLC 2 via the Modbus/TCP interface. As a result, PLC 2 sends a falsified actuator command to the plant while still receiving an intact and legitimate SDU from PLC 1. Because the PROFINET frames are not altered in transit, Wireshark does not capture any SDU-level modification—only the physical consequence of the tampering on the actuator state. Figures 14–16 show the malicious command injection, its impact on the RTC1 SDU, and the resulting actuator behavior.

The dataset recorded in MSSQL captures the resulting inconsistency between the legitimate YB+ command issued by PLC 1 and the unchanged sensor feedback b0, which remains active despite the expected movement. This mismatch forces the controller to enter state D2 in the GEMMA sequence, reporting a marking-actuator fault even though no physical malfunction exists. The SQL log and the PROFINET captures exhibit perfect temporal alignment between the injected DO value (applied internally in PLC 2) and the anomalous process state, confirming that the testbed accurately records both the physical and logical consequences of the attack.

https://doi.org/10.3390/app16073533

These results demonstrate that output-level manipulation can disrupt the control cycle without modifying the PROFINET SDU or triggering protocol-level alarms, thereby complementing the findings reported for Scenario A1.

5.3. Combined Interpretation of A1 and A2

These observations align with the communication-level measurements presented in Section 4.6. The MitM manipulation does not affect L3 latency or packet delivery (Section 4.6.1) and does not disturb the RTC1 cycle timing (Section 4.6.2). The throughput analysis (Section 4.6.3) shows only small variations—within ±0.3%—consistent with normal binning effects and the micro-timing introduced by the relay, none of which interfere with the cyclic I/O exchange. Taken together, the results demonstrate that the attacker can modify the logical perception of the process without altering PROFINET timing behavior or violating protocol-level validity checks.

Overall, the hybrid testbed enables simultaneous observation of the physical domain—represented by the digital twin—and the logical domain through PROFINET captures and MSSQL Server logging. This integrated perspective provides clear evidence of the effectiveness of the attack. It highlights the need for complementary validation mechanisms in industrial systems, as SDU-level manipulation can remain invisible to both the controller and the operator while still inducing incorrect process execution.

The impact of the MitM module becomes most evident at the process-execution level, where each attack mode (idle, DI, DO) interacts differently with the behavior of the marking system. Although the idle mode does not alter the semantic execution of the process, it represents the stage in which the attacker is connected to the communication path while maintaining fully valid PROFINET cyclic traffic. In contrast, the DI and DO manipulation modes introduce changes at the application layer, despite the cyclic communication remaining structurally and temporally valid. Across all scenarios, RTC1 frames preserve their timing and maintain consistent IOPS/IOCS indicators, preventing the controller from detecting manipulation at the communication layer. However, only the DI and DO modes modify the logical behavior of the process—either by altering perceived sensor states (DI manipulation) or by forcing actuator commands (DO manipulation)—while remaining indistinguishable from valid RTC1 cyclic traffic.

6. Discussion

The results of scenario A1 demonstrate that selectively modifying the SDU in RTC1 frames enables an attacker to introduce logically valid yet physically incorrect process states without producing measurable deviations in the evaluated PROFINET cyclic exchange indicators. This behavior is consistent with the PROFINET I/O specification [24], in which the validity of a telegram depends on the structural integrity of IO Data Objects, the correctness of offsets, and valid IOPS/IOCS indicators, rather than on any semantic verification of payload meaning. As long as the frame structure and timing remain correct, the PLC accepts the modified values as legitimate inputs.

The experimental evidence confirms that coherent SDU manipulation is sufficient to influence process behavior without affecting traffic availability or PROFINET timing parameters. Such manipulation enables the attacker to force incorrect GEMMA transitions or generate fictitious sensor activations while maintaining valid cyclic communication, thereby making the manipulation undetectable by standard PROFINET RTC1 integrity mechanisms. This highlights a critical vulnerability: the controller's perception of the process can be altered while all network-level indicators remain nominal.

Scenario A2 complements these findings by demonstrating that output-level manipulation can disrupt the physical behavior of the process without modifying the PROFINET

https://doi.org/10.3390/app16073533

SDU exchanged between PLC 1 and PLC 2. Because the attacker injects the forged actuator command directly into the PLC 2 process image, the RTC1 frames observed on the network remain structurally valid, temporally coherent, and indistinguishable from legitimate cyclic traffic. This makes the attack invisible to communication-layer monitoring while still forcing deterministic deviations in the actuator sequence, leading PLC 1 to misinterpret the plant state and transition to GEMMA fault mode (D2). The contrast between unmodified network traffic and incorrect actuator behavior further illustrates the limitations of relying solely on PROFINET cyclic-integrity indicators for anomaly detection.

The scenario also underscores the importance of reducing the attack surface in industrial networks. In the testbed, the exposure of TCP port 502—associated with a Modbus/TCP service under attacker control—can be identified through basic network-scanning techniques. Although the MitM attack presented here does not rely on this service, its presence illustrates how unnecessary services or misconfigurations may facilitate reconnaissance in real deployments. Measures such as disabling unused ports, enforcing network segmentation, and restricting administrative services are essential to minimize intrusion opportunities.

A key contribution of this work is the integration of three complementary data sources: the physical state from the digital twin, PROFINET traffic captured via port mirroring, and structured logs stored in MSSQL. This combination provides full traceability of the attack, from its manifestation at the communication layer to its propagation through the PLC logic. The ability to correlate physical behavior, network traffic, and control execution offers a comprehensive understanding of the consequences of SDU-level manipulation and strengthens the reliability of the analysis.

This study, however, presents certain limitations. The analysis targets a limited set of digital signals, a single attack configuration, and exclusively RTC1 traffic. More complex architectures, heterogeneous devices, higher-performance modes such as RTC3, and additional manipulation strategies remain unexplored. Likewise, the testbed does not evaluate redundancy mechanisms, PROFIsafe, or integrity layers that may affect the detectability of SDU manipulation.

Despite these limitations, the results show that semantic manipulation of the SDU constitutes a relevant threat in PROFINET I/O systems. The hybrid testbed demonstrates how SDU-level attacks can be introduced, how they propagate through the control logic, and how they remain unnoticed by routine RTC1 structural-integrity checks. This provides a solid foundation for future research on detection methods, semantic-validation techniques, and mitigation strategies against MitM attacks in industrial environments.

Regarding the research questions posed in Section 1.2, the experimental results confirm that (RQ1) the proposed hybrid testbed provides the flexibility required to reproduce realistic PROFINET I/O communication and safely execute controlled MitM attacks; (RQ2) such attacks do not introduce measurable deviations in RTT, throughput, or RTC1 cycle timing and therefore remain stealthy with respect to these communication-layer indicators at the communication layer; and (RQ3) the testbed is capable of generating reproducible baseline and attack datasets suitable for analysis, anomaly-detection research, and training activities in industrial cybersecurity.

Taken together, these insights articulate the broader implications of SDU-level semantic manipulation, clarify the practical limitations of the current implementation, and outline the conditions under which the proposed approach can be generalized to other communication architectures and industrial-control settings.

7. Conclusions and Future Work

This work demonstrates that a Man-in-the-Middle adversary intercepting RTC1 traffic between two PROFINET controllers can selectively modify process semantics without

https://doi.org/10.3390/app16073533

altering cyclic-exchange timing or triggering protocol-level validity checks. Scenario A1 shows that injecting altered digital-input (DI) values induces incorrect GEMMA transitions in the consuming PLC, generating logically valid states that do not correspond to any real physical condition. Scenario A2 further confirms that manipulating actuator-command signals at the digital-output (DO) level can force deterministic deviations in the execution of the process without modifying the PROFINET SDU exchanged between PLC 1 and PLC 2. In both scenarios, the attacks preserve the structural integrity of the RTC1 frames—including IO Data Object organization and IOPS/IOCS semantics—rendering the manipulation indistinguishable from normal cyclic communication.

The combined use of a hybrid digital-twin-based testbed, PROFINET traffic captured via port mirroring, and structured MSSQL-based logging enabled detailed correlation between the physical state of the process, the logical perception of the controller, and the real-time manipulation introduced by the MitM module. This integration provides full traceability of the attack as it propagates from the communication layer to the control logic and offers a comprehensive, non-invasive platform for studying semantic manipulation attacks in PROFINET I/O systems. The methodology is reproducible and compatible with industrial environments where high-fidelity experimentation must not disrupt operational hardware.

As future work, the validated testbed will be used for the systematic generation of datasets aimed at advancing research in industrial cybersecurity for PROFINET-based systems. The infrastructure developed in this study supports the controlled reproduction of diverse SDU manipulation patterns. It enables the creation of labeled datasets that integrate physical (digital twin), logical (PROFINET RTC1/RTC3), and control-level (GEMMA and MSSQL logs) information. These datasets will serve as a foundation for evaluating anomaly-detection mechanisms based on machine-learning and advanced sequential-analysis methods. Beyond demonstrating the feasibility of both DI- and DO-level attacks, the proposed testbed opens new research avenues focused on the detection, analysis, and mitigation of Man-in-the-Middle threats in industrial cyber-physical systems.

These results consolidate the relevance of the proposed testbed as a scientific instrument for analyzing semantic-level attacks in industrial communication systems. They demonstrate its value as a foundation for generating high-quality industrial datasets and for supporting the future development of intelligent detection, validation, and protection mechanisms for PROFINET-based industrial cyber-physical systems.

Author Contributions: Conceptualization, J.V.M.-F.; methodology, J.V.M.-F.; software, J.V.M.-F.; validation, J.E.S.G., N.B., and Á.H.; formal analysis, J.V.M.-F.; investigation, J.V.M.-F.; data curation, J.V.M.-F.; writing—original draft preparation, J.V.M.-F.; writing—review and editing, J.V.M.-F., J.E.S.G., N.B., and Á.H.; visualization, J.V.M.-F.; supervision, J.E.S.G., N.B., and Á.H.; funding acquisition, J.E.S.G. and Á.H. All authors have read and agreed to the published version of the manuscript.

Funding: This publication is part of the AI4SECIoT project ("Artificial Intelligence for Securing IoT Devices"–C032.23), funded by the National Cybersecurity Institute (INCIBE), derived from a collaboration agreement signed between the National Institute of Cybersecurity (INCIBE) and the University of Burgos. This initiative is carried out within the framework of the Recovery, Transformation and Resilience Plan funds, financed by the European Union (Next Generation), the project of the Government of Spain that outlines the roadmap for the modernization of the Spanish economy, the recovery of economic growth and job creation, for solid, inclusive and resilient economic reconstruction after the COVID19 crisis, and to respond to the challenges of the next decade.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

https://doi.org/10.3390/app16073533

Data Availability Statement: The raw data supporting the conclusions of this article will be made available by the authors on request.

Conflicts of Interest: The authors declare no conflicts of interest. The funders had no role in the design of the study; in the collection, analyses, or interpretation of data; in the writing of the manuscript; or in the decision to publish the results.

Abbreviations

The following abbreviations are used in this manuscript:

AR	Application Relation (PROFINET)
ARP	Address Resolution Protocol
CER	Critical Entities Resilience Directive
CPS	Cyber–Physical System
CRA	Cyber Resilience Act
DCP	PROFINET Discovery and Configuration Protocol
DI	Digital Input
DO	Digital Output
DoS	Denial of Service
DT	Digital Twin
EAP-TLS	Extensible Authentication Protocol—Transport Layer Security
EW	Engineering Workstation
GEMMA	Guide for the Study of Machine Operating Modes
HMI	Human–Machine Interface
ICMP	Internet Control Message Protocol
ICPS	Industrial Cyber Physical System
ICS	Industrial Control System
ICT	Information and Communication Technologies
iDevice	PROFINET I/O Device mode (Siemens PLC)
IEC 62443	Industrial Cybersecurity Standards (IEC 62443 series)
I/O	Input/Output
IOCS	Input Output Consumer Status
IOPS	Input Output Provider Status
IP	Internet Protocol
ISO	International Organization for Standardization
IT	Information Technology
kB/s	Kilobytes per second
L3	Layer 3 (Network Layer)
MAC	Media Access Control (Ethernet hardware address)
MitM	Man in the Middle
MSSQL	Microsoft SQL Server
Modbus TCP	Modbus Transmission Control Protocol
NIS2	Network and Information Security Directive 2
OB1	Organization Block 1 (main PLC execution cycle)
OSI	Open Systems Interconnection
OT	Operational Technology
PAC	Programmable Automation Controller
PLC	Programmable Logic Controller
PROFIBUS	Process Field Bus
PROFINET I/O	PROcess Field NETwork Input/Output
PROFIsafe	PROFIBUS/PROFINET Safety Communication Profile
PUT/GET	Siemens S7 read/write communication service
PPS	Packets Per Second
RTC1	Real Time Class 1 (cyclic real-time communication)

https://doi.org/10.3390/app16073533

RTC3	Real Time Class 3 (Isochronous Real Time)
RTT	Round Trip Time
SCADA	Supervisory Control and Data Acquisition
SD	Standard Deviation
SDU	Service Data Unit
S7comm	Siemens S7 communication protocol
SSMS	SQL Server Management Studio
SYN	TCP Synchronize flag (used to initiate a connection)
TCP	Transmission Control Protocol
TIA Portal	Totally Integrated Automation Portal
TSAP	Transport Service Access Point
UDP	User Datagram Protocol
VPN	Virtual Private Network

References

1. Serpanos, D. The Cyber-Physical Systems Revolution. *Computer* **2018**, *51*, 70–73. [CrossRef]
2. Javid, M.; Haleem, A.; Suman, R. Digital Twin Applications toward Industry 4.0: A Review. *Cogn. Robot.* **2023**, *3*, 71–92. [CrossRef]
3. Stranahan, J.; Soni, T.; Heydari, V. Supervisory Control and Data Acquisition Testbed Vulnerabilities and Attacks. In *Proceedings of the IEEE SoutheastCon, Huntsville, AL, USA, 11–14 April 2019*; IEEE: Piscataway, NJ, USA, 2019; pp. 1–6. [CrossRef]
4. Sinha, A.; Patel, S.S.; Kumar, A.; Vyas, O.P. Exploiting Vulnerabilities in the SCADA Modbus Protocol: An ICT Reliant Perspective. In *Advanced Network Technologies and Intelligent Computing*; Woungang, I., Dhurandher, S.K., Pattanaik, K.K., Verma, A., Verma, P., Eds.; Communications in Computer and Information Science; Springer: Cham, Switzerland, 2022; Volume 1534, pp. 94–108. [CrossRef]
5. Plager, A.; Olexa, E.; Gardner, D.; Torres, B.; Hays, A.; Farraj, A. A Study of SCADA System Vulnerabilities and Man-in-the-Middle Threats in Substation Operations. In *Proceedings of the 2025 IEEE Texas Power and Energy Conference (TPEC), National Harbor, MD, USA, 10–11 February 2025*; IEEE: Piscataway, NJ, USA, 2025; pp. 1–5. [CrossRef]
6. Qian, J.; Du, X.; Chen, B.; Qu, B.; Zeng, K.; Liu, J. Cyber–Physical Integrated Intrusion Detection Scheme in SCADA System of Process Manufacturing Industry. *IEEE Access* **2020**, *8*, 147471–147481. [CrossRef]
7. Sun, S.; Haque, K.A.; Huo, X.; Sahu, A.; Goulart, A.; Davis, K. Anomaly Detection in Cyber-Physical Systems Using Long-Short Term Memory Autoencoders: A Case Study with Man-in-the-Middle (MiTM) Attack. In *Proceedings of the 2025 IEEE Texas Power and Energy Conference (TPEC)*; IEEE: Piscataway, NJ, USA, 2025; pp. 1–6. [CrossRef]
8. Houkan, A.; Sahoo, A.K.; Gochhayat, S.P.; Sahoo, P.K. Artificial Intelligence Approach to Intrusion Detection in Industrial Control Systems with Real-World Dataset Generation and Model Evaluation. *Discov. Artif. Intell.* **2025**, *5*, 307. [CrossRef]
9. Walz, A.; Niemann, K.-H.; Göppert, J.; Fischer, K.; Merklin, S.; Ziegler, D.; Sikora, A. PROFINET Security: A Look on Selected Concepts for Secure Communication in the Automation Domain. In *Proceedings of the 2023 IEEE 21st International Conference on Industrial Informatics (INDIN)*; IEEE: Piscataway, NJ, USA, 2023; pp. 1–8. [CrossRef]
10. Niemann, K.-H.; Waldeck, B.; Esslinger, T. PROFINET–Future OT Security Requirements: What Do NIS2, CER, CRA and IEC 62443 Require? *atp Mag.* **2025**, *67*. [CrossRef]
11. Göppert, J.; Sikora, A. Evaluation of the Secure PROFINET Application Relation Establishment Performance. In *Proceedings of the 2024 IEEE 22nd International Conference on Industrial Informatics (INDIN)*; IEEE: Piscataway, NJ, USA, 2024; pp. 1–8. [CrossRef]
12. Mutlu, H.; Akyürek, N.; Korçak, Ö. PROFINET Controller on Cloud. In *Proceedings of the 24th International Symposium INFOTEH-JAHORINA*; IEEE: Piscataway, NJ, USA, 2025; pp. 1–6. [CrossRef]
13. Noorizadeh, M.; Shakerpour, M.; Meskin, N.; Unal, D.; Khorasani, K. A Cyber-Security Methodology for a Cyber-Physical Industrial Control System Testbed. *IEEE Access* **2021**, *9*, 16239–16253. [CrossRef]
14. Srivastava, A.; Zhao, J.; Moore, K.; Anagnostou, E. Development of Cyber Physical Security Simulation Testbed in RTDS Using Modbus Communication. In *Proceedings of the IEEE Power & Energy Society General Meeting (PESGM)*; IEEE: Piscataway, NJ, USA, 2024; pp. 1–6. [CrossRef]
15. Chen, B.; Pattanaik, N.; Goulart, A.; Butler-Purry, K.L.; Kundur, D. Implementing Attacks for Modbus/TCP Protocol in a Real-Time Cyber-Physical System Test Bed. In *Proceedings of the 2015 IEEE Power & Energy Society Innovative Smart Grid Technologies Conference (ISGT), Washington, DC, USA, 18–20 February*; IEEE: Piscataway, NJ, USA, 2015; pp. 1–6. [CrossRef]
16. Erceylan, G.; Akbarzadeh, A.; Gkioulos, V. Leveraging Digital Twins for Advanced Threat Modeling in Cyber Physical Systems Cybersecurity. *Int. J. Inf. Secur.* **2025**, *24*, 151. [CrossRef]
17. Lucchese, M.; Salerno, G.; Pugliese, A. A Digital Twin Based Approach for Detecting Cyber–Physical Attacks in ICS Using Knowledge Discovery. *Appl. Sci.* **2024**, *14*, 8665. [CrossRef]

https://doi.org/10.3390/app16073533

18. Kosugi, A.; Teranishi, K.; Kogiso, K. Experimental Validation of the Attack Detection Capability of Encrypted Control Systems Using Man in the Middle Attacks. *IEEE Access* **2024**, *12*, 10535–10547. [CrossRef]
19. Yılmaz, E.N.; Sayan, H.H.; Üstünsoy, F.; Gönen, S.; Karacayılmaz, G. Cyber Security Analysis of DoS and MitM Attacks Against PLCs Used in Smart Grids. In *Proceedings of the 2019 IEEE International Conference on Smart Grids and Smart Cities (ICSGC), Istanbul, Turkey, 25–26 April 2019*; IEEE: Piscataway, NJ, USA, 2019; pp. 36–40. [CrossRef]
20. Raja, D.J.S.; Sriranjani, R.; Arulmozhi, P.; Hemavathi, N. Unified Random Forest and Hybrid Bat Optimization Based Man in the Middle Attack Detection in Advanced Metering Infrastructure. *IEEE Trans. Instrum. Meas.* **2024**, *73*, 2523812. [CrossRef]
21. Obonna, U.O.; Opara, F.K.; Mbaocha, C.C.; Obichere, J.-K.C.; Akwukwaegbu, I.O.; Amaefule, M.M.; Nwakanma, C.I. Detection of Man in the Middle (MitM) Cyber Attacks in Oil and Gas Process Control Networks Using Machine Learning Algorithms. *Future Internet* **2023**, *15*, 280. [CrossRef]
22. PCSimu: Industrial Automation Simulator (Software Version 3.0). 2023. Available online: https://cade-simu.com/pc-simu (accessed on 7 February 2026).
23. TIA Portal: STEP 7 Basic/Professional, STEP 7 Safety Basic/Advanced, WinCC Basic/Comfort/Advanced and WinCC Unified (Version 18.0). 2022. Available online: https://support.industry.siemens.com/cs/attachments/109807109/TIA_Portal_STEP7_Prof_Safety_WINCC_Prof_V18.iso (accessed on 7 February 2026).
24. Hilscher Gesellschaft für Systemautomation mbH. *PROFINET IO Device API—Protocol API.2. DocumentID: DOC190303API V5.8.0.1EN, Version 5.8.0.1*; Hilscher GmbH: Hattersheim am Main, Germany, 2025.

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.

https://doi.org/10.3390/app16073533

B.3. Paper III (CEA Conference).

Banco de pruebas para evaluación de técnicas IA en la detección de anomalías en sistemas ciberfísicos industriales

XX Simposio CEA de Control Inteligente
25-27 de junio de 2025, Huelva

Banco de pruebas para evaluación de técnicas IA en la detección de anomalías en sistemas ciberfísicos industriales

Martín-Fraile, J.V.[a, *], Sierra-García, J.E.[a], Basurto, N.[b]

[a] *Departamento de Digitalización, Ingeniería de Sistemas y Automática, Universidad de Burgos, C/ Hospital de Rey, sn, 09001 Burgos, España.*
[b] *Departamento de Digitalización, Ciencia de la Computación e Inteligencia Artificial, Universidad de Burgos, C/ Hospital de Rey, sn, 09001 Burgos, España.*

To cite this article: Martín-Fraile, J.V., Sierra-García, J.E., Basurto, N., 2025. Testbed for the Evaluation of AI Techniques in Anomaly Detection for Industrial Cyber-Physical Systems. XX Simposio CEA de Control Inteligente, Huelva (Spain), 2025.

Resumen

Los Sistemas Ciberfísicos (CPS) industriales, impulsados por la Industria 4.0, han mejorado la monitorización y el control en tiempo real de procesos productivos, así como la automatización en fábricas inteligentes, lo que ha incrementado la demanda de datos. Esta situación ha acelerado la convergencia entre tecnologías de la información (IT) y operativas (OT), aumentando la exposición de sistemas de control, especialmente los basados en Controladores Lógicos Programables (PLCs), a ciberataques. En este contexto este artículo presenta un Banco de Pruebas para la Detección de Anomalías y Protección de Sistemas Ciberfísicos Industriales (DAyPSCI), una plataforma híbrida, abierta y flexible, orientada a la investigación y formación en ciberseguridad aplicada a Sistemas de Control Industrial (ICS). DAyPSCI permite generar conjuntos de datos (datasets) con información de proceso y del estado del sistema, siguiendo la metodología GEMMA. Además, posibilita la ejecución de ataques controlados sobre dispositivos industriales reales, facilitando el estudio de vulnerabilidades y la validación de mecanismos de detección y protección.

Palabras clave: Sistema ciberfísico, Control industrial, Controlador lógico programable (PLC), Interfaz hombre-máquina (HMI), Detección de anomalía, Ciberseguridad.

Testbed for the Evaluation of AI Techniques in Anomaly Detection for Industrial Cyber-Physical Systems

Abstract

Industrial Cyber-Physical Systems (CPS), driven by the implementation of Industry 4.0, have enhanced real-time monitoring and control of production processes, as well as automation in smart factories, leading to increased data demand. This situation has accelerated the convergence of Information Technology (IT) and Operational Technology (OT), increasing the exposure of control systems—especially those based on Programmable Logic Controllers (PLCs)—to cyberattacks. This article presents the Testbed for Anomaly Detection and Protection of Industrial Cyber-Physical Systems (DAyPSCI), a hybrid, open, and flexible platform designed for research and training in cybersecurity applied to Industrial Control Systems (ICS). DAyPSCI enables the generation of datasets containing both process data and system state information, following the GEMMA methodology (Guide for the Study of Operating Modes of an Automated System). Additionally, it allows the execution of controlled cyberattacks on real industrial devices, facilitating the study of vulnerabilities and the validation of detection and protection mechanisms.

Keywords: Cyber-physical system, Industrial control, Programmable logic controller (PLC), Human-machine interface (HMI), Anomaly detection, Cybersecurity.

1. Introducción

Los sistemas Ciberfísicos (CPS, por sus siglas en inglés) integran procesos computacionales y físicos para automatizar tareas complejas del mundo real (Serpanos, 2018; Hosseinzadeh *et al.*, 2024). Estos sistemas aprovechan los avances tecnológicos en sensores, redes de comunicación y computación embebida especialmente mediante el uso de Controladores Lógicos Programables (PLC). Combinan

*Autor para correspondencia: jvmartin@ubu.es
Attribution-NonCommercial-ShareAlike 4.0 International (CC BY-NC-SA 4.0)

componentes de software y hardware para interactuar con el entorno físico de forma automática y eficiente.

Entre sus funcionalidades destacan la monitorización y control en tiempo real de procesos físicos, como la gestión energética en redes inteligentes o la optimización y automatización de la producción en fábricas inteligentes.

Sin embargo, esta creciente evolución tecnológica no está exenta de riesgos. En el año 2024, el número de ciberataques a sistemas ciberfísicos creció de forma significativa en comparación con el año anterior (Fortinet, 2024; Faelli, 2025). Este aumento evidencia su creciente vulnerabilidad y exposición, y refuerza la necesidad urgente de implementar medidas de ciberseguridad más robustas para proteger los Sistemas de Control Industrial (ICS) y reducir el impacto de las amenazas cibernéticas (Han *et al.*, 2021; Nankya, Chataut, & Akl, 2023).

Según el informe del Equipo de Respuesta a Emergencias Informáticas del Centro Criptológico Nacional CCN-CERT (2024), así como estudios como los de Pospisil et al. (2021), Domínguez *et al.* (2022) y Benka *et al.* (2025), la creciente integración de tecnologías de la información (IT) y las tecnologías operativas (OT) con sistemas como SCADA (Supervisory Control and Data Acquisition), HMI (Human Machine Interface), PLC (Programmable Logic Controllers) y DCS (Distributed Control Systems) ha incrementado significativamente la exposición de los sistemas de control industrial (ICS, Industrial Control Systems) a amenazas cibernéticas.

La implementación de la industria 4.0 ha acelerado esta convergencia IT/OT permitiendo a las empresas mejorar su eficiencia operativa, tomar decisiones basadas en datos en tiempo real y reducir costes mediante la automatización avanzada y el mantenimiento predictivo. Esta integración proporciona una visión holística de los procesos productivos y del negocio, lo que se traduce en una mayor capacidad de adaptación, competitividad y optimización de recursos.

En este artículo se presenta el Banco de Pruebas para la Detección de Anomalías y Protección de Sistemas Ciberfísicos Industriales (DAyPSCI), una plataforma diseñada con fines de investigación y formación en ciberseguridad aplicada a Sistemas de Control Industrial (ICS). El objetivo de DAyPSCI es doble: por un lado, facilitar el desarrollo y validación de nuevos mecanismos de detección y protección frente a amenazas cibernéticas; y por otro, proporcionar un entorno experimental accesible para la docencia en ciberseguridad industrial.

La estructura del resto del artículo es la siguiente. La sección 2 introduce la arquitectura de los CPS industriales. El banco de pruebas se presenta en la sección 3. La sección 4 ejemplifica un caso de uso del banco. La sección 5 explica cómo se pueden generar un dataset con el banco de pruebas. Se presentan resultados en la sección 6. El artículo finaliza con las conclusiones y trabajos futuros.

2. Arquitectura de los CPS industriales

La arquitectura de un CPS clásico generalmente consta de tres capas principales:

- **Capa física:** Incluye todos los sensores y actuadores que interaccionan con el mundo físico. Los sensores captan información de las variables relevantes del proceso y los actuadores ejecutan acciones físicas basadas en órdenes de la capa computacional.
- **Capa de red**: Está formada por la infraestructura que permite la comunicación entre la capa física y la capa de computacional.
- **Capa de computacional**: Procesa los datos recibidos de los sensores, realiza cálculos y toma decisiones para controlar los actuadores. Esta capa puede incluir sistemas de control en tiempo real, análisis de datos y algoritmos de inteligencia artificial (Oks *et al.*, 2022).

La Figura 1, muestra una imagen de dicha arquitectura. En aplicaciones industriales, como los CPS basados en PLCs, la capa computacional se implementa mediante el propio PLC, que actúa como el componente principal encargado del procesamiento de datos y la ejecución de tareas de control.



Figura 1: Arquitectura de un CPS Industrial

A nivel de implementación, diversos autores han adoptado la Arquitectura de Referencia Empresarial de Purdue (PERA), como es el caso de Alrumaih, Alenazi, AlSowaygh, Humayed y Alablani (2023), así como Hosseinzadeh *et al.* (2024). Esta arquitectura también ha sido incorporada por empresas líderes en soluciones de ciberseguridad, como Fortinet.

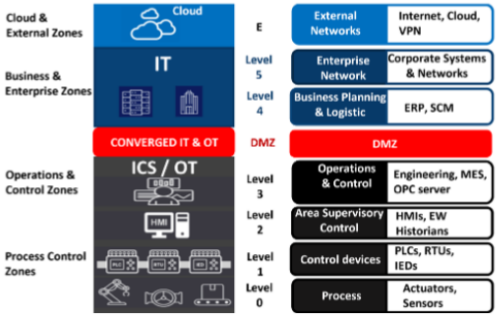


Figura 2: Arquitectura ICS basada en modelo PERA

La arquitectura PERA es un marco para organizar los ICS en zonas o niveles definidos (Alrumaih *et al.*, 2023), creando una estructura jerarquizada separando equipos IT/OT y facilitando la implementación de los requisitos de seguridad establecidos por la norma ISA/IEC 62443-2-1:2024 (International Society of Automation [ISA] & Comisión Electrotécnica Internacional [IEC], 2024). En la Figura 2, se pueden apreciar las zonas, con sus distintos niveles, separadas por la región convergencia entre tecnologías también designada como zona desmilitarizada (DMZ, por sus siglas en inglés), cuya misión es la de permitir a la organización conectarse de manera segura a redes con diferentes requerimientos de seguridad y de este modo proteger la zona empresarial y la de control.

3. Banco de Pruebas para evaluación de técnicas IA en la detección de anomalías en CPS industriales

En general los bancos de prueba de CPS industriales se clasifican, según los componentes que integran en los siguientes tipos (Hosseinzadeh *et al*. 2024):

- **Virtuales**: Son soluciones basadas exclusivamente en software lo que les dota de gran flexibilidad y escalabilidad.
- **Físicos**: Incluyen componentes reales, como sensores, actuadores, dispositivos de red, controladores, etc., es decir son un reflejo del ICS real.
- **Híbridos**: Combinan componentes virtuales y físicos. Son adecuados para probar sistemas integrados industriales sin un elevado coste del equipo de pruebas.

El banco de pruebas presentado en este artículo se enmarca dentro de la categoría híbrida porque combina componentes físicos y virtuales.

3.1. Arquitectura

El banco de pruebas presenta una arquitectura abierta y flexible basada en el modelo PERA, lo que permite recrear condiciones propias de un entorno industrial real. Esto es posible gracias a que la mayoría de sus componentes son equivalentes a los presentes en sistemas automatizados reales, con la excepción del proceso industrial físico (sensores, actuadores, piezas, etc.), el cual se implementa mediante un gemelo digital.

Un gemelo digital es una representación virtual de objetos, procesos o sistemas que existen y operan en tiempo real (Javid, Haleem, & Suman, 2023). Para simular condiciones de comunicación industrial más realistas, este gemelo digital emplea un PLC como interfaz Profinet con el sistema, el cual está conectado a un switch industrial. De esta forma, el PLC de nivel 1 puede enviar órdenes y recibir señales de entrada del proceso a través del PLC de nivel 0, como si estuviera controlando un proceso físico mediante Profinet I/O. Además, el sistema incorpora una interfaz HMI que permite al operador interactuar con el proceso virtual. Esta versión inicial contempla únicamente los niveles correspondientes a la zona de control del proceso, como se muestra en la Figura 3. No obstante, su arquitectura modular y escalable permite la incorporación de nuevos componentes, lo que facilita la expansión hacia otras zonas y niveles del modelo PERA.



Figura 3: Arquitectura de red del Banco

Con el objetivo de utilizar el banco de pruebas para la evaluación de técnicas de inteligencia artificial (AI) en la detección de anomalías y la protección de sistemas ciberfísicos, y así contribuir a mejorar su resiliencia, se ha incorporado una base de datos gestionada mediante Microsoft SQL Server 2022. Esta base de datos permite la generación y almacenamiento de conjuntos de datos (datasets) para su posterior análisis. Aunque por defecto se encuentra alojada en la estación de ingeniería, su ubicación no constituye un requisito obligatorio, ya que puede ser desplegada en otros entornos compatibles.

Asimismo, el banco permite la conexión de un PC de monitorización para la captura de tramas y paquetes de comunicación mediante herramientas como Wireshark u otras similares. También se contempla la posibilidad de conectar un equipo con Kali Linux, lo que facilita la ejecución y evaluación de diversas técnicas de ataque, con fines de prueba y validación de mecanismos de defensa.

3.2. Componentes

La Tabla 1 muestra los componentes que conforman la versión inicial del Banco de pruebas:

Tabla 1: Componentes del Banco de pruebas.

Nº	Dispositivo	Modelo	Descripción
1	PS	6EP1333-4BA00	Fuente de alimentación
2	PLC	6ES7512-1CK01-0AB0	PLC
3	PC		Gemelo digital
4	PLC	215-1BG40-0XB0	Interfaz profinet I/O
5	Switch	208-0BA00-2AC2	Switch industrial
6	HMI	6AV2128-3GB06-0AX1	Panel HMI
7	PC		Estación de ingeniería

4. Contextualización de un caso de uso

En esta sección se presenta un caso de uso del banco de pruebas con el fin de mostrar su aplicabilidad. Para ello, se ha elegido un proceso industrial muy sencillo, como es un sistema de marcaje de piezas. El objetivo de este sistema es el de alimentar con piezas el marcador, realizar la operación de marcaje, y posteriormente expulsar las piezas procesadas por el sistema.

4.1. Gemelo digital de un sistema de marcaje de piezas

Con el fin de disponer de información de sensores, piezas, pre-accionadores y actuadores, es decir datos de proceso comparables a los de un proceso real, junto con su correspondiente visualización animada se ha empleado la aplicación *PCSimu* (PCSimu, 2023). Esta herramienta permite el diseño y visualización 2D y ha sido utilizada para generar el gemelo digital del sistema de marcaje.

Como se muestra en la Figura 4, el sistema está compuesto por tres actuadores (A, B y C), cada uno con sus respectivos sensores magnéticos de posición (a0, a1, b0, b1, c0 y c1), un generador de piezas, y dos sensores de presencia (B1 y B2), ubicados en las zonas de alimentación y marcaje, respectivamente. Los actuadores están configurados como cilindros de doble efecto, controlados mediante electroválvulas 5/2 de tipo biestable.



Figura 4: Gemelo digital.

La aplicación *PCSimu* también permite ajustar las velocidades de avance y retroceso de los actuadores, lo que contribuye a una simulación más realista del comportamiento dinámico del proceso. Además, como se muestra en la Figura 3, el gemelo digital se conecta al PLC 2, que actúa como interfaz Profinet I/O (PNIO), mediante su conexión al puerto 1 del switch industrial. Esta configuración permite el intercambio de información entre el gemelo digital y el PLC 1 a través de comunicación Profinet I/O, reproduciendo así una arquitectura de periferia descentralizada típica de entornos industriales modernos.

4.2 Interfaz HMI

Para interactuar con el sistema ICS, se ha implementado una interfaz como la mostrada en la Figura 5. A través de esta interfaz, el operador puede ajustar parámetros como el **tiempo de actuación** de la matriz sobre la pieza, el **número de golpes** del marcador por pieza y la **cantidad de piezas** a marcar en cada lote. Además, el selector **MAN/AUT** permite elegir entre el modo manual y el modo automático de funcionamiento de la instalación. En modo manual, los pulsadores **AV_x** y **RE_x** (donde *x = A, B o C*), con contacto **NO** (Normally Open), permiten al operador emitir órdenes de avance y retroceso a los actuadores, si fuera necesario.

La interfaz también incorpora otros pulsadores como **marcha**, **paro**, **rearme**, **reset** y **ack**, cuya función es emitir las siguientes órdenes: inicio del modo automático, parada al final del ciclo, rearme de la instalación, reseteo del lote de piezas y confirmación de eventos solicitados, respectivamente.

Asimismo, se incluye un **pulsador de parada de emergencia** con contacto **NC** (Normally Closed, por sus siglas en inglés). En cuanto a la señalización, el sistema cuenta con una **baliza luminosa** compuesta por cuatro pilotos de colores: **rojo**, **verde**, **azul** y **amarillo**, que proporcionan información visual sobre el estado de las salidas digitales.

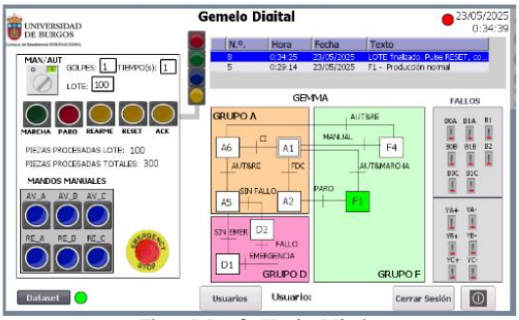


Figura 5: Interfaz Hombre Máquina.

Como ya se ha comentado anteriormente el equipo está diseñado también con fines educativos y por ello incorpora la metodología **GEMMA** (Guía de Estudio de los Modos de Marchas y Paradas, de un sistema automatizado). El **grupo A** incorpora procedimientos de parada. **A1**: parada en el estado inicial, **A2**: parada pedida a fin de ciclo, **A5**: preparación para la puesta en marcha tras fallo y **A6**: puesta del sistema en el estado inicial. El **grupo F**, proceso en funcionamiento incorpora **F1**: producción normal y **F4**: marchas de verificación sin orden o modo manual. Finalmente, el **grupo D** proceso en defecto contempla **D1**: para de emergencia y **D2**: diagnóstico y/o tratamiento de los defectos.

La gestión de la parada de emergencia y el rearme se debe considerar la parte principal de un programa de autómata (Piedrafita, 2004). Por ello, para facilitar al operador la operación de rearme está dotado entre otros con dicho pulsador y de un sistema de avisos que lo guía e informa del estado de la instalación con objeto de facilitar al operador la operación del rearme de la instalación.

Otro de los objetivos es la detección de anomalías en el sistema y en ese sentido está dotado con los interruptores **B0x**, **B1x**, **Yx+** e **Yx-** (*x = A, B y C*) que permiten simular fallos en los sensores (B) y en los preaccionadores, es decir en las electroválvulas (Y) que gobiernan a los actuadores. Esto permite contemplar fallos o anomalías como ocurre en los sistemas reales y añadir esta información en la generación del dataset.

4.3. Proyecto TIA Portal

Para llevar a cabo la configuración, programación y puesta en marcha del banco de pruebas, se generó un proyecto utilizando la herramienta TIA Portal (Totally Integrated Automation Portal) en su versión 18, desarrollada por el fabricante Siemens (TIA Portal, 2022). La Figura 6 muestra la configuración de red realizada dentro de dicho proyecto.



Figura 6: Configuración de red en TIA Porta V18.

Una vez realizada la configuración del hardware del proyecto el siguiente paso fue abordar la programación del PLC 1, cuya función principal es controlar el proceso de marcado. Esta programación se llevó a cabo utilizando programación estructurada, empleando los siguientes lenguajes definidos por la norma IEC-61131-3: diagrama de contactos KOP, diagrama de funciones secuenciales S7-Graph y lenguaje de control estructurado SCL.

5. Creación de un conjunto de datos

Como se mencionó anteriormente el banco de pruebas está equipado con una base de datos que registra todos los eventos programados que ocurren en el sistema, lo que le permite la generación de conjuntos de datos (datasets) para su posterior análisis. Las variables recogidas se describen en la Tabla 2 indicando el tipo de datos y una breve descripción con el significado de la información que registra.

Tabla 2: Campos del dataset.

Variable	Tipo	Descripción
Timestamp	Time	Hora del evento.
Duración	Int	Unidades en milisegundos.
Sensores	Boolean	B0_x0, B1_x1, x=a, b y c, B1 y B2
Actuadores	Boolean	YA+, YA-, YB+, YB-, YC+, YC-.
Anomalías	Int	0-Sin anomalía, 1-Sensores, 2-Actuadores, 3-Múltiple.
GEMMA	Int	Modos: 0-A1, 1-F1, 2-A2, 3-D1, 4-F4, 5-A5, 6-A6, 7-D2.
Ataques	Int	0-Sin ataque, 1- DoS, 2-MITM.

5.1 Captura de datos de proceso

Para validar la información contenida en el conjunto de datos, relativa a las comunicaciones con la periferia descentralizada, se incluye en la Figura 7 una captura de Wireshark que muestra los datos de proceso correspondientes a las entradas, transmitidos mediante el protocolo Profinet IO. En dicha captura, el valor hexadecimal 0x55 indica que los actuadores se encuentran en su posición inicial y que el sistema dispone de una pieza en el alimentador (véase Tabla 2).

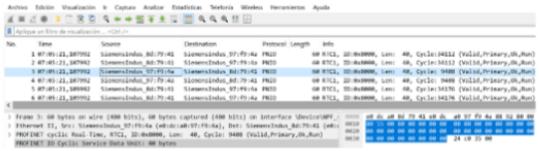


Figura 7: Datos de proceso de entradas.

5.2 Etiquetado del dataset

Con el objetivo de disponer de datos adecuados para su evaluación mediante técnicas de AI orientadas a la detección de anomalías y la protección de sistemas ciberfísicos, se han generado dos conjuntos de datos conforme a la información recogida en la Tabla 2: uno correspondiente a un proceso sin anomalías y otro con anomalías introducidas de forma controlada.

- **Sin anomalías**: Este conjunto recoge la información del banco de pruebas configurado con un lote de 100 piezas, aplicando un golpe por pieza y un tiempo de apriete de un segundo por unidad (véase Figura 5).
- **Con anomalías**: En este caso, la configuración del banco es idéntica a la anterior. No obstante, se han introducido de forma controlada 12 fallos en las electroválvulas que gobiernan los actuadores.

6. Resultados

A partir del conjunto de datos generado sin anomalías, se realizó un análisis exploratorio utilizando la Transformada Rápida de Fourier (FFT) con el objetivo de clasificar los datos en intervalos correspondientes a eventos de origen real o ciclos de proceso. Esta técnica permite identificar patrones periódicos subyacentes en los datos, derivados de la naturaleza cíclica de los procesos de producción automatizados (Sai, *et al.*, 2023).

Aplicando esta metodología, se realizó la reconstrucción temporal de las señales de ambos conjuntos de datos y a continuación se muestrearon a intervalos uniformes de 10 ms. Las señales resultantes de los datos recogidos con anomalías se presentan en la Figura 8 y en ella claramente se observan las 12 anomalías.

A continuación, se aplicó la Transformada Rápida de Fourier (FFT) a las señales binarias de sensores y actuadores, del conjunto de datos sin anomalía identificándose la frecuencia dominante, en el espectro de frecuencias que se representa en la siguiente Figura 9.

Para el conjunto de datos generado con anomalías, también se realizó la FFT de las mismas señales y su espectro se muestra en la Figura 10.

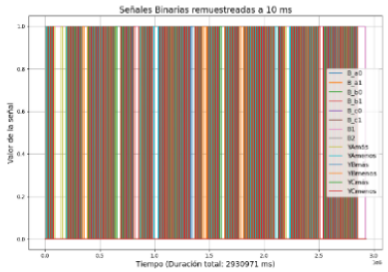


Figura 8: Señales de datos de proceso con anomalías.

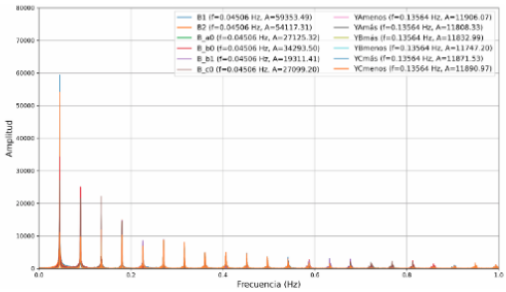


Figura 9: FFT de datos de proceso sin anomalías

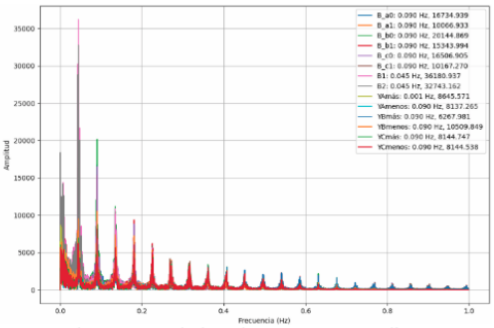


Figura 10: FFT de datos de proceso con anomalías

5.1 Análisis de resultados

El tiempo total empleado en el marcado del lote programado del conjunto de dados sin anomalía fue de **2.223.568 milisegundos**, es decir, aproximadamente **37,06 minutos**.

De forma complementaria, en la Figura 9 se identifica un pico dominante en la frecuencia de **0,04506 Hz**. A partir de esta frecuencia, es posible calcular el tiempo de ciclo de producción utilizando la siguiente expresión:

$$t_{Ciclo} = \frac{1}{f_{pico}} \quad (1)$$

Aplicando la ecuación (1), se obtiene un tiempo de ciclo aproximado de **22,2 segundos**. Al dividir el tiempo total de marcado del lote por este valor, se obtiene el número de ciclos, lo que permite validar el tamaño del lote ajustado: 100 piezas.

En cuanto al conjunto de datos con anomalías, se observa prácticamente la misma frecuencia dominante **0.045 Hz**, indicando igual tiempo de ciclo de producción dominante, si bien aparecen nuevas frecuencias debido a los fallos registrados y también se puede ver en la Figura 8 como el tiempo total de marcado del lote fue de unos **48,85 minutos**, lo que hace que el tiempo de ciclo medio de las piezas con fallo sea superior al del lote sin anomalías, aproximadamente **29,3 segundos** en como era de esperar debido a los fallos inducidos en las electroválvulas.

7. Conclusiones

Este artículo presenta un banco de pruebas para la Detección de Anomalías y Protección de Sistemas Ciberfísicos Industriales (DAyPSCI), el cual integra componentes industriales de última generación y adopta una arquitectura híbrida, abierta y flexible.

Se ha mostrado un caso de uso basado en el gobierno de un gemelo digital de un proceso de marcado, utilizando la metodología GEMMA, y se ha validado su utilidad para la generación de conjuntos de datos destinados a análisis posteriores. En dichos conjuntos se ha incorporado información sobre el estado del proceso, conforme a esta metodología de diseño.

Asimismo, se aplicó un análisis exploratorio de datos y la metodología propuesta por Sai, Gram y Bauernhansl (2023) a un conjunto de datos sin anomalías, logrando identificar el tiempo de ciclo de producción y el número de piezas procesadas con resultados satisfactorios.

También se realizaron capturas de datos de proceso mediante comunicación Profinet I/O, obteniendo resultados positivos. Esta información permitirá, en trabajos futuros, disponer de datos de red industrial para aplicar técnicas de inteligencia artificial orientadas a la detección y protección frente a ciberataques en sistemas ciberfísicos industriales.

Agradecimientos

Esta actividad se lleva a cabo en ejecución del Proyecto Estratégico "Inteligencia Artificial para la Securización de Dispositivos IoT" (C032.23), fruto de un convenio de colaboración suscrito entre el Instituto Nacional de Ciberseguridad (INCIBE) y la Universidad de Burgos. Esta iniciativa se realiza en el marco de los fondos del Plan de Recuperación, Transformación y Resiliencia, financiadas por la Unión Europea (Next Generation), el proyecto del Gobierno de España que traza la hoja de ruta para la modernización de la economía española, la recuperación del crecimiento económico y la creación de empleo, para la reconstrucción económica sólida, inclusiva y resiliente tras la crisis de la COVID19, y para responder a los retos de la próxima década.

Referencias

Alrumaih, T. N. I., Alenazi, M. J. F., AlSowaygh, N. A., Humayed, A., & Alablani, I., 2023. Cyber resilience in industrial networks: A state of the art, challenges, and future directions. Journal of King Saud University - Computer and Information Sciences, 35(9), 101781. https://doi.org/10.1016/j.jksuci.2023.101781.

Benka, D., Horvath, D., Spendla, L., Gaspar, G., & Stremy, M., 2025. Machine learning-based detection of anomalies, intrusions, and threats in industrial control systems. IEEE Access, 13, 12502–12514. https://doi.org/10.1109/ACCESS.2025.3530902.

CCN-CERT. 2024. CCN-CERT IA-04/24: Ciberamenazas y tendencias. Edición 2024. Centro Criptológico Nacional. https://www.ccn-cert.cni.es/es/informes/informes-ccn-cert-publicos/7274-ccn-cert-ia-04-24-ciberamenazas-y-tendencias-edicion-2024/file.html.

Domínguez, M., Fuertes, J. J., Prada, M. A., Alonso, S., Morán, A., & Pérez, D., 2022. Design of platforms for experimentation in industrial cybersecurity. Applied Sciences, 12(13), 6520. https://doi.org/10.3390/app12136520.

Faelli, M., 2025. Informe global de ciberataques del primer trimestre de 2025 de Check Point Software. Check Point Research. https://www.itsitio.com/seguridad/informe-global-de-ciberataques-del-primer-trimestre-de-2025-de-check-point-software/.

Fortinet, 2024. 2024 State of Operational Technology and Cybersecurity Report. https://www.fortinet.com/content/dam/fortinet/assets/reports/report-state-ot-cybersecurity.pdf.

Han, S., Lee, K., Cho, S., & Park, M., 2021. Anomaly detection based on temporal behavior monitoring in programmable logic controllers. Electronics, 10(10), 1218. https://doi.org/10.3390/electronics10101218.

Hosseinzadeh, S., Voutos, D., Barrie, D., Owoh, N., Ashawa, M., & Shahrabi, A., 2024. Design and development considerations of a cyber physical testbed for operational technology research and education. Sensors, 24(12), 3923. https://doi.org/10.3390/s24123923.

International Society of Automation & Comisión Electrotécnica Internacional, 2024. ISA/IEC 62443-2-1:2024 - Seguridad para sistemas de automatización y control industrial – Parte 2-1: Requisitos del programa de seguridad para propietarios de activos de IACS. https://www.isa.org/products/ansi-isa-62443-2-1-2024-security-industrial-automa.

Javid, M., Haleem, A., & Suman, R., 2023. Digital Twin applications toward Industry 4.0: A review. Cognitive Robotics, v. 3, p. 71-92.. https://doi.org/10.1016/j.cogr.2023.04.003.

Nankya, M., Chataut, R., & Akl, R., 2023. Securing industrial control systems: Components, cyber threats, and machine learning-driven defense strategies. Sensors, 23(21), 8840. https://doi.org/10.3390/s23218840.

Oks, S. J., Jalowski, M., Lechner, M., Mirschberger, S., Merklein, M., Vogel-Heuser, B., & Möslein, K. M., 2022. Cyber-Physical Systems in the Context of Industry 4.0: A Review, Categorization and Outlook. Information Systems Frontiers, 26, 1731–1772. https://doi.org/10.1007/s10796-022-10252-x.

PCSimu, 2023. PCSimu: Simulador de automatización industrial (Versión 3.0) [Software]. https://cade-simu.com/pc-simu.

Piedrafita, R., 2004. Ingeniería de la automatización industrial. Ra-Ma.

Pospisil, O., Blazek, P., Kuchar, K., Fujdiak, R., & Misurec, J., 2021. Application perspective on cybersecurity testbed for industrial control systems. Sensors, 21(23), 8119. https://doi.org/10.3390/s21238119.

Sai, B. K., Gram, J., & Bauernhansl, T., 2023. Information-based preprocessing of PLC data for automatic behavior modeling. Procedia CIRP, v. 120, p. 565-571.https://doi.org/10.1016/j.procir.2023.09.038.

Serpanos, D., 2018. The Cyber-Physical Systems Revolution. IEEE Computer, 51(3), 70–73. https://doi.org/10.1109/MC.2018.1731058.

TIA Portal, 2022. TIA_Portal: STEP 7 Basic/Professional, STEP 7 Safety Basic/Advanced and WinCC Basic/Comfort/Advanced and WinCC Unified (Versión 18.0) [Software]. https://support.industry.siemens.com/cs/attachments/109807109/TIA_Portal_STEP7_Prof_Safety_WINCC_Prof_V18.iso.