On the Challenges of Holistic Intrusion Detection in ICS

Stefan Lenz[1], Julia Raab[1], Benedikt Holzbach[1], Deniz Köller[1], Sotiris Michaelides[1], Martin Henze[1,2]
[1]Security and Privacy in Industrial Cooperation, RWTH Aachen University • [2]Cyber Analysis & Defense, Fraunhofer FKIE • {lenz, michaelides, henze}@spice.rwth-aachen.de, {julia.raab, benedikt.holzbach, deniz.koeller}@rwth-aachen.de

Abstract. Past attacks against industrial control systems (ICS) show that adversaries often target both the ICS network *and* the physical process to achieve potential catastrophic impact. To secure ICS, intrusion detection systems promise timely uncovering of such adversaries. However, as these detection mechanisms typically focus on isolated characteristics of ICS (e.g., packet timings), multiple detection systems have to be deployed in parallel, complicating their operation in practice. In this work, to spur discussion and further research, we present challenges encountered during our research towards a holistic intrusion detection system aiming to cover all dimensions of an ICS.

1 Holistic Intrusion Detection in ICS

By connecting industrial control systems (ICS) to the Internet, these safety critical systems are exposed to sophisticated attackers [3, 6, 7]. As ICS consist of low-resource legacy devices that are difficult to replace and not capable of complex security operations, retro-fitting attack detection into ICS has been a major focus of recent research [8, 19]. Since the behavior of ICS is tightly coupled with a deterministic physical process, typical industrial intrusion detection systems (IIDS) model benign behavior (e.g., packet timings [10] or process states [17]) to later identify anomalies. However, as attacks materialize in the network *and* the physical dimension of ICS, comprehensive detection must cover all aspects of ICS [20]. Thus, we argue that IIDS should aim for a holistic detection approach, covering multiple dimensions simultaneously as depicted in Fig. 1.

One method to achieve holistic detection is to use multiple one-dimensional detectors in parallel [20]. Although such ensembles promise great detection performance *in theory*, they are difficult to deploy, since each detector requires sufficient training data as well parameter optimization. Furthermore, individual detectors risk to miss interplay between different dimensions and it is unclear how to combine alerts of individual detectors in a sensible manner [20].

To overcome these challenges, we set out to address the need for a *practical* holistic detection method. Instead of combining different detectors specializing in individual aspects, we strive to cover the complete benign behavior of an ICS in one holistic model. During this pursuit, we encountered several challenges while exploring different approaches. In this paper, we discuss these challenges regarding process state discretization (§2), IIDS parameterization (§3), and collecting sufficiently good training data (§4).

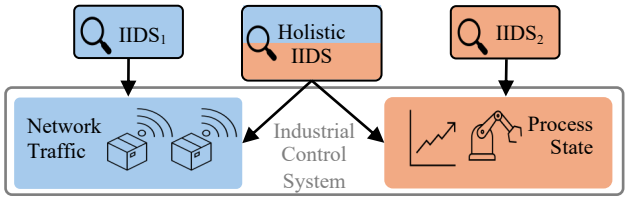

Figure 1: Holistic industrial intrusion detection systems (IIDS) aim to monitor both network behavior and physical process state, while traditional IIDS focus on one.

2 Challenge 1: Discretization

To create a holistic IIDS capturing the complete ICS in a single model, we explored the use of *process mining*, which has shown promising results in modeling physical process state [2, 12] as well as network traffic [15]. Indeed, initial experiments showed that such a detector can detect attacks targeting both these dimensions, while performing *on par* with other IIDSs that capture only one dimension. However, as process mining relies on discrete labels (e.g., bins) to incorporate the physical state into the model, the fundamental *challenge of discretization* arises to assign such labels to continuous ICS data (e.g., temperatures or fill levels).

Although many discretization mechanisms, such as statistical methods or unsupervised learning approaches (e.g., clustering), solve this problem, the quality of the resulting labels/clusters is highly dependent on the distribution of the underlying data [1, 4, 16]. To choose the optimal discretization method for an IIDS using process mining, we compared multiple discretization methods based on their resulting *detection performance* as summarized in Tb. 1. Our experiments show that, even while utilizing benchmarking datasets (with posterior knowledge that is not achievable in any practical scenario), this task is not trivial, as the best discretization method (green in Tb. 1) depends on the considered performance metric.

Take Away While discretization mechanisms promise well-defined labels for process data, in the context of intrusion detection, identifying the "best" approach depends on the considered detection metric.

Name	Metric								
	Accuracy	Precision	Recall	F0.1	F1	eTaP	eTaR	eTaF0.1	eTaF1
Jenkspy	0.68	0.23	0.15	0.23	0.21	0.23	0.52	0.24	0.32
kMeans	0.68	0.24	0.16	0.23	0.21	0.24	0.52	0.24	0.32
Quantiles	0.77	0.71	0.00	0.03	0.00	0.69	0.14	0.66	0.23

Table 1: The best algorithm (green) to discretize process data not only depends on data distribution but also on the detection performance metric.

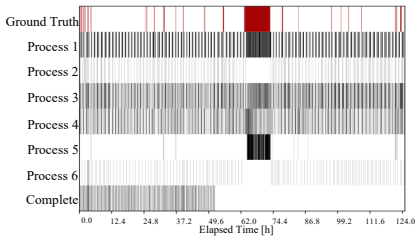

Figure 2: Alerts (black) of LLM detection on SWaT (red) are sensitive to process data variance. Monitoring subprocesses reduces processing overhead.

Medium	Inter-Arrival Times		
	L1 [ms]	L2 [ms]	L3 [ms]
Wired	46.06 (0.37)	46.06 (0.37)	100.00 (0.02)
Wireless^{+}	46.06 (0.37)	46.06 (0.37)	100 (0.12)
Wireless^{-}	45.97 (1.47)	45.97 (1.56)	100 (0.67)

Table 2: Differences in inter-arrival times in wired and good ($^{+}$) / bad ($^{-}$) wireless channels show the need for detection to adapt to dynamic behavior.

3 Challenge 2: Parameterization

A holistic IIDS, e.g., based on process mining, requires multiple additional parameters (e.g., length of the ICS cycle in packets) that must be set manually, resulting in a need for hard-to-come-by expert knowledge. To address the resulting *challenge of parameterization*, we explored different "closed-box" modeling approaches (e.g., based on LLMs), which promise no parameterization at all. The goal of such approaches is to gain "adequate" detection performance while utilizing the reasoning capabilities of LLMs. Exemplary experiments on the SWaT [5] dataset (cf. Fig. 2) indicate that an LLM-based IIDS struggles to monitor the complete process in a holistic manner, due to excessive resource demands, even exceeded the capabilities of the high-performance cluster node (96GB VRAM) we utilized.

Therefore, we focused on individual subprocesses in further experiments. While the smaller scope reduces computational complexity, the detector provided nearly-constant alerts, rendering attacks unrecognizable (except for Process 5 which has comparatively low variance). Still, due to its closed-box nature, it is impossible to comprehend the detector's reasoning behind alerts [14], making discerning between true and false alarm challenging. Together with the high resource demands, these issue prevent a practical application of a parameterless approach as a holistic IIDS.

Take Away Although addressing the challenge of parameterization, monolithic closed-box detectors face challenges in deployment, resource usage, and interpretability, preventing their use for holistic IIDS.

4 Challenge 3: Good Training Data

To meet increasing demands for flexibility and efficiency, ICS adopt new technologies such as wireless communication [11]. These more dynamic systems exaggerate the *challenge of gathering good training data* in IIDS research [18].

To investigate this issue, we built an ICS simulation enabling us to change the communication medium without altering the process behavior [9]. There, we analyze a popular timing-based IIDS [10], which uses deterministic inter-arrival timings (i.e., the time between two packets) to discern between benign and anomalous behavior for three communication scenarios: *Wired* (as a baseline), *Wireless*$^{+}$ (a wireless channel in perfect condition), and *Wireless*$^{-}$ (a wireless channel experiencing substantial disturbance).

Tb. 2 shows the mean (μ) inter-arrival times and standard deviation (σ) for three links from our simulated ICS. First, these results show that a wireless channel in perfect conditions can achieve the same link quality as the wired medium, although the variance in some links may be higher (Tb. 2 L3). The wireless medium under disturbance, however, cannot provide such link quality. Although the mean inter-arrival time is similar or even lower than in the other scenarios, σ is substantially higher for all links. This behavioral change can trigger false alerts in timing-based IIDS, which cannot distinguish a distressed channel from an attack.

As the quality of a wireless channel can change anytime (e.g, due to noise), the model of the IIDS might not reflect the actual (communication) behavior of the ICS, thus either resulting in false alarms or missed attacks. Additionally, current research [13] also suggests that ICS behavior might not be as deterministic as often assumed. Thus, IIDS research should also focus on mechanisms that can cope with these dynamic environments, especially for holistic IIDS.

Take Away Even though the collection of suitable training data is a known challenge for intrusion detection, modern, e.g., wireless, ICS with more dynamic behavior further complicate this challenge.

5 Conclusion

To secure ICS against multi-faceted threats targeting both physical process *and* communication, we pursue a holistic detection approach which covers both dimensions of an ICS simultaneously. In this paper, we report on our research process and discuss challenges—regarding process state discretization (§2), parameterization (§3), and training data (§4)– we faced to push such a holistic IIDS towards practical viability. Ultimately, our findings identify ample research opportunities, especially highlighting opportunities in more "dynamic" detection mechanisms capable of handling characteristics of future ICS.

Acknowledgements The research underlying this publication has in parts been funded by the German Federal Ministry of Research, Technology and Space (BMFTR) under funding reference number 16KIS2409K (6GEM+). Computations were performed with computing resources granted by RWTH Aachen University under project thes2001. The authors are responsible for the content of this publication.

References

[1] Wil van der Aalst. "Academic View: Development of the Process Mining Discipline". In: *Process Mining in Action: Principles, Use Cases and Outlook*. 2020.

[2] Filipe Apolinário et al. "FingerCI: Writing industrial process specifications from network traffic". In: *International Journal of Critical Infrastructure Protection* 47 (2024).

[3] Lennart Bader et al. "Comprehensively Analyzing the Impact of Cyberattacks on Power Grids". In: *IEEE EuroS&P*. 2023.

[4] Chris Fraley and Adrian E. Raftery. "How Many Clusters? Which Clustering Method? Answers via Model-Based Cluster Analysis". In: *The Computer Journal* 41.8 (1998).

[5] Jonathan Goh et al. "A Dataset to Support Research in the Design of Secure Water Treatment Systems". In: *CRITIS*. 2017.

[6] Kevin E. Hemsley and Dr. Ronald E. Fisher. *History of Industrial Control System Cyber Incidents*. Tech. rep. Idaho National Laboratory, USA, 2018.

[7] Eric D. Knapp. "Industrial Cybersecurity History and Trends". In: *Industrial Network Security*. 2024.

[8] Olav Lamberts et al. "SoK: Evaluations in Industrial Intrusion Detection Research". In: *Journal of Systems Research* 3.1 (2023).

[9] Stefan Lenz et al. "Security Implications of 5G Communication in Industrial Systems". In: *Proceedings of the 11th ACM Cyber-Physical System Security Workshop (CPSS)*. 2026. DOI: 10.1145/3775042.3807886.

[10] Chih-Yuan Lin, Simin Nadjm-Tehrani, and Mikael Asplund. "Timing-Based Anomaly Detection in SCADA Networks". In: *CRITIS '21*. 2018.

[11] Sotiris Michaelides et al. "Secure Integration of 5G in Industrial Networks: State of the Art, Challenges and Opportunities". In: *Future Generation Computer Systems* 166 (2025).

[12] David Myers et al. "Anomaly Detection for Industrial Control Systems using Process Mining". In: *Computers & Security* 78 (2018).

[13] Franka Schuster and Hartmut König. "Questioning the Myth: Investigating ICS Traffic Homogeneity from an Anomaly Detection Perspective". In: *Critical Information Infrastructures Security*. Vol. 15549. 2025.

[14] Heinrihs Kristians Skrodelis and Andrejs Romanovs. "Explainable hybrid intrusion detection for SCADA/ICS: a review and research agenda". In: *Frontiers in Computer Science* Volume 8 (2026).

[15] Christian Wakup and Jörg Desel. "Analyzing a TCP/IP-Protocol with Process Mining Techniques". In: *Business Process Management Workshops*. 2015.

[16] Marc Wegmann et al. "A review of systematic selection of clustering algorithms and their evaluation". In: *CoRR* abs/2106.12792 (2021).

[17] Konrad Wolsing et al. "Can Industrial Intrusion Detection Be SIMPLE?" In: *ESORICS '22*. 2022.

[18] Konrad Wolsing et al. "Deployment Challenges of Industrial Intrusion Detection Systems". In: *Computer Security. ESORICS '24 International Workshops*. 2024.

[19] Konrad Wolsing et al. "IPAL: Breaking up Silos of Protocol-dependent and Domain-specific Industrial Intrusion Detection Systems". In: *RAID '22*. 2022.

[20] Konrad Wolsing et al. "One IDS is not enough! Exploring Ensemble Learning for Industrial Intrusion Detection". In: *ESORICS '23*. 2023.