Implementation and Detection of Novel Attacks to the PLC Memory on a Clean Water Supply System

Andres Robles-Durazno[1] and Naghmeh Moradpoor[2] and James McWhinnie[3] and Gordon Russell[4] and Inaki Maneru-Marin[5]

[1] Edinburgh Napier University, Scotland, UK
a.roblesdurazno@napier.ac.uk
[2] Edinburgh Napier University, Scotland, UK
n.moradpoor@napier.ac.uk
[3] Edinburgh Napier University, Scotland, UK
j.mcwhinnie@napier.ac.uk
[4] Edinburgh Napier University, Scotland, UK
g.russell@napier.ac.uk
[5] Edinburgh Napier University, Scotland, UK
40291354@live.napier.ac.uk

Abstract. Critical infrastructures such as nuclear plants or water supply systems are mainly managed through electronic control systems. Such systems comprise of a number of elements, such as programmable logic controllers (PLC), networking devices, and actuators. With the development of online and networking solutions, such electronic control systems can even be managed online. Even though network connected control systems permit users to keep up to date with system operation, it also opens the door to attackers taking advantages of such availability. In this paper, a novel attack vector for modifying PLC memory is proposed, which affects the perceived values of sensors, such as a water flow meter, or the configuration of actuators, such as a pump. In addition, this attack vector can also manipulate control variables located in the PLC working memory, reprogramming decision making rules. To show the impact of the attacks in a real scenario, a model of a clean water supply system is implemented in the Festo rig. The results show that the attacks on the PLC memory can have a significant detrimental effect on control system operations. Further, a mechanism of detecting such attacks on the PLC memory is proposed based on monitoring energy consumption and electrical signals using current-measurement sensors. The results show the successful implementation of the novel PLC attacks as well as the feasibility of detecting such attacks.

Keywords: Industrial Control Systems, Cyber Attacks, Clean Water Supply System, PLC Memory Attack, Clean Water Supply System.

1 Introduction

The evolution of Industrial Control Systems has improved the application of computer-based management systems in industrial settings. For instance, in water industries, the technology has improved the reliability and quality of water services, but as a result it has increased the likelihood of targeted cyber events that could lead to disruption in the water supply. Currently, Industrial Control Systems (ICS) are facing new threat vectors design to extract sensitive information or disrupt operations. One of the biggest recent attacks occurred in May 2017, using the WannaCry ransomware [1] that affected a considerable number of computers running Windows operating system across the globe. This attack affected not only desktop computers, but industrial and social infrastructure facilities as well. For instance, Renault, Nissan and Honda were forced to suspend operations because their facilities were infected. In another example, the ExPetr (Petya) attack [1], which was discovered in Jun 2017, affected power sector companies, transport industry and even the Chernobyl radiation monitoring station. Such attacks are becoming increasingly more sophisticated, and the risk of disrupting industry operations if growing.

Although cyber-attacks to critical infrastructures, such as water plants, have increased globally [2], growing awareness of the risk does not necessarily result in companies implementing better security protocols and safer systems. When cyber-attacks are identified publically, the focus is frequently on security breaches in industries such as banking and retail [3]. Although, according to a number of reports, cyber-attacks on vital infrastructure such as electrical grids and water distribution systems have increased considerably. For instance, on November 21th, 2001, hackers gained access to a clean water utility in Springfield, USA and destroyed a pump [4]. The attackers stole the access credentials by first breaking into a computer belonging to the utility's SCADA software vendor. The control system under attack kept turning on and off, resulting in the burnout of a water pump. According to the forensic report, the hackers may have had access to the water plant two months prior the attack. This attack resulted in about 56000 people without water.

This paper proposes a novel attack vector to the PLC memory and a mechanism of detecting this type of attack, by monitoring the energy consumption and other electrical signals. To validate this approach, a model of an un-interrupted clean water supply system was constructed in the Festo rig [5]. This paper is organized as follows. Section 2 describes the related work in the field. Section 3 gives a brief overview of the PLC operation. Sections 4 refers to the testbed used to conduct this research. Section 5 describes the attacks performed to the PLC memory. Section 6 proposes a new method of attack detection by analyzing energy traces. Section 7 indicates the results obtained. Section 8 presents the conclusions.

2 Related Work

In this section, existing work related to anomaly detection techniques for SCADA systems are discussed. Detecting attacks is challenging, as in particular attacks change over time and such attacks may be using previously unknown attack methods. For this reason the authors have focused on applying machine learning approaches to the task, with the goal of making such detection easier and more effective.

In [6], a behavior-based attack detection and classification scheme for a Secured Water Treatment (SWaT) system is proposed using machine learning algorithms. SWaT is an operational scaled down water treatment plant with six main processes, though they studied intrusions against only one of the processes. Here, Best-First Tree (BFTree) shows the best results in terms of precision and accuracy of detection and classification in comparison with their other eight selected machine learning algorithms. They used 18 attacks based on exploiting 10 different issues in three different subsystems to build the model to evaluate their selected nine machine learning algorithms. The three places that their attacks occurred were inflow into the process (4 attack types in total), outflow from the process (2 attack types in total), and the water level of the tank (4 attack types in total). Their attacks were based on the model proposed by [7] with the aim to mislead the PLC by providing false sensor or actuator information. For instance, for the attack on inflow, one of the attacks changed the operating value of the flow indicator sensor to above the normal operating range, which falsified input flow rate and gave a wrong impression to the PLC that the relevant sensor was faulty. In one of the outflow attacks, the value of the main pump's status was set to "closed". This made the PLC turn on the backup pump while the main pump was still running, which can damage pumps or burst pipes. Likewise, for an attack on the water level aspect of the tank, the value of the water level sensor was changed to below the normal minimum level, which in turn made the outflow and filtration process stop.

In [8], a Support Vector Machine (SVM) was employed, and a proposed an Intrusion Detection System (IDS) with a discriminant model to detect cyber-attacks on Industrial Control Systems (ICS) was presented. Their proposed model was based on a communication profile analysis which considers only packet intervals and packet length. Their testbed contained two water tanks prepared with control devices and controlled automatically. For their experiments, they created two datasets that included penetration test dataset and normal dataset each consisting of 10 sub datasets. While the former was constructed during the period of a four-stage penetration test, over which malicious and benign packets are labelled based on their source IP address, the latter was built throughout the normal operation of the system. For each dataset, packet interval and packet length was captured. For their penetration tests, they used the Metasploit Framework (Rapid7) and then Wireshark to capture packets. In their results they identified a significant difference between attack packets and normal packets in terms of packet intervals and packet lengths.

In [9], big data analysis and behaviour observation techniques were employed for cyber-attack detection within a simulated critical infrastructure: a pressurized water cooled nuclear reactor. The simulated system includes a water source, two water tanks, a condenser, a reactor, a generator, acid tank, and emergency coolant. In their simulation, each component has a corresponding observer to extract physical information about behavior and construct two datasets: one with a smaller and one with a larger number of features and events. They constructed features by taking the maximum, minimum, mean and median of water tank level, steam output, and energy creation for 32 mechanical components and 9 system components, which is sampled at 4Hz (4 times every second) for a 24-hour simulation. After specifying the constraints for the water tank, steam output, and energy creation, they observed the minimum and maximum levels regularly. If the levels recorded are lower or higher than the expected minimum or maximum values then the system behavior identified it as abnormal, otherwise it is identified as normal behavior. They then used five supervised data classifiers to detect attacks. Addressing their captured results, in the initial evaluation, the classifiers were able to produce a reasonably good accuracy, which in turn was significantly increased in their second evaluation by increased the number of the events as well as the number of the features captured per event.

In [10], unsupervised machine learning algorithms for anomaly detection in water treatment systems was examined using a dataset collected through the same SWaT testbed as [6], which was a scaled-down raw water purification plant. They applied and compared two unsupervised argorithms: Deep Neural Network (DNN) including a layer of Long Short-Term Memory (LSTM) architecture followed by feedforward layers of multiple inputs and outputs, and a one-class Support Vector Machine (SVM) in which DNN performs slightly better than one-class SVM in general. They used logs from the testbed that were available online [2] which contained benign and malicious events collected from network traffic, in addition to the data collected from all 51 sensors and actuators available in SWaT over eleven days of continuous operations. This included seven days of continuous normal operation in addition to four days of 36 attack scenarios representative of typical network-based attacks on Cyber-Physical Systems (CPSs). Generally, their attacks were based on hijacking and modifying network packets through the data communication link of the SWaT network, allowing for sensor data and actuator signals to be manipulated before reaching the PLCs, pumps, and valves.

In [11], an unsupervised Recurrent Neural Networks (RNN) for anomaly detection was proposed using a dataset collected through the SWaT. This considered only a single process (Process 1) out of six available processes of the testbed. Addressing their captured results, they were able to detect the majority of attacks with low false positive rate using their proposed RNN approach. The SWaT dataset is available online [2] and includes benign and malicious logs captured from the SWaT testbed [6]. The dataset collected over seven days of continuous normal operation and four days of malicious operation during which 36 attack scenarios were conducted. The attacks have been simulated by network packet hijacking and manipulation which results in sensor data and actuator signals to be manipulated before reaching the PLCs, pumps, and valves. For

the data pre-processing stage, they treated all P1's sensors and actuators as a numeric attribute and normalized each feature by subtracting the mean and scaling to the unit variance given.

In [12], an unsupervised SCADA data-driven approach to detect integrity attacks on a simulated Water Distribution System (WDS) was proposed. This was based on k-nearest neighbour technique and includes two stages of: 1) automatic identification of consistent and inconsistent state of SCADA systems and 2) automatic extraction of proximity-based rules from the identified states to detect inconsistent states. They compared their proposed unsupervised learning approach with three other approaches, two of which are based on unsupervised learning and one is based on semi-supervised learning. Their proposed unsupervised approach shows better performance in terms of detection accuracy and efficiency. For their attack scenarios, they conducted man-in-the-middle attacks to manipulate process parameters such as: water flow, water pressure, water demand, water level, valve status, valve setting, pump status, and pump speed in the WDS server. To simulate a realistic scenario, the WDS server reads and controls these process parameters in response to message commands from a filed device. In total they used three datasets: 1) a publicly available dataset referred to as DUWWTP [13] which comes from the daily measuring of sensors in an urban waste water treatment plant and 2) their two simulated datasets. The DUWWTP includes 527 observations of 38 data nodes in which 14 of them are labelled as inconsistent. The simulated datasets include 10,500 observations from 23 data nodes for which 100 events are labelled as inconsistent.

Attacks documented in previous publications [6][10][11] employed attacks based on packet hijacking, manipulation and injection, as well as old techniques like man-in-the-middle attacks such as in [12], all of which could be detected by Network Intrusion Detection/Prevention Systems (IDS/IPS). In this paper the authors propose a different, novel attack targeting PLC memory. Additionally, a unique approach to detect attacks is proposed for PLC memory based on monitoring the energy consumption of the system.

3 PLC Operation

The operation of a PLC is straightforward. It makes decisions based on the program coded within it by a user. PLCs operate by running a scan cycle and repeat this many times per second [14]. Fig 1 represents the PLC scan cycle. It runs checks on the hardware and software for faults when it starts. Afterwards, the PLC runs a three step process: input scan, program execution and output scan [15].

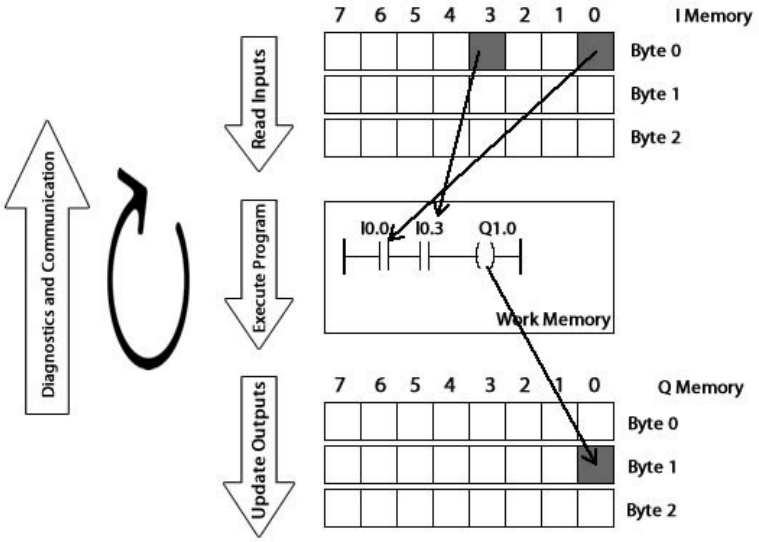

Fig. 1. PLC Scan Cycle

3.1 Input Scan

In this scan, the PLC takes a snapshot of the inputs and determines the state of the devices connected. Then it saves this information in a data table to use in the next step. This speeds up subsequent processing and maintains consistency in cases where an input changes in the period from start to the end of the program [15].

3.2 Execute Program

After getting the information from the inputs, the PLC executes a program, one instruction at a time, using only the memory copy of the inputs. In addition, during the program execution, the PLC may require information allocated in the working memory, such as Process Variables (PV).

3.3 Output Scan

The outputs will be updated when the execution of the program ends, using the temporary values in memory. The PLC updates the status of the outputs by writing to the memory locations associated with each output [15].

4 Testbed

A model of an uninterrupted clean water supply system using the Festo rig [5] was implemented to support the investigation. To simulate clean water demand of a small

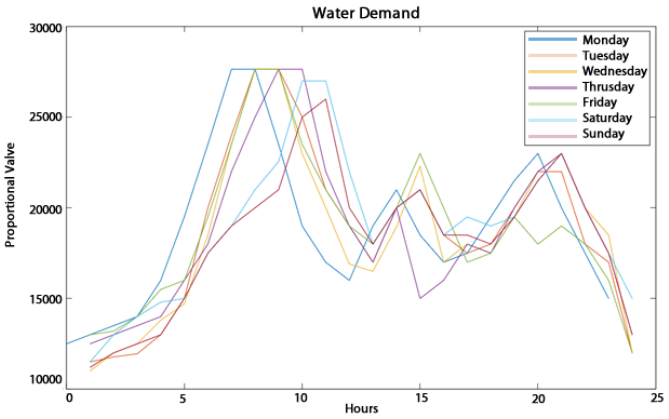

Fig. 2. Water demand models.

town, a water demand model was constructed which is represented in Fig 2. The X-axis represents 24 hours of a day and the Y-axis represents the value applied to the space of memory addressed to the proportional valve. It can be argued that the simulation only represents one week and in some cases the water demand might variate depending on various factors over longer periods. For instance, water demand during the summer might be higher than during the winter, or even during the holidays. However, for experimental purposes the water demand model ignores such variances. Fig 3 represents the network implementation where the attacker is connected to the control network along with the PLC and Supervisor Console. The sensors and actuators, which are connected to the PLC inputs and outputs, are shown in Fig 4.

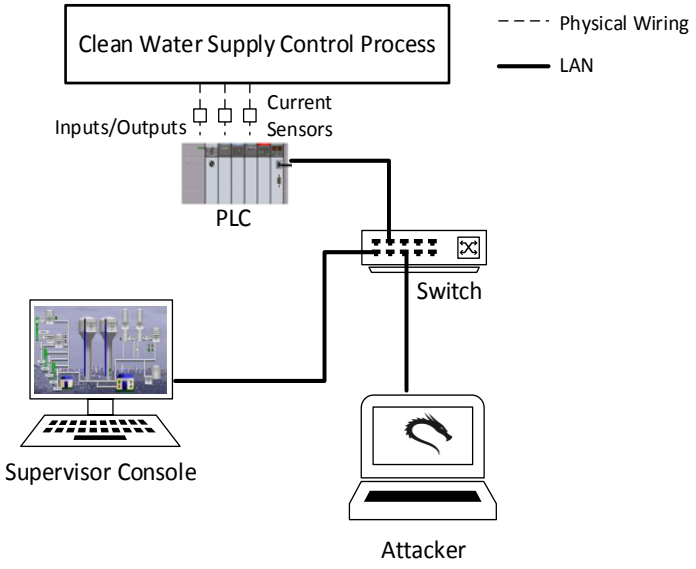

Fig. 3. Testbed implementation.

4.1 Normal Operation

An uninterrupted clean water supply system was modelled. In this model, the water reservoir is assumed to be a continuous naturally filtered water supply that that is pumped up to the supply tank. Fig 4 represents the control process diagram implementation. In normal operation tank B101 simulates the clean water supply which is pumped up to the reservoir tank B102. The valve V102 simulates a non-return valve; therefore, the water does not return when the pump is not operating. The flow in the pipe is measured by a flowmeter FIC/B102. The ultrasonic sensor LIC/B101 placed at the top of tank B102 measures the water level. In this model, a cascade control aiming to maintain the required water level SetPoint (SP) is implemented. The ultrasonic sensor reading is the process variable PV for the cascade outer loop and the flowmeter reading is used as the process variable for the cascade inner loop. The proportional valve V106 simulates the water demand for a small town. Finally, the water returns to tank B101.

5 PLC Memory Corruption

The previous section showed a brief summary of the PLC operation, and in this section three novel PLC memory corruption attacks are introduced.

5.1 Attack to the PLC inputs

In this attack, the aim is to overwrite the bytes of memory assigned to the external sensors in the PLC. This attack affects the control process because the PLC uses wrong readings while executing the internal code that will define the PLC outputs. Crafted ISO 8073/X.224 COTP packets are sent over the network to the PLC. This attack is repeated constantly as the PLC updates the input table at the start of each cycle. For instance, in Fig 4 the control system maintains the tank B102 level at a determined setpoint by controlling the pump 101.

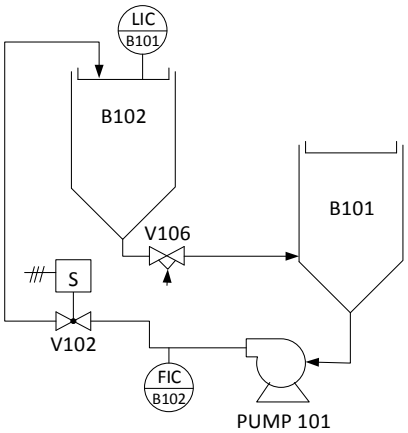

Fig. 4. Clean water supply system control diagram

The control process is configured to maintain the reservoir tank setpoint at 6 litres. In the first scenario, the attacker tampers with the memory used to hold data from the ultrasonic sensor, so that the tank appears to contain 1 litre of water. The reaction of the control system is to achieve the setpoint by speeding up the pump, which results in overflowing the B102 tank.

5.2 Attack to the PLC Outputs

In this attack, the aim is to overwrite the memory associated with the outputs in the PLC. The same attack logic is employed as in the previous attack. However, the main difference is in the devices targeted during this attack. The aim to overwrite the memory associated with the pump with the intention of increasing its speed to overflow tank B102 showed in Fig 4. To achieve this, crafted packets are injected with values that require the pump work at 100% of its capacity. This is reminiscent of the Stuxnet [16] attack discussed earlier where the attack increased and decreased the values injected in the PLC Output memory. In this paper, we plan to increase and reduce the speed of the pump at the time to emulate overflowing and emptying the tank B102.

5.3 Attacks to the PLC working memory

As discussed before, the PLC is composed of different memory elements, such as work memory, retentive memory, etc. The work memory contains the PLC code that is executed at runtime [14]. Process variables such as the setpoint are allocated in this part of the PLC memory. In this scenario, the memory associated with the setpoint variable is modified. Two different values can be controlled, namely the high and low setpoints for a system. In the attack the setpoint value is modified by steady increasing and decreasing it. This is similar to the attack proposed in [17] and [18]. However, they proposed a theoretical approach while in this work it was physically implemented.

6 Mechanisms of attack detection

The detection of attacks in control systems can be seen from different points of view because some of these systems are connected to corporate networks; as a result, they face targeted attacks like Stuxnet [16] and common attacks such as buffer overflow, SQL injection and more. In this paper, attack detection is implemented by monitoring the signal of the ultrasonic sensor and the flowmeter.

In the author's previous work, the feasibility of detecting cyber attacks was demonstrated in a control process by monitoring the energy consumption of the pump [19]. In the first attempt, this was by monitoring the energy consumption of the ultrasonic sensor and the flowmeter; however, these sensors consume a low amount of energy which makes monitoring the energy consumption difficult. For that reason, current sensors

were used instead at the PLC Inputs for the ultrasonic sensor and the flowmeter. This means, for instance, it should detect any change to the setpoint variable by monitoring the input signal of the ultrasonic sensor. This signal is represented in volts (v).

7 Results

The results of the energy consumption records in the pump and voltage signals of the ultrasonic sensor and flowmeter are shown in Fig 5. The shaded areas show each one of the attacks performed and they are explained as follows. Table 1 summarizes the attacks performed to the clean water supply system, in addition it also includes the shadow area that is represented in Fig5.

Table 1. Summary of attacks to the PLC Memory

Shadow Area	PLC Memory	Device Attacked	Results
1	Working Memory	- Setpoint	- Signal in the ultrasonic sensor increases/decreases. - Energy consumption in the pump increases/decreases.
2	PLC Input	- Ultrasonic Sensor Reading	- Signal in the ultrasonic sensor increases/decreases. - Energy consumption in the pump increases/decreases.
3	PLC Output	- Pump Speed	- Signal from the ultrasonic sensor increases/decreases. - Energy consumption in the pump increases/decreases.
4	PLC Input	- Flowmeter Reading	- Energy consumption in the pump decreases.
5	PLC Input/output	- Pump Speed - Ultrasonic Sensor	- System unstable.

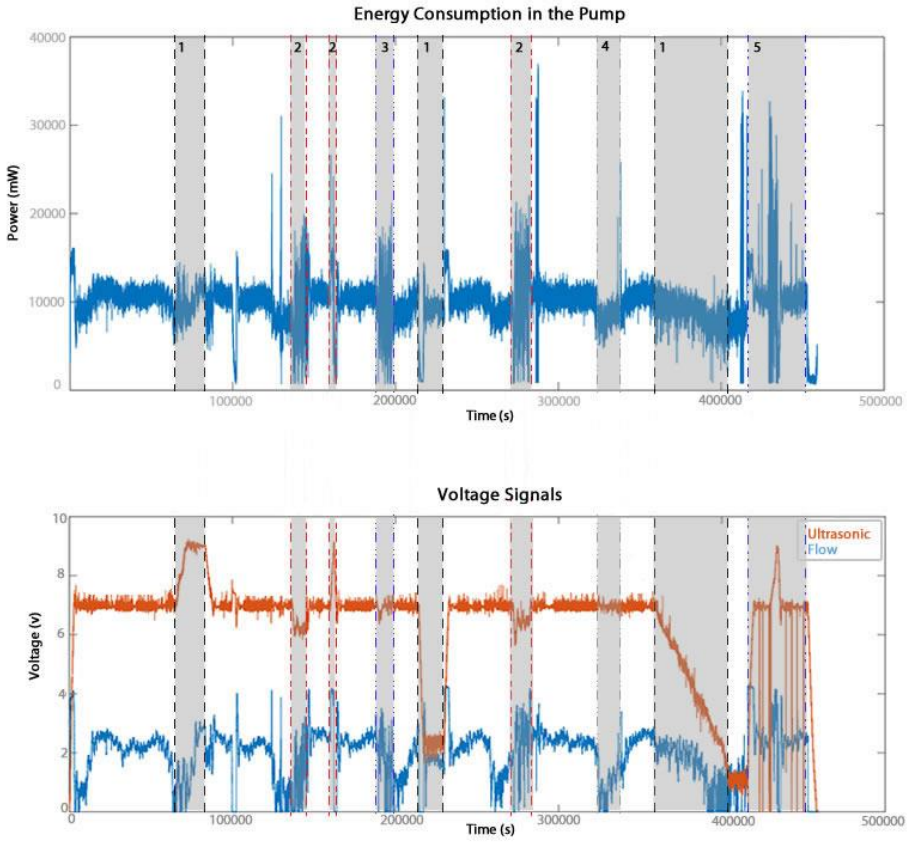

Fig. 5. Energy and voltage signals.

Attack on the PLC working memory. When the setpoint suffers a significant change, for instance from 6 litres to 9 litres, the pump starts speeding up to achieve the new setpoint, as a result, it consumes more energy than expected. In Fig 5, the first shadow area from the left-hand side which marked with 1 shows this sudden change. The energy consumption of the pump might indicate that an attack is happening in the control system, although it does not indicate in which part of the control system it is happening. On the other hand, the ultrasonic sensor signal allows monitoring the water level in the reservoir tank, and when the setpoint increases, the ultrasonic sensor signal also increases. Thus, the energy monitored in the pump correlated with the ultrasonic sensor signal indicating the parts of the control system that have been affected. It can be argued that the PLC permits collecting and monitoring the energy of the ultrasonic sensor, as a result, it could allow detecting the increment in the setpoint using normal rules. However, in this scenario, detection does not rely on the information provided by the PLC

because it might itself be compromised. For that reason, external and independent current sensors are preferred.

In the second scenario, when the setpoint goes from six litres to two litres the pump stops working, and the ultrasonic sensor signal decreases. In Fig 5, the second shadow area from the left-hand side marked with 1 shows this change in the ultrasonic sensor and the pump. Until the setpoint is reached the energy consumption in the pump falls to zero along with the flow in the pipe. Afterwards, the pump starts working again to maintain the new setpoint. However, the energy used by the pump is lower than the normally used.

In the third scenario, the intention is to avoid the sudden change in the energy consumption of the pump and so the attack decreases the setpoint by one over time until the tank is empty. In Fig 5, the third shadow area from the left-hand side marked with 1 shows this attack. It can be seen that the energy consumption of the pump does not suffer a considerable change. However, the signal in the ultrasonic sensor starts to decrease indicating that the tank starts to empty.

Attack to the PLC Inputs. In this case the memory associated with the ultrasonic sensor signal is overwritten by injecting a value that represents eight litres. As a result, the pump slows down and the reservoir tank is reduced to one litre of water. In Fig 5, the first shadow area marked with 2 shows this attack. In the second attack, memory was overwritten with a value that represented one litre of water, and as a result, the pump started working at 100% of its capacity resulting in a rapid increase of water and an imminent overflow of the reservoir tank. This attack can be seen in the second shadow area marked with 2 in Fig 5. It should be considered that the response of the control system to attacks depends on how it is designed and implemented. For this testbed a cascade control system was implemented. It worth mentioning that a different technique such as Proportional Integral (PI) controller might produce different results to this attack.

In the next attack, the PLC memory addressed to the flowmeter was overwritten with a value that represents 4.1 litres per minute, which is the maximum value that the pump provides. This attack does not affect the operation of the control system, because this value is used for the inner loop of the cascade control, which is designed to handle high amounts of noise. Meaning that the control system is interpreting the attack as noise. In addition, the flow in this control system is low, which is another reason for the low impact of the attack. In Fig 5, the shadow area marked with 4 shows this attack.

Attack to the PLC Outputs. In this scenario, the aim is to compromise the pump by overwriting the PLC output memory addressed to it. This increased and decreased the pump speed by injecting integer values into the PLC memory. In Fig 5, the shadow area marked with 3 shows that the energy consumption in the pump starts to fluctuate. It can

be concluded that this affected the control system, however, it does not cause a huge impact. Instead, this attack might lead to damage the pump over time and stop completely the process operation.

Multiple attacks. In this scenario, an attack was executed on all the devices that are part of the control system at the same time, as a result, it affected the entire process operation. In Fig 5, the grey area marked with 5 shows the behaviour of the control process. In this attack, the memory associated with the pump was overwritten with a small value aiming to slow down the pump. In addition, random values were injected in the memory associated with the ultrasonic sensor. At the end, the setpoint was modified to zero. During this attack, the control system operation is highly affected. It could be said that performing this type of attacks on real control systems might have the same effect compromising the water supply of a certain population.

8 Conclusions

In this paper, a novel attack vector is proposed for control systems which targets PLC memory corruption in three places: PLC input, PLC output, and PLC working memory. These attacks were demonstrated in a clean water supply system modelled in the Festo Rig to show the impact of the attacks in a real scenario. The execution of these attacks showed that it is possible to disrupt the control system operation by overwriting the memory locations associated to PLC Inputs and Outputs. In addition, attacks were made which modified the setpoint variable located in the PLC working memory. It can be concluded that most published research proposes different types of theoretical attacks. However, practical implementations are needed to measure the impact of those attacks. It can be argued whether the theoretical attacks proposed are applicable in real implementations. Further, this paper extended our previous work by offering a mechanism of detection based on monitoring the energy consumption of the pump and the electrical signal from the ultrasonic sensor and the flowmeter. The results show the feasibility of detecting the attacks performed to the control system by monitoring the energy and voltage parameters. When reducing the setpoint by one until reaching the minimum setpoint, the energy consumption in the pump did not show a considerable change. However, the electrical signal in the ultrasonic sensor is reduced showing that the reservoir tank is emptying. In future work, the feasibility of attack detection in industrial control systems using control engineering techniques will be investigated.

References

1. KASPERSKY, "THREAT LANDSCAPE FOR INDUSTRIAL AUTOMATION SYSTEMS IN H1 2017," 20 May 2018. [Online]. Available: https://ics-cert.kaspersky.com/reports/2017/09/28/threat-landscape-for-industrial-automation-systems-in-h1-2017/.
2. "Secure Water Treatment (SWaT) Dataset," 21 May 2018. [Online]. Available: https://itrust.sutd.edu.sg/research/dataset/.
3. T. H. Morris and W. Gao, "Industrial Control System Cyber Attacks," in Proceedings of the 1st International Symposium on ICS \& SCADA Cyber Security Research 2013, Leicester, 2013.
4. T. Bradley, "Water Utility Hacked: Are Critical Systems at Risk?," PCWorld, 20 November 2011. [Online]. Available: https://www.pcworld.com/article/244359/water_utility_hacked_are_our_scada_systems_at_risk_.html. [Accessed 30 April 2018].
5. FESTO, "MPS® PA Compact Workstation with level, flow rate, pressure and temperature controlled systems," [Online]. Available: http://www.festo-didactic.com/int-en/learning-systems/process-automation/compact-workstation/mps-pa-compact-workstation-with-level,flow-rate,pressure-and-temperature-controlled-systems.htm?fbid=aW50LmVuLjU1Ny4xNy4xOC44ODIuNDM3Ng.
6. J. Khurum Nazir and J. Goh, "Behaviour-Based Attack Detection and Classification in Cyber Physical Systems Using Machine Learning," in Proceedings of the 2Nd ACM International Workshop on Cyber-Physical System Security, Xi'an, China, 2016.
7. S. Adepu and A. Mathur, "An Investigation into the Response of a Water Treatment System to Cyber Attacks," in 2016 IEEE 17th International Symposium on High Assurance Systems Engineering (HASE), 2016.
8. Terai, S. Abe, S. Kojima, Y. Takano and I. Koshijima, "Cyber-Attack Detection for Industrial Control System Monitoring with Support Vector Machine Based on Communication Profile," in 2017 IEEE European Symposium on Security and Privacy Workshops (EuroS PW), 2017.
9. W. Hurst, M. Merabti and P. Fergus, "Big Data Analysis Techniques for Cyber-threat Detection in Critical Infrastructures," in 28th International Conference on Advanced Information Networking and Applications Workshops, 2014 .
10. J. Inoue, Y. Yamagata, Y. Chen, C. Poskitt and J. Sun, "Anomaly Detection for a Water Treatment System Using Unsupervised Machine Learning," in IEEE International Conference on Data Mining Workshops, 2017.
11. J. Goh, S. Adepu, M. Tan and Z. Lee, "Anomaly detection in cyber physical systems using recurrent neural networks," in 2017 IEEE 18th International Symposium on High Assurance Systems Engineering (HASE), 2017.
12. A. Almalawi, X. Yu, Z. Tari, A. Fahad and I. Khalila, "An unsupervised anomaly-based detection approach for integrity attacks on SCADA systems," Computers & Security, p. 94, 2014.
13. A. Frank and A. Asuncion, "UCI machine learning repository," 12 May 2018. [Online]. Available: http://archive.ics.uci.edu/ml.

14. K. Kamel and E. Kamel, Programmable Logic Controllers: Industrial Control, US: McGraw-Hill Professional, 2013.
15. W. Bolton, "Chapter 2 - Input/Output Devices," in Programmable Logic Controllers (Sixth Edition), Boston, Newnes, 2015, pp. 23 - 61.
16. R. Langner, "Stuxnet: Dissecting a Cyberwarfare Weapon," in IEEE Security Privacy, 2011.
17. I. Shames, A. Texeira, H. Sandberg and K. Johansson, "Revealing stealthy attacks in control systems," in 2012 50th Annual Allerton Conference on Communication, Control, and Computing (Allerton), 2012.
18. D. Urbina, J. Giraldo, A. Cardenas, N. O. Tippenhauer, J. Valente, M. Faisal, J. Ruths, R. Candell and H. Sandberg, "Limiting the Impact of Stealthy Attacks on Industrial Control Systems," in Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security, Vienna, 2016.
19. A. Robles-Durazno, N. Moradpoor, J. McWhinnie and G. Russell, "A supervised energy monitoring-based machine learning approach for anomaly detection in a clean water supply system," in CYBER SECURITY 2018, 2018 (Accepted for publication).
20. NORDPOOL, "Market Data Nord Pool," NORDPOOL, [Online]. Available: https://www.nordpoolgroup.com/Market-data1/Power-system-data/Consumption1/Consumption/ALL/Hourly1/?view=table. [Accessed 30 April 2018].
21. R. Gentleman and V. J. Carey, "Unsupervised Machine Learning," in Bioconductor Case Studies, New York, Springer New York, 2008, pp. 137-157.
22. J. Hartigan and M. Wong, "Algorithm AS 136: A K-Means Clustering Algorithm," Journal of the Royal Statistical Society, vol. 28, no. 1, pp. 100-108, 1979.
23. K. Wagstaf, C. Cardie, S. Rogers and S. Schroedl, "Constrained K-means Clustering with Background Knowledge," Proceedings of the Eighteenth International Conference on Machine Learning, p. 577–584, 2001.
24. W. Gao, T. Morris, B. Reaves and D. Richey, "On SCADA control system command and response injection and intrusion detection," in 2010 eCrime Researchers Summit, 2010.
25. Q. Chen, R. Abercrombie and F. Sheldon, "Risk Assessment for Industrial Control Systems Quantifying Availability Using Mean Failure Cost (MFC)," 2015.