2019 IEEE Security and Privacy Workshops (SPW)

Devil in the Detail: Attack Scenarios in Industrial Applications

1st Simon D. Duque Anton
Intelligent Networks Research Group
German Research Center for AI
67663 Kaiserslautern, Germany
Simon.Duque_Anton@dfki.de

2nd Alexander Hafner
Intelligent Networks Research Group
German Research Center for AI
67663 Kaiserslautern, Germany
Alexander.Hafner@dfki.de

3rd Hans Dieter Schotten
Intelligent Networks Research Group
German Research Center for AI
67663 Kaiserslautern, Germany
Hans_Dieter.Schotten@dfki.de

Abstract—**In the past years, industrial networks have become increasingly interconnected and opened to private or public networks. This leads to an increase in efficiency and manageability, but also increases the attack surface. Industrial networks often consist of legacy systems that have not been designed with security in mind. In the last decade, an increase in attacks on cyber-physical systems was observed, with drastic consequences on the physical work. In this work, attack vectors on industrial networks are categorised. A real-world process is simulated, attacks are then introduced. Finally, two machine learning-based methods for time series anomaly detection are employed to detect the attacks.** ***Matrix Profiles*** **are employed more successfully than a predictor** ***Long Short-Term Memory*** **network, a class of neural networks.**

Index Terms—**Cyber Security, Time Series, Machine Learning, Neural Networks, Industrial Control Systems**

I. Introduction

For decades, the industrial domain has been deemed secure due to two reasons: First, the physical separation of networks. Second, each network was created in an application specific fashion, rendering it extremely difficult for an attacker to exploit it [1]. However, the fourth industrial revolution introduced novel use cases that build on interconnectivity and embedded intelligence [2], [3]. While increasing productivity and flexibility and decreasing operational cost and effort, new attack vectors are introduced to industrial systems as well. An increase in attacks on industrial environments can be detected [4]. While industrial networks have been unique in their applications specific nature, the establishment of Commercial Off The Shelf (COTS) hard- and software introduces standardised modules. This makes set up and maintenance much easier, but also drastically increases the effect of vulnerabilities in one of the modules. In order to tackle these problems, cyber security measures have been adapted to industrial scenarios, such as firewalls, anti virus software and intrusion detection tools. However, the characteristics of industrial networks differ from those of home and office networks, motivating the need for adaption of those tools. A deep understanding of these characteristics is required in order to effectively protect industrial networks. In this work, an overview of possible attacks for industrial networks is provided. Attack vectors are analysed and categorised, with an emphasis on industrial network protocols. Furthermore, the simulation of a real-world scenario is presented, as well as attacks on this scenario. The remainder of this work is structured as follows: In Section II, related work is presented. A systematic categorisation of attack scenarios is provided in Section III. The simulated process and the implementation of attacks is described in Section IV and evaluated in Section V. Finally, the findings are discussed in Section VI.

II. Related Work

In this section, related work on classification of industrial cyber attacks is presented. Furthermore, it is grouped with respect to the scope that is addressed by the work in Table I. *Cherdantseva et al.* survey existing risk assessment methods and evaluate their usefulnes with respect to Supervisory Control And Data Acquisition (SCADA) scenarios [5]. *Gao and Morris* discuss the detection of cyber attacks [6]. They focus on signature-based detection for *Modbus*-based communication. In order to evaluate the intrusion detection and to classify it, possible attacks are grouped. A more thorough analysis of attacks on Industrial Control Systems (ICSs) is performed by *Morris and Gao* as well [7]. *Zhu et al.* provide an overview of cyber attacks while considering many dimensions [8]. They compare industrial cyber security to classic IT security. Furthermore, they consider the security objectives of industrial applications and ways they can be attacked. Finally, they present specific attacks on different attack surfaces of an industrial environment. In another work, *Zhu and Sastry* create a taxonomy for SCADA-specific attacks [9]. They present types of attacks and discuss countermeasures. *Fernandez et al.* discuss the development of secure SCADA systems [10]. In doing so, attacks on industrial systems are evaluated with respect to their attack vector. *Fovino et al.* discuss the effects of SCADA attacks on infrastructure [11]. They first assess the potential damages to eventually discuss potential attack types. *Ten et al.* present a vulnerability assessment of SCADA systems [12]. They consider the increasing dependency of industrial and office Information Technology (IT). Furthermore, they classify attacks according to their type to model and

This work has been supported by the Federal Ministry of Education and Research of the Federal Republic of Germany (Foerderkennzeichen 16KIS0932, IUNO Insec). The authors alone are responsible for the content of the paper.

DOI 10.1109/SPW.2019.00040

evaluate attack scenarios in using attack trees in an earlier work [13]. *Cai et al.* analyse the development of SCADA systems, their applications and threats [14]. Furthermore, they discuss standards and guidelines for protecting such systems. *Igure et al.* discuss SCADA security [1]. They analyse attacks, categorise them and extract research challenges. Furthermore, standardisation efforts are addressed. The summary of addressed topics is shown in Table I.

TABLE I
RESEARCH TOPICS COVERED BY THE INDIVIDUAL WORKS

Subject Covered	**Research Work**
Risk Assessment	[5], [11]–[13]
Industrial vs Home- and Office IT	[8]
Attack Vectors	[10]
Security Objectives	[1], [8]
Types of Attacks	[1], [6]–[9], [11]–[13]
Standards and Guidelines	[14]
Applications in SCADA	[14]
Taxonomy	[1], [9]
Intrusion Detection	[6], [9]

Most research is done regarding the types of attacks, i.e. the way an attacker will influence the systems or networks. Risk assessment is a widely regarded topic as well. In risk assessment, the effects of an attack are discussed in a formal manner. The remaining topics are more specific and only addressed by one or two singular works.

III. INDUSTRIAL ATTACKS

In this section, possible ways for an attacker to break into industrial applications are discussed. This is done by looking at past attacks on industrial networks that have extensively been discussed. *Stuxnet*, the attack that came to attention first, has been widely discussed [15]–[18], but also the lesser known successors, such as *Duqu* [17], *Industroyer/Crashoverride* [15], [19], *Flame* [17], *BlackEnergy* [15], [19], *Havex* [15] and *Red October* [17] have received attention. An assessment of attack vectors for industrial companies is done by *Positive Technologies* [20]. They evaluate points of entry and propagation methods in a general fashion which, however, is in accordance to the above-mentioned malware-specific analyses. The first step in attacking industrial environments is commonly the breach of the perimeter. Even though there are occasions where industrial networks are directly connected to the Internet [4], they are commonly separated from public networks. This is an important recommendation in securing industrial networks [20], especially since many industrial network protocols do not contain means for authentication or encryption. This allows easy propagation and participation in communication for an attacker once the network is accessible. If the production network is not reachable from the outside, the corporate network needs to be breached first. According to *Positive Technologies*, 73% of the corporate systems they tested had insufficient protection of their perimeter [20]. Another common attack vector is the human user. Allegedly the *Stuxnet* attack has breached the perimeter by means of a thumb drive that was carelessly used [18]. After breaking the perimeter, the ICS or the field devices respectively have to be taken over. The analysed malware that was tailor-made for industrial targets used properties or vulnerabilities characteristic to the industrial environment they were designed for. Some malware could stay undetected for long periods of time, at least partly due to missing or insufficient security procedures for industrial networks. Implementing robust security for critical parts of production networks is one of the major take aways. Most industrial malware consists of several modules:

- Backdoor,
- loader module, and
- wiper module

The backdoor allows for communication with Command & Control (C&C) servers. Coupled to the backdoor is the loader module that is tasked with uploading the malware modules to perform certain attacks. And lastly, most industrial malware contains a module for wiping the traces of its existence from the infected system. Breaking the perimeter has proven to be possible most of the time. The difficulty of industrial malware lies in the profound knowledge the malware authors needs to have about the targeted systems. This goes for the architecture of the infrastructure as well as for the protocols and devices used. Most devices used in industrial applications are COTS products and can thus be obtained for vulnerability analysis, so that exploits can be written and re-used. However, to successfully break a process by abusing system parameters, the intent of devices as well as the structure of the process needs to be known. These attacks are hardest to detect, as the attacker can conceal them as irregularities or normal behaviour. The attacks that are implemented and evaluated in this work are such attacks. They are based on the assumption of a successful breach of perimeter and take over of a Programmable Logic Controller (PLC) which subsequently shows malicious behaviour.

In summary, any attack of an industrial application first needs to break the perimeter. Then it needs to move laterally towards the control system or target device. Finally, the malicious intent has to be carried out. During each of these steps, the attack can be discovered by different means. Breaking the perimeter should be observed by IT-based security means. Lateral movement is in the domain of **siem!** (**siem!**)-systems. Detecting attacks in the context of an industrial process is the final method to discover misbehaviour.

IV. PROCESS ENVIRONMENT AND ATTACK SCENARIOS

In this section, the process under investigation as well as the implemented attacks are discussed. First, the real-world process is described and transferred to a simulation. After that, the attack scenarios and their implementations are presented.

A. Process Environment

The process this work is based upon has been used to generate data for industrial intrusion detection already [21]. It is shown in Figure 1. The process environment consists of

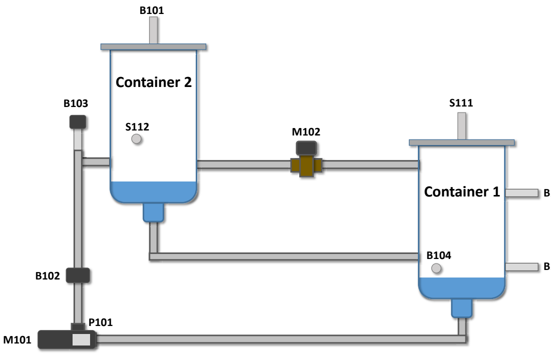

Fig. 1. Schematic Overview of the Process Environment

two water containers, *Container 1* and *Container 2*. Water is pumped with pump *P101*, driven by a DC motor *M101* from *Container 1* to *Container 2* until a threshold is reached. The water level is measured with different sensors, *S111* and *S112*, as well as capacitive sensors *B113* and *B114*. Additionally, A vane sensor measuring the flow of liquid between *Container 1* and *Container 2*, *B102*, and a *PT100* temperature sensor, *B104*, are used. To release water from *Container 2*, a solenoid valve, *M102*, is employed. An exemplary behaviour of this process is shown as a time-series in Figure 2. A selection of process parameters, all of them sensor outputs, during normal operation is shown. This operation has been performed on real implementation of the scenario that was used to create the simulation analysed in this work. For this work, the environment described above has been extended to five instances. They are simulated with real-world hardware, i.e. *Siemens S7-1500* PLCs and *PiXtend* extension boards for *Raspberry Pis*. Five *Raspberry Pis* with a *PiXtend*-board each are used to simulate the process, controlled by a PLC each. The process information is collected on a central Human Machine Interface (HMI). In order to obtain realistic data, the simulation has been developed to mimic the real scenario as good as possible. For communication, *OPC UA* [22] is used. It provides encrypted, authenticated, easy and platform-independent communication and consists of an information model including communication capabilities. In this scenario, a master-slave concept is followed with regular polling of the devices by the HMI.

B. Attack Scenarios

Two scenarios have been implemented and evaluated in this work. They are loosely coupled to the categorisation of *Morris and Gao* [7]. In this work, all attacks are a kind of *Response and Measurement Injection Attack*. In creating the data set for evaluation, one of the five PLCs shows malicious behaviour for five minutes after 15 minutes of normal operation for each attack. The use case is an attacker having breached the perimeter, bridged the air gap and used well-engineered malicious code to disrupt the process. The aim of this scenario is to detect malicious behaviour on field level. This area is currently not well-developed, intrusion detection on field level is a growing field with a short history. The behaviour regarding flow and water level of *Container 1* of the malicious process is shown in Figure 3. Since a specific application use case is discussed in this work, the attacks are domain-specific. However, as discussed in Section III, the general concept of this kind of attack can be generalised to most processing units.

1) Open Valve Attack: In the first attack scenario, the valve *M102* is opened even though water is pumped from *Container 1* to *Container 2*. This leads to an increased time it takes for *Container 2* to be filled up to the desired level. The HMI indicates the valve as open. This attack starts at packet 4,200 and ends around packet 4,800 in Figure 3. As in the process, the valve is not supposed to be opened, this is an identifier of the attack, making it trivial to detect. In order to better evaluate the methods presented in this work, it is not used as an input variable.

2) Stealth Attack: The second attack scenario is implemented in a stealthier fashion. As in attack scenario 1, the valve is opened invalidly. However, the sensor still indicates a closed valve. This leads to an unexpected decrease in filling speed of *Container 2* and an increased emptying once the container is filled. This attack starts at packet 6,500 and ends at the end of the trace in Figure 3.

V. Evaluation

In this section, the methods to detect the discussed attacks are presented and evaluated on the data set. As input values for the anomaly detection, the water flow as well as the water level of *Container 1* are used. They could easily be extended, but proved to be the most expressive variables.

A. Matrix Profiles

Time series-based anomaly detection has proved to be highly effective in industrial intrusion detection [23]. As the process is expected to produce regular sensors and actuator values, deviations of a time series representation of those values should be detectable. In this work, *Matrix Profiles* are used to analyse the data sets. *Matrix Profiles* were introduced by *Yeh et al.* in 2016 and provide a mean to determine the similarity of sequences in a time series to other sequences [24]. In order to employ *Matrix Profiles*, only one hyper-parameter needs to be set, the window size m. It is robust against changes and provides sensible results for a variety of lengths. However, each time window needs to contain a set of values that has a standard deviation that is not zero. Thus, m needs to be chosen in a way that no window in the water flow values only contains zero flow, as shown in Figure 4. In this figure, the water flow, as well as the water level of *Container 1*, are shown in combination with their respective *Matrix Profiles*. The *Matrix Profile* determines the minimal distance of any windowed sequence of length m from any other sequence of length m. In terms of anomaly detection, a high minimal distance represents an outlier, as the corresponding window does not look like any other. In Figure 4, the normal behaviour of the process is shown, with m as 300. Even though m proved to be robust in the evaluation, auto-correlation [25] was employed in

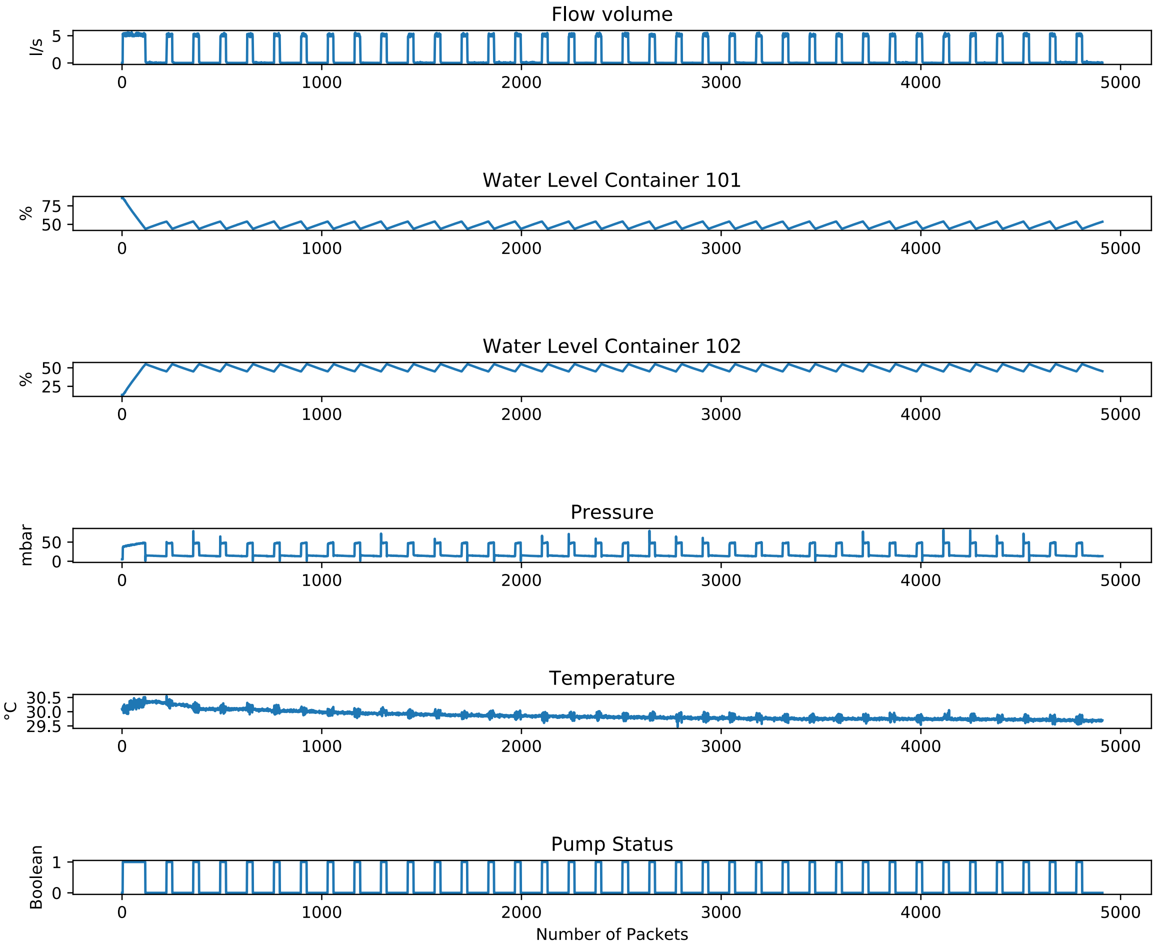

Fig. 2. Normal Process Behaviour

order to find a sensible value. The auto-correlation function of the normal process is shown in Figure 5. The peak at around 150 seconds indicated periodic behaviour. In our experiments, two measurements per second were performed, thus a window size of 300 packets was chosen as m. In Figure 4, the *Matrix Profiles*, named *Min. Dist.*, are small, except for the beginning. The settling of the process is a unique event, thus the high minimal distance. The process was monitored on one of the PLCs that did not exhibit malicious behaviour. This was introduced to another PLC and resulted in the behaviour shown in Figure 6. Both attacks show significant peaks in the minimal distances around both attacks as described in Subsection IV-B. It is noteworthy that the transitions from normal to malicious behaviour, and vice versa, are the events in the time series that are unique and thus result in an increased minimal distance. If an attack has a characteristic signature that is repeated more than once, it is not detected as an anomaly anymore, as another instance with the same characteristic is found. To counter this effect, *Matrix Profiles* can be adapted in a way that they are employed continuously. This shows promising results [23]. This approach indicates any change in behaviour. To mitigate disruptions due to alerts, natural changes in processes can be integrated into *Matrix Profiles* with an extension [26]. Furthermore, *Matrix Profiles* can be used to analyse meta data, providing good results as well [23].

B. Long Short-Term Memory

Many types of Recurrent Neural Networks (RNNs) tend to neglect long-term dependencies in the decision making. In order to keep such information, *Hochreiter and Schmidhuber* proposed a novel kind of RNNs in 1997 to overcome the vanishing gradient problem [27]. This kind of RNN is called *Long Short-Term Memory* (*LSTM*). In this work, an *LSTM* with an input layer consisting of 350 units, two hidden layers

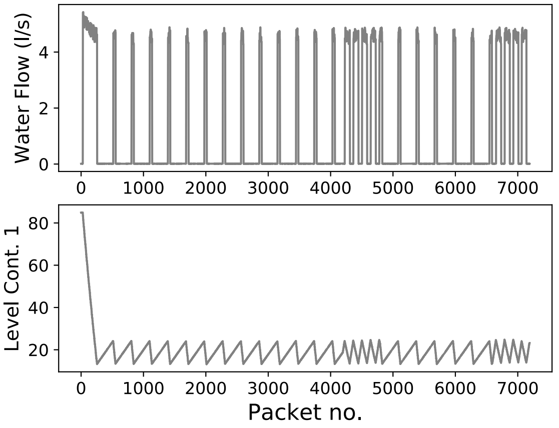

Fig. 3. Malicious Process Behaviour

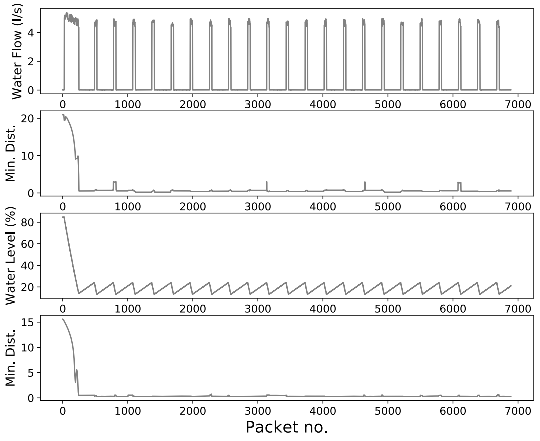

Fig. 4. Time Series and *Matrix Profiles* of Normal Behaviour

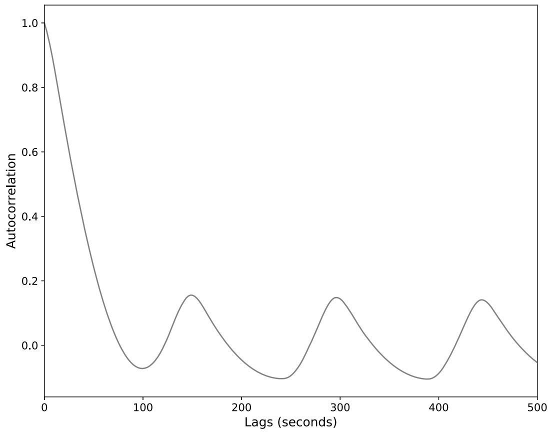

Fig. 5. Auto-correlation of Normal Time Series

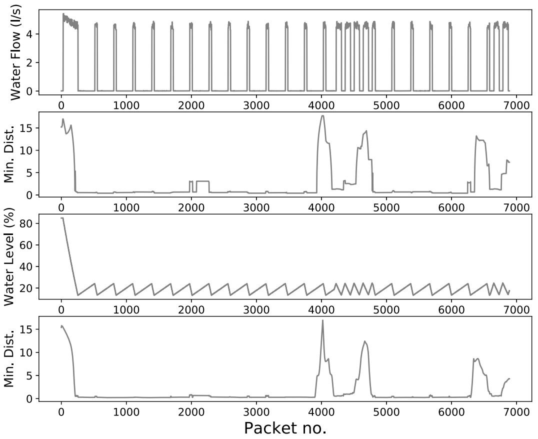

Fig. 6. Time Series and *Matrix Profiles* of Malicious Behaviour

with 350 and 250 units respectively and a dense output-layer is employed. The input length is 300 as this is the minimal periodicity of the data. The learning rate was set to 0.001. An hour of process activity was used to train it in 25 iterations. No attacks were contained in the data. After that, the hour of process activity conducted by the infected PLC was used as the testing data set. In order to monitor anomalousness, a value was predicted by the neural network and compared to the real value. The distance between those values was calculated, a high value indicating an anomalous instance. The result of the *LSTM* is shown in Figure 7. The first row shows the real values of water flow, the second row the predicted values and the third row the absolute error. The fourth row shows the real values of the water *Container 1*, the second row the predicted values and the third row the absolute error. It can be seen that the *LSTM* closely follows the process behaviour. Unfortunately, this also includes the attacks. They are predicted as part of the process by the *LSTM*, making detection of the attacks difficult. Only the frequency of the periodic error behaviour changes, however, values that clearly indicate attacks would enhance the detection probability.

VI. Discussion

In this work, we discussed the attack scenarios in industrial environments. From these scenarios, a use case was derived and implemented. After that, attack scenarios where introduced to the scenario. Two time series-based anomaly detection methods were employed to detect the attacks. *Matrix Profiles* performed satisfactorily, detecting the attacks easily. Only one robust hyper-parameter and no supervised training make it easy to use and transfer between application domains. The *LSTM* approach did not work well, it predicted the attack behaviour as well as the normal behaviour. This behaviour can

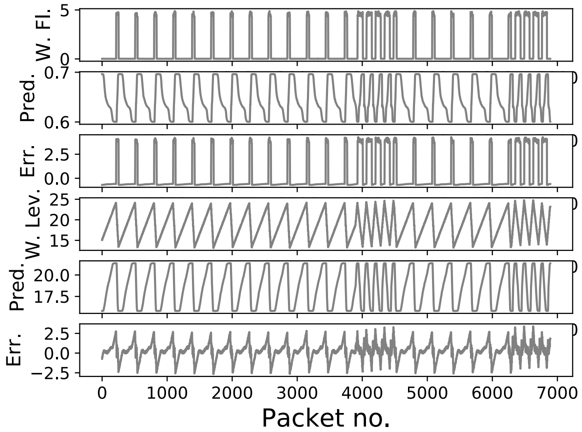

Fig. 7. Time Series, Predictions and Errors Based on *LSTM*s

derive from over-fitting, as regular time series have a tendency to teach neural networks to learn certain patterns, but not to generalise. Context information [28] or methods of machine learning-based classification [29] might address the issue as well.

A. An Epilogue on Sophisticated Industrial Attacks

One of the major features of sophisticated industrial attacks such as *Stuxnet* is the masquerading of any indicators for misbehaviour. However, if no trace of malicious behaviour is simulated, it simply cannot be detected. For the sake of clarity in this work, only attacks with distinctive characteristics were used, so that detection was possible. After attacks such as *Stuxnet* propagated into the industrial domain, side-channel detection, e.g. acoustic signals, would be required if standard field busses were used. *Langner* claims that any engineer with experience in the area would have told something was amiss easily by the sound of the turbines. Unfortunately, such side-channels are hard to simulate. However, there are works creating data sets of real-world applications including side-channel sensor measurements so that they can be used to detect attacks [21].

References

[1] V. M. Igure, S. A. Laughter, and R. D. Williams, “Security issues in SCADA networks,” *Computers & Security*, no. 25, pp. 498–506, 2006.

[2] “Study on communication for automation in vertical domains,” 2017, 3GPP TR 22.804, V1.0.0.

[3] S. Plaga, N. Wiedermann, S. Duque Anton, S. Tatschner, H. D. Schotten, and T. Newe, “Securing future decentralised industrial IoT infrastructures: Challenges and free open source solutions,” *Future Generation Computer Systems*, vol. 93, pp. 596–608, 2019.

[4] S. Duque Anton, D. Fraunholz, C. Lipps, F. Pohl, M. Zimmermann, and H. D. Schotten, “Two decades of SCADA exploitation: A brief history,” in *2017 IEEE Conference on Application, Information and Network Security (AINS)*, November 2017, pp. 98–104.

[5] Y. Cherdantseva, P. Burnap, A. Blyth, P. Eden, K. Jones, H. Soulsby, and K. Stoddart, “A review of cyber security risk assessment methods for SCADA systems,” in *Computers & Security*, no. 56, 2016.

[6] W. Gao and T. H. Morris, “On cyber attacks and signature based intrusion detection for modbus based industrial control systems,” *Journal of Digital Forensics, Security and Law*, vol. 9, no. 1, 2014.

[7] T. H. Morris and W. Gao, “Industrial control system cyber attacks,” in *Proceedings of the 1st International Symposium for ICS & SCADA Cyber Security Research*, 2013, pp. 22–29.

[8] B. Zhu, A. Joseph, and S. Sastry, “A taxonomy of cyber attacks on SCADA systems,” in *Proceedings of the 2011 International Conference on Internet of Things and 4th International Conference on Cyber, Physical and Social Computing*, ser. ITHINGSCPSCOM. Washington, DC, USA: IEEE Computer Society, 2011, pp. 380–388.

[9] B. Zhu and S. Sastry, “SCADA-specific intrusion detection / prevention systems : A survey and taxonomy,” 2010.

[10] E. B. Fernandez, J. Wu, M. Larrondo-Petrie, and Y. Shao, “Designing secure SCADA systems using security patterns,” in *43rd Hawaii International Conference on System Sciences*, January 2010, pp. 1–8.

[11] I. N. Fovino, A. Carcano, and A. Trombetta, “An experimental investigation of malware attacks on SCADA systems,” in *International Journal of Critical Infrastructure Protection*, 2009, pp. 139–145.

[12] C.-W. Ten, C.-C. Liu, and G. Manimaran, “Vulnerability assessment of cybersecurity for SCADA systems,” in *IEEE Transactions on Power Systems*, vol. 23, no. 4, November 2008, pp. 1836–1846.

[13] ——, “Vulnerability assessment of cybersecurity for SCADA systems using attack trees,” in *IEEE Power Engineering Society General Meeting*, June 2007, pp. 1–8.

[14] N. Cai, J. Wang, and X. Yu, “SCADA system security: Complexity, history and new developments,” in *The IEEE International Conference on Industrial Informatics (INDIN)*, July 2008.

[15] Dragos, “Chrashoverride - analysis of the threat to electric grid operations,” Dragos Inc., Tech. Rep. 2.20170613, 2016.

[16] R. Langner, “To kill a centrifuge,” The Langner Group, Tech. Rep., November 2013.

[17] N. Virvilis and D. Gritzalis, “The big four - what we did wrong in advanced persistent threat detection?” in *2013 International Conference on Availability, Reliability and Security*, September 2013, pp. 248–254.

[18] J. R. Lindsay, “Stuxnet and the limits of cyber warfare,” *Security Studies*, vol. 22, no. 3, pp. 365–404, 2013.

[19] A. Cherepanov, “Win32/Industroyer - a new threat for industrial control systems,” ESET, Tech. Rep., June 2017.

[20] Positive Technologies, “Industrial companies - attack vectors,” Positive Technologies, Tech. Rep., 2018.

[21] S. Duque Anton, M. Gundall, D. Fraunholz, and H. D. Schotten, “Implementing scada scenarios and introducing attacks to obtain training data for intrusion detection methods,” in *International Conference on Cyber Warfare and Security (ICCWS)*, 2019.

[22] IEC, “Opc unified architecture - part 1: Overview and concepts,” International Electrotechnical Commission (IEC), Technical Report, 2016. [Online]. Available: https://webstore.iec.ch/publication/25997

[23] S. Duque Anton, L. Ahrens, D. Fraunholz, and H. D. Schotten, “Time is of the essence: Machine learning-based intrusion detection in industrial time series data,” in *IEEE International Conference on Data Mining Workshops (ICDMW)*. IEEE, 2018.

[24] C.-C. M. Yeh, Y. Zhu, L. Ulanova, N. Begum, Y. Ding, H. A. Dau, D. F. Silva, A. Mueen, and E. Keogh, “Matrix profile i: All pairs similarity joins for time series: A unifying view that includes motifs, discords and shapelets,” in *2016 IEEE 16th International Conference on Data Mining (ICDM)*, December 2016, pp. 1317–1322.

[25] J. A. Gubner, “Probability and random processes for electrical engineers,” 2006.

[26] Y. Zhu, M. Imamura, D. Nikovski, and E. Keogh, “Matrix profile vii: Time series chains: A new primitive for time series data mining,” in *IEEE International Conference on Data Mining (ICDM)*, November 2017, pp. 695–704.

[27] S. Hochreiter and J. Schmidhuber, “Long short-term memory,” *Neural Computation*, vol. 9, no. 8, pp. 1735–1780, November 1997.

[28] S. Duque Anton, D. Fraunholz, S. Teuber, and H. D. Schotten, “A question of context: Enhancing intrusion detection by providing context information,” in *13th Conference of Telecommunication, Media and Internet Techno-Economics (CTTE-17)*, 2017.

[29] S. Duque Anton, S. Kanoor, D. Fraunholz, and H. D. Schotten, “Evaluation of machine learning-based anomaly detection algorithms on an industrial Modbus/TCP data set,” in *Proceedings of the 13th International Conference on Availability, Reliability and Security (ARES)*. ACM, 2018.