A Framework for Determining Robust Context-Aware Attack-Detection Thresholds for Cyber-Physical Systems

Downloaded from: <https://research.chalmers.se>, 2023-04-26 12:27 UTC

Citation for the original published paper (version of record):

Aoudi, W., Almgren, M. (2021). A Framework for Determining Robust Context-Aware Attack-Detection Thresholds for Cyber-Physical Systems. ACM International Conference Proceeding Series.
<http://dx.doi.org/10.1145/3437378.3437393>

N.B. When citing this work, cite the original published paper.

A Framework for Determining Robust Context-Aware Attack-Detection Thresholds for Cyber-Physical Systems

WISSAM AOUDI, Chalmers University of Technology, Sweden

MAGNUS ALMGREN, Chalmers University of Technology, Sweden

Process-aware attack detection plays a key role in securing cyber-physical systems. A process-aware detection system (PADS) identifies a baseline behaviour of the physical process in cyber-physical systems and continuously attempts to detect deviations from the baseline attributed to malicious modifications in the process operation. Typically, a PADS triggers an alarm whenever the detection score crosses a fixed and predetermined threshold. In this paper, we argue that in the context of securing cyber-physical systems, relying on a single fixed threshold can undermine the effectiveness of the PADS, and propose a context-aware framework for determining two-dimensional thresholds that enhance the sensibility and reliability of such detection systems by rendering them more robust to false detection. In addition, we propose an algorithm, out of many possible, within this framework as a practical example.

CCS Concepts: • **Security and privacy** → **Intrusion detection systems**; *Systems security*.

Additional Key Words and Phrases: cyber-physical systems, attack detection, threshold, process-aware defense

ACM Reference Format:

Wissam Aoudi and Magnus Almgren. 2021. A Framework for Determining Robust Context-Aware Attack-Detection Thresholds for Cyber-Physical Systems. In *2021 Australasian Computer Science Week Multiconference (ACSW '21)*, February 1–5, 2021, Dunedin, New Zealand. ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/3437378.3437393>

1 INTRODUCTION

Cyber-physical systems (CPSs) are a growing type of systems that supervise, monitor, and control physical processes. CPSs are characterized by the interaction between cyber objects (e.g., computing and communication tools) and physical objects (e.g., sensors and actuators) to control a process, often autonomously or with minimal human intervention [15]. These systems are typically employed in manufacturing, transportation, water networks, energy management, and many other systems [8], where they often take part in executing safety-critical tasks [19].

With the advances in communication technology, embedded systems, and cloud computing, CPSs are increasingly connected to open networks to enable remote functionalities for operators and stakeholders. Consequently, these systems are becoming targets to malicious schemes [14], some of which are reportedly capable of crippling critical infrastructure [4]. In response to the arising threats, process-aware detection systems (PADSs) that leverage the deterministic nature of CPS dynamics have proven viable in detecting malicious manipulation at the process level by monitoring field devices and control networks [6, 13]. The basic idea behind process-level monitoring is that it is rather hard for malicious actors to subvert the target physical process without causing changes in the process dynamics that are detectable by process-aware defense mechanisms.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2021 Copyright held by the owner/author(s). Publication rights licensed to ACM.

Manuscript submitted to ACM

Process-aware attack-detection systems, also referred to as physics-based attack-detection (PBAD) systems in some contexts [17], consist of two phases: an offline identification phase and an online detection phase. In the identification phase, the PADS identifies a baseline behaviour that the CPS is expected to exhibit under normal operating conditions. Subsequently, in the detection phase, a detection metric measures the degree of conformity of current system observations to the identified baseline in an attempt to detect deviations presumably attributed to malicious modifications in the system operation.

Several approaches to identifying the baseline behaviour in CPSs have been proposed in the literature. Linear dynamic state-space models are used in [7, 12, 17, 18] to create a characterization of the process based on physical invariants; the said models are then used to predict the expected next state of the CPS. Hadžiosmanović et al. [11] propose the use of autoregression on historical sensor measurements to identify a predictive model to predict the expected next sensor measurement. Spectral decomposition of time series of sensor measurements are applied in [1–4] to capture the deterministic variability in the physical process using ideas from a time-series exploratory analysis technique known as singular spectrum analysis. Giraldo et al. [10] propose the use of estimation functions incorporating historical sensor measurements and control actions to estimate individual sensors in decentralized process-control environments.

Once the baseline behaviour of the process is identified, the online detection phase is initiated, wherein a detection metric is used to measure the extent to which the most recent observations of the process conform to the baseline. Different metrics are used (e.g., χ^2 statistic, non-parametric CUSUM statistic, Euclidean distance, etc.) to produce a real-valued detection score at every iteration during the online detection phase. In all of the mentioned methods, this detection score is constantly compared to a fixed threshold such that whenever the score crosses the threshold, the PADS issues an alert to the system operators indicating a potential attack on the CPS. Contrary to our contribution, the threshold is fixed and predetermined based on the relative behaviour of the detection score over a validation period.

In a CPS, different machines communicate in a controlled environment, therefore the process exhibits a regular and fairly deterministic behaviour [5]. However, as no PADS is perfect and no data source is flawless and fully representative, it is likely that the anomaly score would repeatedly cross the threshold moderately and for short time intervals during normal system operation, thereby triggering false alarms. Unlike other systems that may tolerate a certain rate of false alarms in favor of accurate and early detection of anomalies, false detection is highly costly in CPSs since they rely primarily on availability and continuity of operation. That is, the necessary avoidance of downtime makes it rather hard to convince system operators to discontinue the system operation and investigate every alert triggered by the PADS if many alerts are likely to be false. Yet, a true anomaly that goes undetected may prove even more costly for industrial stakeholders as it can potentially have a devastating effect on their systems.

In this paper, we argue that in a CPS context, although thresholds as decision boundaries are desirable since they are conceptually simple and computationally efficient, basing the decision of issuing alerts solely on a single fixed threshold can undermine the effectiveness of the PADS due to the high cost of false alarms. Therefore, in Section 2 we propose a framework for determining two-dimensional attack-detection thresholds for CPSs that are less susceptible to false detection, and we present an algorithm based on the proposed framework in Section 3, which we experimentally evaluate in Section 4, and discuss evaluation results in Section 5.

2 PROPOSED FRAMEWORK

Our proposed framework is context-aware and rests on two pillars: two-dimensional thresholds and actionability of alerts. The flowchart in Figure 1 describes the workings of the framework. Table 1 summarizes how our framework compares favourably with the single-threshold approach in terms of requirements.

Fig. 1. A flowchart describing the workings of the proposed framework.

Approach	Simple	Accurate	Efficient	Context-aware	Low FP
Single threshold	✓	✓	✓	✗	✗
Our framework	✓	✓	✓	✓	✓

Table 1. Comparison between classical single-threshold approach and our framework in terms of requirements.

2.1 Context-Awareness

It is important to realize that a PADS is essentially an anomaly detection system (ADS) that is tailored to the specific application as a security monitor for physical processes in cyber-physical systems. The distinction between a PADS and other ADSs, or between any two ADSs in different domains for that matter, resides in the individual characteristics and dynamics of the system they monitor, and by extension, in the difference in the behaviour of the detection scores and in handling alerts.

Since our framework is context-aware, it takes into account the expected or typical reaction of a CPS to malicious manipulations in its operation. Due to the sequential nature of its operation, in a truly anomalous situation, a CPS is expected to exhibit one of two typical behaviours: a *surge* behaviour or a *chronic* behaviour.

DEFINITION (SURGE BEHAVIOUR). *In a surge-behaviour scenario, the CPS dynamics gradually and increasingly depart from the baseline behaviour, in which case an efficient PADS would produce scores that gradually increase in magnitude until they are significantly higher than the norm.*

DEFINITION (CHRONIC BEHAVIOUR). *In a chronic-behaviour scenario, the CPS exhibits instability that is moderate in intensity but persistent, in which case an efficient PADS would produce scores that are moderately higher than the norm but persist above the fixed threshold for some time.*

The surge behaviour should be distinguished from drastic abrupt changes that can take place in other kinds of systems (e.g., a sudden increase in network traffic due to a flooding attack). Such a behaviour is atypical in the context of process-level monitoring in CPSs where changes usually build up cumulatively (e.g., gradually increasing water level to overflow a water tank in a water-purification process or manipulating process variables to increase pressure to dangerous levels in a chemical process) leading to a gradual increase in the detection score. The chronic behaviour is more likely in stealth-attack situations wherein attackers try to maintain their malicious manipulations within the valid range of process measurements.

2.2 Assumptions

Before we explain the concepts underlying our framework, we put forth the assumptions we make in this paper.

Assumption 1: The PADS we abstractly refer to throughout the paper is a sequential algorithm, or a set of algorithms, that initially identify a baseline behaviour of a CPS. In an online phase, the PADS then monitors the CPS by iteratively producing a real-valued detection score measuring how far the current system state is from the baseline. If the computed score meets certain conditions, the PADS raises an alert.

Assumption 2: The PADS we abstractly refer to throughout the paper is efficient in the sense that the identified baseline captures the normal dynamics of the monitored CPS fairly accurately and that the anomaly score reacts proportionately to the departure of the system state from the baseline.

Assumption 3: The data processed by the PADS is time-dependent raw process data, mainly sensor measurements, control commands, and actuator signals. Anomaly detection systems that process other types of CPS data, e.g. inter-arrival times of event-based packets [16], do not necessarily benefit from our framework.

2.3 Two-Dimensional Thresholds

Setting a fixed threshold may seem simplistic but it is in fact an intuitive and effective decision boundary, which is probably why it is a common practice. When the anomaly score computed by an efficient PADS maintains an average value below a certain level for a sufficiently long validation period, it is plausible to assume that the monitored CPS may be behaving abnormally when that level is crossed.

We acknowledge the simplicity and efficacy of using a fixed threshold but we argue that it is insufficient for making a reliable decision because it is unlikely that the baseline determined in the identification phase captured all the variability in the data. One could increase the threshold level (e.g., a multiple of standard deviations from the mean of the detection scores) to reduce the likelihood of false positives but only at the risk of increasing the chances of true anomalies going undetected (false negatives) and increasing the time-to-detection in the case of a successful detection of suspicious behaviour. Managing this trade-off between false positives, false negatives, and early detection proves challenging when there is only one variable to tune. We therefore propose adding a second dimension to threshold determination to make the detection more robust to false positives without sacrificing accuracy. The second dimension is simply an additional condition that the anomaly score is verified against once the first condition is met (crossing the fixed threshold) to decide on whether the detected change is a true or false anomaly. The property that the second condition is checked

Fig. 2. A depiction of the oblique threshold mechanism constructed dynamically and initiated by the attestation subroutine upon violation of the flat threshold.

only after the first condition has been met significantly limits the complexity added by the second dimension since, as we assumed the PADS is efficient (see Assumption 2, Section 2.2), the anomaly score is not expected to cross the fixed threshold all too often. Crossing the fixed threshold sets off the execution of what we refer to as the “attestation subroutine”, which verifies if the second condition is satisfied.

DEFINITION (ATTESTATION SUBROUTINE). *Upon crossing the fixed threshold, the mechanism for verifying the anomaly scores against the second condition consists of executing an attestation subroutine for a limited time window during which the algorithm attempts to attest that the detected change is indeed an anomaly.*

Whenever the first condition is met, for as long as the detection score is above the fixed threshold, the main purpose of the attestation process is to either substantiate or refute the suspicion of an anomalous behaviour while taking into account the expected behaviour of the detection-score series under different process conditions.

2.4 Actionability of Alerts

Although changes in the anomaly score that only meet the first condition (i.e., first-degree changes) are not considered as true alerts, they could provide a valuable insight into both the dynamics of the monitored CPS and the PADS itself. For instance, if these changes occur frequently, it could mean that not enough data were incorporated in the process

Algorithm 1 An example of a two-dimensional threshold algorithm.

Required: Anomaly scores s_t , flat threshold θ , slack variable ϵ , and parameters for the oblique threshold (α, δ) .

Outcome: Register a weak alert if the flat threshold is crossed and raise an actionable alert if, in addition, the oblique threshold is crossed.

```

1:  $\theta^+ \leftarrow \theta + \epsilon$ 
2:  $\theta^- \leftarrow \theta - \epsilon$ 
3: crossed  $\leftarrow$  FALSE
4: for  $t \leftarrow$  current_score_index do
5:   if  $s_t \geq \theta^+$  then
6:     if crossed == FALSE then
7:       Register weak alert
8:        $\tau \leftarrow t$ 
9:       crossed  $\leftarrow$  TRUE
10:    end if
11:     $\gamma_t = \theta - \frac{\alpha}{\delta} (t - \tau - \delta)$ 
12:    if  $s_t \geq \gamma_t$  then
13:      Raise actionable alert
14:    end if
15:  else if  $s_t < \theta^-$  & crossed == TRUE then
16:    crossed  $\leftarrow$  FALSE
17:  end if
18:   $t \leftarrow$  next_score_index
19: end for
```

of identifying the baseline behaviour, or that the dynamics of the CPS have shifted slightly due to, e.g., a change in configuration parameters, in which cases it may prove wise to re-evaluate the baseline. Therefore, in our proposed framework, the PADS issues two types of alerts: *weak* alerts and *actionable* alerts.

DEFINITION (WEAK ALERT). *Whenever the PADS detects first-degree changes (first condition is met) in the monitored process, it issues a “weak” alert, which is a non-actionable alert that is registered together with metadata describing the underlying event for later analysis.*

DEFINITION (ACTIONABLE ALERT). *Whenever the PADS detects second-degree changes (both conditions are met) in the monitored process, it issues an “actionable” alert that solicits the immediate attention of the system operators.*

3 A TWO-DIMENSIONAL THRESHOLD ALGORITHM

We now present an algorithm that implements the concepts underlying our proposed framework in Section 2. The algorithm implements two thresholds, raises both weak and actionable alerts according to which threshold is crossed, and the second threshold takes into account the typical score behaviour. We stress that the presented algorithm is just one example of algorithms that can fit our proposed framework.

To detect first-degree changes, the algorithm uses the classical flat threshold θ as an initial decision boundary for the anomaly score that is computed iteratively during the online phase using the detection metric. The second dimension consists of an oblique threshold γ_t describing a line with negative slope determined by two points, where t is the index of the anomaly score. The first point lies a vertical distance α from the flat threshold θ and the second point lies on the flat threshold and is a horizontal distance δ from the first point (see Figure 2), where α and δ are two parameters to be

Fig. 3. Slacked thresholds are used to reduce the sensitivity of the flat threshold θ towards minor and momentary breaches.

determined based on historical behaviour of the anomaly-score series. The oblique threshold is thus described by the line equation

$$\gamma_t = \theta - \frac{\alpha}{\delta} (t - \tau - \delta) \quad (1)$$

for $t \in [\tau, \tau + \delta]$, where τ is the index of the first anomaly score that crossed θ .

Hence, from the moment τ the anomaly score s_t crosses the flat threshold θ , a non-actionable alert is registered and the attestation subroutine execution starts. The attestation subroutine checks at every iteration $t = \tau, \tau + 1, \dots, \tau + \delta$ if the score s_t crosses the oblique threshold γ_t , such that if $s_t \geq \gamma_t$, an actionable alert is raised. The execution of the attestation subroutine ends at index $\tau + \delta$ or prematurely when the anomaly score falls back below θ .

A pseudocode of the algorithm is presented in Algorithm 1. The algorithm requires the anomaly scores s_t as soon as they are received sequentially from the process, the value of the flat threshold θ , a slack variable ϵ , and the parameters for the oblique threshold (α, δ) . As output, the algorithm ensures that a weak alert is registered together with relevant metadata as soon as the flat threshold is crossed and an actionable alert is raised if, in addition, the oblique threshold is crossed. The algorithm starts by defining two slacked thresholds θ^+ and θ^- (line 1–2) by adding and subtracting the slack variable ϵ from the flat threshold θ respectively. The purpose of the slacked thresholds (see Figure 3) is to reduce the sensitivity of θ towards minor and momentary breaches in order to avoid reporting useless alerts.

Fig. 4. Using data from the TE process and anomaly scores computed by PASAD, this figure demonstrates the workings of our two-dimensional threshold algorithm in a surge-behaviour case.

On line 5, the first condition is checked. If this condition is true, the attestation subroutine starts (line 6). The boolean variable `crossed` is used to ensure that the weak alert is registered once and not for every new score during the subroutine, and also to initialize τ to the first anomaly score index since it is used as a constant in the line equation on line 11. The second condition is checked on line 12. So long as this condition is true, the algorithm will keep raising actionable alerts. The attestation subroutine ends prematurely when the score falls back below the fixed threshold (line 15–17).

4 EVALUATION

We evaluate our two-dimensional threshold algorithm using a PADS proposed in [4] called PASAD, and using process data generated from the Tennessee-Eastman (TE) process control [9], which is a simulation model that simulates a real plant-wide chemical process. Both the PADS code and the TE process data used in the experiments are publicly available.¹

PASAD is a PADS that monitors sensor measurements from physical processes to detect cyberattacks on CPSs. In the identification phase, PASAD defines the baseline behaviour by identifying a low-dimensional signal subspace that

¹<https://github.com/mikeliturbe/pasad>

Fig. 5. Using data from the TE process and anomaly scores computed by PASAD, this figure demonstrates the workings of our two-dimensional threshold algorithm in a chronic-behaviour case.

captures the deterministic variability of the process recorded in time series of measurements. In the online detection phase, PASAD applies the Euclidean distance between current measurement vectors and a cluster of baseline vectors as a detection metric to measure the degree of departure of the system dynamics from the baseline behaviour, producing a detection score for each new sensor measurement.

Figures 4 and 5 demonstrate the workings of our proposed two-dimensional threshold algorithm in a surge-behaviour case and a chronic-behaviour case respectively using data from the TE process. As shown in the figures, a weak alert is generated every time the detection score crosses the flat threshold, and upon execution of the attestation subroutine, the oblique thresholds are constructed dynamically from preset parameters (θ, δ, γ) and the index τ of the first breach of the flat threshold by the anomaly score. As we stated in Section 3, the attestation subroutine ends prematurely as soon as the anomaly score falls back below θ (the second vertical line between the two thresholds). Note that the part of the oblique line segment with the faded colour is displayed in the figures for an illustration purpose only; effectively, the line no longer exists. At this point, the algorithm goes back to checking the first condition only and sets off the execution of the attestation subroutine with a new τ upon detection of another first-degree change. Finally, actionable alerts are raised when the oblique threshold is crossed during the attack period (shaded region).

5 DISCUSSION

As we stated in Section 1, the low tolerance to false alerts in CPS environments makes it imperative to reason about attack-detection thresholds beyond a single and static decision boundary. The concepts of two-dimensional thresholds and actionability of alerts that we introduced in our framework and their effects on reducing false positives are depicted in Figures 4 and 5. In these experiments, we used a real PADS to produce anomaly scores using data from the popular TE process. The results suggest that our practical approach reduces the susceptibility of PADS to false alerts without significantly increasing the complexity of the detection procedures.

The reason for choosing oblique lines as the second threshold dimension (instead of e.g., a second flat threshold) is to detect chronic anomaly score behaviours, where the score does not surge to very high values but persists above the first threshold for some time. Observe that the oblique threshold is a sequence of δ threshold levels γ_t linearly decreasing in value during the attestation subroutine as t varies from τ to $\tau + \delta$. The idea is that in a chronic-behaviour situation, although the score is not expected to cross the early higher levels, it will inevitably cross one of the lower levels later in the sequence as it persists above the first threshold (see Figure 5).

The classification of alerts into actionable and non-actionable alerts, where the latter are documented for offline analysis, is an important component of the proposed framework. As we mentioned in Section 2.4, the dynamics of the monitored system may shift slightly over time due to, for example, configuration updates or adding new components to the CPS. One implication of such a shift in dynamics, also known as the concept drift, is that the baseline behaviour identified by the PADS is no longer as representative as it was upon deployment, the fact that increases the likelihood and frequency of weak alerts. Offline analysis of the weak alerts may help identifying such trends early on and taking corrective measures accordingly.

6 CONCLUSION

PADSs, which are process-aware defense mechanisms that monitor physical processes, have proven to be viable for detecting cyberattacks on cyber-physical systems. In this paper, we argued that relying on single fixed thresholds for raising alerts can undermine the effectiveness of PADSs since the environments they are designed for have low tolerance to false alerts. Then, we presented a context-aware framework for constructing attack-detection thresholds for CPSs that are less susceptible to false detection, as well as a two-dimensional threshold algorithm based on the framework. Finally, we demonstrated the efficacy of our approach through experiments using a PADS called PASAD and publicly available data from the Tennessee-Eastman process.

ACKNOWLEDGMENTS

The research leading to these results has been partially supported by the Swedish Civil Contingencies Agency (MSB) through the projects “RICS” and “RIOT”, the Vinnova-funded project “KIDSAM”, and the European Community’s Horizon 2020 Framework Programme through the “UNITED-GRID” project under grant agreement 773717.

REFERENCES

- [1] Magnus Almgren, Wissam Aoudi, Robert Gustafsson, Robin Krah, and Andreas Lindhé. 2018. The Nuts and Bolts of Deploying Process-Level IDS in Industrial Control Systems. In *Proceedings of the 4th Annual Industrial Control System Security Workshop* (San Juan, PR, USA) (ICSSS '18). Association for Computing Machinery, New York, NY, USA, 17–24. <https://doi.org/10.1145/3295453.3295456>
- [2] Wissam Aoudi and Magnus Almgren. 2020. A Scalable Specification-Agnostic Multi-Sensor Anomaly Detection System for IIoT Environments. *International Journal of Critical Infrastructure Protection* 30 (2020), 100377. <https://doi.org/10.1016/j.ijcip.2020.100377>

- [3] Wissam Aoudi, Albin Hellqvist, Albert Overland, and Magnus Almgren. 2019. A Probe into Process-Level Attack Detection in Industrial Environments from a Side-Channel Perspective. In *Proceedings of the Fifth Annual Industrial Control System Security (ICSS) Workshop* (San Juan, PR, USA) (ICSS). Association for Computing Machinery, New York, NY, USA, 1–10. <https://doi.org/10.1145/3372318.3372320>
- [4] Wissam Aoudi, Mikel Iturbe, and Magnus Almgren. 2018. 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* (Toronto, Canada) (CCS '18). Association for Computing Machinery, New York, NY, USA, 817–831. <https://doi.org/10.1145/3243734.3243781>
- [5] Alvaro Cardenas, Saurabh Amin, Bruno Sinopoli, Annarita Giani, Adrian Perrig, Shankar Sastry, et al. 2009. Challenges for Securing Cyber Physical Systems. In *Workshop on future directions in cyber-physical systems security*, Vol. 5. Citeseer.
- [6] Alvaro Cardenas and Santa Cruz. 2019. Cyber-Physical Systems Security Knowledge Area. *The Cyber Security Body Of Knowledge (cybok)* (2019).
- [7] Alvaro A. Cárdenas, Saurabh Amin, Zong-Syun Lin, Yu-Lun Huang, Chi-Yen Huang, and Shankar Sastry. 2011. Attacks against Process Control Systems: Risk Assessment, Detection, and Response. In *Proceedings of the 6th ACM Symposium on Information, Computer and Communications Security* (Hong Kong, China) (ASIACCS '11). Association for Computing Machinery, New York, NY, USA, 355–366. <https://doi.org/10.1145/1966913.1966959>
- [8] Hong Chen. 2017. Applications of Cyber-Physical System: a Literature Review. *Journal of Industrial Integration and Management* 02, 03 (10 2017), 1750012. <https://doi.org/10.1142/S2424862217500129>
- [9] J.J. Downs and E.F. Vogel. 1993. A Plant-Wide Industrial Process Control Problem. *Computers & Chemical Engineering* 17, 3 (1993), 245 – 255. [https://doi.org/10.1016/0098-1354\(93\)80018-I](https://doi.org/10.1016/0098-1354(93)80018-I) Industrial challenge problems in process control.
- [10] Jairo Giraldo, David Urbina, CheeYee Tang, and Alvaro A. Cardenas. 2020. The More the Merrier: Adding Hidden Measurements to Secure Industrial Control Systems. In *Proceedings of the 7th Symposium on Hot Topics in the Science of Security* (Lawrence, Kansas) (HotSoS '20). Association for Computing Machinery, New York, NY, USA, Article 3, 10 pages. <https://doi.org/10.1145/3384217.3385624>
- [11] Dina Hadžiosmanović, Robin Sommer, Emmanuele Zamboni, and Pieter H. Hartel. 2014. Through the Eye of the PLC: Semantic Security Monitoring for Industrial Processes. In *Proceedings of the 30th Annual Computer Security Applications Conference* (New Orleans, Louisiana, USA) (ACSAC '14). Association for Computing Machinery, New York, NY, USA, 126–135. <https://doi.org/10.1145/2664243.2664277>
- [12] Y. Hu, Hong Li, H. Yang, Yuyan Sun, L. Sun, and Z. Wang. 2019. Detecting Stealthy Attacks Against Industrial Control Systems Based on Residual Skewness Analysis. *EURASIP Journal on Wireless Communications and Networking* 2019 (2019), 1–14.
- [13] F. Khorrami, P. Krishnamurthy, and R. Karri. 2016. Cybersecurity for Control Systems: A Process-Aware Perspective. *IEEE Design Test* 33, 5 (2016), 75–83. <https://doi.org/10.1109/MDAT.2016.2594178>
- [14] Marina Krotofil, Klaus Kursawe, and Dieter Gollmann. 2019. *Securing Industrial Control Systems*. Springer International Publishing, Cham, 3–27. https://doi.org/10.1007/978-3-030-12330-7_1
- [15] Marina Krotofil, Jason Larsen, and Dieter Gollmann. 2015. The Process Matters: Ensuring Data Veracity in Cyber-Physical Systems. In *Proceedings of the 10th ACM Symposium on Information, Computer and Communications Security* (Singapore, Republic of Singapore) (ASIA CCS '15). Association for Computing Machinery, New York, NY, USA, 133–144. <https://doi.org/10.1145/2714576.2714599>
- [16] Chih-Yuan Lin and Simin Nadjm-Tehrani. 2019. Timing Patterns and Correlations in Spontaneous SCADA Traffic for Anomaly Detection. In *22nd International Symposium on Research in Attacks, Intrusions and Defenses (RAID 2019)*. USENIX Association, Chaoyang District, Beijing, 73–88. <https://www.usenix.org/conference/raid2019/presentation/lin>
- [17] Raul Quinonez, Jairo Giraldo, Luis Salazar, Erick Bauman, Alvaro Cardenas, and Zhiqiang Lin. 2020. SAVIOR: Securing Autonomous Vehicles with Robust Physical Invariants. In *29th USENIX Security Symposium (USENIX Security 20)*. USENIX Association, 895–912. <https://www.usenix.org/conference/usenixsecurity20/presentation/quinonez>
- [18] David I. Urbina, Jairo A. Giraldo, Alvaro A. Cardenas, Nils Ole Tippenhauer, Junia Valente, Mustafa Faisal, Justin Ruths, Richard Candell, and Henrik Sandberg. 2016. Limiting the Impact of Stealthy Attacks on Industrial Control Systems. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security* (Vienna, Austria) (CCS '16). Association for Computing Machinery, New York, NY, USA, 1092–1105. <https://doi.org/10.1145/2976749.2978388>
- [19] J. Wu, S. Luo, S. Wang, and H. Wang. 2019. NLES: A Novel Lifetime Extension Scheme for Safety-Critical Cyber-Physical Systems Using SDN and NFV. *IEEE Internet of Things Journal* 6, 2 (2019), 2463–2475. <https://doi.org/10.1109/JIOT.2018.2870294>