A Methodology to Measure the "Cost" of CPS Attacks: Not all CPS Networks are Created Equal

Martin Rosso, Luca Allodi, Emmanuele Zambon, Jerry den Hartog
Department of Mathematics and Computer Science
Eindhoven University of Technology
Eindhoven, The Netherlands
{*m.j.rosso, l.allodi, e.zambon.n.mazzocato, j.d.hartog*}*@tue.nl*

***Abstract*—Cyber-Physical Systems (CPS) are (connected) computer systems used to monitor and control physical processes using digital control programs. Cyberattacks targeting CPS can cause physical impact with potentially devastating consequences. While some past attacks required expert CPS knowledge (e.g., Stuxnet), other attacks could be implemented by anyone, solely with pure IT knowledge. Understanding what causes these differences is essential in effectively defending CPS, however, as of now, there is no way of qualifying let alone quantifying them. In this paper, we first define a notion of (non-monetary) attack "cost" focusing on the required CPS-specific attacker knowledge. We then identify several context factors that may influence this cost and, finally, provide a methodology to analyze the relation between attack cost and CPS-context using past cyberattacks. To validate the methodology in a reproducible way, we apply it to publicly reported CPS incidents with physical impact. Though this constitutes only a small set of attacks, our methodology is able to find correlations between context factors and the attack cost, as well as significant differences in context factors between CPS domains.**

***Index Terms*—attacker capabilities, attacker cost, attacker knowledge, cyber-physical system, industrial control system**

1. Introduction

Cyber-Physical Systems (CPS) bridge the digital and physical world by monitoring and controlling physical processes using sensors, actuators, and digital control programs. As a result, cyberattacks targeting these systems can impact the physical world. However, the necessary expertise or knowledge attackers need to have at their disposal to implement such attacks and the extent of physical impact they obtain varies widely. On the one hand, we observe sophisticated cyberattacks like Stuxnet, a state-sponsored attack tailored to one specific CPS instance, an uranium enrichment facility, designed and implemented by teams of domain experts [1]. On the other hand, we see reports of intruders with no particular education or expertise performing actions, potentially even without awareness of the consequences, such as a 2021 incident where an intruder found a graphical interface for a drinking water treatment facility in Florida exposed over TeamViewer[1] and then arbitrarily changed some values on the graphical user interface [2].

1. remote desktop control software, https://www.teamviewer.com/

It is unsurprising that mounting cyberattacks against an air-gapped top-secret uranium enrichment facility requires more CPS knowledge (and related skills) compared to opportunistic cyberattacks changing values on a graphical user interface exposed on the Internet. Yet these two (and other) real-world examples raise the question of how to capture and explain such differences, motivating the need for more systematic studies characterizing differences and similarities between CPS attacks, particularly which context characteristics influence attack "cost", and whether these context characteristics are influenced by the domain or economical sector of the target CPS.

Currently, we observe two major obstacles that hamper the transfer of security knowledge or best-practices between CPS domains and sectors:
(1.) We observe many standards and best practices to be domain or sector specific. For example, there are many standards tailored to power generation, transport and distribution, and the distribution of gas and fuels [3], including a dedicated ISO 27019[4]. Standards, risk assessment, best-practices, and cyber-threat intelligence (CTI), are tailored to and thus only shared within a single domain (cf. [4]–[7]). Here, the lack of an objective way to characterize, measure, and compare CPS characteristics across CPS-installations and -domains prevents cross-domain information exchange.
(2.) Because there are many common pitfalls, trying to transfer ideas and results across CPS domain-boundaries is prone to errors and oversights. For example, we observe academic publications transferring security solutions from one CPS domain into another but due to the different context between the domains, the *practical applicability* of the proposed solution is drastically limited or even futile in the target domain and context [8], [9].

Research focusing on CPS domain similarities and differences can thus become an enabling factor to improve the general state of CPS security across multiple domains and thus improve the defensive capabilities of CPS operators. Our work provides a first step in this direction, focusing on the following research questions:

(Q1) *What type of knowledge do CPS-attackers need?*
(Q2) *Which characteristics of the attacked CPS influence this required attacker knowledge?*
(Q3) *How do these CPS characteristics differ across CPS domains?*

For Q1, we look at the CPS capabilities that attackers have demonstrated in past attacks (i.e., knowledge and the

skill to effectively apply it). In contrast to CPS capabilities[10], IT-attack capabilities are often readily available e.g., in the form of malware- or access as a service (e.g., [11], [12]). As such, our notion of "cost" only captures *additional* (cf. [13]) CPS-specific capabilities (*Cap*) an attacker needs, be it to attack systems themselves or offer a service to others. Relating this "cost" to context factors, capturing specific characteristics of the attacked CPS and its environment, is the focus of Q2. We score the context factors (*Ctx*) of attacked CPS and check for correlation with the attacker capabilities that were required. With Q3, we focus on consistent similarities and differences in context between CPS-domains.

We define and motive the capabilities and context factors in Section 2 and 3 respectively, before presenting our methodology for measuring and relating them in Section 4. Results are presented in Section 5. Interpretation of our results (Section 6) shows that our methodology is capable of comparing CPS installations across domains and identifying the CPS capabilities required to carry out an attack in different domains. We observe that attackers need relatively little CPS capabilities to obtain physical impact in building automation systems (BAS) and that, in general, the majority of CPS-specific knowledge is needed for ATT&CK Tactics *Data Collection* (and interpretation), and to *Impair Process Control*, while Tactics like *Initial Compromise*, *Execution*, or *Privilege Escalation* rarely require any CPS-specific knowledge. Further, industrial control systems (ICS) and BAS installations show significantly distinct *Ctx* characteristics, with BAS systems scoring significantly lower in most of them, ranging from technical (e.g., network protection) to procedural (e.g., legislation and regulation). Our research artifacts are publicly available in an artifact repository [14].

2. Attack Cost & Attacker Capabilities (*Cap*)

In this section, we address how to capture CPS attack "cost". After analyzing aspects of cost and illustrating why widely established cost metrics do not fit our use case in Section 2.1, we introduce Attacker-Capabilities (***Cap***) by leveraging the MITRE ATT&CK Framework in Section 2.2.

Knowledge is considered a dominant factor in cost. For example, Green, Krotofil, and Abbasi reason about the importance of "process comprehension", i.e., the ability of attackers to know and understand the physical process that they try to manipulate. In their work, they claim most attackers "would [not] have adequate knowledge and resources to achieve targeted operational process manipulation" [15]. Allodi and Etalle consider whether an attacker, after having successfully compromised the computer network connected to a CPS, is willing and able to assume control over the CPS to eventually aim for physical impact. They conclude that economical attackers often have no incentive to do so, due to the additional risk, costs, required knowledge, and required resources [13].

In the IT domain, cybercrime and attacker tools became more sophisticated and proficient over time, thus reducing the technical knowledge required by attackers [16], [17]. Eventually, malware and access to compromised systems became commercial products ("as-a-Service") in underground forums and markets [11], [12]. As of 2020, there is no offer of CPS-attacks or Malware-as-a-Service in cybercrime markets according to Dodson, Beresford, and Thomas. The authors attribute this to the high development costs and unclear or non-existing customer demand [10]. As a result, attackers can outsource the parts of their attack dealing with IT systems, but need to conduct CPS-specific parts of the attack themselves.

There is, however, research focusing on how attackers can reduce the required (CPS-specific) capabilities, e.g., by standardizing tools similar to how IT attacks evolved over time. Green, Derbyshire, Krotofil, *et al.* present a software tool that uses pattern recognition on compiled PLC programs to automatically identify PLC programming software libraries used to produce the code [18]. This way, a vulnerability found in a PLC software library can be exploited by attackers who do not have the necessary background to decompile and understand compiled PLC code. Esquivel-Vargas, Castellanos, Caselli, *et al.* also rely on pattern recognition. They present a method to automatically identify parts of the control process that are likely easy to manipulate while still maximizing impact [19].

From literature, we extract that CPS attackers need CPS-specific capabilities to obtain physical impact and that as of 2020, there is no wide-scale offer that allows attackers to outsource CPS-specific components of attacks. To answer Q1 (and Q2), we need a CPS attack cost metric that must capture all the required knowledge and skills for the entirety of all CPS-specific steps of an attack and not just the "cost" of individual phases (e.g., only the initial exploit of a Windows workstation used by CPS-technicians).

2.1. Background: Attack Cost Metrics

A prominent example of the numerous approaches to measure attack cost is the Common Vulnerability Scoring System (CVSS). CVSS captures, among other factors, how difficult a vulnerability is to exploit, as well as the expected severity of impact on the target system. As CVSS focuses only on the exploit, its score does not adequately estimate attack *cost* in an environment where the exploit does not incur the majority of this cost. For example, numerous CVEs [2] describe lack of authentication or authentication bypass vulnerabilities in CPS, equipment exploitable over the network. The NIST National Vulnerability Database assigns them a CVSS 3.x score of 9.0 or higher ("Critical"), with only one exception rated 8.8 ("high"), as an attacker can easily, without other prerequisites, obtain privileged access over the network to a process control device. Even though there is a multitude of vulnerable CPS equipment connected to the Internet [20]–[22], public reports do not seem to suggest that these systems are systematically targeted by cyberattacks aiming to obtain physical impact. One possible explanation for this mismatch between high CVSS scores and low exploitation rate is that exploit complexity, as measured by CVSS, does not capture the full complexity of a cyberattack against CPS. In fact, it is commonly believed that to obtain physical impact after obtaining access to a CPS, the

2. e.g., CVE-2016-5815, CVE-2019-18250, CVE-2021-22779, CVE-2022-30319, CVE-2022-33139 & CVE-2022-45789

attacker requires expert knowledge of the physical process [13], [15]. Because CVSS does not capture the entirety of the attack and does not include the required attacker knowledge, it does not fulfill our requirements.

Rocchetto and Tippenhauer create threat actor profiles covering, among others, their motivation, their available background knowledge, and available resources, as well as their determination or willingness to actively overcome defense mechanisms deployed at the target [23]. Tommasini, Rosso, Zambon, *et al.* extend their work and compare results obtained for ICS with BAS [24]. However, the categorization is focused on attacker profiles, i.e., it tries to capture more generally which *type* of attacker typically have which type of knowledge or resources at their disposal. Furthermore, their knowledge-dimensions focus on offensive IT-knowledge and insider-specific system knowledge, but do not capture whether *CPS-specific* knowledge is necessary. As a result, their framework is not suitable to analyze and compare the attacker knowledge required for individual cyberattacks.

Other approaches focus on modelling the economics of cybercrime, including monetarian expenses, earnings, and studying attacker decision making for profit maximization [11], [13], [25].

2.2. Cost Metric for Attacker-Capabilities (*Cap*)

As mentioned, CPS attack cost hinges on CPS-specific knowledge and skills an attacker needs to cause physical impact, i.e. Attacker-Capabilities (*Cap*). To be able to capture *Cap* in a structured way, our metric builds on the MITRE ATT&CK Framework. This framework provides a list of ICS Tactics, each with associated Techniques. Analogue to ATT&CK for ICS, *Cap* is constructed from 12 sub-dimensions, one for each ICS Tactic. Values for each *Ctx* sub-dimension are computed as the maximum rating over all relevant ATT&CK ICS Techniques implemented by the attacker. Finally, *CAPABILITY* denotes the highest required *Cap* throughout the entire attack (i.e., the maximum over all Tactics). While the latter summarizes the attacker knowledge and skills required for the entirety of the attack, the per Tactic cost reveals in which attack phases that knowledge is needed. Here, an implemented Technique implying *Process-Mapping Knowledge*, i.e., clear understanding of how this particular control system manipulates this particular physical process, is rated as `high` $(+1)$ while one that involves only *Technical Knowledge* and/or *Process Knowledge* is rated as `medium` (0). `low` (-1) indicates no CPS-specific knowledge is needed at all. A detailed description of these three types of knowledge is provided in Appendix A.2.2.

3. CPS-Context (*Ctx*)

In this section, we introduce the notion of CPS-Context (***Ctx***) as a way to characterize a specific CPS installation. After we summarize literature analyzing factors that influence attack cost and motivate the notion of CPS-Context in Section 3.1, we then extract relevant *Ctx*-factors (or dimensions) from literature that together form the *Ctx* of a CPS in Section 3.2.

Ortiz, Rosso, Zambon, *et al.* [26] compared three CPS network captures from different sectors and observed notable differences among them. Notably, the three physical processes have different requirements on process control and as a result, the deployed CPS differ in network architecture and utilization of network protocols.

However, to the best of our knowledge, there is no methodology to quantify or qualify observable characteristics of CPS deployments to compare them across domains. In fact, most publications, technical standards, or best practices only focus on isolated aspects, e.g., technical security [4], [27], [28], physical process [15], [21], or legal aspects [4], [27]. As a result, there is no comprehensive summary of relevant CPS characteristics, and no overview of the impact of interdisciplinary factors on (cyber)security.

To measure *Ctx*, we require a metric capable of capturing the technical and non-technical *context* that characterizes an operational CPS. The metric needs to be relevant across all different types of CPS, and it needs to be abstract enough to apply to all CPS domains, but granular enough to allow for comparison. Further, we expect that changes in *Ctx* should have some impact on the *Cap* required to attack CPSs with different *Ctx*-factors (and obtain physical impact).

3.1. Background Literature

Over the years, many publications and reports investigated different aspects of CPS, including e.g., CPS technical security or legal aspects.

The 2018 ENISA [29] report on "Good Practices for Security of IoT", is a rare publication naming multiple different factors that determine *Ctx*. The report, however, is too generic to be used as-is to define specific *Ctx*-factors and derive clear rater instructions. We are not aware of any other work presenting a similar overview of CPS specific context factors.

Even though not focused on CPS specifically, there are multiple publications in the field of "Science of Security", aiming to measure the impact of policies, standards, or implementation choices on "security". The publications of Herley and Oorschot and Stolfo, Bellovin, and Evans both describe the need for metrics capable of measuring security and the difficulties associated to create such metric [30], [31]. One of the named difficulties is caused by metrics relying on underlying assumptions that are not always applicable in the real world [30], [31].

3.2. CPS-Context Metric

We need a metric capturing technical and non-technical *Ctx* factors that characterize an operational CPS. The ENISA report [29] fulfills this requirement, but lacks in details. Thus, we extend their *Ctx*-factors and definitions using information from academic publications, international standards, and CPS vendors' product descriptions. We briefly present the eight *Ctx*-factors in Table 1, including references to literature used to extend the *Ctx*-factors presented in the ENISA report. To preempt possible misconceptions, we briefly present the factors we believe to be not self-explanatory. *Device Rarity* is the inverse of "market penetration" and captures devices, protocols, and software libraries. *Contractor Independence* relates to various types of dependencies ranging from

	Ctx	ENISA Security Challenges and Threats [29, ch 2.2 & fig 8] and other sources
1	Physical Protection	physical attack threats [29, fig 8] & [27], [4, ch. 11], [32, tab 1], [28, tab. 1]
2	Device Protection	vulnerable components, legacy industrial control systems, unused functionalities, security updates, secure product lifecycle [29, ch 2.2] & [4], [27], [32, ch. 2.6 f.] MITRE ATT&CK, [28, tab 1]
3	Network Protection	increased connectivity, IT/OT convergence, legacy industrial control systems, insecure protocols [29, ch 2.2] & MITRE ATT&CK , [32, ch 2.4, 2.5, 2.7], [28, tab. 1]
4	Device Rarity	[33]–[35] [36]–[39] [40], [41], [32, ch 2.6.1], MITRE ATT&CK
5	Contractor Independence	supply chain complexity [29, ch 2.2] & [4], [27]
6	Process Complexity	management of processes [29, ch 2.2] & [15], [18], [19], [42]
7	Safety Monitoring	safety aspects [29, ch 2.2] & [43], [44], [45, ch 3.3.1]
8	Legislation & Regulation	legal threats [29, fig 8] & [4], [27]

TABLE 1. OVERVIEW OF LITERATURE RELATED TO CTX-FACTORS

Value	Description
`low`	Context factor does not apply, is not implemented, implemented ineffectively, poorly, or with low maturity.
`medium`	Context factor is implemented according to best-practices; implementation is somewhat effective.
`high`	Context factor is implemented effectively, above and beyond what is considered best-practice or mandatory.

TABLE 2. SUMMARY OF CTX SCORING VALUES

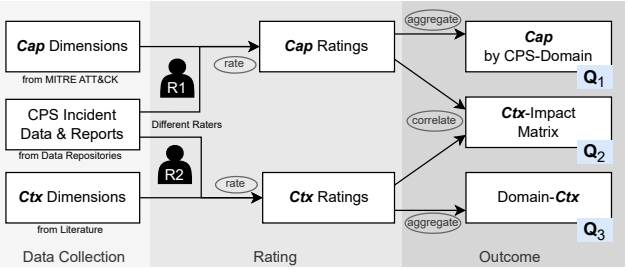

Figure 1. Methodology

software updates to outsourcing device configuration and dependencies on the physical process. *Safety Monitoring* includes all safety-related procedures, ranging from human observations by means of a process expert, to physical interlocks or automated Safety Instrumented Systems (SIS). Lastly, *Legislation & Regulation* is focused on the physical process and captures, among others, availability or quality guarantees of the process (e.g., availability of drinking water or critical infrastructure in general). Regulation related to the safety and security of networked or "smart" devices is not the primary scope.[3] We refer the reader to Appendix A.3 for a detailed explanation of all *Ctx*-factors.

Even though our eight interdisciplinary *Ctx*-factors cover a wide range of aspects, we do not claim completeness, as there may be other (measurable) factors not captured by our characteristics.

Similar to *Cap*, we use a coarse rating schema with a three-value scale: `high` (+1), `medium` (0), and `low` (−1). The metric assigns a value to each *Ctx* factor, based on whether the factor applies to the CPS instance or was correctly and effectively implemented in the respective CPS instance. Table 2 provides a short summary, for details we refer to Appendix A.3. Finally, the *CONTEXT* score for a CPS instance is computed as the average score of all *Ctx*-factors thus ensuring that each factor contributes, without making explicit assumption about the weights of the individual factors.

4. Methodology & Execution

In the previous sections we defined the notions of attack cost and context. Here we describe how to rate those for a dataset of past CPS cyberattacks. We then use the ratings to compare required attacker capabilities (Q1), the impact of context on cost (Q2) and the differences in context between CPS domains (Q3). An overview of our approach is given in Figure 1. We now describe the three phases of our approach in more detail. To illustrate the approach we also already give some details of its application to ICS and BAS systems discussed in Section 5.

4.1. Phase 1: Data Collection

Having already described *Cap* and *Ctx* in the previous two sections, here we address the collection of real-world CPS incidents to obtain historic information about CPS attacks and the respective context of the attacked CPS at the time of attack. To filter out CPS incidents not in scope of our work, we derive the following inclusion criteria from our research questions:

1) the attackers must have obtained physical impact;
2) the physical impact is caused by the attackers, i.e., is not self-inflicted by the victim (e.g., impact is caused by the defender taking responsive measures); [4]
3) incidents must be real-world cyber-attacks (as e.g., security research, penetration tests, or proof of concepts would distort information about the impact);
4) incidents must be relatively recent, i.e., past 2010.[5]

Gathering the required data can be non-trivial, especially if relying purely on publicly available information. Collections of historic CPS incident data can be obtained from public databases [47]–[49]. While public attack databases provide curated lists of incidents, they often do not provide sufficient information to extract the required *Cap* and *Ctx*. By searching the Internet for publicly available reports, additional information can be obtained. However, reports often target a non-technical audience and thus sometimes lack the necessary details or clarity. Occasionally, reports contain contradictory information.

3. Legislation on safety and security of "smart" and IoT devices, e.g., the upcoming EU Cyber Resilience Act, is likely to have an impact on Ctx, more particularly on *Device Protection* (Table 1).

4. We would except from this rule if the attacker forced or tricked the victim to obtain the desired impact (cf. false sensor reporting attacks) but no such case was observed in our datasets.

5. Evans wrote in a Cisco report, that 2008–2009 marks the beginning of IoT. Assuming that attacks changed over time, combined with older reports often lacking necessary details, we chose 2010 as cutoff [46].

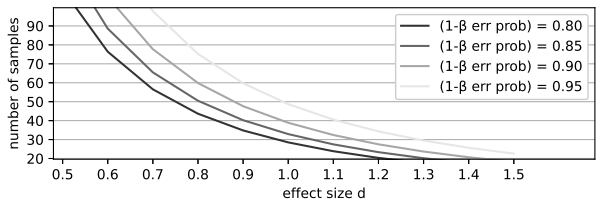

Figure 2. A priori Power Analysis

These issues can be resolved by cross-referencing information across multiple reports and evaluating plausibility of each new piece of information obtained from a report. As there is no universally agreed structured way to report such information, extraction is a manual process. More general information, such as applicable standards and best or typical practices, can help assess *Ctx*, especially if no specific information indicates a CPS deviates from these.

Statistical Power.

Statistical power analysis allows to reason about the number of samples required to observe an effect with a specific probability or certainty. Following literature describing best practices for exploitative studies [50], we suggest $\alpha = 0.05$ and $(1 - \beta$ err. prob.$) = 0.8$. By doing so, one can determine the relation between the number of required samples and detectable effect size. Considering that the proposed metrics for *Cap* and *Ctx* work on an integer scale, we suggest to target an effect size of $d = 1.0$ or lower. Figure 2 shows that, with these settings, an application of our approach should aim at 29 samples or more for a good chance of finding such effects. Significant results may still be observed with fewer data points but the chance of not finding existing effects grows.

Implementation. Implementing our approach, we gather data from the two most comprehensive incident databases of their kind: the "OT cyber-attack Database" (OT-CAD) [47] with 127 ICS incidents, and the BAS attack database [24] with 26 incidents from building automation and smart city. After applying the inclusion criteria (presented earlier in Section 4.1), a total of 25 incidents remain in scope, resulting in an estimated statistical power around 75% (assuming effect size $d = 1.0$). Both databases are still maintained, created with a similar methodology, and aim at exhaustively listing publicly known cyberattacks for their respective domain and scope. We thus deem them comparable for our purposes. While there are other databases available (e.g., RISI [49], last updated 2014), none of the public datasets match the selected repositories in quality or quantity. The inclusion of non-public datasets is possible and would probably lead to statistically stronger results, however to not hamper reproducibility we do not do so in this work.

4.2. Phase 2: Rating

Rating must follow a consistent, clearly defined rating process to create repeatable results. For this purpose, Appendixes A.2.3 and A.3.3 provide codebooks with rating instructions for *Cap* and *Ctx* respectively. Given the expected limitations caused by the data collection challenges discussed above, rating should be limited to a coarse scale, as indicated in Sections 2 and 3. When there is insufficient information to infer a score, the corresponding value should remain unassigned ($null$).

To minimize (the risk of) rater bias and the impact of subjectivity, we split the rating of *Cap* and *Ctx* over two independent raters; one *Cap*-rater (Rater R1 in Figure 1) and one *Ctx*-rater (Rater R2 in Figure 1). Raters should not confer, but data collection and rating do not have to be temporally separated; the raters can search the Internet for additional references when deriving scores.

Rating of *Cap*. To determine *Cap*, the rater assigns a value to each of the ATT&CK Tactics, depending on what type of CPS knowledge (technical, process, or process-mapping knowledge) is required to apply the respective Tactic for each cyberattack, as defined in Section 2.

The rater should consider the greater context when assigning values, e.g., the act of *Collection* may not necessarily require CPS-knowledge, however the interpretation of the collected data may require or even provide the attacker with *process-mapping* knowledge. In these cases, the rater should include the implied capabilities and opt for the higher rating. Note that knowledge available to the attacker at the end of the attack is rated, so e.g., process-mapping knowledge obtained as a result of Collection is explicitly included in the rating.

Because our *Cap*s are based of ATT&CK Tactics, we expect any cybersecurity expert or researcher familiar with general concepts of ICS and the ATT&CK Framework to produce comparable Attacker-Capability ratings. Khalil, Bahsi, and Korõtko confirm that (CPS) attack modelling does not require extensive CPS knowledge of the target system [51, ch 2.3].

In our application, the *Cap* rating was performed by a cyber- and ICS-security researcher with $4+$ years of experience (Rater R1), following the instructions presented here, in Section 2 and in Appendix A.2.3.

Rating of *Ctx*. To determine the *Ctx*, a rater scores all eight *Ctx* dimensions for each incident, according to the criteria defined in Section 3 and the instructions provided in the codebook (Appendix A.3.3). To reliably rate *Ctx*, extensive experience across multiple CPS domains and deployments is necessary.

For our implementation, we asked a domain expert (Rater R2) to perform our *Ctx* rating according to the instructions presented in Section 3 and Appendix A.3.3. Having $10+$ years of experience as university researcher in CPS security, as well as $10+$ years as entrepreneur and senior R&D for a major vendor of Intrusion Detection Systems for CPS, our rater has seen many CPS-deployments across different domains and countries.

Rating quality. Given a codebook with rating instructions, raters are likely to assign comparable scores (even though subjective), making human error and imperfect information during rating the biggest threats [52]–[54]. These effects can be reduced by asking multiple raters to rate the same metric, compute agreement scores, and perform conflict resolution and following general best practices [54].

Following best practices, we publish all ratings and inter-rater agreement scores (Cohen's Kappa) for transparency, as well as our codebook as part of the Methodology (cf. Appendices A.2.3 and A.3.3). To compute an inter-rating agreement score for our application, where we

only had one rater per metric, we asked both raters to also rate a small random sample (5 incidents) for the respective other metric, performing conflict resolution and checking at least moderate agreement was reached. The inter-rater agreement on *CAPABILITY* (cf. Section 2) is > 0.705, indicating moderate agreement [55][56, tab 3]. Despite Rater R1 not having as much experience as we would like to see in a context rater, the inter-rater agreement score for the *CONTEXT* (cf. Section 3) is ≈ 0.644, indicating moderate agreement [56, Table 3]. These scores indicate sufficient quality of the rating process, as well as sufficient clarity of the rating instructions and codebook.

4.3. Phase 3: Outcome

The final phase of our approach aims at answering the three research questions.

Question Q1. We focus on the total *CAPABILITY* score to determine the required *Cap*s for CPS in general and for each CPS domain separately. To test whether attacks against different CPS domains require significantly different *Cap*s, we use the Mann-Whitney U test ($\alpha = 0.05$). A high variance implies that some CPS require extensive CPS *Cap*s, while others can be attacked with basically no CPS-specific knowledge at all. Observing clear differences in mean and distribution of *Cap*s when comparing two CPS domains, implies that *Cap* requirements significantly differ between the two domains. The individual *Cap* sub-dimension scores can help explain what causes these differences.

Based on broader intuition in the subject area, we expect to observe a wide distribution of required *Cap*s for CPS attacks. For BAS, we expect that attackers often do not require any CPS-knowledge at all [24], [57], while we expect that attacks against ICS, especially ICS in the power and electricity sector (ICS_{elec}) require process-mapping knowledge (`high` / $+1$) [15]. We test this expectation against the alternative hypothesis that BAS and ICS share identical capability requirements.

Question Q2. To analyze which *Ctx* factors influence the required *Cap*s and how, we compute a *Context-Impact Matrix*, expressing the correlation between *Cap* and *Ctx*. We use Spearman's rank correlation coefficient, skipping missing ($null$) values. Because of how the *Ctx*-factors were defined, we expect only positive correlation (i.e., higher *Ctx* rating is expected to lead to higher *Cap* requirements). A high correlation (close to 1) between a *Ctx*-factor and a *Cap*-dimension suggests (but does not yet prove[6]) that the maturity of a specific *Ctx*-factor increases the minimum required *Cap* of an attack Tactic.

As *Ctx* characteristics capture many security related aspects considered in literature, we expect the overall *CONTEXT* score and some of the individual characteristics to show correlation to at least some of the *Cap* scores, even with a limited sample size. Al-Sada, Sadighian, and Oligeri suggest that patterns will emerge, e.g., that high correlation with e.g., *Collection* may imply high correlation with other Tactics [58].

Question Q3. To characterize and analyze similarities and differences in *Ctx* across domains, we aggregate and compare the *Ctx* ratings for each CPS domain. To test the hypothesis that two domains differ in *Ctx*, we use the Mann-Whitney U test with significance level $\alpha = 0.05$ and alternative hypothesis that they do not differ in *Ctx*.

In accordance with academic and expert opinion [57], we expect BAS to receive lower *Ctx* scores, compared to ICS. Further, we expect ICS_{elec}, as a subset of ICS, to be more mature than other ICS implementations.

6. correlation does not imply causation; see Section 6.3.

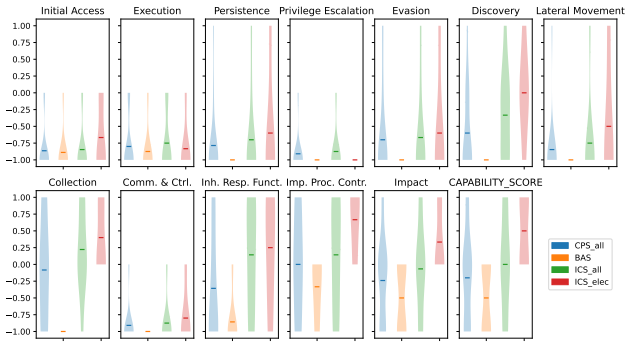

Figure 3. Required Attacker-Capabilities (*Cap*) by Domain

5. Results

In this section we give the results of applying our approach to the BAS and ICS domains. An interpretation of these results is provided in Section 6.1.

Our final dataset holds 25 incidents with physical impact between 2010 and 2020. Of these, 10 come from the extended BAS domain and 15 from the ICS domain. OTCAD further divides the ICS incidents into 6 cyberattacks from the sector of electrical power generation and distribution, 4 from water and waste water, as well as metals (2), oil and gas (1), and automotive (1). For one ICS incident, OTCAD does not specify the sector [47].

5.1. Attacker-Capabilities (Q1)

We observe `low` dominating the required *Cap* in Figure 3, indicating that most steps of a CPS attack do not require any CPS *Cap*. In total, 9/25 attacks obtained `low` for all capability sub-dimensions, in 12/25 the highest score is `medium`, and only four attacks received at least one `high` rating. These `high` scores are predominantly obtained in the four ATT&CK Tactics *Collection*, *Inhibit Response Function*, *Impair Process Control*, and *Impact*.

Figure 3 presents the *Cap* ratings by CPS domain, including the mean and distribution of *Cap* ratings, across all *Cap* sub-dimensions (i.e., ATT&CK Tactics). According to the definitions in Section 2.2, -1 implies no CPS-knowledge was required, 0 implies that at least *technical* or *process* knowledge was required, and $+1$ implies that *process-mapping* knowledge was necessary. In general, during most steps of a CPS attack, low or no CPS-specific *Cap*s are required, visible by the low value with low distribution in Figure 3 (blue). Across all CPS domains, the first seven capabilities and *Comm. & Ctrl.* score below -0.5 (average), indicating attackers require next to no CPS knowledge for achieving the related ATT&CK Tactics, regardless of the CPS domain. Especially Tactics *Initial Access*, *Privilege Escalation*, and *Comm & Ctrl.*

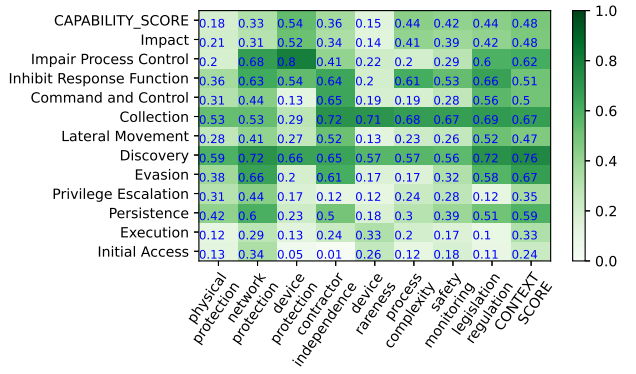

Due to *Ctx*-factor definitions, only positive correlation is expected.

Figure 4. *Ctx*-Impact Matrix (cf. Q1)

can be performed with offensive IT-knowledge only. The figure also shows *Persistence*, *Evasion*, *Discovery*, and *Lateral Movement* only requiring CPS-specific *Cap* for few attacks, while *Collection*, *Inh. Resp. Funct.* and *Impact* show a wide range from `low`, i.e., no CPS-specific knowledge to `high`, i.e., *process-mapping* knowledge.

Domain-specific differences emerge, when looking at the BAS and ICS domains separately (Figure 3, orange and green). All BAS received `low` (-1) for 7/12 *Cap*-dimensions.[7] BAS-specific *Cap* for *Inh. Resp. Funct.*, *Imp. Proc. Contr.*, and *Impact*, range from `low` to `medium` scores, implying that attackers require *at most technical* or *process*-specific CPS knowledge for these Tactics. The ICS domain, on the other hand, is characterized by a wide distribution ranging from `low` to `high` for the same Tactics. This implies most attacks require *some* CPS knowledge, be it *technical* or *process* knowledge, with some attacks requiring *process-mapping* knowledge. ICS incidents from the power and electricity domain (ICS_{elec}, red in Figure 3) always mark the higher end of ICS scores, implying that attacks in that specific ICS (sub)sector are linked to higher *Cap* than the average ICS (ICS_{all}).

The *CAPABILITY* difference between BAS and ICS_{elec} is significant with respect to our threshold (Mann-Whitney U test, $stat = 7.5$, $p \approx 0.009$). The *CAPABILITY* difference between BAS and ICS, however, is not significant ($stat = 47.5$, $p \approx 0.102$). The power analysis presented in Section 4, together with Figure 3, implies that this may be due to a lack of statistical power, i.e., the number of samples being too small to reliably detect the effect ($\Delta 0.5$ between the two means).

In summary, CPS-specific knowledge is not equally important for all ATT&CK Tactics. Especially in BAS, attacks often do not require any CPS-specific knowledge. In ICS, where some Tactics do require CPS-knowledge, the required *Cap* vary drastically, implying that even in ICS some attacks do not require any CPS knowledge at all, while others require even *process-mapping* knowledge. We see those attacks requiring *process-mapping* knowledge predominantly in the electricity (sub-)domain.

5.2. *Ctx*-Impact Matrix (Q2)

The *Ctx*-Impact-Matrix, presented in Figure 4, shows Spearman correlation between the required *Cap* for each attack and the *Ctx* ratings for the attacked CPS deployments. The aggregated values *CONTEXT* (from *Ctx*) and *CAPABILITY* (from *Cap*) are moderately correlated (0.48), hinting that *Ctx* and *Cap* are related. We note *Evasion* to be notably correlated with *network protection* and *contractor independence*. *Discovery* and *Collection* both show correlation with almost all *Ctx* dimensions; we especially note correlation between *Collection* and *device rarity*, *process complexity*, and *safety monitoring*. Lastly, *Inhibit Response Function* also shows notable correlation with *process complexity*, *safety monitoring*, and *legislation & regulation*. There is generally lower correlation between *Ctx* and "early" ATT&CK Tactics, for which we previously found low *Cap* requirements (Figure 3); i.e., *Initial Access*, *Execution*, *Privilege Escalation*, as well as *Persistence* and *Lateral Movement*. Column-wise, *network protection*, *contractor independence* and *legislation & regulation* show correlation with most *Cap*-dimensions.

5.3. *Ctx* Differences by Domain (Q3)

Most CPS were rated either `medium` or `low` *Ctx*, with few `high` ratings mostly in *physical protection*, *network protection*, *process complexity*, and *safety monitoring*. The vast majority of `high` ratings stem from ICS deployments; only three `high` ratings were given to BAS, all in *device rarity* due to obscure, non-"typical" BAS or smart-city devices like industrial fridges or radio-based emergency siren systems. BAS deployments received a `low` score in most cases (66/80 assigned values). In particular, *contractor independence*, *safety monitoring*, and *legislation & regulation* score exclusively `low`, with the exception of one *null* value. This indicates that many of the attacked BAS networks did not follow even the most basic security recommendations or best practices. In contrast, the predominant value in the ICS domain is `medium` (71/120), followed by `high` (28/120), indicating that many ICS systems follow at least basic best-practices. Notably, every ICS setup has at least one `high` value.

Looking at Figure 5, *Ctx*-differences between the CPS domains BAS and ICS (orange and blue) become apparent. Without exception, ICS showed notably higher *Ctx* than BAS, often with difference of more than 1.0 points between the averages. Especially in *physical protection*, *process complexity*, and *safety monitoring*, ICS deployments often obtained high scores, implying that ICS are often shielded from physical access, run complicated control operations, and make use of extensive mechanism to ensure safety even outside of normal operating parameters (e.g., Safety Instrumented Systems). In direct contrast, the physical processes in BAS are often relatively simple, the lack of immediate threat makes extensive *safety monitoring* unnecessary, and there are fewer legislative regulations. The low score in *contractor independence* is the result of outsourcing and many sub- and service-contractors being involved in the operation of a modern

7. Tactics: Persistence, Privilege Escalation, Evasion, Discovery, Lateral Movement, Collection, Comm & Ctrl.

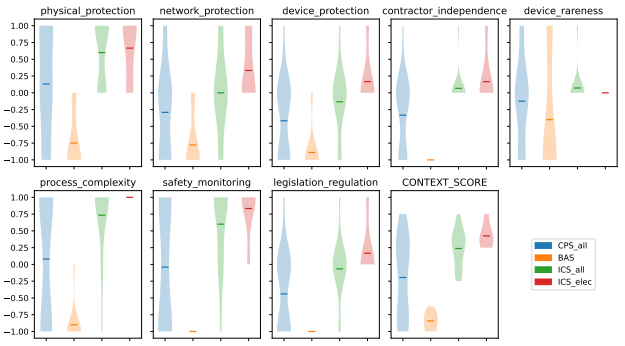

Figure 5. CPS-Context (*Ctx*) by Domain

building. As a subdomain of ICS, ICS_{elec} is responsible for many of the high ICS-*Ctx* scores. Only in *device rareness*, ICS_{elec} scores lower than ICS_{all}, implying some form of standardization or limited set of device vendors capable of creating the necessary operational equipment, resulting in somewhat homogeneous deployments. Finally, the high score in *process complexity* is result of the high complexity of operating a distributed power grid, balancing offer and demand, while maintaining a stable frequency.

The *CONTEXT* between BAS and ICS_{all} differs ($\Delta 1.15$) significantly (Mann-Whitney U test, $stat = 0.0$, $p \ll 0.001$). Between BAS and ICS_{elec}, the difference ($\Delta 1.27$) is also significant ($stat = 0.0$, $p \approx 0.001$). As expected, the difference between ICS_{all} and ICS_{elec} ($\Delta 0.19$) is non-significant ($stat = 61.5$, $p \approx 0.207$); there is no reason to believe that the two are distinct.

In summary, *Ctx* varies drastically across CPS installations, with stark differences between CPS-domains. Most BAS implementations violate basic recommendations or implement them with low maturity. In contrast, ICS often follow basic best practices, some implement even extensive measures at high maturity levels. The power and energy sector rates high, even amongst ICS, but this does not yield a statistically relevant difference.

6. Discussion and Conclusion

6.1. Interpretation of Results

***Cap* in Different Domains (Q1).** In Section 1, we presented an example of attackers obtaining physical impact without any CPS knowledge. Our results (see Figure 3) confirm that not all attacks against CPS require extensive background knowledge in process control or the underlying physical process. Often, the attacker can suffice with "standard" IT-techniques. This partially deviates from previous arguments, e.g., by Green, Krotofil, and Abbasi [15].

Where necessary, *Cap* are mostly used in later steps of the ATT&CK Framework, e.g., to *Inhibit Response Function*, and *Impair Process Control*. Further, *Collection* plays a key role in obtaining *process-mapping* knowledge during the attack.

Attacks against BAS required notably less CPS-knowledge than attacks against ICS; many BAS devices are attacked with common offensive IT knowledge, e.g., using denial of service attacks, interrupting network connectivity, restarting or bricking the device itself, or exploiting an exposed Web-UI. This does not require CPS-specific *Cap*; it could imply that attackers may be unaware of the alternative ways to obtain physical impact in BAS. High *Cap* scores mostly occur for sophisticated ICS attacks (including e.g., attacks against the Ukrainian power grid) where attackers knew which operation on which point would lead to specific impact in the physical world (i.e., *process-mapping* knowledge) and which safeguards had to be circumvented.

***Ctx*-Impact Matrix (Q2).** Figure 4 showed *Collection* and *Inhibit Response Function* correlating with *process complexity*. It is reasonable to assume that understanding the *process-mapping* is harder for complex CPS; fully understanding the control logic and process gets challenging and more data needs to be collected. Being specific to the CPS deployment, this knowledge can only be obtained through device *Discovery* and data *Collection* on site and the interpretation requires both *technical* and *process* experience. As the complexity of the program logic increases with the complexity of the processes, it becomes harder for an attacker to identify a (minimal) set of actions that keep the CPS from recovering during the attack (i.e., the difficulty of *Inhibit Response Function* increases). Likewise, *safety monitoring* influences *Collection* and *Inhibit Response Function*; To find ways to avoid triggering the safety system or suppress its response, attackers need to anticipate how a human supervisor or a pre-programmed Safety Instrumented System (SIS) will react to deviations from the desired state and thus must perform *Collection*.

In summary, high scores in *process complexity* and *safety monitoring* indicate that attackers need higher *Cap*s and more time & effort to perform *Collection* and to *Inhibit Response Function*. The latter is necessary, to ensure that neither the regular process control, nor a human operator or safety system interfere when the attacker starts to *Impair Process Control*.

***Ctx* Differences by Domain (Q3).** That BAS, on average, performs worse than ICS on all *Ctx* characteristics aligns with academic publications [24] and expert opinions [57]. In general, BAS networks are, by design, more open and easier to access, both physically and over the network [57]. This is reflected in low values for *physical protection* and *network protection*. Most BAS functions are not safety critical and do not have hard time- or quality requirements [57], so there is no need for elaborated fail-safe mechanisms, human supervision, or legislative regulation. Lastly, many BAS applications run on standard Linux-based controllers and are rarely patched [57] making successful cyber-attacks re-usable.

Unsurprisingly, traditional ICS networks are, in comparison to BAS, better protected, monitored by human operators, and have defined fail-safe mechanisms (such as Safety Instrumented Systems). In contrast to BAS, most ICS systems are deployed in non-public spaces or enclosed in locked closets. We expect most ICS systems operate on separate networks that are not directly accessible from the Internet or IT-network. This is supported by our findings (see e.g., Figure 5). Especially for critical infrastructure, the availability and quality of delivered services or goods is enforced by law. In addition, whenever human lives or the environment are at risk, there is some form of regulation to ensure minimum safety standards.

Overall, we observe that not all CPS networks are created equal; some, especially in the BAS domain, are

susceptible to simple and low-effort attacks in which the attacker can obtain physical impact regardless of their CPS *Cap*. We also observe notable *Ctx* differences between the domains and saw that *Ctx* and *Cap* are correlated. While the differences between ICS and BAS are most notable, we also observe minor but explainable differences between ICS and ICS_{elec}.

6.2. Implications for Research and Practice

Research. While there are studies analyzing the conditions under which cyber-attacks are economically viable for different threat actors (e.g., when the expected gain exceeds the estimated cost of a cyber-attack) [59], [60], we are unaware of similar research focused on the CPS domain. Our work shows *how* CPS (especially BAS and ICS) differ in *Ctx* and, as a result, face different threats and attackers. Our results qualitatively confirm general intuitions and observations made in previous literature [24] and gives directions for more detailed research into the causal relationship that is necessary. The importance of insight into *Ctx* is underlined by unawareness of the impact of *Ctx*-differences between domains in prior literature. For instance, Wendzel, Kahler, and Rist propose countermeasures to covert channels in BACnet [8]. Modern BAS-devices support a multitude of protocols including common IT-protocols like TLS, DNS, HTTP(s), and communicate with cloud services. In a BAS-Context, attackers thus have no incentive to use covert channels to stay hidden, but can utilize "normal" IT-protocols instead. As another example, Kaur, Tonejc, Wendzel, *et al.* present a traffic-normalizer to protect BACnet devices from malformed network traffic [9]. While many ICS-devices have limited protocol stacks and several reports of malformed packets crashing ICS devices exist, BAS-controllers also implement multiple application-layer protocols and often run web servers making other, likely easier, (IT-)attacks possible. The *Ctx* and the resulting *Cap* drastically reduces the effectiveness of a sophisticated traffic-normalizer in the BAS-Context.

Practice. We show that, in line with [24], different types of CPS are exposed to different threats and attackers. Especially for BAS, minimal *Cap* are required, meaning that the number of potential attackers who can obtain physical impact on BAS is likely significantly higher than for ICS. The low *Cap* threshold can make BAS devices economic targets, even for unskilled attackers. To keep BAS systems secure, it is vital for BAS owners to understand the profile and capabilities of a typical BAS attacker, as these do not coincide with threat models made for ICS. Similarly, ICS operators from the electricity and energy domain cannot solely rely on best-practices for general CPS, as they would undercut the domain-specific *Ctx* and protection levels.

Our results also imply that the security of some CPS, especially BAS, can be improved with common and well-established IT-security techniques at comparably low costs, especially by restricting network access and defining fail-safe mechanisms.

While our results cannot and should not be used to derive concrete actions for individual CPS deployments, they can help create general recommendations and best practices, as well as help implementing generic best practices to a specific domain. Best practices should not be applied blindly across domains but instead need to be tailored to the target domain and CPS. Our results can help with this domain-transfer and thus allow for informed and economic decisions when designing secure CPS.

6.3. Study Limitations & Threats to Validity

Even though very little data is publicly available, we obtained significant results (threshold $\alpha \leq 0.05$). Additional data may allow for higher rating granularity and to measure smaller effects.

Limitations. Even though *Cap* and *Ctx* capture relevant dimensions, we cannot claim that they capture all relevant problem dimensions. Further, the design of our study cannot examine causation between *Ctx* and *Cap*, thus we cannot be certain that the domain characteristics *caused* the differences in required attacker capabilities; we recommend further studies on this subject.

Internal validity. We obtained dimensions for *Ctx* and *Cap* from relevant literature. Values for *Ctx* and *Cap* are assigned by different raters to avoid confounding variables and the inter-rater agreement score suggests that the rating criteria allow for reproducible results.

External validity. Despite the scarcity of available data points, we showed significant differences and strong trends between domains. However, the analyzed public incidents may not be representative of the overall attack landscape. The influence of individual factors may change over time and other factors, not captured by *Cap* and *Ctx*, may exist. Even though cross-referencing information from different articles can increase data quality for the rating process, false information in reports remains a threat to validity. Lastly, we only evaluate *Ctx* for CPS that were compromised, the obtained results may not be representative for all (non-compromised) CPS.

6.4. Future Research

Besides applying our methodology to future CPS incidents, we consider three topics of interest for future research. First, research that focuses on individual *Ctx*-factors, their impact on security, and whether there is a causal relationship between *Ctx* and required *Cap*. Secondly, a longitudinal study to analyze how *Ctx* and *Cap* developed and changed over time. Lastly, we see value in applying our methodology to private incident datasets to determine how representative publicly available incident information is, and whether the selection of publicly available incidents of attacks are biased.

6.5. Ethical Considerations

As we only use public information and do not process any personal data, our university policies do not require IRB-approval. The authors believe that the public availability of this paper and associated research artifacts does not pose a threat to existing CPS and has no immediate use to attackers. Information about incidents is collected manually from the incident database and public online resources (i.e., no use of automated tooling or scrapers).

Acknowledgements

This research was funded by the Dutch Research Council (NWO), grant number 628.001.032 (DEPICT) and NWA.1160.18.301 (INTERSECT). We thank Daniel dos Santos from Forescout and Stash Kempinski (TU/e) for sharing feedback and domain expertise. We thank the anonymous reviewers for their helpful comments and pointers to relevant literature. As part of the open-report model followed by the Workshop on Attackers & CyberCrime Operations (WACCO), all the reviews for this paper are publicly available at https://github.com/wacco-workshop/WACCO/tree/main/WACCO-2024.

References

[1] N. Falliere, L. O. Murchu, and E. Chien, "W32.stuxnet dossier," Symantec, Report, version 1.4, 2011, Archived. [Online]. Available: https://web.archive.org/web/20200224233811/http://www.symantec.com/content/en/us/enterprise/media/security_response/whitepapers/w32_stuxnet_dossier.pdf (visited on 01/25/2023).

[2] C. Bing, "Hackers try to contaminate florida town's water supply through computer breach," 2021. (visited on 05/13/2024).

[3] "Protecting industrial control systems, Annex iii," European Union Agency for Cybersecurity, Report, 2011. [Online]. Available: https://www.enisa.europa.eu/publications/annex-iii/@@download/fullReport.

[4] ISO, Standard ISO/IEC 27019:2017, Oct. 2017.

[5] Ponemon Institute LLC, "Third annual study on exchanging cyber threat intelligence: There has to be a better way," Report, Nov. 2017. [Online]. Available: https://www.ponemon.org/local/upload/file/2017%20Inflobox%20Report%20V6.pdf (visited on 11/01/2023).

[6] J. Soldatos, J. Philpot, and G. Giunta, Eds., *Cyber-Physical Threat Intelligence for Critical Infrastructures Security, A Guide to Integrated Cyber-Physical Protection of Modern Critical Infrastructures*. Now Publishers, 2020. DOI: 10.1561/9781680836875.

[7] M. Van Horenbeeck, M. Davidson, B. Grobauer, *et al.*, "Actionable information for security incident response," European Union Agency for Cybersecurity, Report, Nov. 2014. DOI: 10.2824/38111.

[8] S. Wendzel, B. Kahler, and T. Rist, "Covert channels and their prevention in building automation protocols: A prototype exemplified using BACnet," in *IEEE Conference on Green Computing and Communications*, 2012. DOI: 10.1109/GreenCom.2012.120.

[9] J. Kaur, J. Tonejc, S. Wendzel, and M. Meier, "Covert channels and their prevention in building automation protocols: A prototype exemplified using BACnet," in *SEC: ICT Systems Security and Privacy Protection*, 2015. DOI: 10.1007/978-3-319-18467-8_41.

[10] M. Dodson, A. R. Beresford, and D. R. Thomas, "When will my plc support mirai? the security economics of large-scale attacks against internet-connected ics devices," in *APWG Symposium on Electronic Crime Research (eCrime)*, Nov. 2020. DOI: 10.1109/ecrime51433.2020.9493257.

[11] M. Campobasso and L. Allodi, "Impersonation-as-a-service: Characterizing the emerging criminal infrastructure for user impersonation at scale," in *SIGSAC CCS Proceedings*, 2020. DOI: 10.1145/3372297.3417892.

[12] P. H. Meland, Y. F. F. Bayoumy, and G. Sindre, "The ransomware-as-a-service economy within the darknet," *Computers & Security*, vol. 92, 2020. DOI: 10.1016/j.cose.2020.101762.

[13] L. Allodi and S. Etalle, "Towards realistic threat modeling: Attack commodification, irrelevant vulnerabilities, and unrealistic assumptions," in *Workshop on Automated Decision Making for Active Cyber Defense*, 2017. DOI: 10.1145/3140368.3140372.

[14] M. Rosso, *Data supplementary to the publication "a methodology to measure the "cost" of cps attacks: Not all cps networks are created equal"*, Eindhoven University of Technology. DOI: 10.4121/2cd80dea-9eea-4dfb-b0d3-3e72c0e47804. [Online]. Available: https://gitlab.tue.nl/sec-lab/cps-security/cps-knowledge (visited on 05/13/2024).

[15] B. Green, M. Krotofil, and A. Abbasi, "On the significance of process comprehension for conducting targeted ICS attacks," in *Workshop on Cyber-Physical Systems Security and PrivaCy*, Nov. 3, 2017. DOI: 10.1145/3140241.3140254.

[16] A. Alazab, J. Abawajy, M. Hobbs, R. Layton, and A. Khraisat, "Crime toolkits: The productisation of cybercrime," in *2013 12th IEEE International Conference on Trust, Security and Privacy in Computing and Communications*, 2013. DOI: 10.1109/TrustCom.2013.273.

[17] L. Ablon, M. C. Libicki, and A. A. Golay, "Markets for cybercrime tools and stolen data, Hackers' bazaar," 2014. [Online]. Available: https://www.rand.org/content/dam/rand/pubs/research_reports/RR600/RR610/RAND_RR610.pdf (visited on 03/01/2024).

[18] B. Green, R. Derbyshire, M. Krotofil, W. Knowles, D. Prince, and N. Suri, "Pcaad: Towards automated determination and exploitation of industrial systems," *Comput. & Secur*, 2021. DOI: 10.1016/j.cose.2021.102424.

[19] H. Esquivel-Vargas, J. H. Castellanos, M. Caselli, N. O. Tippenhauer, and A. Peter, "Identifying near-optimal single-shot attacks on ICSs with limited process knowledge," in *ACNS 2022: Applied Cryptography and Network Security*, LNCS,volume 13269, 2022. DOI: 10.1007/978-3-031-09234-3_9.

[20] J. Wetzels, D. Dos Santos, and M. Ghafari, "Insecure by design in the backbone of critical infrastructure," in *Cyber-Physical Systems and Internet of Things Week*, 2023. DOI: 10.1145/3576914.3587485.

[21] Forescout Technologies Inc., "OT:ICEFALL," 2022. [Online]. Available: https://www.forescout.com/resources/ot-icefall-report/ (visited on 02/28/2024).

[22] W. Xu, Y. Tao, and X. Guan, "The landscape of industrial control systems (ICS) devices on the internet," in *2018 International Conference On Cyber Situational Awareness, Data Analytics And Assessment (Cyber SA)*, 2018. DOI: 10.1109/CyberSA.2018.8551422.

[23] M. Rocchetto and N. O. Tippenhauer, "On attacker models and profiles for cyber-physical systems," in *ESORICS*, 2016. DOI: 10.1007/978-3-319-45741-3_22.

[24] M. Tommasini, M. Rosso, E. Zambon, L. Allodi, and J. Hartog, "Characterizing building automation system attacks and attackers," in *2022 IEEE EuroS&PW*, 2022. DOI: 10.1109/EuroSPW55150.2022.00020.

[25] R. Anderson, C. Barton, R. Böhme, *et al.*, "Measuring the cost of cybercrime," in *The Economics of Information Security and Privacy*. Springer, 2013, pp. 265–300, ISBN: 9783642394980. DOI: 10.1007/978-3-642-39498-0_12.

[26] N. Ortiz, M. Rosso, E. Zambon, J. den Hartog, and A. Cardenas, "From power to water: Dissecting SCADA networks across different critical infrastructures," Mar. 2024. (visited on 03/11/2024).

[27] ISO, Standard ISO/IEC 27002:2022, Feb. 2022.

[28] ICS-CERT, "Ics-cert monitor, November/december 2017," U.S. DHS, Tech. Rep., 2017. [Online]. Available: https://www.cisa.gov/sites/default/files/Monitors/ICS-CERT_Monitor_Nov-Dec2017_S508C.pdf (visited on 07/13/2023).

[29] European Network and Information Security Agency (ENISA), *Good practices for security of Internet of things in the context of smart manufacturing*. Nov. 2018. DOI: 10.2824/851384.

[30] C. Herley and P. C. van Oorschot, "Sok: Science, security and the elusive goal of security as a scientific pursuit," in *Symposium on Security and Privacy*, 2017. DOI: 10.1109/SP.2017.38.

[31] S. J. Stolfo, S. M. Bellovin, and D. Evans, "Measuring security," *IEEE Secur. Priv.*, vol. 9, 2011. DOI: 10.1109/MSP.2011.56.

[32] ICS-CERT, "Recommended practice: Improving industrial control system cybersecurity with defense-in-depth strategies," U.S. DHS, Tech. Rep., Sep. 2016. [Online]. Available: https://www.cisa.gov/sites/default/files/recommended_practices/NCCIC_ICS-CERT_Defense_in_Depth_2016_S508C.pdf (visited on 07/13/2023).

[33] X. Qin, M. Rosso, A. A. Cardenas, S. Etalle, J. den Hartog, and E. Zambon, "You can't protect what you don't understand: Characterizing an operational gas SCADA network," in *IEEE SPW*, 2022. DOI: 10.1109/SPW54247.2022.9833864.

[34] Y. Yang, K. McLaughlin, T. Littler, S. Sezer, B. Pranggono, and H. F. Wang, "Intrusion detection system for iec 60870-5-104 based SCADA networks," in *IEEE Power & Energy Society General Meeting*, Jul. 2013. DOI: 10.1109/PESMG.2013.6672100.

[35] X. Qin, K. Mai, N. Ortiz, K. Koneru, and A. A. Cardenas, "Cybersecurity and resilience for the power grid," in *Resilient Control Architectures and Power Systems*. Wiley, Dec. 3, 2021, ch. 13. DOI: 10.1002/9781119660446.ch13.

[36] J. R. Clark and W. L. Davis, "A human capital perspective on criminal careers," *J. of Appl. Bus. Res.*, vol. 11, 1995. DOI: 10.19030/jabr.v11i3.5860.

[37] N. Kshetri, "The simple economics of cybercrimes," *Secur. & Priv.*, vol. 4, 2006. DOI: 10.1109/MSP.2006.27.

[38] K. Huang, M. Siegel, and S. Madnick, "Systematically understanding the cyber attack business: A survey," *Comput. Surv.*, 2018. DOI: 10.1145/3199674.

[39] M. Garcia, A. Bessani, I. Gashi, N. Neves, and R. Obelheiro, "Analysis of operating system diversity for intrusion tolerance," *Softw.: Pract. and Exp.*, vol. 44, 2013. DOI: 10.1002/spe.2180.

[40] S. Forrest, A. Somayaji, and D. Ackley, "Building diverse computer systems," in *Workshop on Hot Topics in Operating Systems*, 1997. DOI: 10.1109/hotos.1997.595185.

[41] S. A. Hofmeyr and S. Forrest, "Archit. for an artif. immune syst.," *Evolutionary Computation*, vol. 8, 2000. DOI: 10.1162/106365600568257.

[42] L. Garcia, F. Brasser, M. H. Cintuglu, A. Sadeghi, O. A. Mohammed, and S. A. Zonouz, "Hey, my malware knows physics! attacking PLCs with physical model aware rootkit," in *24th NDSS*, 2017. (visited on 02/20/2023).

[43] IEC, Standard IEC TR 61508:2010, Apr. 2010.

[44] A. Di Pinto, Y. Dragoni, and A. Carcano, "TRITON: The first ics cyber attack on safety instrument systems," Nozomi Networks, Tech. Rep., 2018, at Black Hat USA 2018. [Online]. Available: https://i.blackhat.com/us-18/Wed-August-8/us-18-Carcano-TRITON-How-It-Disrupted-Safety-Systems-And-Changed-The-Threat-Landscape-Of-Industrial-Control-Systems-Forever-wp.pdf (visited on 03/20/2024).

[45] German Federal Office for Information Security (BSI), "The state of it security in germany 2014," Tech. Rep. BSI-LB15503e, Mar. 2, 2015. [Online]. Available: https://www.bsi.bund.de/SharedDocs/Downloads/EN/BSI/Publications/Securitysituation/IT-Security-Situation-in-Germany-2014.pdf (visited on 07/13/2023).

[46] D. Evans, "The internet of things: How the next evolution of the internet is changing everything," Whitepaper, Apr. 2011. [Online]. Available: https://www.cisco.com/c/dam/en_us/about/ac79/docs/innov/IoT_IBSG_0411FINAL.pdf (visited on 05/15/2024).

[47] S. Kempinski, *Operational technology cyber attack database: OTCAD*, Secura B.V., Sep. 24, 2021. [Online]. Available: https://github.com/SecuraBV/OTCAD (visited on 09/16/2022).

[48] M. Tommasini and M. Rosso, *Data supplementary to the paper: "characterizing building automation system attacks and attackers"*, Eindhoven University of Technology, Jul. 2022. DOI: 10.4121/19617243.

[49] *The repository of industrial security incidents: RISI*. [Online]. Available: https://www.risidata.com/Database (visited on 09/16/2022).

[50] C. C. Serdar, M. Cihan, D. Yücel, and M. A. Serdar, "Sample size, power and effect size revisited: Simplified and practical approaches in pre-clinical, clinical and laboratory studies," *Biochem. Med. (Zagreb)*, vol. 31, 2021, PMID 33380887.

[51] S. M. Khalil, H. Bahsi, and T. Korõtko, "Threat modeling of industrial control systems: A systematic literature review," *Comput. & Secur.*, vol. 136, 2024. DOI: 10.1016/j.cose.2023.103543.

[52] P. Johnson, R. Lagerström, M. Ekstedt, and U. Franke, "Can the common vulnerability scoring system be trusted? a bayesian analysis," *Transactions on Dependable and Secur. Comput.*, vol. 15, 2018. DOI: 10.1109/TDSC.2016.2644614.

[53] M. Höst, B. Regnell, and C. Wohlin, "Using students as subjects — a comparative study of students and professionals in lead-time impact assessment," *Empir. Softw. Eng.*, 2000. DOI: 10.1023/a:1026586415054.

[54] J. Belur, L. Tompson, A. Thornton, and M. Simon, "Interrater reliability in systematic review methodology: Exploring variation in coder decision-making," *Sociol. Methods Res.*, 2018. DOI: 10.1177/0049124118799372.

[55] M. Tavakol and R. Dennick, "Making sense of cronbach's alpha," *Int. J. Med. Educ.*, vol. 2, 2011. DOI: 10.5116/ijme.4dfb.8dfd.

[56] M. L. McHugh, "Interrater reliability: The kappa statistic," *Biochem. Med. (Zagreb)*, 2012, PMID 23092060.

[57] D. dos Santos, C. Speybrouck, and E. Costante, "Cybersecurity in building automation systems (BAS)," Forescout, Whitepaper, version 08_20, 2020. [Online]. Available: https://www.forescout.com/resources/bas-research-report-the-current-state-of-smart-building-cybersecurity-2/ (visited on 02/14/2023).

[58] B. Al-Sada, A. Sadighian, and G. Oligeri, "Analysis and characterization of cyber threats leveraging the MITRE ATT&CK database," *Access*, vol. 12, 2024. DOI: 10.1109/ACCESS.2023.3344680.

[59] W. Shim, L. Allodi, and F. Massacci, "Crime pays if you are just an average hacker," in *Conference on Cyber Security*, 2012. DOI: 10.1109/CyberSecurity.2012.15.

[60] K. Huang, M. Siegel, and S. Madnick, "Systematically understanding the cyber attack business: A survey," *Comput. Surv.*, vol. 51, 2018. DOI: 10.1145/3199674.

[61] D. Eidemuller, *Nuclear power explained*. Springer Nature, Aug. 2021. DOI: 10.1007/978-3-030-72670-6.

[62] National Research Council, *Digital Instrumentation and Control Systems in Nuclear Power Plants*. National Academies Press, Apr. 1997. DOI: 10.17226/5432.

[63] K. E. Herold, R. Radermacher, and S. A. Klein, *Absorption chillers and heat pumps*, 2nd ed. CRC press, Feb. 11, 2016, ISBN: 978-1-4987-1435-8.

[64] M. Conti, D. Donadel, and F. Turrin, "A survey on industrial control system testbeds and datasets for security research," *Commun. Surv. & Tutor.*, vol. 23, 2021. DOI: 10.1109/COMST.2021.3094360.

[65] A. P. Mathur and N. O. Tippenhauer, "SWaT: A water treatment testbed for research and training on ics security," in *Workshop on Cyber-physical Systems for Smart Water Networks*, 2016. DOI: 10.1109/CySWater.2016.7469060.

[66] C. M. Ahmed, V. R. Palleti, and A. P. Mathur, "WADI: A water distribution testbed for research in the design of secure cyber physical systems," in *Workshop on Cyber-Physical Systems for Smart Water Networks (CySWater)*, 2017. DOI: 10.1145/3055366.3055375.

[67] M. Krotofil, A. Isakov, A. Winnicki, D. Gollmann, J. Larsen, and P. Gurikov, "Rocking the pocket book: Hacking chemical plants for competition and extortion," Whitepaper, Aug. 2015, Blackhat. [Online]. Available: https://www.blackhat.com/docs/us-15/materials/us-15-Krotofil-Rocking-The-Pocket-Book-Hacking-Chemical-Plant-For-Competition-And-Extortion-wp.pdf (visited on 02/17/2023).

[68] H. Esquivel-Vargas, M. Caselli, G. Laanstra, and A. Peter, "Putting attacks in context: A building automation testbed for impact assessment from the victim's perspective," in *DIMVA*, 2020. DOI: 10.1007/978-3-030-52683-2_3.

[69] C. J. Deloglos, C. R. Elks, and A. Tantawy, "An attacker modeling framework for the assessment of cyber-physical systems security," in *SAFECOMP Conference*, 2020. DOI: 10.1007/978-3-030-54549-9_10.

[70] Y. Roumani, J. K. Nwankpa, and Y. F. Roumani, "Examining the relationship between firm's financial records and security vulnerabilities," *Int. J. of Information Management*, 2016. DOI: 10.1016/j.ijinfomgt.2016.05.016.

[71] M. E. Byvaikov, E. F. Zharko, N. E. Mengazetdinov, A. G. Poletykin, I. V. Prangishvili, and V. G. Promyslov, "Experience from design and application of the top-level system of the process control system of nuclear power-plant," *Automation and Remote Control*, vol. 67, 2006. DOI: 10.1134/s0005117906050067.

A. Appendix

A.1. List of Cyber-security Incidents

Src.	ID	Incident Name
BAS	c14a63df-...-e2f934ecb6cf	DDos to Finnish heating...
BAS	a90bc67e-...-d5764bebe268	Hack on security cameras
BAS	5e72ec7f-...-065f447b1f86	Stadium SEA Games cam...
BAS	1ac2cba7-...-7fde513eb813	Ransomware in hotel lock...
BAS	2c01e5c1-...-d4c1419e2503	Hack on heating system in...
BAS	ca6f2dfa-...-d13f4a3bc879	BrickerBot IoT botnet
BAS	aae330f9-...-18ce93589705	Silex IoT botnet
BAS	d8daa229-...-348ebc62e190	Hack on Dallas emergency...
BAS	e7572a67-...-c742ee53dc7f	EMS of government facilit...
BAS	8877d596-...-91551ac55d18	KNX-based smart building...
ICS	0c7720ef-...-ba08313293ae	Malware Targets Uranium...
ICS	6a8726c8-...-1f50d95ed1f1	Steel plant infected with...
ICS	1cb89c03-...-91b77aa9c77d	Irainian [sic] cyber attack...
ICS	a578cd69-...-981b4399184d	TGB water station hacks
ICS	d35b7f52-...-19b7820bc5a7	Control systems device...
ICS	bd9478c1-...-f2ad4e228e51	German Steel Mill Cyber...
ICS	6f9c6c61-...-6ec1ae360470	Russian-Based Dragonfly...
ICS	c8e89e94-...-faefdf8396aa	Ukraine BlackEnergy 3...
ICS	13d2b0de-...-a77a04f3d1fd	Kemuri Water Company
ICS	ef2fc8dc-...-54365be5cf1f	Industroyer/CRASHOVE...
ICS	298e4a0b-...-a2263f5cb29e	TRITON attack on Saudi...
ICS	8d4c80a7-...-ecf60c8d62c7	Dragonfly 2.0 attacks
ICS	0187d507-...-a91bd1f7ac5c	Cyber attack on U.S. Pow...
ICS	f3e5f21c-...-8c797f3566c6	Honda EKANS ransom...
ICS	e24ac06d-...-19479e65075d	Cyber attack on Israels...

TABLE 3. LIST OF INCLUDED CPS INCIDENTS FROM BAS [48] AND ICS [47]; FOR MORE INFORMATION, SEE ARTIFACT REPOSITORY [14]

A.2. Attacker-Capabilities (*Cap*)

A.2.1. Dimensions. We define our own attack cost metric, focused on the type of required attacker knowledge, i.e., the type of knowledge an attacker needs to bring or obtain during the cyberattack, to be able to successfully obtain physical impact at the end. We call this the "(minimum) required *Cap*".

To be able to capture *Cap* in a structured way, we build our metric on top of the ATT&CK Framework. Each dimension of our attack cost schema is made of an ATT&CK for ICS Tactic. The value of an ATT&CK Tactic is determined by rating all Techniques used by the attackers as part of their attacks. Finally, all *Cap* can be summarized in a total *CAPABILITY* score.

A.2.2. Values. For each step (i.e. ATT&CK Tactic) of the attack, we want to measure the CPS-knowledge required to implement the respective attack step. For this purpose, we distinguish between three types of CPS-knowledge with associated level and the *Cap*-score captures the maximum level of CPS-knowledge used;

No CPS knowledge. There are attack steps or even whole attacks for which the attacker did not have any CPS knowledge or experience. The example given in the Introduction is such a case. When it is proven, or very likely, that an attacker did not have or did not use any CPS-specific knowledge, we encode this with the value `low` (-1). Please note, that we do not try to capture the overall offensive cybersecurity skills used by the attacker, meaning that even highly experienced IT attacker activities could obtain a low score.

Technical knowledge comprises conceptional knowledge on how CPS work and familiarity with the device vendors, devices and hardware, controllers, software, human-machine interfaces (HMI), how control logic works, how control programs are written and deployed, device and point configuration, how points are configured, network layout, and network protocols. While possibly vendor or product(family) specific, it does not require understanding of the physical process under control. As such, it is not CPS domain specific. CPS technical knowledge may be obtained from product manuals, experience built up working with CPS, or, if applicable, open standards.

The mentioned articles [18], [19] on automation of vulnerability discovery rely on technical knowledge. Independent of the physical process, they solely depend on the device and software/library versions; the former, for instance, scans the PLC memory layout for use of libraries with known vulnerabilities. Examples of attacks that required, among others, technical knowledge are the masquerading and evasion techniques used by Garcia, Brasser, Cintuglu, *et al.* [42] or Stuxnet [1] as developing these required detailed knowledge of the PLCs including the internal memory layout of the PLC and how the control program is invoked on the PLC. In case of Stuxnet, the attackers also knew operating conditions that would cause severe damage to the equipment [1]. Rocchetto and Tippenhauer capture technical CPS knowledge as *system knowledge* and further break it down into knowledge of *source code*, (network) *protocols*, and *credentials* [23]. Attacks relying solely on technical CPS knowledge can be

reused across sectors and domains, as long as the target CPS is composed of devices from the same device-family and runs similar software versions. When technical CPS knowledge was shown or required, we encode this with the value `medium` (0).

Process knowledge comprises intricacies of the physical process. Operators use it to supervise a process and to steer needed corrective measures (e.g., using an HMI). Process knowledge is valid within a single domain or sector. For example, an operator in a water purification facility is familiar with the individual process steps and knows when to add flocculation agents or disinfectants (e.g., Chlorine), in which amounts, and for how long these agents need to remain in the water. The operator does not need to know how the control program is implemented or what network protocols are used; no technical CPS knowledge is required. Being tied to the domain, this knowledge does not easily translate to other domains or physical processes: For example, expert knowledge in water purification does not imply expert knowledge in the operation of drinking water distribution networks.

Garcia, Brasser, Cintuglu, *et al.* use process knowledge (having full testbed specifications) to understand the physical consequences of specific actions and to generate false measurements to hide their activities [42]. The developers of Stuxnet had detailed knowledge of the Uranium enrichment process, including, e.g., the frequency to operate centrifuges [1]. Some process knowledge is publicly available, mostly based on underlying (physical, chemical, mechanical, ...) principles. However, availability decreases with increasing process complexity and specificity, e.g., details on operating nuclear power plants [61], [62] or Uranium enrichment facility are scarcer than explanations and models of heat pumps and fridges [63].

To better understand and evaluate the physical impact of cyberattacks, numerous publications delved into process-based anomaly detection and the development of physics- and simulation-in-the-loop testbeds or honeypots [64]. Examples include the Secure Water Treatment (SWaT) and Water Distribution (WADI) testbeds [65], [66], the simulation of a chemical process [67], or a building automation testbed [68].

When process knowledge was shown or required, we encode this with the value `medium` (0).

Process-mapping knowledge constitutes how the technical setup manipulates the physical process. It concerns the exact devices, setpoints, and protocol messages involved in obtaining a specific physical impact, e.g., to open a specific valve for a certain duration, releasing this much Chlorine into the water. It requires both basic technical and process knowledge at least.

This knowledge is specific to a single setup, site, or plant (or a few, if multiple sites are built identically) and thus cannot be reused in other CPS deployments. It has to be collected on-site. While this knowledge is available to some (but not all) process engineers (read: insider), with enough time, it can also be obtained during an attack by an attacker with at least moderate technical CPS knowledge, general process knowledge and a way to extract information from the respective CPS deployment. Sources can be information repositories, project files on operator workstations, or the collection of screenshots of HMIs and recordings of the commands sent over the network while operators interact with the human-machine interface (HMI). ATT&CK for ICS mentions possible ways to obtain this type of knowledge as part of *Collection*. [8]

Green, Krotofil, and Abbasi and Deloglos, Elks, and Tantawy respectively write about the difficulty of obtaining what they call "process comprehension" [15] and the iterative way in which attackers try to obtain this process-mapping knowledge [69]. Qin, Rosso, Cardenas, *et al.* reverse-engineered partial process-mapping knowledge by passively observing management traffic in a SCADA network [33]. Several other publications, while not addressing its acquisition, do use process-mapping knowledge. Garcia, Brasser, Cintuglu, *et al.*, for example, know which setpoints to manipulate to execute their attack, and which sensor readings to forge to remain undetected [42]. Process-mapping knowledge is used in many of the highly sophisticated targeted CPS attacks, including Stuxnet [1], Triton, and Black Energy 3/Ukraine Power Grid. A `high` (+1) encodes that *process-mapping* knowledge (and thus also technical and process knowledge) is involved or acquired.

A.2.3. Rater Instructions.

Initial Access

ATT&CK Tactic TA0108 captures techniques attackers can use to obtain access to a CPS system.

- **`low`** attacker obtains access using IT-knowledge and tools, e.g., by exploiting a Windows computer in the CPS network or an exposed webserver running on a controller: Exploit Public-Facing Application (T0819).
- **`medium`** attacker demonstrates *technical* CPS knowledge e.g., exploiting CPS network protocols, or utilizing CPS-specific control software. It is unclear, how *process* knowledge can help an attacker obtain initial access.
- **`high`** attacker demonstrates insights into process mappings. Techniques Rogue Master (T0848) and Supply Chain Compromise (T0862) can indicate *process mapping* knowledge.

Execution

ATT&CK Tactic TA0104 captures techniques attackers can use to obtain the ability to execute commands or assume (partial) control over a device.

- **`low`** attacker does not show any CPS-specific capabilities. Execution can, for example, be obtained via shell commands (Command-Line Interface, T0807) or by running arbitrarily interacting with a GUI (Graphical User Interface, T0823).
- **`medium`** attacker demonstrates *technical* CPS knowledge. For example, Hooking (T0874) or Modify Controller Tasking (T0821) require knowledge over the target platform. It is unclear, how *process* knowledge can help an attacker to achieve Execution.
- **`high`** attacker demonstrates insights into the process mappings. It is unclear how *process-mapping*

8. In particular, *Data from Information Repositories* (T0811), *Point & Tag Identification* (T0861), *Adversary-in-the-Middle* (T0830), and *Screen Capture* (T0852)

knowledge can help an attacker to achieve Execution.

Persistence

ATT&CK Tactic TA0110 captures techniques attackers can use to persist their access to a device.

low attacker does not show any CPS-specific capabilities. Hardcoded Credentials (T0891) may be an example that allows to obtain persistence without CPS-specific knowledge.

medium attacker demonstrates *technical* CPS knowledge, e.g., using Module Firmware (T0839) or Project File Infection (T0873). It is unclear how *process* knowledge can help an attacker to achieve Persistence.

high attacker demonstrates insights into process mappings. Depending on how the attacker implements Project File Infection (T0873), *process-mapping* knowledge can be utilized.

Privilege Escalation

ATT&CK Tactic TA0111 captures techniques attackers can use to extend or elevate their permissions.

low attacker does not show any CPS-specific capabilities. Exploitation for Privilege Escalation (T0890) may be implemented without CPS-specific knowledge.

medium attacker demonstrates *technical* CPS knowledge, e.g., Hooking (T0874). It is unclear how *process* knowledge can help an attacker to achieve Privilege Escalation.

high attacker demonstrates insights into process mappings. It is unclear how *process-mapping* knowledge can help an attacker to achieve Privilege Escalation.

Evasion

ATT&CK Tactic TA0103 captures techniques attackers can use to hide their presence or disguise their activities.

low attacker does not show any CPS-specific capabilities. Techniques Masquerading (T0849) or Rootkit (T0851) can be implemented without CPS-specific knowledge.

medium attacker demonstrates *technical* CPS knowledge, e.g., building a Rootkit (T0851) may require *technical* ICS knowledge, or *process* knowledge, e.g., some cases of Spoof Reporting Message (T0856) may require *process* knowledge.

high attacker demonstrates insights into process mappings. Advanced cases of Spoof Reporting Message (T0856) may require *process-mapping* knowledge, especially when the attacker needs to spoof a complex physical process.

Discovery

ATT&CK Tactic TA0102 captures techniques attackers can use to obtain information about other network devices.

low attacker does not show any CPS-specific capabilities. Techniques Network Sniffing (T0842) Remote System Discovery (T0846) can be implemented without CPS-specific knowledge, e.g., using common IT tools like tcpdump or nmap.

medium attacker demonstrates *technical* CPS knowledge. For example, Techniques Network Sniffing (T0842) Remote System Discovery (T0846) can be implemented in ways that require *technical* knowledge, e.g., when the attacker passively sniffs ICS-specific protocols and extracts remote system information from domain-specific protocols. It is unclear how *process* knowledge can help an attacker to achieve Discovery, however this type of knowledge may help an attacker during Discovery e.g., because the attacker may know what to look out for.

high attacker demonstrates insights into process mappings. It is unclear how *process-mapping* knowledge can help an attacker to achieve Discovery.

Lateral Movement

ATT&CK Tactic TA0109 captures techniques attackers can use to laterally move between devices and networks.

low attacker does not show any CPS-specific capabilities. Many common IT techniques apply, including e.g., Default Credentials (T0812), Lateral Tool Transfer (T0867), or Remote Services (T0886) using common (IT-typical) protocols (e.g., Telnet, SSH, or RDP).

medium attacker demonstrates *technical* CPS knowledge, e.g., by implementing Remote Services (T0886) using a domain-specific protocol. It is unclear how *process* knowledge can help an attacker to achieve Lateral Movement.

high attacker demonstrates insights into process mappings. It is unclear how *process-mapping* knowledge can help an attacker to achieve Lateral Movement.

Collection

ATT&CK Tactic TA0100 captures techniques attackers can use to collect information about devices and physical processes. If the attacker does not have access to insider knowledge, Collection is necessary to obtain *process-mapping* knowledge of a CPS.

low attacker does not show any CPS-specific capabilities. Some common IT techniques apply, e.g., Screen Capture (T0852).

medium attacker demonstrates *technical* or *process* knowledge. While it seems plausible that an attacker obtains these types of knowledge during Collection, it seems more likely that an attacker could utilize such knowledge to speed-up Collection e.g., because they know which pieces of information are relevant and how to obtain them.

high attacker demonstrates or obtains *process-mapping* knowledge. Most, if not all Collection techniques aim to extend *process-mapping* knowledge, three arbitrarily selected techniques are Screen Capture (T0852), Point & Tag Identification (T0861), and I/O Image (T0877).

Command and Control

ATT&CK Tactic TA0101 captures techniques attackers can use to implement command and control capabilities.

low attacker does not show any CPS-specific capa-

bilities. Some common IT techniques apply, e.g., usage of a Standard Application Layer Protocol (T0869).

medium attacker demonstrates *technical* CPS knowledge, e.g., by using a Standard Application Layer Protocol (T0869) or a Commonly Used Port (T0885) that is unique to the soft- and hardware used in the target CPS. It is unclear how *process* knowledge can help an attacker to achieve Command and Control.

high attacker demonstrates *process-mapping* knowledge. It is unclear how *process* knowledge can help an attacker to achieve Command and Control.

Inhibit Response Function

ATT&CK Tactic TA0107 captures techniques attackers can use to impair the ability of the control system to maintain control over the physical process.

low attacker does not show any CPS-specific capabilities. Some common IT techniques apply, e.g., Denial of Service (T0814) or Device Restart/Shutdown (T0816).

medium attacker demonstrates *technical* CPS knowledge e.g., with Activate Firmware Update Mode (T0800). Alternatively, an attacker might use Alarm Suppression (T0878), Block Command Message (T0803), or Block Reporting Message (T0804) using *technical* ICS knowledge of the used vendors, protocols, and software. It is unclear how *process* knowledge can help an attacker to achieve Inhibit Response Function.

high attacker demonstrates *process-mapping* knowledge. Techniques like Manipulate I/O Image (T0835) and Modify Alarm Settings (T0838) require the attacker to have detailed understanding of the impact of the current device configuration on the physical world and how to manipulate it to their advantage.

Impair Process Control

ATT&CK Tactic TA0106 captures techniques attackers can use to alter or manipulate normal execution of the physical process in favor of the attacker.

low attacker does not show any CPS-specific capabilities. An attacker without any CPS-specific knowledge may be able to implement Unauthorized Command Message (T0855) or Modify Parameter (T0836) e.g., using an exposed web GUI by arbitrarily changing values without understanding the consequences.

medium It is unclear how *technical* CPS knowledge can help an attacker to achieve Impair Process Control. An attacker with awareness of *process* knowledge may implement Spoof Reporting Message (T0856) or Unauthorized Command Message (T0855) to either trick the CPS or operators into performing reactions or by directly sending commands to the controllers that manipulate the process in a desired way. Any manipulation of a physical process done with an intended goal, implies *process* knowledge.

high attacker demonstrates *process-mapping* knowledge. Techniques like Manipulate Modify Parameter (T0836) or Unauthorized Command Message (T0855) imply *process-mapping* knowledge, if the attacker directly addresses and manipulates actuators in a targeted way (i.e., the attacker knows that manipulating a specific point has a certain impact on the physical process or process control system).

Impact

ATT&CK Tactic TA0105 captures techniques attackers can use to obtain impact on the control system and the physical process.

low attacker does not show any CPS-specific capabilities. Techniques like Loss of Availability (T0826) or Denial of View (T0815) can be implemented with IT-techniques only, e.g., by bricking devices or using network DDoS.

medium attacker demonstrates *technical* CPS knowledge, e.g., by abusing device-, vendor-, or software-specific functionalities e.g., to obtain Loss of Availability (T0826) or Loss of Control (T0827). One example is resetting devices or setting passwords. An attacker may require *process* knowledge to implement Manipulation of Control (T0831) in a meaningful way, as one needs understanding of the process to manipulate it.

high attacker demonstrates *process-mapping* knowledge. Advanced implementations of Manipulation of Control (T0831) and Manipulation of View (T0832) may require *process-mapping*. One example being Stuxnet that not only sent specific commands to a specific set of actuators to destroy centrifuges, but also manipulated the reported values, implying that the developers knew exactly how these points relate to the physical process.

A.3. CPS-Context (*Ctx*)

A.3.1. Dimensions. To capture *Ctx* we extract eight characteristics from the ENISA [29] report on "Good Practices for Security of IoT" and an exploratory literature survey of academic publications, international standards, and CPS vendors' product descriptions. Each context characteristic is expected to have a security impact, i.e., a change is expected to result in a change of attacks and required *Cap*. Even though our eight interdisciplinary *Ctx* factors cover a wide range of aspects, we do not claim completeness, as there may be other (measurable) factors not captured by our characteristics.

Physical Protection refers to measures taken to restrict (unauthorized) physical access to equipment and protect devices from physical tampering. It includes, among others, whether devices are located in a publicly accessible or access-restricted area, whether the devices are physically enclosed, and whether (and if so, how) access restriction is enforced. The ISO 27002:2022 standard dedicates an entire chapter exclusively to security controls related to physical access [27, ch. 7] and ISO 27019:2017 further extends this specifically for process control systems in the energy sector [4, ch. 11]. The U.S. Industrial Control Systems Cyber Emergency Response Team (ICS-CERT) links

physical access to CPS cybersecurity [32, ch. 2.3] and states that "maturity varies for site facilities based on criticality/culture" [32, tab. 1]. According to the *ICS-CERT Monitor* 2017 report, ICS-CERT assessments found physical access control to be among the top five most common weaknesses in CPS in 2014, 2016, and 2017 [28, tab. 1].

Device Protection concerns on-device cybersecurity measures such as user authentication, running secure software, and host-based security monitoring. Both ISO 27001:2022 and ISO 27019:2017 list security controls regarding access control, vulnerability and patch management, and implementation of (configuration) backup policies [4], [27]. The U.S. Homeland Security ICS-CERT lists "host security", including patch management, and (host-based) "security monitoring" as important factors of CPS cybersecurity [32, ch. 2.6 f.] and identify "known vulnerabilities" as major risk factor [32]. ATT&CK for ICS mentions possible mitigation strategies to increase the level of device protection.[9] According to the 2017 *ICS-CERT Monitor* report, weaknesses related to Access Control and User Authentication appeared in the top four most common weaknesses in 2014, 2015, 2016, and 2017: Often passwords are weak, identical among all devices, and shared among all CPS staff, bypassing security best-practices [28, tab. 1].

Network Protection concerns network architecture and segregation (e.g., air-gapping), effective implementation using e.g., VLANs, firewalls, as well as (network) security monitoring and intrusion detection systems. ATT&CK for ICS[10] and the ICS-CERT both recommend best practices for network architecture, -security mechanisms, and security monitoring [32, ch 2.4, 2.5, 2.7]. The 2017 *ICS-CERT Monitor* report lists network "boundary protection" as the most common weakness every single year, starting from 2014 [28, tab. 1].

Device Rarity captures how many types of devices are present and how common they are; whether devices and software are from well-established, globally operating brands or smaller regional CPS vendors (see [33]). There are also geographical differences. For example, the SCADA protocol DNP3 is historically more prominent in Northern America, while IEC-104 (and its predecessor IEC-101) is more prominent in Europe [34], [35]. Qin, Rosso, Cardenas, *et al.*, describe a European gas distribution network (running IEC-104) in which all remote terminal units (RTU) were exclusively manufactured by a local Dutch vendor [33]. In contrast, the BACnet protocol or the Tridium Niagara controllers are globally present.

To maximize profits, cyber criminals re-use tools, infrastructure, and knowledge [36]–[39]. Assuming comparable effort, it is economical to target devices with higher market penetration, as the predicted return on investment is higher [36]–[39]. Numerous studies link diversity to resilience [40], [41] and the operational risk of a monoculture with a singular market leader. One study even found a correlation between company profits and the number of vulnerabilities in their products [70]. Best practices suggest patch and vulnerability management [32, ch 2.6.1] to limit attackers' abilities to re-use tools and ATT&CK proposes measures to identify and prevent re-use of known vulnerabilities.[11]

Contractor Independence measures how many (external) parties are regularly interacting with the CPS and how much the process depends on external factors, providers, or resources. In some domains, stakeholders are commonly distinct parties (e.g., manufacturer, CPS-operator, -integrator, -maintainer, -owner). Outsourced maintenance can make ownership and (security) responsibilities unclear and introduces remote access functionalities that would otherwise not be necessary. Standard ISO 27002:2022 highlights the relationship between supply chains or external contractors, and the associate risk to information security. [27, ch. 5.19 et seq.]. Controls specific to the electricity sector are listed in ISO 27019:2002 [4, ch. 15].

Process Complexity considers the complexity of the physical process and control program, and whether information about the process is publicly accessible. Complexity varies a lot; from simple IoT devices like a smart industrial fridge to complex systems comprised of multiple interacting sub-systems like a nuclear power plant [71][62, Fig. 1-1]). With increasing process complexity, learning the *process-mapping* becomes more difficult, requires better *process* and *technical* knowledge, as well as more time and resources to collect and *understand* the obtained data. As such, the lack of "process comprehension" [15] makes it harder for attackers to obtain meaningful and targeted physical impact. In fact, many academic publications consider full process comprehension by default [15], [18], [19], [42].

Safety Monitoring regards checks for critical process states. With human supervision, a process engineer can manually coordinate corrective actions. In high-risk contexts, an independent control system, the so-called Safety-Instrumented System (SIS), automatically detects critical process states and immediately executes the appropriate fail-safe mechanisms [43]. However, safety systems are not present in all CPS. In BAS, for example, dedicated safety systems are uncommon, as failure of the "normal" system often has very little safety impact.[12]

ATT&CK for ICS mentions additional protection layers for hazard scenarios and non-digitally controlled safety mechanisms to contain the (physical) impact of cyberattacks.[13] The TRITON malware became famous as the first cyberattack that targeted and manipulated the SIS of a CPS [44]. The German BSI describes a 2014 cyberattack against a German Steel Mill "led to the uncontrolled shutdown of a blast furnace, leaving it in an undefined state and resulting in massive damage" [45, ch 3.3.1], very likely triggering the SIS. In another case, a very simple cyberattack was noted by the process operator and real physical impact was inhibited by the process operator implementing corrective actions. [14]

9. Update Software (M0951), Operating System Configuration (M0928), Password Policies (M0927)

10. Network Segmentation (M0930), Network Intrusion Prevention (M0931), Filter Network Traffic (M0937), Network Allowlists (M0807)

11. Vulnerability Scanning (M0916) and Threat Intelligence Program (M0919)

12. In most buildings, safety-critical features work independently of the BAS and do not require the BAS to be operational (e.g., an emergency exit door can be monitored via BACnet but in state of emergency, the door can always be unlocked, physically, regardless of the state of the BAS).

13. Safety Instrumented Systems (M0812) and Mechanical Protection Layers (M0805)

14. https://www.wired.com/story/oldsmar-florida-water-utility-hack/

Legislation and Regulation captures the presence of authority-imposed requirements or standards that form an extrinsic motivation to guarantee properties of the physical process (e.g., availability or quality). Regulations pose minimal requirements that have to be achieved, often regardless of economic interests of the responsible company or organization. Critical infrastructure such as drinking water purification and distribution of water and of natural gas are examples of regulated domains where many countries specify quality and availability requirements for the delivered resource. For example, the Netherlands regulates disruptions to drinking water distribution (expected to be) longer than 24 hours by law.

Standards ISO 27002:2022 and ISO 27019:2022 have dedicated sections on compliance and legal requirements [4], [27]. Especially ISO 27019 mentions compliance and legal obligations of the physical process in context of cybersecurity.

Based on the cited literature, we conjecture these characteristics positively correlate with the CPS capabilities required for attacks with physical impact.

A.3.2. Values. As we require a simple scoring system, we define three categorical values, `high`, `medium`, and `low` and assign them numerical values $+1, 0,$ and -1 respectively.

High value. The `high` value ($+1$) is used to indicate that a *Ctx* factor is implemented or present correctly and completely, at a high maturity level. This often implies the presence of regular checks and enforcement. For technical *Ctx* factors, this means that the CPS operator implemented these correctly, for external factors like e.g., legislation and regulation, this implies that the CPS operator needs to comply with strict regulations that are enforced and successfully achieve the intended purpose.

Medium value. A `medium` value (0) depicts implementation of a *Ctx* factor following basic or minimum-level best-practices or implementations.

Low value. The `low` value (-1) indicates that a *Ctx* factor is not implemented at all, or implementation does not have any notable effect, e.g., because measures are easily circumvented. Examples are e.g., a lack of any form of network segregation (network protection), no or default credentials (device protection), the absence of any safety system (safety monitoring) or the absence of any effective supervision from a public authority (legislation and regulation). This value is assigned for *Ctx* characteristics that are implemented below general best-practices.

A.3.3. Rater Instructions.

Physical Protection

Depicts whether a device is easily accessible in a physical sense or not.

- **`low`** devices in public space or easily accessible to the public; *Example:* building automation, gas distribution street closets, wireless networks, ...
- **`medium`** devices not easily accessible (physically) e.g., because they are located in a restricted area. *Example:* factory assembly line, elevator control room, gas/power/water distribution stations, ...
- **`high`** devices not easily accessible, attacker would have to overcome multiple layers of security before obtaining physical access. *Example:* power plant

Device Protection

Depicts whether the application(s) running on the device are resilient against cyber attacks or not.

- **`low`** devices with default passwords, outdated and known to be vulnerable software libraries. *Example:* IoT-cameras with known authentication bypass, unpatched controllers, ...
- **`medium`** authentication is required, devices have (non-default) passwords. There are no exposed vulnerabilities than can be exploited easily. *Example:* (secure) remote access gateways for BAS
- **`high`** devices known to be difficult to compromise e.g., because they are hardened or designed for public-facing scenarios. Authentication with strong password protection, no exploitable software bugs, etc. *Example:* Secure elements

Network Protection

Depicts whether a device is easily accessible by means of a cyberattack (i.e., over the network) or not.

- **`low`** devices facing the public Internet or any other insecure network without proper firewalling or security measures. *Example:* unprotected or default-password remote management access is exposed over the Internet
- **`medium`** devices are not directly facing the Internet, or the exposed protocols are well restricted and authentication is required. Devices and networks are properly segregated using firewalls and an attacker needs to implement MITRE ATT&CK Tactic Lateral Movement to successfully compromise or interact with the device. *Example:*factory
- **`high`** network is properly segmented and there are trust boundaries. Attackers would need to show skills in Lateral Movement to move between individual network segments (without being detected) or get access to the device in another way. *Example:* power-plants

Device Rarity

Depicts how "local" or "global" the market is, with respect to software and hardware. Some devices like CCTV IP cameras typically run a Linux kernel and are typically the same all over the world, devices like Siemens S7 are globally available but less common, BACnet and KNX are only common in specific regions, and when looking at critical infrastructures we expect to see local vendors.

- **`low`** devices and network protocol implementations are very common across the globe. One can expect that many other devices are identical or run the same or very similar software libraries and/or protocol stacks. *Example:* IoT IP-cameras running Linux kernel, standard web server, ...
- **`medium`** devices or software libraries are not common or de-facto standard on an international scale. *Example:* railway ICs devices, factory automation, power grid RTUs, ...
- **`high`** devices and software libraries are very re-

gional, manufactured from national or local companies. *Example:* RTUs for gas or water distribution networks.

Contractor Independence

Depicts how much the process relies on external stakeholders (example: in building automation there is the network architect, the integrator, the physical maintenance contractor, the network provider the client who runs the building, potentially camera/elevator/... operators). This concept includes service contractors to operate the CPS, peering-partners in case of a distributed CPS, and also dependencies for the physical process itself.

low there are many dependencies on contracted parties, as a result ownership is not always clear, change requests take a long time and are implemented by externals.*Example:* building automation networks.

medium there are a few, well selected contracted parties, the roles are clear and the associated risks are well-documented and acceptable. Some interruptions are expected when then contractors fail to meet agreements but the physical process is expected to revert to normal within a reasonable time. *Example:* district heating.

high all core functionalities and competencies are in-house or, if they come from a third party, timely available, strongly supervised and monitored. The process is expected to continue with acceptable restrictions in case dependencies fail or try to sabotage the process. *Example:* power-plants.

Process Complexity

Depicts how complex the controlled physical process is and therefore how complex it is to learn it (i.e., the process). This metric also includes how much of this information is given from the labels on the HMI or sniffing the network traffic.

low the process is fairly simple and comparably easy to learn (e.g., because it is limited to a single device), even for somebody with no control system familiarity (but security/networking experience/knowledge). *Example:* fridge, CCTVs, ...

medium some effort is needed to learn to understand the physical process and how the physical process is implemented on this instance ("process mapping"), as the process is not trivial (e.g., because it is distributed or contains many related in- and outputs). Valid for "subsystems". *Example:* building automation network for an entire building, one production belt in a factory, a specific sub-system of e.g., a battleship

high ICS-technical and ICS-process background knowledge is required to be able to understand the process in a reasonable amount of time (i.e., within months). Understanding the system without proper documentation requires a high amount of effort. *Example:* stuxnet Iranian uranium enrichment facility, entire power plant, airplanes, battleships, ...

Safety Monitoring

Whether a dedicated safety system is in place, either in form of a human supervisor or an automated system monitoring the physical process and taking over process control to avoid desaster.

low there is no such safety monitoring system in place. *Example:* building automation network

medium a safety-system of some form is in place, ensuring that the monitored system or process cannot reach critical states or notify the operator for intervention, or automatically run fail-safe operations. Systems can, but do not have to be, fully mechanical. *Example:* emergency pressure release valves, airplane safety-systems notifying the pilot immediately

high a safety system is in place and properly configured. The system automatically applies safe-default and fail-safe operation when needed, making worst-case scenarios virtually impossible. *Example:* nuclear power plant, railway interlocking

Legislation and Regulation

Whether a dedicated authority oversees the physical process and/or safety, and how strict the regulatory boundaries are.

low No or very few legislative boundaries related to safety or the physical process. *Example:* building automation

medium Legislative boundaries exist and require the operator of the CPS and physical process to comply with a set of regulations e.g., related to process safety or process continuity. *Example:* human safety in manufacturing

high The process and safety-aspects are strictly supervised and monitored. Deviations (e.g., accidents, failures, cyberattacks, ...) are generally not acceptable and have to be disclosed and analyzed. *Example:* aviation, railway, critical infrastructure (energy, water distribution)

Full rater instructions are included in the data artifact associated with the paper.

A.3.4. Rating Example. We illustrate the rating process with an example for the "KNXlock" incident from the BAS domain. From public reports, it is known that an IP-KNX-gateway connected the KNX building network, supposed to be separated from other networks, to the public Internet, resulting in `low` *Network Protection*. All KNX devices follow a known standard and protocol (KNX), thus have `low` *Device Rarity*. Usually, building networks are maintained and operated by external contractors (cf. *Contractor Independence*), the available reports do not give reason to believe that this was different in this case, thus resulting in another `low` value. Typically, Building Automation networks have comparably simple processes (cf. *Process Complexity*) and do not have dedicated *Safety Monitoring* system.

This partial example shows that, even though public reports lack many details, the quality and quantity of available information is sufficient to perform Context rating.