

IPAL: Breaking up Silos of Protocol-dependent and Domain-specific Industrial Intrusion Detection Systems

Konrad Wolsing
Fraunhofer FKIE
RWTH Aachen University
konrad.wolsing@fkie.fraunhofer.de

Antoine Saillard
RWTH Aachen University
Fraunhofer FKIE
antoine.saillard@rwth-aachen.de

Eric Wagner
Fraunhofer FKIE
RWTH Aachen University
eric.wagner@fkie.fraunhofer.de

Martin Henze
RWTH Aachen University
Fraunhofer FKIE
henze@cs.rwth-aachen.de

ABSTRACT

The increasing interconnection of industrial networks exposes them to an ever-growing risk of cyber attacks. To reveal such attacks early and prevent any damage, industrial intrusion detection searches for anomalies in otherwise predictable communication or process behavior. However, current efforts mostly focus on specific domains and protocols, leading to a research landscape broken up into isolated silos. Thus, existing approaches cannot be applied to other industries that would equally benefit from powerful detection. To better understand this issue, we survey 53 detection systems and find no fundamental reason for their narrow focus. Although they are often coupled to specific industrial protocols in practice, many approaches could generalize to new industrial scenarios in theory. To unlock this potential, we propose IPAL, our industrial protocol abstraction layer, to decouple intrusion detection from domain-specific industrial protocols. After proving IPAL's correctness in a reproducibility study of related work, we showcase its unique benefits by studying the generalizability of existing approaches to new datasets and conclude that they are indeed not restricted to specific domains or protocols and can perform outside their restricted silos.

CCS CONCEPTS

• **Security and privacy** → **Intrusion detection systems**; • **Networks** → **Cyber-physical networks**; *Network monitoring*.

KEYWORDS

Industrial Intrusion Detection, IDS, CPS, ICS, Industrial Protocols

ACM Reference Format:

Konrad Wolsing, Eric Wagner, Antoine Saillard, and Martin Henze. 2022. IPAL: Breaking up Silos of Protocol-dependent and Domain-specific Industrial Intrusion Detection Systems. In *25th International Symposium on Research in Attacks, Intrusions and Defenses (RAID 2022)*, October 26–28, 2022, Limassol, Cyprus. ACM, New York, NY, USA, 16 pages. <https://doi.org/10.1145/3545948.3545968>



This work is licensed under a Creative Commons Attribution International 4.0 License.

RAID 2022, October 26–28, 2022, Limassol, Cyprus
© 2022 Copyright held by the owner/author(s).
ACM ISBN 978-1-4503-9704-9/22/10.
<https://doi.org/10.1145/3545948.3545968>

1 INTRODUCTION

Digitized industries play an essential role in today's society across various domains such as water treatment and distribution, power delivery, chemical processing, or manufacturing [59]. To realize their functionality, industries rely on (near) real-time communication of process values and commands using a large pool of specialized industrial protocols, such as Modbus or EtherNet/IP [57]. The increasing reliance of these protocols on the Internet Protocol (IP) to foster automation, remote control, and optimized processes, however, also moves traditionally air-gapped industrial networks closer to cybersecurity threats [59]. Indeed, a growing number of cyber incidents with detrimental environmental damage and risk to human life [55, 116] highlight the importance and urgency of adequately protecting industrial communication networks.

As a non-intrusive, retrofittable, and cheaply deployable security solution, Intrusion Detection Systems (IDSs) offer a great additional layer of defense that is already well-established in traditional environments, such as office or data-center networks. While these traditional IDSs (e.g., Zeek [94]) offer a certain benefit to industrial networks, they also reach their limitations: Stealthy attacks differing just enough from normal operations to cause serious harm are hardly detectable by them [114]. Meanwhile, industrial protocol's rather deterministic and predictable nature offers unique opportunities for *process-aware* or *semantic* intrusion detection [49]. Hence, Industrial IDSs (IIDSs) leveraging these regularities are necessary to protect industrial networks. Consequently, a large research community has gathered [34, 49, 58, 65, 84, 90, 98, 114], proposing specialized solutions for a wide variety of industrial scenarios.

However, although different industrial domains exhibit similar predictable communication and process patterns, research on IIDSs is highly tailored to specific domains and communication protocols. While technical and engineering reasons for such tight coupling might exist, we find this strong interdependence rather surprising and contrary to intuition. From an IIDS perspective, all industrial protocols essentially exhibit the same distinct characteristics: They primarily exchange sensed data and commands using a small set of well-defined communication patterns [121]. Additionally, while specifics of benign and malicious behavior change between scenarios, these can mostly be trained prior to deployment.

Although there exists no fundamental conceptual reason for the dependence of IIDSs on specific industrial protocols, IIDSs still specialize to few protocols, limiting their applicability to single

scenarios. Moreover, besides contrary claims [33, 46, 61, 68], it is seldom tested whether a specific IIDS even works outside the precise scenario and industrial protocol it was developed for. Consequently, the overall progress in research on IIDSs is slowed down due to niche solutions that do not improve on prior work from other industrial fields, wasting tremendous advancements in detecting cyber attacks on industrial systems.

Contributions. To break up the prevalent research silos of protocol-dependent and domain-specific IIDSs and ultimately transfer the advancements of IIDS solutions to new industrial scenarios, we make the following contributions in this paper:

- Beginning with a survey of 53 IIDSs from related work, we identify the characteristics of industrial communication they rely on to detect anomalies and attacks. Thus, we *theoretically* identify the potential for protocol-independent and domain-agnostic intrusion detection across existing IIDSs (Sec. 3).
- To turn these results into *practice*, we introduce the IPAL industrial protocol abstraction layer. IPAL sits between industrial protocols and IIDSs, captures all aspects relevant for intrusion detection in a unified representation, and thus facilitates protocol-independent and widely applicable IIDSs (Sec. 4).
- We showcase the practical *applicability and correctness* of IPAL in a reproducibility study of eight IIDSs from related work. We show that IPAL provides all required information, and we simultaneously reproduce scientific results as an independent party, an important but often neglected step in research (Sec. 5).
- We illustrate the various *benefits for IIDS research* of IPAL in three case studies. With IPAL, we are able to compare different IIDSs, even from opposing research branches, study to what extent they generalize to new scenarios, and, for the first time, transfer the advancements of IIDSs to new protocols and domains (Sec. 6).

Availability Statement. To let IPAL become the foundation of future work on industrial intrusion detection, we make our source code available. This includes tools to transcribe industrial protocols [43] and existing datasets into IPAL [41], as well as IIDSs implementations from our reproducibility study [42] (to the extent permitted by licenses or authors).

2 INDUSTRIAL INTRUSION DETECTION

As a foundation for our work, we begin by recapitulating the core idea behind intrusion detection and what makes it special in industrial networks. Based on this, we identify key limitations of industrial intrusion detection research that prevent the broad application of improvements achieved for isolated industrial scenarios.

2.1 Intrusion Detection

Intrusion (or anomaly) detection is the art of automatically uncovering cyber attacks or other suspicious activity by passively monitoring a system’s behavior, e.g., w.r.t. communication [110, 119]. The core assumption behind Intrusion Detection Systems (IDSs), which implement intrusion detection, is that cyber attacks lead to distinctively different (anomalous) system behavior than benign activity. As their passive nature allows retrofitting them to existing deployments easily, IDSs are popular to complement preventive security measures, i.e., integrity or authentication.

Traditionally, IDSs were designed for office, server, or data-center networks with largely fluctuating traffic and usage patterns [108]. Even though this makes it challenging to define benign activity comprehensively, there is a decent understanding of how individual attacks are performed on these networks, as attacks and attack patterns are typically used across many different networks, e.g., by widely-used malware. Thus, standard methods to reliably detect such (known) attacks are signature or rule-based approaches, which compare user and network behavior against a list of attack indicators, with well-established implementations such as YARA [117], Suricata [44], Zeek [94] (previously Bro [106]), and Snort [100].

2.2 Promises of Industrial Intrusion Detection


Intrusion detection is particularly attractive for *industrial* networks, as retrofitting other (active) security measures often results in costly hard- or software modifications, major downtime, or the need to redo acceptance tests. W.r.t. attack and system behavior, however, we observe exactly opposite conditions for intrusion detection in industrial networks: While rather deterministic industrial tasks lead to more predictable communication patterns [57], attacks are often more subtle and scenario-specific, severely limiting predictability and thus signature or rule-based approaches.

Most importantly, attacks against industrial networks are often specifically designed for a single or few targets and hardly generalize across domains (as, e.g., malware does), making it hard to catch them by predefined rules. Furthermore, industrial settings are susceptible to subtle attacks [114, 126] where attackers send legitimate traffic (as allowed by rules) but at the “wrong” time to trip the system into a dangerous state. Thus, IDSs designed for typical IT networks cannot simply be transferred to industrial settings [128].

Still, industrial scenarios provide unique opportunities to detect intrusion and anomalies through different means, e.g., the tight coupling of industrial processes such as water treatment, power delivery, chemical processing, or manufacturing, with the physical environment via sensors and actuators. To this end, *process-aware* or *semantic* IIDSs incorporate information from the physical and/or communication level to uncover highly specialized and subtle attacks and anomalies. Therefore, IIDS approaches typically “learn” a model of benign behavior for a specific facility and alert deviations from this model. Such models may incorporate additional system knowledge, e.g., that certain actuator combinations never occur, resulting in an IIDS not requiring individual attack knowledge.

2.3 Limitations of Current IIDS Research

Given these promises of IIDSs to detect sophisticated and safety-critical attacks on industrial networks and processes, a plethora of different research fields for IIDSs in various industrial domains has been established [49, 90, 114], including industrial control systems [34, 58, 65], intra-vehicular communication [84], or water treatment [98]. Across these various research fields, we find that individual IIDSs are tailored to specific domains and communication protocols. Consequently, advancements made for one scenario often cannot generalize to other industrial protocols or domains, limiting their widespread use, slowing down innovation, and the overall progress in securing industrial networks. We identify two core problems that lead to these limitations of current IIDS research.



(a) The current state-of-the-art requires a tailored implementation of IIDSs for each protocol. (b) Protocol-independent IIDSs would allow separating protocol parsing and IDS development.

Figure 1: To overcome the limitations of current IIDS research, protocol-independent IIDSs promise to break up the strong interdependence between IIDS approaches and the specific industrial protocol and domain they operate in.

First, compared to traditional computer networks, we observe a huge *heterogeneity* in industrial domains, each with unique and custom communication protocols for legacy reasons. Thus, current IIDSs generally focus on one specific domain and industrial protocol combination, presenting yet another (novel) approach in isolation without building on top of prior research findings. E.g., two similar approaches independently propose to use Probabilistic Suffix Trees for Modbus [123] and IEC-104 [79], highlighting that work is done twice across different industrial scenarios.

Second, IIDS research across different industrial domains suffers from an *evaluation bias*, where vast inconsistencies in used evaluation methodologies make comparisons across publications hardly possible [49, 78, 114]. Moreover, the use of only a few datasets (e.g., SWaT [50]) or many private ones steers the collective research efforts in a direction where solutions are optimized and biased towards specific scenarios. This problem is already known for traditional IDSs [92] but is even more severe in industrial settings, where systems expose very narrow yet domain-specific behavior [111].

Consequently, although different approaches claim broader applicability (without proof) [33, 46, 61, 68], the actual generalizability of proposed IIDSs to new settings within the same domain or across domains remains mostly unexplored [36, 78]. Thus, the various research streams in the area of IIDSs are indeed disjoint. As a result, the overall research field is slowed down, as new approaches are designed for niches without improving on previous research and are biased to datasets from specific scenarios. It thus remains unknown to what extent the tremendous advancements in the detection of cyber attacks on industrial systems can or cannot be generalized and transferred from their isolated niches to a broader scale and thus sustainably increase security for many industrial networks.

3 THE CASE FOR PROTOCOL-INDEPENDENT INDUSTRIAL INTRUSION DETECTION

Currently, the IIDS research landscape is scattered across heterogeneous industrial domains (cf. Sec. 2.3), leading to a situation as depicted in Fig. 1a, where IIDSs are developed in isolation for specific industrial protocols and domains. Consequently, research on IIDSs is limited w.r.t. generalizability, evaluation scenarios, widespread use, and coherence between research streams.

To overcome these limitations, in this paper, we make a case for *protocol-independent IIDSs*, as shown in Fig. 1b, where IIDSs do not

have to care about the underlying protocol anymore while simultaneously being transferable to various scenarios with little to no additional effort. Our proposal for protocol-independence in IIDSs is motivated by two observations: First, for the purpose of intrusion detection, essentially all industrial protocols show similar functionality, i.e., the exchange of sensor readings and actuator commands. Second, in principle, the underlying detection approaches of IIDSs work similarly as they each know or learn a model of a specific industrial process’s behavior and detect anomalies by comparing their models to the monitored physical process.

Given these observations, we postulate that a large portion of existing work is not necessarily bound to the specifics of the industrial domain and protocol they are designed for and may generalize well to other domains and protocols. Thus, we are convinced that tackling the diversity in industrial protocols is key to breaking up the isolated silos in the IIDSs research landscape and hence allows the wide-ranging application of IIDSs across domains and protocols.

To realize this goal, we first highlight the benefits of protocol-independent IIDSs for the research community (Sec. 3.1) before we survey 53 state-of-the-art IIDS approaches to analyze their potential for generalizing across industrial domains and protocols (Sec. 3.2).

3.1 Benefits of Protocol-Independent IIDSs

Although IIDSs are specifically tailored to a single industrial protocol and domain nowadays, this does not imply that their benefits are necessarily restricted to a single combination. *Protocol-independent IIDSs* (cf. Fig. 1b) instead promise to operate detached from specific industrial protocols or domains and thus exhibit several advantages.

First, this would allow applying IIDSs to multiple industrial protocols within the same domain. E.g., power grids that historically rely on different protocols such as IEC-104 (covered by IIDSs [40, 79–81]) or DNP3 (covered by [46, 96]) would no longer require separate IIDSs. Moreover, if the underlying industrial protocol could be exchanged, the potential to generalize detection methodologies across domains not initially thought of would become directly accessible.

Also, by designing IIDSs with protocol-independence in mind, developers no longer have to extract relevant data from a specific industrial protocol in a time-consuming and error-prone process. Instead, this process can be left to domain experts who know the ins and outs of specific network protocols, which has already proved beneficial for prominent intrusion detection datasets, e.g. SWaT [50] providing pre-processed process data from network traffic.

Finally, it becomes easier to reproduce results, validate findings in different scenarios, and fairly compare IIDSs based on common metrics, thus constituting to sound scientific experiments [49, 113]. Protocol-independence allows to efficiently validate results from prior work and expand upon them, a step that is often neglected in research [18, 113], especially when the original scenario is unavailable, flawed, or too narrow in scope.

Consequently, decoupling IIDSs from the underlying industrial protocol promises to address the transferability of IIDSs and their generalizability to new scenarios. It can also streamline future research efforts on IIDSs by enabling extensive evaluations, validations, or removing potential biases with regard to specific protocols or domains. However, so far, it remains open how big the potential for realizing protocol-independent IIDSs actually is.

3.2 Potential for Protocol-Independent IIDSs

To identify the potential for protocol-independent IIDSs, we first study to which extent existing IIDSs can generalize to new industrial protocols and domains, which is directly related to the protocol and process information they operate on. However, while existing surveys identify the need for generalizing IIDSs [36], they do not investigate whether existing work shows potential in this regard. Bridging this gap, we survey 53 IIDSs covering a wide range of detection methods, industrial protocols, and domains. We specifically focus on identifying which information they use for detection.

As our goal is to identify the potential for protocol-independent IIDSs, we specifically target *process-aware* or *semantic* IIDSs (leveraging knowledge about the physical process), as these are mostly independent of the used protocol. Contrary, we exclude IIDSs leveraging specific characteristics of physical communication [71, 72].

Survey Structure. In the following, we structure and discuss our survey, as depicted in Tab. 1, along two broad categories of process-aware and semantic intrusion detection: *Communication-based* approaches (Sec. 3.2.1) which incorporate characteristics of industrial network traffic and *process state-aware* approaches (Sec. 3.2.2) restricted to the current (physical) state of an industrial setting.

For each approach, we identify the input format, i.e., network packets (P) or aggregated process states (S). We complement this with information an IIDS operates on, i.e., (i) timing information of exchanged packets, (ii) knowledge about communication partners and type of communication, as well as (iii) which values are communicated, also including the process state. We identify whether additional packet features are required (e.g., packet lengths), whether an externally provided process model is required, as well as whether it trains on anomalous (A) or benign-only (B) traffic. Finally, we list the number of evaluation scenarios, the number of IIDSs a novel approach compares to, the intended industrial domain and protocol, as well as the availability of source code and evaluation dataset.

3.2.1 Communication-based Intrusion Detection. Different streams of research propose to leverage the inherent communication patterns found across industrial domains. We identify three categories of how regularities caused by the periodic distribution of, e.g., new sensor measurements, can be utilized for anomaly detection. First, irregularities in *timings* indicate potential malicious activity. Secondly, the order of message *sequences* also exposes repeating patterns that can be learned to identify expected normal communication. Finally, a class of *classifier-based* approaches grade each packet as benign or anomalous based on a set of predefined features.

Our analysis shows that all three categories are used across different industrial domains. Analyzing the type of information these IIDSs operate on, we observe that timing- and sequence-based approaches mostly rely on just four data types: timing information, communicating entities, message types, and often detailed information about transmitted data. In addition, classifier-based approaches often incorporate the exchanged values. Rarely do such IIDSs use further packet features, e.g., packet sizes or checksums for their classification. Still, none of the approaches take advantage of protocol-specific information, such that, in theory, they could be employed for other protocols and scenarios (for Chu et al. [31], this claim cannot be validated as their automatic feature extraction approach could potentially result in the use of arbitrary features).

Overall, communication-based approaches are not based on the particular structure and behavior of the underlying industrial protocol. Instead, despite IIDSs being tailored to specific protocols, they nearly exclusively use protocol-agnostic information to detect anomalies. Hence, in theory, virtually all considered communication-based IIDSs work protocol-independent and show potential for deployment in a wide range of industrial domains.

However, such potential for protocol-independence remains mostly unexplored in practice. As a rare exception, Lin et al. evaluate their IIDS [81], checking for irregular timings between similar messages, for three industrial protocols. Still, they focus on a single domain (power grids) and thus do not examine generalizability. Generally, despite theoretical generalizability to other domains and protocols, communication-based IIDSs are not compared against each other or evaluated across different domains. Consequently, researchers cannot determine which IIDSs show the best performance and should be used as a benchmark for novel proposals, improved in future research, or deployed in practice.

3.2.2 Process State-aware Intrusion Detection. In contrast to the communication-based approaches, process state-aware IIDSs solely base on (time-)series of industrial system’s physical states. A state comprises the combined process information of all sensors and actuators aggregated over multiple packets to assess whether a process’ physical state indicates an anomaly over time. Thus, IIDSs from this category do not rely on packet-specific information and already naturally abstract from many industrial protocol’s characteristics. To identify attacks, process state-aware approaches leverage the repetitiveness and predictability of physical processes, e.g., Programmable Logic Controllers (PLCs) controlling a pump to keep a water reservoir’s fill level within certain bounds [39].

While process state-aware IIDSs operate on a more restricted set of information than communication-based IIDSs, their detection methodologies show more diversity, which we broadly classify along four categories. The first category of approaches defines *critical states* of an industrial system, e.g., through externally provided system states by domain experts, and raises an alarm if such a state is reached. Complementary, the second category of IIDSs attempts to *predict* future states based on past observations, raising an alarm if the behavior significantly deviates from the prediction. The third category of approaches takes a more local view by investigating *correlations* between an individual or a small set of sensor readings over time, searching for indicators of anomalous activities, such as outliers from learned clusters. Finally, as for network-based approaches, a category of approaches train *classifiers* to distinguish between genuine and anomalous behavior based on process state.

The majority of approaches are evaluated on public datasets, but essentially split across merely two datasets (SWaT [50] and TEP [35]), suggesting an evaluation bias. Only Aoudi et al. evaluate their IIDS [16] on both of these datasets. Thus, our analysis of process state-aware approaches’ evaluation methodologies again indicates that most publications neither compare their approach to related work sufficiently nor evaluate on multiple datasets. This observation is surprising since process state-aware IIDSs are especially interesting in terms of transferability, as their underlying process representation contains information available across many industrial protocols and domains.

	Publication(s)	Year	Detection Methodology	Representation Timing	Comm. Entities	Message Type	Comm. Values	Process State	Packet Features	External Model	Trains on	Open Source	Public Dataset	Eval. Scenarios	Compares to Domain(s)	Protocol(s)	
Communication-based Approaches (Sec. 3.2.1)	Timings	Valdes et al. [115]	2009	Flow Periodicity	P	●	●	○	○	○	●	○	○	1	0	🏠	Modbus
		Barbosa et al. [19]	2012	Periodicity	P	●	●	○	○	○	○	○	○	1	0	🏠	-
		Ponomarev et al. [93]	2016	Telemetry	P	●	●	○	○	○	○	○	○	1	0	🏠	Modbus
		Lin et al. [81]	2017	Inter arrival time	P	●	●	●	○	○	○	○	○	3	0	🏠	Modbus, IEC-104, S7
		Lin et al. [80]	2019	Inter-arrival time	P	●	●	○	○	○	○	○	○	1	0	🏠	IEC-104
	Sequences	Goldenberg et al. [51]	2013	DFA	P	●	●	●	○	○	○	○	○	1	0	⚙️	Modbus
		Yoon et al. [123]	2014	PST	P	●	●	●	○	○	○	○	○	2	0	🏠	Modbus
		Caselli et al. [23, 24]	2015	DTMC	P	●	●	●	○	○	○	○	○	1	0	🏠	Modbus
		Ferling et al. [40]	2018	DTMC	P	●	●	●	○	○	○	○	○	1	0	🏠	IEC-104
		Lin et al. [79]	2018	PST	P	●	●	●	○	○	○	○	○	2	0	🏠	IEC-104
Classifiers	Yun et al. [124]	2018	Nearest Neighbor	P	●	●	○	○	○	○	○	○	1	0	-	-	
	Shang et al. [103]	2016	SVM	P	●	●	●	○	○	○	○	○	1	0	🏠	Modbus	
	Feng et al. [38]	2017	LSTM	P	●	●	○	○	○	○	○	○	1	1	🏠	Modbus	
	Perez et al. [83]	2018	SVM, RF, BLSTM	P	●	●	○	○	○	○	○	○	1	0	🏠	Modbus	
	Anton et al. [14]	2019	RF, SVM	P	●	●	○	○	○	○	○	○	1	0	🏠	Modbus	
	Chu et al. [31]	2019	NN	P	○	○	○	○	○	○	○	○	1	3	🏠	Modbus	
	Maglaras et al. [85]	2019	OCVM	P	●	○	○	○	○	○	○	○	1	0	-	Modbus	
DIDEROT [96]	2020	Decision Tree, DNN	P	●	●	○	○	○	○	○	○	1	0	🏠	DNP3		
Process State-aware Approaches (Sec. 3.2.2)	Critical States	Carcano et al. [21]	2011	Language	S	○	○	○	○	○	○	○	○	1	0	🏠	Modbus
		Almalawi et al. [11]	2014	Outlier detection	S	○	○	○	○	○	○	○	○	3	0	🏠	Modbus
		Kiss et al. [70]	2015	GMM	S	○	○	○	○	○	○	○	○	1	0	🏠	-
		Pan et al. [91]	2015	Common Paths	S	●	○	○	○	○	○	○	○	1	0	🏠	Modbus, IEEE C37.118
		SysDetect [66]	2015	Frequent Itemsets	S	○	○	○	○	○	○	○	○	1	0	⚙️	-
		Kong et al. [63, 73, 74]	2016	Temporal logic	S	○	○	○	○	○	○	○	○	2	0	🏠	-
		Adepu et al. [2]	2016	Invariants	S	●	○	○	○	○	○	○	○	1	0	🏠	EtherNet/IP
		Feng et al. [39]	2019	Invariants	S	○	○	○	○	○	○	○	○	2	2	🏠	EtherNet/IP
	Monzer et al. [86]	2019	Rules	S	●	○	○	○	○	○	○	○	1	0	🏠	Modbus	
	Das et al. [33]	2020	Data Analysis	S	●	○	○	○	○	○	○	○	1	0	🏠	EtherNet/IP	
	Behavior Prediction	Hadziosmanović et al. [53]	2014	Autoregression	S	●	○	○	○	○	○	○	○	2	0	🏠	Modbus
		Caselli et al. [23]	2015	DTMC	S	●	○	○	○	○	○	○	○	1	0	🏠	EtherNet/IP
		Ahmed et al. [6]	2017	Kalman Filter	S	●	○	○	○	○	○	○	○	1	0	🏠	-
		PASAD [15, 16]	2018	PCA	S	○	○	○	○	○	○	○	○	3	1	🏠	Modbus EtherNet/IP
		Choi et al. [28]	2018	Control Invariants	S	●	○	○	○	○	○	○	○	1	0	🏠	-
		Myers et al. [88]	2018	Petri-nets	S	●	○	○	○	○	○	○	○	1	0	⚙️	S7
		Kravchik et al. [75]	2018	Neural Networks	S	○	○	○	○	○	○	○	○	1	3	🏠	EtherNet/IP
		TABOR [82]	2018	TA, BN	S	●	○	○	○	○	○	○	○	1	2	🏠	EtherNet/IP
Anton et al. [12]		2019	Matrix Profiles	S	●	○	○	○	○	○	○	○	1	1	🏠	EtherNet/IP	
HybTester [25]		2019	Hybrid-Automata	S	○	○	○	○	○	○	○	○	1	0	🏠	EtherNet/IP	
Correlations	Kim et al. [68]	2019	Neural Networks	S	●	○	○	○	○	○	○	○	1	1	🏠	EtherNet/IP	
	Denque Anton [13]	2020	Matrix Profiles	S	●	○	○	○	○	○	○	○	1	0	🏠	EtherNet/IP	
	SAVIOR [95]	2020	Physical Invariants	S	●	○	○	○	○	○	○	○	2	1	🏠	MAVLink	
	Krotofil et al. [76]	2015	Entropy Analysis	S	●	○	○	○	○	○	○	○	1	0	🏠	-	
	Alippi et al. [10]	2016	Hidden Markov Model	S	●	○	○	○	○	○	○	○	3	0	🏠	-	
	Aggarwal et al. [3]	2018	Hidden Markov Model	S	○	○	○	○	○	○	○	○	2	0	🏠	-	
	Hau et al. [54]	2019	Statistics	S	●	○	○	○	○	○	○	○	1	0	🏠	-	
	NoiSense [4]	2020	Noise Fingerprinting	S	○	○	○	○	○	○	○	○	2	0	🏠	EtherNet/IP	
Classifiers	ProcessSkew [8, 9]	2020	Noise Fingerprinting	S	●	○	○	○	○	○	○	○	1	0	🏠	EtherNet/IP	
	Nader et al. [89]	2014	SVDD, KPCA	S	●	○	○	○	○	○	○	○	2	0	🏠	-	
	Junejo et al. [64]	2016	Machine-learning	S	●	○	○	○	○	○	○	○	1	0	🏠	EtherNet/IP	
	Inoue et al. [61]	2017	SVM, DNN	S	○	○	○	○	○	○	○	○	1	0	🏠	EtherNet/IP	
	Chen et al. [26]	2018	SVM	S	○	○	○	○	○	○	○	○	1	0	🏠	EtherNet/IP	
	AADS [1]	2019	DNN	S	○	○	○	○	○	○	○	○	1	3	🏠	EtherNet/IP	
	Anton et al. [12]	2019	OCVM, Isolation Forest	S	○	○	○	○	○	○	○	○	1	0	🏠	EtherNet/IP	
FALCON [102]	2020	LSTM+ML	S	●	○	○	○	○	○	○	○	1	2	🏠	EtherNet/IP		

🏠: Power Grid 🌊: Water Treatment 🏠: Gas ⚙️: Manufacturing 🏠: Medical 🏠: Chemical 🚰: Water Distribution 🚚: Logistics 🏠: Synthetic Setup
 S: State, P: Packet ●: yes, ○: partial, ○: no A: Anomalous, B: Benign ●^S: SWaT [50], ●^W: WADI [7], ●^T: TEP [35] ●^M: Morris [87]

Table 1: Our survey of 53 intrusion detection approaches confirms the heterogeneity across the industrial research landscape. While IIDSs operate on few information types, they are mostly developed in isolated silos and seldomly compare to existing research. Moreover, we observe similar detection methodologies across different protocols and domains, indicating an enormous potential for realizing protocol-independent IIDSs to transfer achievements to a broader scale of industries.




Figure 2: To achieve protocol independence in IIDS research, IPAL separates the detection methods from the underlying industrial network protocol with an abstract representation.

3.2.3 Meta-analysis of Current IIDS Research. Our analysis of 53 IIDSs confirms our initial observation from Sec. 2.3 that the IIDS research landscape indeed suffers from *heterogeneity*, and the overall progress in this research field is thus unnecessarily slowed down. Further confirming the problem of prevalent *evaluation bias*, we find that IIDSs are compared to a median of zero (average 0.38) existing IIDSs. This phenomenon can partly be explained by the exclusive use of proprietary evaluation datasets (47% of the proposals), a lack of published code or artifacts (75%), and single evaluation scenarios (75%). Similarly, concerning dataset diversity, the median number of evaluated datasets is 1 (average 1.32), and 52% (13 of 25) evaluating on public datasets did so exclusively on SWaT [50].

Laying out a path towards protocol-independent IIDSs as a remedy for current limitations of IIDS research (cf. Sec. 2.3), we observe that IIDS models do not depend on specific communication protocols but rather on the knowledge of which information is exchanged at which time between whom. Furthermore, two-thirds (66%) operate on a state of the underlying process only, completely oblivious of how and when this information is exchanged. Thus, our survey indicates a huge potential for realizing protocol-independent IIDSs.

4 PROPOSING IPAL – AN INDUSTRIAL PROTOCOL ABSTRACTION LAYER

Given the enormous potential for realizing protocol-independent IIDSs (cf. Sec. 3.2) and thus capitalize on their manifold benefits (cf. Sec. 3.1), we set out to turn this vision into reality. To this end, we propose IPAL, our design of an *industrial protocol abstraction layer* to decouple intrusion detection from domain-specific industrial communication protocols. As shown in Fig. 2, we derive IPAL from the knowledge gathered in our survey (Sec. 4.1). We discuss how the landscape of industrial protocols can be transcribed into IPAL (Sec. 4.2) and how IIDSs implemented on top of IPAL help transfer their potential to new scenarios (Sec. 4.3). Finally, we discuss the limitations of introducing an abstraction layer (Sec. 4.4).

4.1 Designing IPAL

IPAL, our design of an *abstract representation* for industrial protocols, lays the foundation to design and study protocol-independent IIDSs by unifying industrial communication. During our comprehensive survey of IIDSs across various industries and communication protocols, we identified inherent similarities with regard to the required data (cf. Sec. 3.2). Leveraging this knowledge, we deduce

IPAL	Description
id	Unique identifier for IPAL messages within a dataset
timestamp	Time corresponding to this IPAL message
length	Length of the industrial protocol’s layer
malicious	Label for training and evaluation (benign or anomalous)
source	Sender of the network packet (e.g., IP and port)
destination	Receiver of the network packet (e.g., IP and port)
message_type	Fine-granular industrial packet’s message type
activity	Abstracted message type (e.g., request or response)
responds_to	List of related IPAL messages this message responds to
process_data	Arbitrary number of process variables and their values

Table 2: IPAL captures ten features in an abstract representation to support current IIDSs across multiple domains.

the information captured by IPAL, as shown in Tab. 2. First, process state-aware IIDSs (cf. Sec. 3.2.2) only operate on a very restricted set of information, encompassing only timings and information about exchanged state information. Meanwhile, communication-based IIDSs (cf. Sec. 3.2.1) operate on a broader set of information but still only rely on packets with (direct) influence on the physical process, such that other packets (e.g., diagnostic data or TCP handshakes) can be ignored. As detailed in the following, ten features suffice to represent industrial network packets in a common format while preserving all information required by process-aware IIDSs.

Meta Data. The first observation in our survey regarding the abstraction is the incorporation of metadata across all IIDS categories. Thus, IPAL includes a packet’s timestamp and length. In addition, a unique identifier (*id*) is included to disambiguate packets. Since several IIDSs require labeled training data differentiating benign and malicious packets, IPAL includes the *malicious* field.

Addressing Information. Besides meta data, communication-based IIDSs require features to identify the *communicating entities*. Thus, IPAL includes the *source* and the *destination* of a single packet. These fields are arbitrary strings that, in most cases, represent an IP-port combination, but can be extended, e.g., by adding Modbus’s unit identifier field, to further disambiguate devices. Likewise, *destination* can remain empty for broadcast protocols.

Message Identification. One relevant aspect found in our survey is that communication-based IIDSs leverage *message types* of industrial packets on different levels of abstraction (cf. Sec. 3.2.1). Some IIDSs [23, 24] analyze the sequence of specific, reoccurring messages, and further IIDSs [81] differentiate between requests or responses. Therefore, IPAL captures a fine-granular *message_type*, as well as a more generic *activity*. Here, four activities (requests, commands, and their respective answers) suffice [121], while the *message_type* closely matches the differentiation schemes used by the underlying protocols. Lastly, the *responds_to* field lists the *ids* of all IPAL packets a given message is a response to.

Process Data. The most important feature used by both types of IIDSs is the process state. In IPAL, the field *process_data* collects all process variables as well as their current values as communicated within a single packet. Depending on the transcribed packet, *process_data* thus can range, e.g., from a single temperature to an array of sensor readings and actuator commands. As process state-aware IIDSs demand a snapshot of the entire system’s state, we support state aggregation over several IPAL messages to realize a similar extraction as, e.g., through active polling [45].

A few rare cases exist where additional features (e.g., checksum validity tests [31]) are included in the decision-making process of otherwise process-aware IIDSs. These checks are, however, already covered by traditional IDS (e.g., Zeek [94]) and do not contribute to information about the physical process monitored by process-aware IIDSs, and can thus be skipped to avoid redundant work. Consequently, the ten features of IPAL suffice to preserve all information required by the process-aware IIDSs surveyed in Tab. 1.

4.2 Transcribing Industrial Protocols

IPAL, our abstract representation of industrial network traffic, promises to decouple process-aware intrusion detection from underlying industrial protocols. Yet, to realize protocol-independent IIDSs, real-world network traffic must be transferred into this abstract representation. As shown in Fig. 2, we designed a *transcriber* [43] that automates this conversion. While integrating specific industrial protocols into IPAL requires a deep protocol understanding, this work needs to be performed only once.

We have already incorporated nine industrial protocols (Modbus, EtherNet/IP, MQTT, IEC-104, DNP3, S7, IEC 61850-GOOSE, IEC 61162-450, and NMEA 0183) used in popular datasets. Thus, future IIDS research can focus on modeling industrial processes, while a broad set of datasets from different industrial scenarios is readily available for evaluation through IPAL [41]. In the following, we discuss challenges faced while implementing our transcriber.

Most features of IPAL can be directly extracted from network packets. However, certain protocols require special attention. E.g., in ModbusTCP, multiple sub-devices may operate behind a single TCP connection. Thus, source and destination need to include not only the IP and port but also the device’s unit identifier. Also, while Modbus’s function code directly corresponds to the message type, the activity has to be derived from the traffic direction, as function codes for requests and responses are identical.

A more challenging aspect is the extraction of the actual process values from industrial communication because of their encoding. To this end, transcribers can contain custom rules to extract and post-process process values from observed communication, similar to StreamPipes [125]. Such rules are small code snippets that allow, e.g., to interpret two 16-bit registers (the biggest data type in Modbus) as a single 32-bit float. Thus, we can integrate the necessary flexibility to define the interpretation of potentially multiple related variables and even directly annotate the interpreted value with a more descriptive and human-readable name.

4.3 Developing IIDSs with IPAL

To leverage its benefits, we provide a framework to implement protocol-independent IIDSs on top of IPAL [42]. We observed in our survey (cf. Sec. 3.2) that IIDSs are either based on communication traffic or on the global process state. As shown in Fig. 2, for communication-based IIDSs, our framework passes each packet to the IIDS after abstracting it into an IPAL message.

For process state-aware IIDSs, we need to aggregate multiple packets over time to obtain the current state of the entire process. Thus, our framework simply caches the process values included within the most recent IPAL messages. We considered different state aggregation methods, but simply outputting the most recent process

values in regular intervals (e.g., each second) yields the best results and allows us to convert EtherNet/IP traffic included in a newer run of the SWaT testbed (without attacks) to a state that closely matches the pre-processed state information provided by the dataset authors. Furthermore, as some common datasets (e.g., SWaT [50]) already provide a state representation, our framework optionally allows restoring these pre-computed states directly.

Our framework enables the realization of communication-based and process state-aware IIDSs on IPAL messages. IPAL already supports nine industrial protocols from various domains. Thus, IIDSs developed for or adjusted to IPAL can be evaluated for various scenarios based on current and future datasets, accelerating research on IIDSs and moving them closer to deployment (cf. Sec. 3.1).

4.4 Limitations of Protocol Abstraction

Any form of abstraction involves the risk of introducing an information loss, as crucial features may be discarded. In our case, IIDSs taking advantage of a feature not covered by our abstraction would likely not be compatible with IPAL. However, considering that IPAL’s features are derived from an extensive survey of 53 scientific publications (cf. Sec. 3.2), we expect the number of currently developed IIDSs that cannot be realized on top of IPAL to be low.

Still, our choice of features captured in IPAL could potentially narrow the scope of future research. However, for dominantly used datasets, we already observe a restriction in available information to pre-processed data, e.g., only providing timing and process-state information (cf. SWaT in Tab. 1), presumably to facilitate easy utilization of the dataset. Contrary to intuition, the adoption of IPAL could even lead to the availability of more features by encouraging the release of unprocessed packet captures, as IPAL provides the tooling to easily process raw data according to researchers’ needs. In this vein, the IPAL framework is publicly available [42, 43] such that the extracted feature set can be extended in the future if deemed necessary. A similar trend has been observed for Netflow [32], whose abstraction revolutionized network monitoring and introduced newer versions as the need for more features arose.

From a different perspective, protocol abstraction comes with the potential risk of introducing transcription or abstraction inaccuracies affecting the final detection performance. Yet, as we show in the following (Sec. 5), such inaccuracies are generally not introduced by IPAL, and the IIDS performance is independent of the use of our abstraction layer. Consequently, the value of IPAL homogenizing a split-up research landscape far outweighs the potential drawbacks arising from introducing an abstraction layer (so early) for the growing industrial intrusion detection research field.

5 PRACTICAL APPLICABILITY OF IPAL

We derived IPAL from the requirements of existing IIDSs and discussed its theoretical applicability for industrial network protocols. Now, we complement this viewpoint by showing IPAL’s *practical applicability and correctness* by performing a reproducibility study of eight IIDS approaches from previous work and reproducing their evaluation results on top of IPAL. Thereby, we also contribute to reproducing scientific research results as an independent party. While deemed extremely important [18], this is often a challenging and tedious task with little reward and thus rarely performed [36, 97].

Category	IIDS	Code	Dataset	Result
Communication	Inter arrival time [81]	○	●	✓
	DTMC [40]	●	●	(✓)
	RF [83]	●	●	✓
	SVM [83]	●	●	✓
	BLSTM [83]	●	●	✓
Process state	PASAD [16]	●	●	✓
	Seq2Seq-NN [68]	●	●	✓
	TABOR [82]	●	●	✓

●: available ●: dataset ○: attacks ○: no (✓) no comparison to original paper

Table 3: We prove the applicability and correctness of IPAL by successfully reproducing eight IIDSs on top of IPAL. These cover all categories from the survey and were selected by the availability of code and evaluation datasets.

IIDS selection criteria. To conduct a comprehensive reproducibility study, we set out to reproduce at least three representative approaches for each of the two major categories in our survey (cf. Tab. 1), covering all features provided by IPAL. Since we re-evaluate existing IIDSs, we rely on the availability of the original evaluation dataset or tools for attack generation. Also, as re-implementing IIDSs from scratch can be challenging, we focus on open source approaches and only implement IIDSs ourselves if necessary. Thus, our selection was driven by the availability of code or datasets.

Note that we reproduce evaluation results and thus do not invent new attack methodologies. This implies that the reproducibility results are not necessarily comparable, as the authors used fairly different evaluation procedures. Still, we will show in Sec. 6 how IPAL enables conducting a side-by-side evaluation of different IIDSs.

In the following, we briefly argue why we selected each IIDS, summarize its core idea, provide re-implementation and evaluation details, and compare our results to the original publication. Tab. 3 summarizes our IIDS selection and reproduction results.

5.1 Inter-arrival Time (Communication)

IIDSs from the communication-based category utilize data on a per-packet basis and consider features such as communicating entities and message types. An ideal representative for this category is the inter-arrival time (IAT) approach by Lin et al. [81], which utilizes periodic traffic patterns in industrial protocols such as Modbus, S7, and IEC-104 to detect, e.g., packet-injection attacks. To this end, the approach measures the *mean* inter-arrival time, i.e., the elapsed time between two packets of the same type (e.g., requests and responses), as well as the maximum temporal deviation (*range*) between packets with the same content and checks for timing violations.

As no implementation of the inter-arrival time approach was available, we re-implemented it based on the paper [81]. To reproduce the results, we use the only publicly available dataset from the original work [81], which evaluates S7 traffic [105]. As only the attack-free dataset was available, we performed an identical 1/10 train-test dataset split and injected “malicious” packets according to the authors’ description. Our reproduced results on top of IPAL (cf. Tab. 4) closely match those of the original paper. We attribute the remaining differences to randomness and minor implementation uncertainties w.r.t. to attack generation. Overall, IPAL provides everything needed to re-implement and reproduce this IIDS.

Attack	Model	Metric	Original		IPAL	
			Request	Response	Request	Response
Flooding	Mean	TPR	99.90	99.90	99.98	99.97
		FPR	0.01	0.20	0.01	0.01
		ODR	100.00	100.00	100.00	100.00
	Range	TPR	59.10	56.40	68.21	65.48
		FPR	0.80	1.10	0.98	1.20
		ODR	100.00	100.00	100.00	100.00
Injection	Mean	TPR	96.20	96.60	98.95	98.95
		FPR	0.01	0.20	0.01	0.01
		ODR	100.00	100.00	97.37	97.37
	Range	TPR	100.00	99.50	99.47	99.47
		FPR	0.80	1.10	1.01	1.21
		ODR	100.00	100.00	97.37	97.37
Prediction	Mean	TPR	0.10	0.20	4.60	4.92
		FPR	0.00	0.20	0.01	0.01
		ODR	99.80	99.50	95.21	95.23
	Range	TPR	90.60	91.00	93.69	93.59
		FPR	0.80	1.10	0.97	1.28
		ODR	99.80	99.50	95.21	95.23

TPR: True Positive Rate, FPR: False Positive Rate, ODR: Overall Detection Rate (in %)

Table 4: Our re-implemented IAT IIDS closely resembles the original detection rates (cf. Tab. 2 in [81]). Slight deviations result from minor differences in the attack generation.

5.2 DTMC (Communication)

As an additional communication-based IIDS, we consider an approach relying on Discrete-time Markov Chains (DTMCs) modeling the sequence of network packets for a single connection [40]. This approach leverages a core property of industrial networks: Devices such as PLCs perform several communication steps in series, e.g., reading values from remote devices and adjusting setpoints later on, resulting in periodic packet sequences repeating in similar order.

We adapted the DTMC approach’s publicly available source code [30] to operate on top of IPAL. While a full reproducibility study is infeasible without the original dataset, we can still validate whether the original implementation and the version on top of IPAL produce identical results when presented with the same input. To this end, we generated our own IEC-104 network trace using a simulation framework for power distribution grids [56] and the attack tool by the authors to add equivalent attacks to our trace [29]. We performed the same procedure as in the original paper [40], i.e., a 50/50 train/test split, applying the attack tool ten times for each attack type and counting the number of state and transition violations compared to the trained DTMC. While not comparable to the original evaluation, IPAL correctly preserves all information resulting in identical attack coverage compared to the original DTMC implementation, as shown in Tab. 5.

5.3 Classifiers (Communication)

Complementing the IAT and the DTMC approach, IIDSs using machine-learning classifiers may operate on additional information such as the packet length. Furthermore, they additionally require traces of past attacks for training. Thus we selected the work of Perez et al. [83] as a representative for this category. In their work, they compare three different machine learning algorithms against each other (Random Forest (RF), Support Vector Machine (SVM), and Bidirectional Long Short Term Memory (BLSTM)).

Reduction Type	Anomaly Type	Copy [%]			Remove [%]			Swap [%]		
		0.1	1	10	0.1	1	10	0.1	1	10
none	state	0	0	0	0	0	0	0	0	0
	transition	0	15.5	38.5	0	2	14	0	20.5	46.6
overlapping	state	0	0	0	0	0	0	0	0	0
	transition	0	8.5	19	0	1.5	6.5	0	11.5	22
all	state	1	1	1	1	1	1	1	1	1
	transition	0	9.5	22	0	1.5	6.5	0	13	27

Table 5: Adapting the DTMC approach to IPAL yields identical detection results in the number of state and transition validations. We do not show the identical values twice here.

	Paper [83]	Original [99]	IPAL
SVM	94.36 %	94.36 %	96.10 %
RF	99.58 %	99.52 %	99.73 %
BLSTM	98.40 %	98.12 %	96.78 %

Table 6: Measuring the accuracy of the public source code (Original) resembles the original paper results. Also, on top of IPAL, we achieve equivalent accuracy most of the time.

For our reproducibility study, we adapted the implementations of the machine-learning classifiers published by the authors [99] to operate on top of IPAL. Furthermore, we use the same evaluation dataset [87], hyper-parameters, and pre-processing steps as in the original paper [83]. In Tab. 6, we compare the original results of the paper with our replication of the original evaluation as well as the same evaluation on top of IPAL. All three approaches perform equivalently, and the minor deviations can be inferred from randomly selecting a different train and test dataset. Thus, for these three approaches and the previous two communication-based IIDSs, we showed that IPAL correctly preserves all relevant information to achieve equivalent attack detection performance.

5.4 PASAD (Process state)

In contrast to the thus far considered IIDSs, Process-Aware Stealthy Attack Detector (PASAD) [16] detects structural changes in data series and thus represents the process state-aware category. PASAD’s core idea is that legit data series span a vector space, and the distance to the mean of all training vectors indicates an anomaly. An alarm is raised if the distance (departure score) exceeds the maximum observed distance during training (threshold).

Two implementations of PASAD are available, the authors’ version in Matlab [62] and a re-implementation in Python [97]. We realized PASAD for IPAL based on the Python version and used the original Matlab implementation to generate reference data. PASAD was trained and evaluated visually on output logs of the TEP [35], parts of the SWaT datasets [50], and private data of a real water treatment facility. We verify our implementation on the public TEP datasets and exemplarily discuss one example in Fig. 3. The results obtained on top of IPAL are identical to the original implementation [62], and both detect the attack simultaneously. By reproducing these results, we show that IPAL is able to preserve all information required by a process state-aware IIDSs such as PASAD.




Figure 3: PASAD’s original departure scores (yellow, cf. Fig. 4 in [16]) match ours (blue). Both have identical thresholds of ~ 0.065 and detect the attack simultaneously after $\sim 3.58h$.




Figure 4: Seq2Seq-NN’s original anomaly scores (yellow, cf. Fig. 3 in [68]) match our reproduced implementation (blue). Both implementations emit alarms simultaneously.

5.5 Seq2Seq-NN (Process state)

Extending on the idea of PASAD, the approach by Kim et al. [68] requires knowledge about *multiple* actuators and sensors to predict the expected future states of an industrial system. More specifically, sequence-to-sequence Neural Networks (Seq2Seq-NN), known from language translation, are provided with the recent history of system states (e.g., the last 99 seconds) and predict the next state. Then, the distance between the predicted and actual state is calculated, and the error aggregated over time, resulting in an anomaly score. If the anomaly score exceeds a threshold, the IDS raises an alarm.

For our Seq2Seq-NN implementation on top of IPAL, we adapted its publicly available version [67], extracted the original anomaly scores as a comparison baseline, and used the authors’ pre-trained models for the evaluation on IPAL. We exemplarily compare the original implementation and the one on top of IPAL w.r.t. their ability to detect an attack on stage 3 of SWaT [50] in Fig. 4. There the anomaly score is identical, and the attack is detected simultaneously, showing IPAL correctly providing all information for Seq2Seq-NN.

5.6 TABOR (Process state)

Complementing PASAD and Seq2Seq-NN, TABOR focuses on one sensor and models the impact of actuators on this sensor [82]. To this end, TABOR combines three distinct detection approaches into one IIDS: An out-of-alphabet check alerts if unseen data or data out of regular boundaries is observed, a timed automaton captures the sequence and duration of linear sensor value segments to enable anomaly detection, and a Bayesian network couples sensor segments with actuator states to identify unexpected combinations.

We were able to retrieve the implementation of TABOR from the authors for our re-implementation on top of IPAL. TABOR was evaluated on the SWaT dataset [50], which was divided into smaller logically coherent units (models). Here, we concentrate on reproducing model 1 as its intermediate results in the original paper [82]

Model 1	Original				IPAL			
	FP	TP	CP	PS	FP	TP	CP	PS
TABOR	0	9	7.85	1889	0	9	7.09	1765
TA	7	26	87.48	189516	7	26	86.3	188829
BN	0	5	1.95	629	0	5	1.27	504
OOA	0	5	5.90	1260	0	5	5.82	1261

Table 7: TABOR, on top of IPAL, detects the identical nine attack scenarios (TP) with slightly different coverage percentage (CP) and penalty scores (PS) (cf. Tab. 2 in [82]).

ease verifying the correctness of our implementation. We achieved equivalent results for our version of TABOR on top of IPAL (cf. Tab. 7). Our re-implementation scores are identical in true positives (TP) and false positives and only deviates marginally within the coverage percentage (CP) and penalty scores (PS), probably due to minor temporal deviations in the alarms. We confirmed manually that the alerts overlap between the original and our implementation. Concluding, IPAL provides the necessary information to realize even complex, combined process state-aware IIDSs.

5.7 Summary and Lessons Learned

To showcase the practical applicability and correctness of IPAL, we performed a reproducibility study of eight IIDSs covering the two IIDS categories with at least three representatives (cf. Tab. 1). By (re-)implementing a diverse set of IIDSs on top of IPAL, we have shown that IPAL provides all information required by various types of IIDSs. By reproducing previous work’s results, we have shown that IPAL operates correctly and neither reduces functionality nor detection rates (cf. Sec. 4.4). Consequently, IPAL provides a reliable foundation for protocol-independent IIDSs, thus capitalizing on the benefits of an abstract representation (cf. Sec. 3.1).

As indicated in Tab. 3, successfully reproducing others’ research requires the availability of the original implementation. Only in rare cases does an evaluation dataset suffice. While we were able to successfully reproduce seven IIDS approaches (and even achieve fully identical results for two of them), we could not reproduce the evaluation results of the DTMC approach [40] since the evaluation dataset was not available to us. Still, as the original source code and attack tool were available, we could use another dataset to show that our re-implementation on top of IPAL produces exactly the same results as the original implementation.

While detection performance is of utmost interest, adding abstraction layers, such as IPAL, introduces overheads. For a fast IIDS such as DTMC, we exemplarily compared the execution time against the original implementation. IPAL yields an overhead of just 10.1% for the benefits of increased applicability. This overhead is expected to shrink for more complex IIDSs.

Overall, the considered eight IIDSs not only rely on the requirements derived from our survey (Sec. 4) but also cover four distinct protocols (Modbus, EtherNet/IP, S7, and IEC-104) from four datasets [35, 50, 87, 105]. Besides contributing to the important task of reproducing scientific research results as an independent party [18, 36], we lay the foundation to ease the *protocol-independent* development, testing, and evaluation of IIDSs by providing a rich foundation for broad evaluations.

Scenario	DoS [20]			Advanced [20]			MitM [109]		
	prec.	rec.	f1	prec.	rec.	f1	prec.	rec.	f1
IAT (mean) [%]	99	100	99	0	0	-	100	98	99
IAT (range) [%]	90	100	95	0	62	1	99	31	47
DTMC [%]	100	100	100	100	95	97	0	0	-

Table 8: IPAL enables the comparison of communication-based IIDSs in realistic scenarios. We see that each IIDS detects certain attacks well in practice, while other attacks that should be detectable in theory are not caught.

6 IIDS GENERALIZABILITY STUDY

Our reproducibility study showed that existing IIDSs can be implemented on top of IPAL without loss of functionality or detection capabilities (cf. Sec. 5). What remains to be shown is that IPAL is actually useful, i.e., it addresses the pressing problems of current IIDS research (Sec. 2.3). More precisely, we demonstrate how IPAL helps to transfer an IIDS to a new industrial protocol, how IIDSs generalize to new scenarios, how existing IIDSs compare against each other, and how distinctly communication-based and process state-aware IIDSs detect different attack types. To this end, we perform three case studies (i) for communication-based IIDSs (Sec. 6.1), (ii) for process state-aware IIDSs (Sec. 6.2), and (iii) combining both directions of intrusion detection (Sec. 6.3).

6.1 Generalizing Communication-based IIDSs

The first case study demonstrates IPAL’s ability to transfer communication-based IIDSs to new and realistic scenarios. Therefore, we study the two approaches IAT [81] and DTMC [40] from our reproduction study (Sec. 5). Both were previously not compared against each other and were evaluated on artificial attacks retroactively integrated into network traces. Thus, their effectiveness in realistic scenarios is currently unknown.

To this end, we utilize novel datasets from realistic industrial systems not yet covered by both IIDSs. The first dataset resembles a Denial of Service (DoS) attack against a power grid communicating with the IEC 61850-GOOSE protocol [20]. In the same setting, more advanced attacks are performed based on a combination of message modification, injection, and replaying, comprising our second dataset [20]. Finally, we look at a Machine-in-the-Middle (MitM) attack dataset based on ARP poisoning to manipulate the data within Modbus packets [109]. Note that these datasets already contained these attacks and were specifically designed for evaluating the performance of security solutions. Our transcriber (cf. Sec. 4.2) transparently translates the GOOSE and Modbus network traces into IPAL, obviating the need to modify the existing IIDS implementations. Tab. 8 summarizes the results on these three new datasets.

DoS. All IIDSs detect the packets injected by the DoS attacks to a high degree. For the IAT’s mean model, these results are expected, as this attack is similar to the artificial flooding attack evaluated in the original publication [81]. Surprisingly, even the IAT’s range model achieves good detection rates, even though it performed worse in the artificial scenario (cf. Tab. 4). However, the DTMC achieved the best performance, classifying all packets correctly despite no similar attack being discussed in the original paper outperforming an IIDS specifically designed to detect such attacks.




Figure 5: IPAL allows the side-by-side comparison of three process state-aware IIDSs on industrial datasets and simultaneously enables to understand how IIDSs generalize to similar (WADI) and new domains (HAI).

Advanced. The second scenario considers advanced data manipulation attacks against a power grid similar to the artificial attacks evaluated in both original publications [40, 81]. Still, neither of the IAT models can sufficiently detect these attacks, despite performing well in the artificial attacks originally evaluated. This behavior can be explained by the low number of data manipulations and variable retransmission timings of the GOOSE protocol, leading to a larger variety of benign inter-arrival times. On the other hand, DTMC achieves great detection rates, confirming the results from the original publication in a new industrial domain and protocol.

MitM. Finally, we consider a MitM attack manipulating data received by industrial controllers [109]. While such an attack is a realistic threat in practice, it has neither been considered by IAT nor by DTMC thus far. Still, the IAT’s mean model performs especially well in this scenario without issuing any false alarm. Contrary, DTMC is not able to detect a single malicious packet since this attack does not change the sequence of observed messages.

These results indicate that artificial attacks, retroactively integrated into a dataset [40, 81], do not necessarily carry over to more realistic attack scenarios, even when considering a similar attack behavior. Yet, these new attack types can be detected well in some instances. This observation highlights that IIDSs have to be studied more extensively to understand which attacks they can detect and which not. Additionally, our case study highlights that one IIDS does not necessarily have to be superior in every attack scenario. Thus, cooperative IIDSs, i.e., several detecting specific attack classes, may improve overall performance (cf. Sec. 6.3). Finally, we showed the adaption from one domain or protocol to another scenario, confirming the enormous potential for protocol-independent and domain-agnostic IIDSs as previously identified in Sec. 3.1.

6.2 Generalizing Process state-aware IIDSs

Our second case study focuses on process state-aware IIDSs. Since the generalizability of IIDSs has not been evaluated sufficiently before (cf. Sec. 2.3), we show how IPAL eases the adaption of IIDSs to new domains and thus understand how IIDSs transfer to new scenarios. Furthermore, as Urbina et al. [49] already observed, approaches are insufficiently compared in the literature. E.g., PASAD is evaluated only visually [16], while TABOR defines its own metrics [82], and Seq2Seq-NN counts an attack as detected even 15 minutes after the attack ended [68] (cf. Sec. 5). Therefore, we evaluate these approaches against each other in a detailed comparison.

From our reproducibility study, we have three process state-aware IIDSs (TABOR [82], PASAD [16], and Seq2Seq-NN [68]) at

hand, which were all evaluated on the SWaT dataset [50]. With IPAL, we can apply these to two new datasets (WADI [7] and HAI [104]). Fig. 5 depicts the alerts of all three approaches on these three datasets and Appx. A summarizes further in-depth results.

SWaT. We compare the IIDSs on the first stage of the SWaT dataset, on which they were initially evaluated. Here, we see significant differences between the operation of each IIDS. TABOR detects 9 attacks, most of which are captured by Seq2Seq-NN too. Seq2Seq-NN detects 10 attacks while still having 17 false alarms. PASAD can detect 6 scenarios, some of which were neither detected by TABOR nor by Seq2Seq-NN. Besides detection performance, we also notice differences in how these IIDSs emit alerts. PASAD buffers recent states internally, and Seq2Seq-NN accumulates deviations over time, delaying alerts for both approaches. Contrary, TABOR has great precision w.r.t. the attack time, yet alarms are announced retrospectively due to internal segmentation. Such differentiation only show in an in-depth side-by-side analysis as enabled by IPAL.

WADI. While the WADI dataset [7] is not unknown (cf. Tab. 1), none of the three approaches evaluated it. WADI is similar to SWaT, and we can show how simple transferring these IIDS to a similar domain with IPAL is again considering the first process stage of WADI only. Here, TABOR detects 4 scenarios, detects PASAD only 2 attacks and Seq2Seq-NN 8 scenarios. During the training of PASAD, computational limits were reached due to the larger dataset size, which might contribute to the lower detection quality. Still, the initial results of transferring an IIDS to a new dataset with IPAL are promising without the need to design new IIDSs from scratch.

HAI. After transferring these IIDSs to a similar domain, we now study how they generalize to other industrial settings. Thus, we consider the fairly new HAI dataset [104] based on a hardware-in-the-loop simulator. As some stages of HAI are of a similar domain as SWaT and WADI, we consider the remaining boiler process used for power generation. This dataset is much noisier than SWaT or WADI and less seldomly exhibits clear repetitive patterns challenging all IIDSs. TABOR does not generalize well, as it flags most parts of the dataset as an anomaly, while Seq2Seq-NN detects 2 scenarios with 5 false alarms. PASAD also detects 2 scenarios but with far more false positives (22). Transferring these IIDSs to HAI with IPAL reveals their unexplored potential to generalize to new scenarios.

IPAL enables a detailed look at how different IIDSs compare against each other. We even observed how differently IIDSs emit alerts motivating the use of well-chosen metrics [60]. Besides this in-depth comparison, IPAL allows to apply IIDSs to new datasets quickly, and we showed that Seq2Seq-NN indeed generalizes.




Figure 6: While until now considered as separate research streams, IPAL can pool the detection capabilities of communication-based (IAT & DTMC) and process state-aware IDSs (PASAD, TABOR, and Seq2Seq-NN), as showcased for the WDT dataset [37].

6.3 Cooperative Communication-based and Process state-aware Intrusion Detection

As a final showcase, we explore IPAL’s capability to compare communication-based and process state-aware intrusion detection systems. Currently, both research branches are largely disconnected (cf. Sec. 3.2), and, as shown in Sec. 6.1, a single IIDS may not perform optimally across all attack types, motivating to combine their potentials even within one industrial scenario. For our study, we train two communication-based IIDSs (IAT [81] & DTMC [40]) and three process state-aware IIDSs (TABOR [82], Seq2Seq-NN [68] & PASAD [16]) on the water distribution testbed (WDT) dataset [37]. What sets it apart from other datasets is that it comprises of network attacks (DoS flooding, MitM to manipulate packet data, and TCP scanning), as well as, physical attacks (e.g., pump breakdowns) and is thus well-suited to analyze and compare both IIDS types.

We visualize the results of applying the five IIDSs to the WDT dataset in Fig. 6 (Appx. B provides further in-depth metrics). Similar to Fig. 5, alerts raised by process state-aware IIDSs are visualized as time ranges. In contrast, individual packets flagged as malicious by communication-based IIDSs are marked with \times . We highlight the attack types in the dataset with different colors.

Without further tweaks, IAT range and DTMC do not produce reliable results, as indicated by the continuous emitting of alerts (in grey). Upon closer manual inspection, these “false” alerts can be traced back to either a missing communication sequence in the training data for DTMC or a missing margin of error for a lower-bound check in the IAT range model. Since these alerts (in grey) are easily circumventable, we excluded them from the following evaluation and considered only alerts marked with black \times .

Process state-aware. We expect process state-aware IIDSs to perform best on the physical attack type since they operate on aggregated state information (cf. Sec. 3.2.2). Indeed, three out of the seven attacks detected by PASAD and five out of nine attacks detected by TABOR are of the physical type. Interestingly, Seq2Seq-NN does not detect a single physical attack even though it should be capable of doing so as it analyses the physical state. Meanwhile, network attacks are covered insufficiently by process state-aware IIDS, as only attacks directly aiming at physical consequences, such as some MitM and few DoS attacks, are detected. In total, TABOR and PASAD detect three MitM attacks each.

Communication-based. Like process state-aware IIDSs, communication-based IIDSs detect MitM attacks, despite not having an understanding of the underlying physical process. DTMC detects three MitM instances, IAT-Range five attacks, and IAT-Mean only two. Another observation is that these IIDSs detect attacks early,

likely even shortly before the attack has a real physical impact (which is detected by process state-aware IIDSs). Regarding DoS, all except one instance are detected by IAT and DTMC. Scanning, since typically considered out of scope for semantic IIDSs (cf. Sec. 2.2 and Sec. 3.2), is, as expected, hardly detectable by any approach.

Both research branches benefit from each other through IPAL. Overall, process state-aware and communication-based IIDSs detect overlapping and contradicting flavors of attack types. Thus, they complement each other nicely, either to widen the covered attack types or to realize redundancy for equivalent attacks – better than a single research branch can achieve on its own. Here, IPAL not only facilitates the generalization to another dataset but also, for the first time, enables the interworking of contrary research streams.

7 FURTHER RELATED WORK

Besides research on IIDSs covered extensively in Sec. 3.2, our work draws inspiration from different streams of related research. Existing tools from traditional network monitoring, such as NetFlow [32], lack essential features and are thus not powerful enough for IIDSs. Furthermore, adapting traditional, rule-based IIDSs [94, 100, 107] to industrial domains [22, 27, 46–48, 122] is also closely related to our work. These rule-based IIDSs, however, only detect traditional (known) attacks (cf. Sec. 2.2) and thus only complement process-aware IIDSs that are able to detect even (stealthy) attacks exploiting physical processes to cause harm [49, 114].

Furthermore, past analysis of IIDSs [5, 49, 90, 114], partially with a focus on specific domains [34, 58, 65, 84, 98], already describe problems such as data heterogeneity and evaluation bias in IIDS research (cf. Sec. 2.3). To analyze the ability of existing IIDSs to detect unseen sensor spoofing attacks, the only reproducibility study (considering only model-free state process-based IIDSs) known to us [36] shows significant differences between claimed generalizability and reality. These works, in addition to reports on issues with widely-used datasets [112], motivated us to quantify these problems and ultimately mitigate them through protocol independence.

The problem of protocol heterogeneity in industrial communication is not exclusive to intrusion detection. Proposals to address this issue (e.g., PLC4X [45] and StreamPipes [120, 125]) do, however, only extract insufficient information for IIDSs and consume valuable bandwidth through polling, a limited good in many industrial scenarios [17, 118]. Other work in this context is concerned with the interoperability of devices [52, 69, 77, 127] and proposes translation between different data representations. We do, however, expect these related fields to benefit from IPAL in the future too, e.g., to reduce development costs and improve protocol support.

Lastly, Ryšavý et al. [101] propose a library to ease industrial network data preprocessing for intrusion detection. Since their approach requires IIDS developers to implement custom data extractors for each protocol, this solution does not tackle protocol-independence, nor does it facilitate generalizability. Still, to motivate the library, the authors similarly acknowledge the problem of heterogeneity within the IIDS research landscape [101] and likewise strive to move towards an interconnected IIDS community.

8 CONCLUSION

A growing number of cyber attacks against industrial networks not only inflict substantial financial and environmental damage but even put human lives at risk [55]. Intrusion detection is regarded as a promising complementary protective measure to timely detect such attacks, which is especially well suited for industrial settings due to their repetitive processes and predictable network behavior. Consequently, a large research community gathered around industrial intrusion detection. The produced research advancements are, however, surprisingly scattered, as many solutions are proposed for distinct communication protocols in specific industrial domains, thus hindering their transfer to other industrial domains.

To better understand this phenomenon, we survey 53 scientific IIDSs and identify an unexplored potential: While practically operating in protocol and domain-dependent silos, theoretically, neither do IIDSs operate on the information of certain communication protocols nor are their fundamental detection methodologies specific to individual domains. Consequently, there is a huge potential for *protocol-independent* IIDSs to protect industrial networks across scenarios, not only to improve the research landscape but also to contribute towards widespread real-world deployments.

To unleash this potential, we propose IPAL, our industrial protocol abstraction layer that decouples intrusion detection from domain-specific industrial communication protocols. To this end, IPAL transcribes protocol features relevant for intrusion detection into a common abstract representation. To show the applicability and correctness of our approach, we conducted a reproducibility study of eight IIDSs from related work, proving that IIDSs can indeed be implemented on top of IPAL. Finally, with the ability to transfer IIDSs seamlessly across different industrial protocols and domains, we studied how existing IIDSs generalize to new scenarios. We find that existing works are (i) indeed not confined to specific domains or protocols, but also that (ii) the type of attacks a given system can detect is not sufficiently studied, thus (iii) motivating the idea to cooperatively combine approaches from previously disjoint detection domains. With our work, we lay the foundation to break up protocol-dependence of IIDSs research and enable further studies regarding the generalization and application of IIDSs in new research or real-world scenarios.

ACKNOWLEDGMENTS

Funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany's Excellence Strategy – EXC-2023 Internet of Production – 390621612. The authors would like to thank Olav Lamberts, Stefan Lenz, Tim Nebel, Leonardo Pompe, and Sven Zemanek for their work on contributing to the IPAL transcriber and IIDS framework.

REFERENCES

- [1] Maged Abdelaty, Roberto Doriguzzi-Corin, and Domenico Siracusa. 2020. AADS: A Noise-Robust Anomaly Detection Framework for Industrial Control Systems. In *Information and Communications Security*.
- [2] Sridhar Adepu and Aditya Mathur. 2016. Distributed Detection of Single-Stage Multipoint Cyber Attacks in a Water Treatment Plant. In *Proceedings of the 11th ACM on Asia Conference on Computer and Communications Security (ASIA CCS '16)*. <https://doi.org/10.1145/2897845.2897855>
- [3] Ekta Aggarwal, Mehdi Karimibiuki, Karthik Pattabiraman, and André Ivanov. 2018. CORGIDS: A Correlation-Based Generic Intrusion Detection System. In *Proceedings of the 2018 Workshop on Cyber-Physical Systems Security and Privacy (CPS-SPC '18)*. <https://doi.org/10.1145/3264888.3264893>
- [4] Chuadhry Mujeeb Ahmed, Aditya P. Mathur, and Martin Ochoa. 2020. NoiSense Print: Detecting Data Integrity Attacks on Sensor Measurements Using Hardware-Based Fingerprints. *ACM Trans. Priv. Secur.* 24, 1, Article 2 (2020). <https://doi.org/10.1145/3410447>
- [5] Chuadhry Mujeeb Ahmed, Gauthama Raman MR, and Aditya P. Mathur. 2020. Challenges in machine learning based approaches for real-time anomaly detection in industrial control systems. In *Proceedings of the 6th ACM on Cyber-Physical System Security Workshop*.
- [6] Chuadhry Mujeeb Ahmed, Carlos Murguia, and Justin Ruths. 2017. Model-Based Attack Detection Scheme for Smart Water Distribution Networks. In *Proceedings of the 2017 ACM on Asia Conference on Computer and Communications Security (ASIA CCS '17)*. <https://doi.org/10.1145/3052973.3053011>
- [7] Chuadhry Mujeeb Ahmed, Venkata Reddy Palleti, and Aditya P. Mathur. 2017. WADI: A Water Distribution Testbed for Research in the Design of Secure Cyber Physical Systems. In *Proceedings of the 3rd International Workshop on Cyber-Physical Systems for Smart Water Networks (CySWATER '17)*. <https://doi.org/10.1145/3055366.3055375>
- [8] Chuadhry Mujeeb Ahmed, Jay Prakash, Rizwan Qadeer, Anand Agrawal, and Jianying Zhou. 2020. Process Skew: Fingerprinting the Process for Anomaly Detection in Industrial Control Systems. In *Proceedings of the 13th ACM Conference on Security and Privacy in Wireless and Mobile Networks (WiSec '20)*. <https://doi.org/10.1145/3395351.3399364>
- [9] Chuadhry Mujeeb Ahmed, Jianying Zhou, and Aditya P. Mathur. 2018. Noise Matters: Using Sensor and Process Noise Fingerprint to Detect Stealthy Cyber Attacks and Authenticate Sensors in CPS. In *Proceedings of the 34th Annual Computer Security Applications Conference (ACSAC '18)*. <https://doi.org/10.1145/3274694.3274748>
- [10] Cesare Alippi, Stavros Ntalampiras, and Manuel Roveri. 2017. Model-Free Fault Detection and Isolation in Large-Scale Cyber-Physical Systems. *IEEE Transactions on Emerging Topics in Computational Intelligence* 1, 1 (2017). <https://doi.org/10.1109/TETCI.2016.2641452>
- [11] Abdulmohsen Almalawi, Xinghuo Yu, Zahir Tari, Adil Fahad, and Ibrahim Khalil. 2014. An unsupervised anomaly-based detection approach for integrity attacks on SCADA systems. *Computers & Security* 46 (2014). <https://doi.org/10.1016/j.cose.2014.07.005>
- [12] Simon D. Duque Anton, Anna Pia Lohfink, Christoph Garth, and Hans Dieter Schotten. 2019. Security in Process: Detecting Attacks in Industrial Process Data (CECC 2019). <https://doi.org/10.1145/3360664.3360669>
- [13] S. D. D. Anton and H. D. Schotten. 2020. Intrusion Detection in Binary Process Data: Introducing the Hamming-distance to Matrix Profiles. In *2020 IEEE 21st International Symposium on "A World of Wireless, Mobile and Multimedia Networks" (WoWMoM)*. <https://doi.org/10.1109/WoWMoM49955.2020.00065>
- [14] S. D. D. Anton, S. Sinha, and H. Dieter Schotten. 2019. Anomaly-based Intrusion Detection in Industrial Data with SVM and Random Forests. In *2019 International Conference on Software, Telecommunications and Computer Networks (SoftCOM)*. <https://doi.org/10.23919/SOFTCOM.2019.8903672>
- [15] 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). <https://doi.org/10.1016/j.ijcip.2020.100377>
- [16] 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 (CCS '18)*. <https://doi.org/10.1145/3243734.3243781>
- [17] Frederik Armknecht, Paul Walther, Gene Tsudik, Martin Beck, and Thorsten Strufe. 2020. ProMACs: progressive and resynchronizing MACs for continuous efficient authentication of message streams. In *Proceedings of the 2020 ACM SIGSAC Conference on Computer and Communications Security*.
- [18] Vaibhav Bajpai, Anna Brunstrom, Anja Feldmann, Wolfgang Kellerer, Aiko Pras, Henning Schulzrinne, Georgios Smaragdakis, Matthias Wählisch, and Klaus Wehrle. 2019. The Dagstuhl Beginners Guide to Reproducibility for Experimental Networking Research. *SIGCOMM Comput. Commun. Rev.* 49, 1 (2019).
- [19] R. R. Barbosa, R. Sadre, and A. Pras. 2012. Towards periodicity based anomaly detection in SCADA networks. In *Proceedings of 2012 IEEE 17th International Conference on Emerging Technologies Factory Automation (ETFA 2012)*. <https://doi.org/10.1109/ETFA.2012.6288888>

- //doi.org/10.1109/ETFA.2012.6489745
- [20] Partha P Biswas, Heng Chuan Tan, Qingbo Zhu, Yuan Li, Daisuke Mashima, and Binbin Chen. 2019. A synthesized dataset for cybersecurity study of IEC 61850 based substation. In *2019 IEEE International Conference on Communications, Control, and Computing Technologies for Smart Grids (SmartGridComm)*. IEEE.
- [21] A. Carcano, A. Coletta, M. Guglielmi, M. Masera, I. Nai Fovino, and A. Trombetta. 2011. A Multidimensional Critical State Analysis for Detecting Intrusions in SCADA Systems. *IEEE Transactions on Industrial Informatics* 7, 2 (2011). <https://doi.org/10.1109/TII.2010.2099234>
- [22] Andrea Carcano, Igor Nai Fovino, Marcello Masera, and Alberto Trombetta. 2010. State-Based Network Intrusion Detection Systems for SCADA Protocols: A Proof of Concept. In *Critical Information Infrastructures Security*.
- [23] Marco Caselli, Emmanuele Zambon, and Frank Kargl. 2015. Sequence-Aware Intrusion Detection in Industrial Control Systems. In *Proceedings of the 1st ACM Workshop on Cyber-Physical System Security (CPSS '15)*. <https://doi.org/10.1145/2732198.2732200>
- [24] Marco Caselli, Emmanuele Zambon, Jonathan Petit, and Frank Kargl. 2015. Modeling message sequences for intrusion detection in industrial control systems. In *Critical Infrastructure Protection IX*. Springer.
- [25] John Henry Castellanos and Jianying Zhou. 2019. A Modular Hybrid Learning Approach for Black-Box Security Testing of CPS. In *Applied Cryptography and Network Security*.
- [26] Yuqi Chen, Christopher M. Poskitt, and Jun Sun. 2018. Learning from Mutants: Using Code Mutation to Learn and Monitor Invariants of a Cyber-Physical System. In *2018 IEEE Symposium on Security and Privacy (SP)*. <https://doi.org/10.1109/SP.2018.00016>
- [27] Steven Cheung, Bruno Dutertre, Martin Fong, Ulf Lindqvist, Keith Skinner, and Alfonso Valdes. 2007. Using model-based intrusion detection for SCADA networks. In *Proceedings of the SCADA security scientific symposium*.
- [28] Hongjun Choi, Wen-Chuan Lee, Youssa Aafer, Fan Fei, Zhan Tu, Xiangyu Zhang, Dongyan Xu, and Xinyan Deng. 2018. Detecting Attacks Against Robotic Vehicles: A Control Invariant Approach (CCS '18). <https://doi.org/10.1145/3243734.3243752>
- [29] Justyna Chromik. 2017. <https://github.com/jjchromik/manipulateTraces/>.
- [30] Justyna Chromik. 2019. <https://github.com/jjchromik/intravis>.
- [31] Ankang Chu, Yingxu Lai, and Jing Liu. 2019. Industrial Control Intrusion Detection Approach Based on Multiclassification GoogLeNet-LSTM Model. *Security and Communication Networks* 2019 (2019).
- [32] B. Claise. 2004. *Cisco Systems NetFlow Services Export Version 9*. RFC 3954. RFC Editor. <http://www.rfc-editor.org/rfc/rfc3954.txt> <http://www.rfc-editor.org/rfc/rfc3954.txt>
- [33] Tanmoy Kanti Das, Sridhar Adepu, and Jianying Zhou. 2020. Anomaly detection in Industrial Control Systems using Logical Analysis of Data. *Computers & Security* 96 (2020). <https://doi.org/10.1016/j.cose.2020.101935>
- [34] Derui Ding, Qing-Long Han, Yang Xiang, Xiaohua Ge, and Xian-Ming Zhang. 2018. A survey on security control and attack detection for industrial cyber-physical systems. *Neurocomputing* 275 (2018). <https://doi.org/10.1016/j.neucom.2017.10.009>
- [35] James J Downs and Ernest F Vogel. 1993. A plant-wide industrial process control problem. *Computers & chemical engineering* 17, 3 (1993).
- [36] Alessandro Erba and Nils Ole Tippenhauer. 2020. No Need to Know Physics: Resilience of Process-based Model-free Anomaly Detection for Industrial Control Systems. *CoRR* abs/2012.03586 (2020). [arXiv:2012.03586](https://arxiv.org/abs/2012.03586) <https://arxiv.org/abs/2012.03586>
- [37] L. Faramondi, F. Flammini, S. Guarino, and R. Setola. 2021. A Hardware-in-the-Loop Water Distribution Testbed Dataset for Cyber-Physical Security Testing. *IEEE Access* 9 (2021). <https://doi.org/10.1109/ACCESS.2021.3109465>
- [38] C. Feng, T. Li, and D. Chana. 2017. Multi-level Anomaly Detection in Industrial Control Systems via Package Signatures and LSTM Networks. In *2017 47th Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN)*. <https://doi.org/10.1109/DSN.2017.34>
- [39] Cheng Feng, Venkata Reddy Palleti, Aditya Mathur, and Deepthi Chana. 2019. A Systematic Framework to Generate Invariants for Anomaly Detection in Industrial Control Systems. In *NDSS*.
- [40] Benedikt Ferling, Justyna Chromik, Marco Caselli, and Anne Remke. 2018. Intrusion Detection for Sequence-Based Attacks with Reduced Traffic Models. In *Measurement, Modelling and Evaluation of Computing Systems*.
- [41] Fraunhofer FKIE-CAD. 2022. IPAL - Datasets. https://github.com/fkie-cad/ipal_datasets.
- [42] Fraunhofer FKIE-CAD. 2022. IPAL - Industrial Intrusion Detection Framework. https://github.com/fkie-cad/ipal_ids_framework.
- [43] Fraunhofer FKIE-CAD. 2022. IPAL - Transcriber. https://github.com/fkie-cad/ipal_transcriber.
- [44] Open Information Security Foundation. 2021. Suricata. <https://suricata.io/>
- [45] The Apache Software Foundation. 2020. *PLC4X*. <https://plc4x.apache.org/>
- [46] I. N. Fovino, A. Carcano, T. De Lacheze Murel, A. Trombetta, and M. Masera. 2010. Modbus/DNP3 State-Based Intrusion Detection System. In *2010 24th IEEE International Conference on Advanced Information Networking and Applications*. <https://doi.org/10.1109/AINA.2010.86>
- [47] Wei Gao and Thomas H Morris. 2014. On cyber attacks and signature based intrusion detection for modbus based industrial control systems. *Journal of Digital Forensics, Security and Law* 9, 1 (2014).
- [48] Hamid Reza Ghaeini and Nils Ole Tippenhauer. 2016. HAMIDS: Hierarchical Monitoring Intrusion Detection System for Industrial Control Systems (CPS-SPC '16). <https://doi.org/10.1145/2994487.2994492>
- [49] Jairo Giraldo, David Urbina, Alvaro Cardenas, Junia Valente, Mustafa Faisal, Justin Ruths, Nils Ole Tippenhauer, Henrik Sandberg, and Richard Candell. 2018. A Survey of Physics-Based Attack Detection in Cyber-Physical Systems. *ACM Comput. Surv.* 51, 4 (2018). <https://doi.org/10.1145/3203245>
- [50] Jonathan Goh, Sridhar Adepu, Khurum Nazir Junejo, and Aditya Mathur. 2016. A dataset to support research in the design of secure water treatment systems. In *International conference on critical information infrastructures security*. Springer.
- [51] Niv Goldenberg and Avishai Wool. 2013. Accurate modeling of Modbus/TCP for intrusion detection in SCADA systems. *International Journal of Critical Infrastructure Protection* 6, 2 (2013). <https://doi.org/10.1016/j.ijcip.2013.05.001>
- [52] Isaías González, Antonio José Calderón, and José María Portalo. 2021. Innovative multi-layered architecture for heterogeneous automation and monitoring systems: Application case of a photovoltaic smart microgrid. *Sustainability* 13, 4 (2021).
- [53] Dina Hadziomanović, Robin Sommer, Emmanuele Zambon, and Pieter H. Hartel. 2014. Through the Eye of the PLC: Semantic Security Monitoring for Industrial Processes (ACSAC '14). <https://doi.org/10.1145/2664243.2664277>
- [54] Zhongyuan Hau and Emil C. Lupu. 2019. Exploiting Correlations to Detect False Data Injections in Low-Density Wireless Sensor Networks. In *Proceedings of the 5th on Cyber-Physical System Security Workshop (CPSS '19)*. <https://doi.org/10.1145/3327961.3329530>
- [55] Kevin E. Hemsley and Dr. Ronald E. Fisher. 2018. History of Industrial Control System Cyber Incidents. (2018). <https://doi.org/10.2172/1505628>
- [56] Martin Henze, Lennart Bader, Julian Filter, Olav Lamberts, Simon Ofner, and Dennis van der Velde. 2020. Poster: Cybersecurity Research and Training for Power Distribution Grids – A Blueprint. In *Proceedings of the 2020 ACM SIGSAC Conference on Computer and Communications Security (CCS) - Poster Session*. <https://doi.org/10.1145/3372297.3420016>
- [57] Jens Hiller, Martin Henze, Martin Serror, Eric Wagner, Jan Niklas Richter, and Klaus Wehrle. 2018. Secure low latency communication for constrained industrial IoT scenarios. In *2018 IEEE 43rd Conference on Local Computer Networks (LCN)*. IEEE.
- [58] Yan Hu, An Yang, Hong Li, Yuyan Sun, and Limin Sun. 2018. A survey of intrusion detection on industrial control systems. *International Journal of Distributed Sensor Networks* 14, 8 (2018).
- [59] Abdulmalik Humayed, Jingqiang Lin, Fengjun Li, and Bo Luo. 2017. Cyber-Physical Systems Security—A Survey. *IEEE Internet of Things Journal* 4, 6 (2017). <https://doi.org/10.1109/JIOT.2017.2703172>
- [60] Won-Seok Hwang, Jeong-Han Yun, Jonguk Kim, and Hyoung Chun Kim. 2019. Time-Series Aware Precision and Recall for Anomaly Detection: Considering Variety of Detection Result and Addressing Ambiguous Labeling. In *Proceedings of the 28th ACM International Conference on Information and Knowledge Management*. ACM.
- [61] J. Inoue, Y. Yamagata, Y. Chen, C. M. Poskitt, and J. Sun. 2017. Anomaly Detection for a Water Treatment System Using Unsupervised Machine Learning. In *2017 IEEE International Conference on Data Mining Workshops (ICDMW)*. <https://doi.org/10.1109/ICDMW.2017.149>
- [62] M. Iturbe. 2018. PASAD. <https://github.com/mikeliturbe/pasad>.
- [63] A. Jones, Z. Kong, and C. Belta. 2014. Anomaly detection in cyber-physical systems: A formal methods approach. In *53rd IEEE Conference on Decision and Control*. <https://doi.org/10.1109/CDC.2014.7039487>
- [64] Khurum Nazir Junejo and Jonathan Goh. 2016. 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 (CPSS '16)*. <https://doi.org/10.1145/2899015.2899016>
- [65] Mohamad Kaouk, Jean-Marie Flaus, Marie-Laure Potet, and Roland Groz. 2019. A Review of Intrusion Detection Systems for Industrial Control Systems. In *2019 6th International Conference on Control, Decision and Information Technologies (CoDIT)*. IEEE. <https://doi.org/10.1109/CoDIT.2019.8820602>
- [66] Abdullah Khalili and Ashkan Sami. 2015. SysDetect: A systematic approach to critical state determination for Industrial Intrusion Detection Systems using Apriori algorithm. *Journal of Process Control* 32 (2015). <https://doi.org/10.1016/j.jprocont.2015.04.005>
- [67] J. Kim. 2019. [swat-seq2seq](https://github.com/jukworks/swat-seq2seq). <https://github.com/jukworks/swat-seq2seq/>.
- [68] Jonguk Kim, Jeong-Han Yun, and Hyoung Chun Kim. 2020. Anomaly Detection for Industrial Control Systems Using Sequence-to-Sequence Neural Networks. In *Computer Security*.
- [69] Jun-Sung Kim, Seong Min So, Joong-Tae Kim, Jung-Won Cho, Hee-Jeong Park, Fauzan Hanif Jufri, and Jaesung Jung. 2019. Microgrids platform: A design and implementation of common platform for seamless microgrids operation. *Electric Power Systems Research* 167 (2019). <https://doi.org/10.1016/j.epsr.2018.10.019>

- [70] I. Kiss, B. Genge, and P. Haller. 2015. A clustering-based approach to detect cyber attacks in process control systems. In *2015 IEEE 13th International Conference on Industrial Informatics (INDIN)*. <https://doi.org/10.1109/INDIN.2015.7281725>
- [71] Marcel Kneib and Christopher Huth. 2018. Scission: Signal Characteristic-Based Sender Identification and Intrusion Detection in Automotive Networks. In *Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security (CCS '18)*, 14 pages. <https://doi.org/10.1145/3243734.3243751>
- [72] Marcel Kneib, Oleg Schell, and Christopher Huth. 2020. EASI: Edge-Based Sender Identification on Resource-Constrained Platforms for Automotive Networks. In *NDSS*.
- [73] Z. Kong, A. Jones, and C. Belta. 2017. Temporal Logics for Learning and Detection of Anomalous Behavior. *IEEE Trans. Automat. Control* 62, 3 (2017). <https://doi.org/10.1109/TAC.2016.2585083>
- [74] Zhaodan Kong, Austin Jones, Ana Medina Ayala, Ebru Aydin Gol, and Calin Belta. 2014. Temporal Logic Inference for Classification and Prediction from Data. In *Proceedings of the 17th International Conference on Hybrid Systems: Computation and Control (HSCC '14)*. <https://doi.org/10.1145/2562059.2562146>
- [75] Moshe Kravchik and Asaf Shabtai. 2018. Detecting Cyber Attacks in Industrial Control Systems Using Convolutional Neural Networks. In *Proceedings of the 2018 Workshop on Cyber-Physical Systems Security and Privacy (CPS-SPC '18)*. <https://doi.org/10.1145/3264888.3264896>
- [76] 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 (ASIA CCS '15)*. <https://doi.org/10.1145/2714576.2714599>
- [77] Viacheslav Kulik and Ruslan Kirichek. 2018. The Heterogeneous Gateways in the Industrial Internet of Things. In *2018 10th International Congress on Ultra Modern Telecommunications and Control Systems and Workshops (ICUMT)*. <https://doi.org/10.1109/ICUMT.2018.8631232>
- [78] Dominik Kus, Eric Wagner, Jan Pennekamp, Konrad Wolsing, Ina Berenice Fink, Markus Dahlmans, Klaus Wehrle, and Martin Henze. 2022. A False Sense of Security? Revisiting the State of Machine Learning-Based Industrial Intrusion Detection. In *Proceedings of the 8th ACM Cyber-Physical System Security Workshop (CPSS)*. <https://doi.org/10.1145/3494107.3522773>
- [79] Chih-Yuan Lin and Simin Nadjm-Tehrani. 2018. Understanding IEC-60870-5-104 Traffic Patterns in SCADA Networks. In *Proceedings of the 4th ACM Workshop on Cyber-Physical System Security (CPSS '18)*. <https://doi.org/10.1145/3198458.3198460>
- [80] 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)*. <https://www.usenix.org/conference/raid2019/presentation/lin>
- [81] Chih-Yuan Lin, Simin Nadjm-Tehrani, and Mikael Asplund. 2017. Timing-Based Anomaly Detection in SCADA Networks. In *International Conference on Critical Information Infrastructures Security*.
- [82] Qin Lin, Sridha Adepu, Sicco Verwer, and Aditya Mathur. 2018. TABOR: A Graphical Model-Based Approach for Anomaly Detection in Industrial Control Systems. In *Proceedings of the 2018 Asia Conference on Computer and Communications Security (ASIACCS '18)*. <https://doi.org/10.1145/3196494.3196546>
- [83] R. Lopez Perez, F. Adamsky, R. Soua, and T. Engel. 2018. Machine Learning for Reliable Network Attack Detection in SCADA Systems. In *2018 17th IEEE International Conference On Trust, Security And Privacy In Computing And Communications/ 12th IEEE International Conference On Big Data Science And Engineering (TrustCom/BigDataSE)*. <https://doi.org/10.1109/TrustCom/BigDataSE.2018.00094>
- [84] George Loukas, Eirini Karapistoli, Emmanouil Panaousis, Panagiotis Sarigiannidis, Anatolij Bezemskij, and Tuan Vuong. 2019. A taxonomy and survey of cyber-physical intrusion detection approaches for vehicles. *Ad Hoc Networks* 84 (2019). <https://doi.org/10.1016/j.adhoc.2018.10.002>
- [85] L. A. Maglaras and J. Jiang. 2014. Intrusion detection in SCADA systems using machine learning techniques. In *2014 Science and Information Conference*. <https://doi.org/10.1109/SAL.2014.6918252>
- [86] M. H. Monzer, K. Beydoun, and J. FLAUS. 2019. Model based rules generation for Intrusion Detection System for industrial systems. In *2019 International Conference on Control, Automation and Diagnosis (ICCAD)*. <https://doi.org/10.1109/ICCAD46983.2019.9037882>
- [87] Thomas H Morris, Zach Thornton, and Ian Turnipseed. 2015. Industrial control system simulation and data logging for intrusion detection system research. *7th annual southeastern cyber security summit* (2015).
- [88] David Myers, Suriadi Suriadi, Kenneth Radke, and Ernest Foo. 2018. Anomaly detection for industrial control systems using process mining. *Computers & Security* 78 (2018). <https://doi.org/10.1016/j.cose.2018.06.002>
- [89] Patric Nader, Paul Honeine, and Pierre Beausery. 2014. L_p -norms in One-Class Classification for Intrusion Detection in SCADA Systems. *IEEE Transactions on Industrial Informatics* 10, 4 (2014). <https://doi.org/10.1109/TII.2014.2330796>
- [90] Felix O Olowononi, Danda B Rawat, and Chunmei Liu. 2020. Resilient machine learning for networked cyber physical systems: A survey for machine learning security to securing machine learning for cps. *IEEE Communications Surveys & Tutorials* 23, 1 (2020).
- [91] Shengyi Pan, Thomas Morris, and Uttam Adhikari. 2015. Developing a Hybrid Intrusion Detection System Using Data Mining for Power Systems. *IEEE Transactions on Smart Grid* 6, 6 (2015). <https://doi.org/10.1109/TSG.2015.2409775>
- [92] Ranjit Panigrahi and Samarjeet Borah. 2018. A detailed analysis of CICIDS2017 dataset for designing Intrusion Detection Systems. *International Journal of Engineering & Technology* 7, 3, 24 (2018).
- [93] S. Ponomarev and T. Atkinson. 2016. Industrial Control System Network Intrusion Detection by Telemetry Analysis. *IEEE Transactions on Dependable and Secure Computing* 13, 2 (2016). <https://doi.org/10.1109/TDSC.2015.2443793>
- [94] The Zeek Project. 2021. Zeek. <https://zeek.org/>
- [95] 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)*.
- [96] Panagiotis Radoglou-Grammatikis, Panagiotis Sarigiannidis, George Efstathiopoulos, Paris-Alexandros Karypidis, and Antonios Sarigiannidis. 2020. DIDEROT: An Intrusion Detection and Prevention System for DNP3-Based SCADA Systems. In *Proceedings of the 15th International Conference on Availability, Reliability and Security (ARES '20)*. <https://doi.org/10.1145/3407023.3409314>
- [97] R. Raj. 2020. PyPASAD. <https://github.com/rahulrajpl/PyPASAD/>.
- [98] Daniel Ramotsoela, Adnan Abu-Mahfouz, and Gerhard Hancke. 2018. A survey of anomaly detection in industrial wireless sensor networks with critical water system infrastructure as a case study. *Sensors* 18, 8 (2018).
- [99] Rocio. 2021. ML-NIDS-for-SCADA. <https://github.com/Rocionightwater/ML-NIDS-for-SCADA>.
- [100] Martin Roesch et al. 1999. Snort: Lightweight intrusion detection for networks. In *Lisa*, Vol. 99.
- [101] Ondřej Ryšavý and Petr Matoušek. 2021. A Network Traffic Processing Library for ICS Anomaly Detection. In *7th Conference on the Engineering of Computer Based Systems (ECBS 2021)*. <https://doi.org/10.1145/3459960.3459963>
- [102] Subin Sapkota, A K M Nuhil Mehdy, Stephen Reese, and Hoda Mehrpouyan. 2020. FALCON: Framework for Anomaly Detection in Industrial Control Systems. *Electronics* 9, 8 (2020). <https://doi.org/10.3390/electronics9081192>
- [103] Wenli Shang, Junrong Cui, Ming Wan, Panfeng An, and Peng Zeng. 2016. Modbus Communication Behavior Modeling and SVM Intrusion Detection Method. In *Proceedings of the 6th International Conference on Communication and Network Security (ICCNS '16)*. <https://doi.org/10.1145/3017971.3017978>
- [104] Hyeok-Ki Shih, Woomyo Lee, Jeong-Han Yun, and HyoungChun Kim. 2020. HAI 1.0: HIL-based Augmented ICS Security Dataset. In *13th USENIX Workshop on Cyber Security Experimentation and Test (CSET 20)*. <https://www.usenix.org/conference/cset20/presentation/shih>
- [105] J. Smith. 2016. ICS-pcap. <https://github.com/automayt/ICS-pcap/blob/master/AdditionalCaptures/4SICS-GeekLounge-151022/4SICS-GeekLounge-151022.pcap>.
- [106] Robin Sommer. 2003. Bro: An open source network intrusion detection system. In *Security, E-learning, E-Services, 17. DFN-Arbeitsstagung über Kommunikation-snetze*.
- [107] Robin Sommer, Johanna Amann, and Seth Hall. 2015. Spicy: A unified deep packet inspection framework dissecting all your data. *Technical report, Tech. rep., ICSI, 2015. TR-15-004* (2015).
- [108] Robin Sommer and Vern Paxson. 2010. Outside the closed world: On using machine learning for network intrusion detection. In *2010 IEEE symposium on security and privacy*. IEEE.
- [109] Bruno Sousa, Tiago Cruz, Vasco Pereira, and Miguel Arieiro. 2021. Denial of Service and Man in The Middle attacks in Programmable Logic Controllers. <https://doi.org/10.21227/mewp-g646>
- [110] William Stallings and Lawrie Brown. 2015. *Computer security: principles and practice* (third edition ed.). Pearson.
- [111] Federico Turrin, Alessandro Erba, Nils Ole Tippenhauer, and Mauro Conti. 2020. A Statistical Analysis Framework for ICS Process Datasets. In *Proceedings of the 2020 Joint Workshop on CPS&IoT Security and Privacy (CPSIoTSEC'20)*.
- [112] Federico Turrin, Alessandro Erba, Nils Ole Tippenhauer, and Mauro Conti. 2020. A Statistical Analysis Framework for ICS Process Datasets. In *Proceedings of the 2020 Joint Workshop on CPS&IoT Security and Privacy*.
- [113] Rafael Uetz, Christian Hemminghaus, Louis Hackländer, Philipp Schlipper, and Martin Henze. 2021. Reproducible and Adaptable Log Data Generation for Sound Cybersecurity Experiments. In *Annual Computer Security Applications Conference (ACSAC '21)*.
- [114] 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*.
- [115] A. Valdes and S. Cheung. 2009. Communication pattern anomaly detection in process control systems. In *2009 IEEE Conference on Technologies for Homeland Security*. <https://doi.org/10.1109/THS.2009.5168010>
- [116] Cybersecurity Ventures. 2019. 2019 Official Annual Cybercrime Report.

- [117] VirusTotal. 2021. YARA. <https://virustotal.github.io/yara/>
- [118] Eric Wagner, Jan Bauer, and Martin Henze. 2022. Take a Bite of the Reality Sandwich: Revisiting the Security of Progressive Message Authentication Codes. (2022). <https://doi.org/10.1145/3507657.3528539>
- [119] Jie Wang. 2009. The Art of Intrusion Detection. In *Computer Network Security: Theory and Practice*.
- [120] Patrick Wiener, Philipp Zehnder, and Dominik Riemer. 2020. Managing geo-distributed stream processing pipelines for the IIoT with StreamPipes edge extensions. In *Proceedings of the 14th ACM International Conference on Distributed and Event-based Systems*.
- [121] Konrad Wolsing, Eric Wagner, and Martin Henze. 2020. Poster: Facilitating Protocol-independent Industrial Intrusion Detection Systems. In *Proceedings of the 2020 ACM SIGSAC Conference on Computer and Communications Security*.
- [122] Y. Yang, H. Xu, L. Gao, Y. Yuan, K. McLaughlin, and S. Sezer. 2017. Multidimensional Intrusion Detection System for IEC 61850-Based SCADA Networks. *IEEE Transactions on Power Delivery* 32, 2.
- [123] Man-Ki Yoon and Gabriela F Ciocarlie. 2014. Communication pattern monitoring: Improving the utility of anomaly detection for industrial control systems. In *Proceedings of the NDSS workshop on Security of Emerging Network Technologies (SENT'14)*.
- [124] Jeong-Han Yun, Yoonho Hwang, Woomyo Lee, Hee-Kap Ahn, and Sin-Kyu Kim. 2018. Statistical Similarity of Critical Infrastructure Network Traffic Based on Nearest Neighbor Distances. In *Research in Attacks, Intrusions, and Defenses*.
- [125] Philipp Zehnder, P. Wiener, T. Straub, and D. Riemer. 2020. StreamPipes Connect: Semantics-Based Edge Adapters for the IIoT. In *ESWC*. Springer.
- [126] Sven Zemanek, Immanuel Hacker, Konrad Wolsing, Eric Wagner, Martin Henze, and Martin Serror. 2022. PowerDuck: A GOOSE Data Set of Cyberattacks in Substations. In *Proceedings of the 15th Workshop on Cyber Security Experimentation and Test (CSET)*.
- [127] Feng Zhang, Min Liu, Zhuo Zhou, and Weiming Shen. 2016. An IoT-based online monitoring system for continuous steel casting. *IEEE Internet of Things Journal* 3, 6 (2016).
- [128] Chunjie Zhou, Shuang Huang, Naixue Xiong, Shuang-Hua Yang, Huiyun Li, Yuanqing Qin, Xuan Li 2015. Design and Analysis of Multimodel-Based Anomaly Intrusion Detection Systems in Industrial Process Automation. *IEEE Transactions on Systems, Man, and Cybernetics: Systems* 45, 10.

A PERFORMANCE COMPARISON OF PROCESS STATE-AWARE IIDSS

In Sec. 6.2, we used IPAL to compare three different process state-aware IIDSSs against each other on the SWaT dataset [50]. To study their potential to generalize to new industrial domains, we applied these IIDSSs to two new datasets (WADI [7] and HAI [104]). Our analysis in Fig. 5 depicts the alerts of these three approaches visually, which provides a decent intuition on how these approaches work in general. To additionally confirm our observations objectively and in further detail here, Tab. 9 summarizes different metrics for all evaluated settings from that case study.

Detected attacks counts the number of scenarios overlapping with an alarm and false alarms are all alarms not overlapping with an attack. Penalty score accumulates for all true positives the time in seconds with regions not labeled as attack within the dataset, based on the definition by Lin et al. [82]. The remaining traditional metrics (accuracy, precision, recall, and f1-score) were calculated one by one for each entry within the evaluation dataset. While these metrics may be valuable for randomized datasets, they struggle to infer the quality of time-aware datasets [60], especially since approaches such as PASAD or Seq2Seq-NN detect attacks with a slight delay.

Dataset	SWaT [50]			WADI [7]			HAI [104]			
	IIDS	S	P	T	S	P	T	S	P	T
Det. Attacks	10	6	9	8	2	4	2	2	35	
False Alarms	17	18	0	6	3	0	5	22	23	
Penalty Score	149	9374	1765	153	13177	5792	0	11669	214100	
Accuracy [%]	87.71	87.59	88.33	94.29	81.53	82.31	98.46	90.33	26.68	
Precision [%]	43.25	49.13	68.71	51.92	5.14	30.05	28.93	0.66	1.18	
Recall [%]	4.00	63.90	7.09	13.99	12.60	24.94	4.26	3.79	60.12	
F1-Score [%]	7.32	55.55	12.86	22.04	7.30	27.26	7.42	1.12	2.31	

Table 9: IPAL enables the comparison of process state-aware IIDSSs Seq2Seq-NN (S), PASAD (P), and TABOR (T) on the SWaT dataset and in new scenarios on WADI, and HAI.

IIDS Type	Process state-aware			Communication-based			
	IIDS	S	P	T	IAT Mean	IAT Range	DTMC
Det. Attacks	2	7	9	6	9	7	
False Alarms	3	2	0	2330	3390	59	
Penalty Score	58	694	765	3.67	2.48	0.06	
Accuracy [%]	73.99	67.32	67.22	75.39	75.39	75.40	
Precision [%]	19.44	27.49	30.03	39.25	42.63	80.58	
Recall [%]	0.96	17.96	22.41	0.06	0.09	0.01	
F1-Score [%]	1.83	21.72	25.67	0.12	0.19	0.02	

Table 10: For the first time, IPAL enables the simultaneous use of process state-aware and communication-based IIDSSs. To showcase the capabilities of both research branches, we trained representative IIDSSs study on the WDT dataset [37].

B COOPERATIVE COMMUNICATION-BASED AND PROCESS STATE-AWARE IIDSSs

In Sec. 6.3, we used IPAL to compare three process state-aware IIDSSs (TABOR [82], Seq2Seq-NN [68] & PASAD [16]) against the complementary research branch of communication-based IIDSSs represented by the IAT [81] and DTMC [40] approaches. To study which type of alerts these two IIDSS directions detect best, we applied them to the WDT dataset [37], as it consists of network and physical attacks. As before, we provide further details by summarizing the different metrics for this case study (cf. Fig. 6) in Tab. 10.

For a short description of the used metrics, refer to the previous Appx. A. Note that process state-aware IIDSSs alert over time-ranges while communication-based IIDSSs mark individual network packets. This results in poor performance of all communication-based IIDSSs w.r.t. recall and the F1-score since this type of IIDSS emits short alerts while the metrics favor coverage over the entire attack’s range. Again, as described in Sec. 6.3, the circumventable false alerts from the IAT range and DTMC approach (cf. grey annotations in Fig. 6) are neglected for this evaluation.