

THESIS FOR THE DEGREE OF LICENTIATE OF ENGINEERING

Departure-Based Intrusion Detection
*Process-Level Detection of Stealthy Attacks on
Industrial Control Systems*

WISSAM AOUDI

Department of Computer Science and Engineering
CHALMERS UNIVERSITY OF TECHNOLOGY
Gothenburg, Sweden, 2019

DEPARTURE-BASED INTRUSION DETECTION

WISSAM Aoudi

Copyright ©2019 Wissam Aoudi
except where otherwise stated.
All rights reserved.

Technical Report No 192L
ISSN 1652-876X
Department of Computer Science & Engineering
Division of Networks and Systems
Chalmers University of Technology and Gothenburg University
Gothenburg, Sweden

This thesis has been prepared using L^AT_EX.
Printed by Chalmers Reproservice,
Gothenburg, Sweden 2019.

This thesis is dedicated to my beloved wife.

“When solving a problem of interest, do not solve a more general problem as an intermediate step. Try to get the answer that you need, but not a more general one. It is quite possible that you have enough information to solve a particular problem of interest well, but not enough information to solve a general problem.”

- Vladimir Vapnik

Abstract

Industrial Control Systems (ICS) combine information technology with operation technology to monitor or control physical industrial processes via computer-based programs and often operate on critical infrastructures. As such, compromised or maliciously operated ICS can cause devastating consequences on society at large. To meet efficiency requirements, ICS are becoming increasingly connected to corporate networks and to the Internet, thereby elevating the risk of cyberattacks. Resilient and sustainable highly connected ICS therefore require a serious consideration of proper security measures. Securing ICS solely from an IT perspective, while necessary, proves insufficient because, at the physical layer, the critical process would remain unmonitored and therefore vulnerable to sabotage by the attackers.

The recent years have witnessed an increased interest in process-level intrusion detection where the process network connecting field devices is monitored for malicious behavior. One prominent approach in the literature proposes to build a model of the physical process, which is then used to compare a predicted state with the actual state in the hope of identifying attacks. Building and using a predictive model of the physical process, however, is non trivial, domain specific, and prone to detection inaccuracies due to noise in the process data.

This thesis introduces a novel model-free approach to detecting cyberattacks on ICS by monitoring the process network in real time and deciding when the system operation is departing from normal dynamics. The proposed process-aware stealthy-attack detection mechanism processes raw sensor measurements to capture the dynamics of the underlying control system during a training phase, and then during a detection phase, it measures the extent to which current sensor observations conform with the estimated dynamics. The thesis provides a comprehensive treatment of the introduced method by thoroughly discussing its theoretical basis, proving its efficacy through extensive experiments on various systems, and, finally, demonstrating its applicability to real environments.

Keywords: Intrusion Detection; Industrial Control Systems; Singular Spectrum Analysis; Stealthy Attacks; Departure Detection

Acknowledgment

I am grateful to my supervisor Magnus for being such an inspiring mentor. I would have hardly accomplished this thesis without your generous advice and continuous guidance. Thank you for your honesty, diligence, and kindness!

I would also like to thank MSB for supporting my research, and everyone in the Computer Science and Engineering department at Chalmers, especially the Networks and Systems division, for the stimulating and enriching work environment.

My special thanks go to my beloved wife Malak for the tremendous support throughout my bumpy journey. Thank you for making it all possible!

I extend my gratitude to all other family members and friends for their encouragement, advice, and confidence.

List of Publications

Appended publications

This thesis is based on the following publications:

- [A] **Wissam Aoudi**, Mikel Iturbe, Magnus Almgren,
“Truth Will Out: Departure-Based Process-Level Detection of
Stealthy Attacks on Control Systems”,
*The 25th ACM Conference on Computer and Communica-
tions Security (CCS '18), October 15–19, 2018, Toronto, ON,
Canada.*

- [B] Magnus Almgren, **Wissam Aoudi**, Robert Gustafsson, Robin Krahl,
Andreas Lindhé,
“The Nuts and Bolts of Deploying Process-Level IDS in Indus-
trial Control Systems”,
*The 4th Annual Workshop on Industrial Control Systems Se-
curity (ICSS '18), December 4, 2018, San Juan, Puerto Rico,
USA.*

Contents

Abstract	vii
Acknowledgement	ix
List of Publications	xi
I INTRODUCTION	1
1 Introduction	3
1.1 Scope	4
1.2 Outline	5
1.3 Industrial Control Systems	5
1.4 Imminent Threats to Critical Infrastructures	6
1.5 Intrusion Detection Systems for ICS	7
1.5.1 Motivation	7
1.5.2 Research Challenges	8
1.5.3 Existing Solutions & Limitations	9
1.5.4 Departure-Based Attack Detection	10
1.6 Research Questions	10
1.7 Contributions	11
1.7.1 Paper A — Truth Will Out: Departure-Based Process-Level Detection of Stealthy Attacks on Control Systems	11

1.7.2	Paper B — The Nuts and Bolts of Deploying Process-Level IDS in Industrial Control Systems	12
1.8	Conclusions and Future Directions	12

II PAPERS 17

2	Paper A: <i>Truth Will Out: Departure-Based Process-Level Detection of Stealthy Attacks on Control Systems</i>	21
2.1	Introduction	23
2.2	PASAD: Process-Aware Stealthy Attack Detection	27
2.2.1	Motivation	28
2.2.2	Preliminaries & Notation	30
2.2.3	The Four Steps of PASAD	31
2.2.4	Departure Detection: The Basic Idea	33
2.2.5	Projection onto the Signal Subspace	34
2.2.6	The Isometry Trick	35
2.2.7	Efficiency: Implicit Projection	37
2.2.8	Validation: Visualizing the Departure	38
2.2.9	Choice of Parameters	40
2.2.10	Implementation & Performance	41
2.3	A Framework for Validation	43
2.3.1	Scenario I: The Tennessee-Eastman Process	43
2.3.1.1	Stealth Attacks	45
2.3.1.2	Direct Damage Attacks	46
2.3.2	Scenario II: The SWaT Dataset	47
2.3.3	Scenario III: A Water-Distribution Plant	47
2.4	Experiments & results	48
2.4.1	EXP. I: Detection of Stealth Attacks	50
2.4.2	EXP. II: Detection of Direct Damage Attacks	50
2.4.3	EXP. III: Detection of SWaT Attacks	52
2.4.4	EXP. IV: Comparison with the AR Method	53
2.4.5	EXP. V: Validating the Choice of Threshold	55
2.4.6	EXP. VI: Experimenting with Real Data	56
2.5	Related work	57
2.6	Conclusion	59

3	Paper B: <i>The Nuts and Bolts of Deploying Process-Level IDS in Industrial Control Systems</i>	69
3.1	Introduction	71
3.2	Related Work	72
3.3	Challenges	73
3.4	System Design	75
3.4.1	PASAD	75
3.4.1.1	Learning Phase	76
3.4.1.2	Detection Phase	77
3.4.2	Midbro	77
3.4.2.1	Network Capturing Engine	78
3.4.2.2	Buffering Mechanism	79
3.4.3	Choice of Hardware	80
3.5	Experiments	80
3.5.1	Modbus Communication	81
3.5.2	Controlled Experiment: Local Testbed	82
3.5.3	Live Experiment: The Paper Factory	83
3.5.4	Discussion & Results Overview	84
3.6	Lessons Learned: Guidelines & Recommendations	86
3.6.1	Process Knowledge	86
3.6.2	Signal Data Disruptions	86
3.6.3	Modbus Parsing	87
3.6.4	Buffering	88
3.7	Conclusion	89

Part I

INTRODUCTION

1

Introduction

“Imagination is more important than knowledge.”

– Albert Einstein

Nowadays, Industrial Control Systems (ICS) are everywhere. Societies are becoming growingly dependent on critical infrastructure operated by these types of systems and industries are heading towards increased connectivity to scale and meet efficiency requirements [1]. Leveraging advances in information and communication technologies is paving the way for unprecedented efficacy and flexibility of operation. Connectivity and digitalization of control systems, however, open doors to malicious actors with high motivation and resources to remotely compromise these historically isolated systems [2].

It is no understatement to say that the risk accompanying the rapidly growing connectivity trend in ICS is alarming, as threats to critical infrastructures on which normal societal functioning depends are worryingly escalating. Unlike attacks on IT systems that are often bounded by virtual impact, attacks on ICS are more serious because they can cause physical damage to critical infrastructures [3–9], potentially leading to loss of human lives or large-scale infrastructural chaos [10]. The need to secure these systems is unquestionable and efforts to secure them are increasing, albeit at a relatively modest pace in light of the scale and nature of the looming threats.

Due to the distinctive nature of ICS systems, traditional IT-based security mechanisms are often inapplicable, hence many attempts to detect attacks via direct application of off-the-shelf techniques are doomed to fall short. One approach that has proven viable in recent years proposes to monitor the *process-level* network connecting field devices to detect intrusions [11–17]. State-of-the-art methods, however, have limitations: they involve building complex models as an intermediate step and are mostly domain-specific. A model-free approach that is tailored to the ICS environment is therefore highly desirable.

This thesis presents a novel model-free data-driven detection method that is theoretically sound, efficient, practical, and widely applicable. We present a comprehensive treatment of our proposed solution, named PASAD, which is tailored to cyber-physical systems [18]. In addition, as deploying process-level Intrusion Detection Systems (IDS) in real ICS environments is still an unexplored territory, we discuss how, based on empirical evidence from real-world experiments, PASAD is sufficiently lightweight and fairly easy to deploy on limited-resource hardware [19].

1.1 Scope

The recent proliferation of cyber-physical systems has been accompanied with an ever expanding attack surface. As such, a firm security approach requires the consideration of the various threats and attack vectors. Notwithstanding the importance of *preventing* attacks on such critical systems, this thesis attempts to treat the problem of *detecting* potential attacks on ICS infrastructure and raising an alert upon detection of suspicious behavior.

While the proposed IDS is generally applicable to various kinds of cyber-physical systems, the focus of this thesis is on the ICS subset. Furthermore, while different subsystems may be monitored for intrusions, PASAD operates on the process-level network where the communication between field devices is monitored.

Finally, even at the process level, different data types may be used as input to the IDS. In this work, we consider sensor measurements generated by physical processes to detect deviations in the system behavior. By only requiring to process time series of raw sensor measurements to make decisions on the system's current

state, thereby dismissing the need for system-dependent features, PASAD proves applicable to a wide range of cyber-physical systems.

1.2 Outline

In the remainder of this chapter, we present background material, highlight the problem of increasing security threats to critical infrastructure, and introduce the concept of process-level attack detection. We conclude this chapter with research questions, our contributions, and future directions.

In Chapter 2, we present Paper A—“Truth Will Out: Departure-Based Process-Level Detection of Stealthy Attacks on Control Systems”, in which we present PASAD, a Process-Aware Stealthy-Attack Detection mechanism that monitors the process network in ICS and detects structural changes in the behavior of the underlying physical process.

Finally, in Chapter 3, we present Paper B—“The Nuts and Bolts of Deploying Process-Level IDS in Industrial Control Systems”, in which we discuss our experience and the lessons we learned from deploying an IDS prototype in a real environment.

1.3 Industrial Control Systems

Industrial control systems are cyber-physical systems that enable communication between field devices (actuators and sensors) and controllers in a closed loop fashion to control a physical process (see Figure 1.1). Abstractly considered, closed-loop control systems involve sensors that sense some physical property from the controlled process and communicate the measurements to a controller. Based on the received sensor measurements and on the implemented control logic, controllers send commands to actuators that directly manipulate the physical process to maintain a desired state of operation [1]. The controlled physical process is often sophisticated, cost-sensitive, and high-precision, and ICS are typically found in safety-critical environments. No matter if it is due to failure or malicious acts, undesired changes in the dynamics of these systems may prove highly costly and it is imperative that proper mechanisms are in place to detect them.

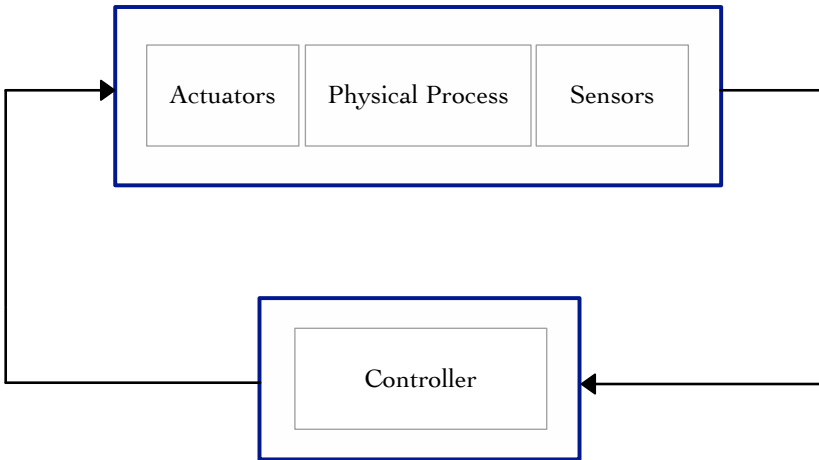


Figure 1.1: An abstract architecture of Industrial Control Systems.

1.4 Imminent Threats to Critical Infrastructures

Controllable “smart” devices are invading our societies. Through digitalization and inter-connectivity, the Industrial Internet of Things (IIoT) is transforming our critical infrastructure and reshaping the cyber landscape into one with much higher destructive potential [10]. The rapidly developing trend of digitalization and connectivity is posing imminent threats to critical infrastructure on which societies highly depend; including health care, transportation, manufacturing, and power distribution to name a few. The recent years have witnessed allegedly state-sponsored cyberattacks suggesting that cyber warfare is looming in the horizon [3, 4, 8, 9]. Unfortunately, the need for meeting efficiency requirements and enabling more controllability and interfacing with industrial assets is overshadowing the thought of resilient and sustainable modernized infrastructure, and cyber adversaries are becoming ever more capable in the process.

Being a double-edge sword, inasmuch as connectivity is increasing industrial efficiency, it is rendering critical systems increasingly vulnerable by introducing single points of failures into systems with high availability and reliability requirements. As a result, cyberattacks on critical infrastructure have potential to throw entire communities into chaos, causing large-scale and far-reaching

consequences on society at large. For instance, a cyberattack on a nation's power grid, which could be launched from anywhere in the world, has been shown capable of depriving thousands of households and facilities of electricity [20].

1.5 Intrusion Detection Systems for ICS

In response to the rising cyber threats in critical infrastructure, considerable effort has recently been devoted by the research community to investigating proper defensive measures. Designing intrusion detection systems suitable for cyber-physical environments has been at the forefront of this effort.

1.5.1 Motivation

It comes as no surprise that IDSs are considered as an important piece of the puzzle since one indispensable step in combating adversarial acts in ICS, or any information system for that matter, is in fact detecting the presence of the attacker.

Intrusion detection has its academic roots in the 1987 work by Denning [21] and has been extensively studied in the context of typical IT systems ever since. In a broad sense, intrusion detection is divided into two main categories: misuse detection and anomaly detection [22]. In misuse detection, traffic patterns that match with predefined so-called attack signatures are flagged as anomalous while all other traffic is considered normal. By contrast, anomaly detection involves creating a baseline from traffic data defining the normal behavior such that all other traffic that deviates from the baseline is considered anomalous.

Defining attack signatures in ICS environments is tricky due to the attacks on ICS being rare, specialized, and targeted at complex system components that are often legacy and proprietary. On the other hand, anomaly-based intrusion detection, although less favorable in IT environments due to intolerably high false-positive rates, proves more adequate for ICS environments by virtue of their highly regular behavior.

In industrial control systems, anomaly-based intrusion detection can be performed in two different layers: the IT layer consisting of supervisory and monitoring information systems, and the

physical layer where sensors, actuators, and controllers collectively control a physical process. Solely monitoring the IT layer for attack-indicating anomalies is problematic because of the fact that completely normal-looking traffic packets may still violate the semantics of the communication protocol at the payload level and deliver process data designed by the attacker to drive the process to an unsafe state [15]. Therefore, a holistic approach to detecting attacks on ICS requires a complementary more in-depth monitoring mechanism capable of detecting attacks *at the process level* as an advanced line of defense.

A process-level attack-detection mechanism monitors process data, such as sensor measurements and control commands, to detect misbehaviors in the physical process. The machine-to-machine communication in ICS process networks produces traffic that is highly deterministic, thereby enabling data-driven methods as a viable approach to attack-detection in these environments. The process-level attack-detection approach is particularly motivated by the *regularity* of ICS behavior, which is a distinctive feature that enables reliably constructing a baseline from historical process measurements and subsequently detecting deviant behavior due to anomalous operation.

1.5.2 Research Challenges

When it comes to equipping ICS with intrusion detection capabilities, there are several challenges to consider.

First, ICS are heterogeneous systems with complex architecture that often lack detailed specifications, hence creating a model of a system is in most cases a tremendous task that is hardly feasible in practice.

Second, process variables by their nature exhibit noisy behavior due to, e.g., vibrations in sensing machinery, which is likely to affect the detection accuracy of the employed detection mechanisms. Therefore, it is imperative that the detection mechanisms are by design insensitive to noise and that the assumption of noiseless communication is dropped when developing these mechanisms.

Third, the diversity of industrial control systems and their widely different application domains require the proposed attack-detection mechanisms to be application-agnostic and to operate independently of the underlying system specifications.

Finally, when it comes to deploying process-level IDS research prototypes in real environments, the fact that ICS incorporate both IT technology and operation technology (OT) poses a hurdle to designing and deploying security mechanisms that take both aspects into account. As OT-based systems are increasingly adopting IT-based solutions, the coordination between IT and OT personnel proves essential for achieving the intended outcome. Furthermore, the strict availability requirements of ICS necessitate the design of lightweight mechanisms to ensure secure real-time operation.

1.5.3 Existing Solutions & Limitations

Most existing approaches to detecting misbehaviors in the physical process propose the use of model-based techniques to model the normal behavior of the process and then detect deviations therefrom [13, 17, 23–25]. While such approaches might prove viable in some cases where a detailed and complete specification of the physical process is at hand, in the real world, it is often the case that the system to monitor is fairly complex and lacks a roadmap to creating a model of the controlled process. Thus, building a model of the physical process requires extensive human effort and domain knowledge, if at all possible [14].

Another disadvantage of model-based techniques lies in the fact that they involve solving a more general problem. Specifically, after presumably modeling the normal behavior of the process, the identified model is subsequently used to *predict the future behavior* of the underlying system, which is then compared to the observed behavior such that large deviations are labelled as potential attacks. Predicting the future behavior based on historical data is a more general problem than detecting misbehaviors, and it is known to be difficult and prone to inaccuracies due to noise in the data. Finally, since models are specific to the environments for which they were identified, model-based techniques prove difficult to generalize.

Approaches that use machine learning and data mining have been considered as well [12, 14, 26–30]. While machine learning methods do not require a model of the physical process, they involve a feature extraction and engineering phase, where *system-dependent* features need to be selected for training. Feature selection is tricky, hard to automate, and finding the best (most representative) features require a great deal of tuning and cross-validation. Moreover,

the fact that features are constructed by combining various process variables and then transformed into high-dimensional feature spaces makes it difficult to identify the whereabouts of the attack and affects the interpretability of the detection results.

1.5.4 Departure-Based Attack Detection

Our proposed solution is model-free, meaning that rather than creating a model to predict the future system behavior, PASAD directly compares the *current* behavior with the historical behavior of the process to detect changes in dynamics.

With PASAD, we introduce the notion of *departure-based* attack detection where a *departure* is a specific type of anomaly that refers to the process dynamics being forced to *depart* from the normal behavior due to potentially malicious structural changes in the stream of sensor measurements. The normal behavior is established in an offline training phase through a mathematical construction that enables the representation of the process dynamics in a noise-reduced geometric space. Thereafter, to detect a departure in the process behavior, PASAD iteratively computes a *departure score* during an online detection phase. Whenever this departure score crosses a predetermined threshold, an alarm is raised to the operators.

1.6 Research Questions

This thesis identifies and contributes to the following research questions.

- RQ1: How to monitor industrial control systems at the physical process level thereby promoting a suitable and tailored security framework?
- RQ2: How to achieve a detection approach that works across diverse and proprietary cyber-physical systems?
- RQ3: Are process-level solutions practical? To which extent are they applicable? What are the nuts and bolts of deploying such research techniques in practice?

1.7 Contributions

We contribute to research questions RQ1 and RQ2 in Chapter 2 (Paper A), where we introduce PASAD as a process-level attack-detection mechanism and discuss how being a model-free technique allows for handling diverse noisy ICS environments.

In Chapter 3 (Paper B), we address research question RQ3 by demonstrating the practicality of PASAD and its deployability in real industrial settings, where we describe our experience and lessons learned from running a live monitoring experiment in a real environment.

Following is a brief overview of the articles appended to this thesis.

1.7.1 Paper A — Truth Will Out: Departure-Based Process-Level Detection of Stealthy Attacks on Control Systems

As a contribution to research questions RQ1 and RQ2, a novel ICS-specific intrusion detection method (PASAD) is introduced. PASAD is an anomaly-based process-level intrusion detection system that monitors ICS process activity in real time to determine whether the system operation is normal or anomalous. Initially, PASAD learns the normal behavior recorded in a time series of sensor measurements through a training phase, during which ideas from a time-series analysis technique known as Singular Spectrum Analysis are applied to extract signal information from process output under normal conditions. Thereafter, the system continuously checks if incoming observations are departing from the normal behavior captured during the training phase.

PASAD is a theoretically sound, purely data-driven, lightweight, model-free mechanism that requires no prior knowledge of the system dynamics. Specifically, rather than creating a model of the physical process to predict future system behavior, PASAD seeks to solve the easier problem of deciding whether present sensor readings are departing from past readings due to a change in the mechanism generating them.

Furthermore, by virtue of its impressive noise-reduction capabilities, PASAD is capable of detecting *slight variations* in the sensor signal. This leads to the possibility of detecting strategic attackers

who may try to hide their stealthy attacks even at the process level.

Finally, we show that PASAD compares favourably with state-of-the-art data-driven techniques and we demonstrate its effectiveness using a simulation platform, data from a physical testbed, and data from a real system.

1.7.2 Paper B — The Nuts and Bolts of Deploying Process-Level IDS in Industrial Control Systems

The evaluation of ICS intrusion-detection methods in the literature seems to have been restricted to simulations and offline analysis of relevant datasets. In an attempt to bridge the existing simulation-based evaluation efforts with the real world by creating a roadmap characterizing potential hurdles to be expected when bringing the systems into a real environment, we take the evaluation of process-level monitoring a step further by running a fully fledged prototype in a real environment to examine the feasibility of the proposed methods in real-world settings. We contribute to research question RQ3 by building a complete system around PASAD, deploying a prototype in an operational paper factory, and describing our experience of running the prototype for 75 days. Finally, we highlight some technical challenges and practical aspects of live process-level monitoring for intrusions in ICS and then propose a set of guidelines and recommendations for both security researchers and practitioners who may consider designing or deploying IDS solutions for control systems.

1.8 Conclusions and Future Directions

Research on securing industrial control systems is gaining momentum in light of the expanding attack surface fuelled by the surging connectivity trend. Monitoring the process-level network for attack-induced changes in dynamics has proven sensible and suitable for these types of systems. This thesis highlighted the viability of this approach and introduced a novel data-driven attack-detection method that is efficient, theoretically sound, lightweight, and scalable. Furthermore, the practicality of the proposed system, and by extension process-level intrusion detection systems, was scrutinized

in real-world settings to investigate their applicability and better understand deployment challenges.

Future work is set to focus on exploring other application areas where PASAD is likely to succeed, such as the vehicular domain and power systems.

Bibliography

- [1] K. Stouffer, J. Falco, and K. Scarfone, “Guide to Industrial Control Systems (ICS) Security,” *NIST special publication*, 2011.
- [2] B. Gregory-Brown, “Securing Industrial Control Systems-2017,” *SANS Institute InfoSec Reading Room*, 2017.
- [3] N. Falliere, L. Murchu, and E. Chien, “W32. Stuxnet Dossier,” *White paper, Symantec Corp., Security Response*, 2011.
- [4] T. Chen and S. Abu-Nimeh, “Lessons from Stuxnet,” *Computer*, 2011.
- [5] R. Lee, M. Assante, and T. Conway, “German Steel Mill Cyber Attack,” *SANS Industrial Control Systems, Tech. Rep.*, 2014.
- [6] M. Abrams and J. Weiss, “Malicious Control System Cyber Security Attack Case Study—Maroochy Water Services, Australia,” *McLean, VA: The MITRE Corporation*, 2008.
- [7] B. Johnson, D. Caban, M. Krotofil, D. Scali, N. Brubaker, and C. Glycer. (2017) Attackers Deploy New ICS Attack Framework “TRITON” and Cause Operational Disruption to Critical Infrastructure. . Last visited 2018-08-01. [Online]. Available: <https://www.fireeye.com/blog/threat-research/2017/12/attackers-deploy-new-ics-attack-framework-triton.html>
- [8] R. Lee, M. Assante, and T. Conway, “Analysis of the Cyber Attack on the Ukrainian Power Grid,” *SANS Industrial Control Systems and E-ISAC, Tech. Rep.*, 2016.
- [9] O. Vukmanovic and S. Jewkes. (2017) Suspected Russia-Backed Hackers Target Baltic Energy Networks. . Last visited 2018-08-01. [Online]. Available: <http://mobile.reuters.com/article/idUSKBN1871W5>

-
- [10] M. Allen and C. Pisani. (2018) Hacking and Cyber Warfare are Top Humanitarian Concerns. . Last visited 2018-08-01. [Online]. Available: https://www.swissinfo.ch/eng/peter-maurer_hacking-and-cyber-warfare-are-top-humanitarian-concerns/43847744
- [11] D. Urbina, J. Giraldo, A. Cárdenas, 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*. ACM, 2016.
- [12] M. Krotofil, J. Larson, and D. Gollmann, “The Process Matters: Ensuring Data Veracity in Cyber-Physical Systems,” in *Proceedings of the 10th ACM Symposium on Information, Computer and Communications Security*, ser. ASIA CCS ’15. ACM, 2015.
- [13] A. Cárdenas, S. Amin, Z.-S. Lin, Y.-L. Huang, C.-Y. Huang, and S. Sastry, “Attacks Against Process Control Systems: Risk Assessment, Detection, and Response,” in *Proceedings of the 6th ACM Symposium on Information, Computer and Communications Security*. ACM, 2011.
- [14] I. Kiss, B. Genge, and P. Haller, “A Clustering-Based Approach to Detect Cyber Attacks in Process Control Systems,” in *Industrial Informatics (INDIN)*, 2015.
- [15] D. Hadžiosmanović, R. Sommer, E. Zambon, and P. H. Hartel, “Through the Eye of the PLC: Semantic Security Monitoring for Industrial Processes,” in *Proceedings of the 30th Annual Computer Security Applications Conference*. ACM, 2014.
- [16] A. Cárdenas, S. Amin, B. Sinopoli, A. Giani, A. Perrig, and S. Sastry, “Challenges for Securing Cyber Physical Systems,” in *Workshop on Future Directions in Cyber-Physical Systems Security*, 2009.
- [17] Y. Liu, P. Ning, and M. Reiter, “False Data Injection Attacks Against State Estimation in Electric Power Grids,” *ACM Transactions on Information and System Security (TISSEC)*, 2011.

- [18] W. Aoudi, M. Iturbe, and M. Almgren, “Truth will out: Departure-based process-level detection of stealthy attacks on control systems,” in *Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security*, ser. CCS ’18. New York, NY, USA: ACM, 2018. [Online]. Available: <http://doi.acm.org/10.1145/3243734.3243781>
- [19] M. Almgren, W. Aoudi, R. Gustafsson, R. Krahl, and A. Lindhé, “The nuts and bolts of deploying process-level ids in industrial control systems,” in *Proceedings of the 4th Annual Industrial Control System Security Workshop*, ser. ICSS ’18. New York, NY, USA: ACM, 2018, pp. 17–24. [Online]. Available: <http://doi.acm.org/10.1145/3295453.3295456>
- [20] P. Polityuk, O. Vukmanovic, and S. Jewkes. (2017) Ukraine’s Power Outage was a Cyber Attack: Ukrenergo. . Last visited 2018-08-01. [Online]. Available: <https://reut.rs/2mPSZqb>
- [21] D. E. Denning, “An intrusion-detection model,” *IEEE Transactions on Software Engineering*, vol. SE-13, no. 2, pp. 222–232, Feb 1987.
- [22] H. Debar, M. Dacier, and A. Wespi, “A revised taxonomy for intrusion-detection systems,” *Annales Des Télécommunications*, vol. 55, no. 7, pp. 361–378, Jul 2000. [Online]. Available: <https://doi.org/10.1007/BF02994844>
- [23] D. Urbina, J. Giraldo, A. Cárdenas, J. Valente, M. Faisal, N. O. Tippenhauer, J. Ruths, R. Candell, and H. Sandberg, “Survey and New Directions for Physics-Based Attack Detection in Control Systems,” National Institute of Standards and Technology, Tech. Rep., 2016.
- [24] A. Mathur and N. Tippenhauer, “SWaT: A Water Treatment Testbed for Research and Training on ICS Security,” in *2016 International Workshop on Cyber-physical Systems for Smart Water Networks (CySWater)*, 2016.
- [25] A. Kerns, D. Shepard, J. Bhatti, and T. Humphreys, “Unmanned Aircraft Capture and Control via GPS Spoofing,” *Journal of Field Robotics*, 2014.

-
- [26] C. Feng, T. Li, and D. Chana, "Multi-Level Anomaly Detection in Industrial Control Systems via Package Signatures and LSTM Networks," in *47th Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN)*. IEEE, 2017.
- [27] P. Nader, P. Honeine, and P. Beauseroy, "Lp-Norms in One-Class Classification for Intrusion Detection in SCADA Systems," *IEEE Transactions on Industrial Informatics*, 2014.
- [28] Y.-j. Xiao, W.-y. Xu, Z.-h. Jia, Z.-r. Ma, and D.-l. Qi, "NIPAD: A Non-Invasive Power-Based Anomaly Detection Scheme for Programmable Logic Controllers," *Frontiers of Information Technology & Electronic Engineering*, 2017.
- [29] K. N. Junejo 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*. ACM, 2016.
- [30] S. Pan, T. Morris, and U. Adhikari, "Developing a Hybrid Intrusion Detection System Using Data Mining for Power Systems," *IEEE Transactions on Smart Grid*, 2015.