

On the Role of Latent Design Conditions in Cyber-Physical Systems Security

Sylvain Frey, Awais Rashid, Alberto Zanutto, Jerry Busby, Karolina Follis
Security Lancaster Research Centre, Lancaster University, United Kingdom
{s.frey, a.rashid, a.zanutto, j.busby, k.follis}@lancaster.ac.uk

ABSTRACT

As cyber-physical systems (CPS) become prevalent in everyday life, it is critical to understand the factors that may impact the security of such systems. In this paper, we present insights from an initial study of historical security incidents to analyse such factors for a particular class of CPS: industrial control systems (ICS). Our study challenges the usual tendency to blame human fallibility or resort to simple explanations for what are often complex issues that lead to a security incident. We highlight that (i) perception errors are key in such incidents (ii) latent design conditions – e.g., improper specifications of a system’s borders and capabilities – play a fundamental role in shaping perceptions, leading to security issues. Such design-time considerations are particularly critical for ICS, the life-cycle of which is usually measured in decades. Based on this analysis, we discuss how key characteristics of future smart CPS in such industrial settings can pose further challenges with regards to tackling latent design flaws.

Categories and Subject Descriptors

Software Engineering [D.2.10]: Design

Keywords

cyber-physical systems, industrial control systems, risk, perception, design

1. INTRODUCTION

The role and impact of users on the security of regular IT systems is a common matter of study in literature. For instance [1, 2] investigate how users’ psychological and cognitive biases affect security features and discuss how better system designs should account for these human characteristics. These works identify users’ perception as a critical aspect, including non-malicious behaviours becoming a threat [2].

Industrial control systems (ICS) differ from pure IT systems that are the focus of such works. Firstly, ICS are

cyber-physical systems (CPS) that combine regular software systems with physical ones that control and operate various sensors and actuators impacting their environment. Secondly, ICS have a greater diversity of human roles around them, e.g., operator, technician, maintainer, engineer, manager, etc. instead of mere “end-user”. Finally, and most importantly, ICS’ typical life cycle is counted in decades: the interval between the design and the end of life of a system can span over 50 years. Such long periods leave little room for evolution, due to stringent availability and safety requirements in a number of infrastructures, some of them critical – power grid, water supply, rail network, etc. The design of such systems is therefore a fundamental step with long-term consequences.

Of course, risk perception has been studied in industrial contexts [4, 14] and safety has been a central concern in such settings [6]. However, the increased connectivity of ICS and the emergent smart CPS settings pose challenges not just for safety but for security – with security lapses inadvertently impacting on safety (as our case studies show). To our knowledge, ours is the first work to undertake a socio-technical analysis of perception errors underpinning security issues in CPS.

In this paper, we start by considering the role of ICS operators’ perception during security incidents. We investigate a corpus of 6 case studies to analyse how operators perceive the system and its various parts before and during an incident. Our investigation reveals that perception errors are central to all the case studies we investigated. We identify and classify perception errors and analyse the possible causes and conditions behind them. We show that, beyond individual operator mistakes, latent design conditions [13] play a fundamental role in shaping perceptions, leading to security issues. Our study offers two key insights:

1. We challenge the idea that humans are necessarily the weak link via which most incidents occur. We show that latent design condition are a key factor that shapes operators’ perception, leading to operational mistakes and incidents. Our classification of perception errors provides additional insights regarding different types of latent design flaws, in terms of system borders, capabilities, observability and controllability.
2. We discuss how fundamental characteristics of future smart CPS deployed in such settings can further complicate the early identification and management of latent design flaws.

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

SEsCPS’16, May 16 2016, Austin, TX, USA

© 2016 ACM. ISBN 978-1-4503-4171-4/16/05...\$15.00

DOI: <http://dx.doi.org/10.1145/2897035.2897036>

Case study	Date	Description
Maroochy water services	2000	Attack on city water system by rogue employee causing major sewage spillage [11].
Aurora	2007	Proof of concept that cyber attacks can cause physical damage to ICS [10].
Stuxnet	2007	Silent destruction of uranium centrifuges by a computer worm [8].
Turkish pipeline	2008	Destruction of a pipeline by a cyber-commando [3].
Metcalf sabotage	2013	Commando performs a physical attack on an electricity transmission substation [7].
German steel mill	2014	Advanced Persistent Threat caused the destruction of an industrial furnace [5].

Table 1: Case studies analysed.

2. STUDY SETTING

We base our analysis on a corpus of 6 public case studies summarised in Table 1. The case studies cover a wide range of systems (single appliances, plants, networks) in different sectors (energy, oil, water, nuclear, chemical). The Aurora case is a controlled experiment rather than an actual attack. Though the Metcalf incident is not a cyber attack, there are key system security issues involved (operator reactions to remote alarms and attackers cutting communication cables).

We utilised publicly available data – incident reports, interviews, media reports – to undertake a qualitative analysis of the root causes underpinning the incidents. For our analysis, we used a combination of grounded theory and incident fault trees [12] to iteratively develop a theory of cause-effect relationships across our corpus of case studies. Critical fault tree elements related to various aspects of each incident were iteratively labelled and aggregated into significant categories – the building and coding of new fault trees reflectively shedding new light on previous incident analyses until a saturation point was reached and no new insights or categories emerged.

Mustering six security-related case studies on cyber-physical systems in industrial settings was a significant effort: public literature on the topic is scarce and scattered. Consequently, some case studies are based on limited data. The Turkish pipeline and Metcalf sabotage cases are based on press articles whose accuracy and reliability cannot be ascertained. The German Steel Mill case has been described in a report which provides little technical detail about the attack. In the absence of better information, the stories provided by these sources must be considered with caution – although the popularity of these cases makes them interesting *per se* as important narratives in today’s security culture in cyber-physical environments.

3. FACTORS IMPACTING CPS SECURITY

3.1 Classification of Perception Errors

Our analysis leads us to 4+1 categories of perception errors (cf. Fig. 1 for illustration and Table 2 for examples):

1. *Perception of system qualities and capabilities:* Operators having a bad perception of the intrinsic (static) qualities and capabilities of a system: its safety, its resilience, etc. They overestimate the capability of their system to withstand faults and failures in certain conditions and to stay in an acceptable state while enduring challenging conditions.
2. *Perception of system boundaries:* Operators having a bad perception of the (static) boundaries of a system, regarding its environment and neighbour systems. Misunderstanding the borders of systems – typically

assuming that a system is isolated from another one whereas it is not, either physically or virtually – leads to underestimating the possible mutual influences and knock-on effects between adjacent systems.

3. *System observability:* Operators incorrectly assuming that (at time t) they have access to a complete enough and accurate enough representation of the current state of their system.
4. *System controllability:* Operators incorrectly assuming that (at time t) they are indeed in control of their system. This perception error is often preceded by an observability error or bad perception of system capabilities.

The first two categories are horizontal concerns, i.e. regarding the system layer itself and other systems around it (see Fig. 1). These perceptions also concern static properties of the system that are not supposed to change significantly in operations: these are therefore related to the design of the system and the operator’s understanding of this design *a priori*. On the other hand the last two categories are vertical concerns, i.e. regarding the interactions between the system and its operator (system layer vs. organisation layer in Fig. 1). These are also dynamic factors: whether or not the operator can observe and control the system depends on the current state of the system, its sensors and actuators, its controllers, the type of attack, etc. These are therefore related to operational contingencies and operator practices *during a particular event*.

A *fifth category (+1)* is that of *Attackers actively tampering* with the “observation & control” link between a system and its operators. A canonical example is Stuxnet: the worm intercepted both sensor and actuator data, replaying harmless patterns impersonating a normal behaviour to the operators while silently tampering with the centrifuge, in a subtle way that would not be attributed to a malicious attack. In other words, this category actively alters the *vertical concerns* with regards to the properties depicted by the *horizontal concerns*.

3.2 Influence of Latent Design

Reason [13] identifies two categories of causes for incidents:

- *Active failures*, which encompass human slips, mistakes and procedure violations that have a direct, temporary impact on a system. These acts can generally be explained by human fallibility, although the conditions in which active failures happen can also play a catalysing role. Due to their unpredictability, active failures can only be detected and identified a posteriori.




Figure 1: Modelling the perception of operators.

- *Latent design conditions*, which stem from prior decisions from system designers and organisation management. These conditions can stay undetected for a long time until they affect the system or its operators. The effects range from pure system vulnerabilities to particular conditions shaping operator perception and leading to active failures. Unlike active failures, latent condition can be identified and mitigated a priori, before any incident happens.

As can be seen in Table 2, all the perception errors identified in the case studies proceed from the latter category: flawed system’s capabilities, porosity in system’s boundaries, vulnerabilities in system’s observability and controllability links. Bad latent design is therefore the main driver behind the key mistakes that ultimately end up in security events.

4. CHALLENGES POSED BY SMART CPS

Latent conditions stem from past design choices that can prove extremely difficult and costly to mitigate, as they are related not only to a single person, but to the systems, structures and organisations around them, and to the integration between them all. There has been lots of existing work in the system safety and dependability domains to tackle such latent design conditions, e.g., [9, 15]. As we can observe in our case studies, such latent conditions also impact security in complex cyber-physical systems. One can argue, and assume, that the best practices developed in the dependability literature can also be utilised to tackle such latent conditions and mitigate their impact on security. However, smart CPS – as they become increasingly prevalent in industrial settings, e.g., Industrial IoT – pose new challenges in this regard:

4.1 Dynamically aggregated nature

Smart CPS environments are expected to be highly dynamic, where different sub-systems or intelligent devices can be aggregated on-the-fly to achieve particular requirements or goals, e.g., in organisations that enable a Bring Your Own Device (BYOD) policy. Such a dynamically aggregated setting changes the formerly, largely, static properties of *system qualities* and *system boundaries* – the horizontal dimension in our model in Fig. 1 – to highly dynamic ones. System

boundaries and neighbouring systems can change often, e.g., monitoring systems may be deployed by a regulatory body during visits to a plant to integrate with existing sensors and actuators. These dynamically aggregated systems would also exhibit differing resilience with regards to security, not to mention the increased connectivity and attack surface that may arise from multiple systems coming together. This, in turn, makes it more challenging to understand the system qualities and boundaries, particularly when the system comes under attack and operators need to take actions to maintain its security and safety.

4.2 Autonomous behaviour

Smart CPS, by their very nature, aim to take over various functions and hide that complexity from the user while augmenting the range of possible behaviours for the system. This impacts the controllability and observability – the vertical dimension in our model in Fig. 1 – as the size of possible state spaces expands and operators need to rely on possibly opaque operational details of the industrial processes and the decision-making logic of their smart controllers. These factors make it difficult for the operator to perceive what is going on and take corrective action – all this becoming particularly challenging if an attacker actively tampers with the operator’s perception. In addition to the induced cognitive load, the attack surface and criticality of smart controllers must be taken into account during the design of such systems.

4.3 Multiple stakeholders

Most current industrial environments, at least those in our case studies, are systems that are normally under the control of a single stakeholder. Already we are seeing the emergence of multi-stakeholder environments in smart grids, intelligent transportation systems and so on. As future factory scenarios come to be implemented, it is not inconceivable to envisage situations where an infrastructure provider offers basic facilities including a building management CPS which are then utilised by multiple other parties for manufacturing and production, each with their own smart CPS interacting with shared sensors and actuators. Such a multi-stakeholder setting impacts all aspects of perception in Fig. 1 as the overall environment is no longer under the observation and control of a single operator: system qualities and boundaries can change frequently as complex inter-dependent and overlapping CPS are deployed by different stakeholders.

4.4 From managing latent design to addressing emergent design

As we discuss above, the design of smart CPS cannot be established early on and then evolve at a slow pace over time. In this context, design is more dynamic and opportunistic rather than a pre-conception in a designer’s mind. This makes it challenging to identify and mitigate latent design conditions as the conditions are not so much latent (they do not arise from a past design decision) but *emergent*. Understanding latent design flaws is still highly relevant in such a context: whether latent or emergent, the four categories of perception errors highlighted in this paper can act as a useful basis to understand and reason about complex CPS configurations in dynamically aggregated multi-stakeholder settings:

- System qualities and capabilities need to become first-class elements of system management interfaces, as

Error Category	Case Study	Details
Capabilities	German steel mill	Lack of fail-safe mechanisms.
Capabilities	Aurora	Safety equipment exploited to attack a system.
Boundaries	German steel mill	Porosity between corporate and control networks.
Boundaries	Turkish pipeline	Porosity between surveillance, control and corporate networks.
Boundaries	Metcalfe sabotage	Intrusion detection system did not monitor outside the perimeter of the station.
Observability	Stuxnet	Vulnerabilities in the monitoring system allowed Stuxnet to take control of it.
Observability	Turkish pipeline	Surveillance camera did not catch intruders and served as an entry point.
Observability	Maroochy water services	Lack of proper monitoring and alarm systems.
Observability	Metcalfe sabotage	Intrusion detection system did not detect the attackers, only bullet impacts.
Controllability	German steel mill	Lack of fail-safe controllers.
Controllability	Stuxnet	Vulnerabilities in the control system allowed Stuxnet to take control of it.

Table 2: Analysis of perception error causes.

operators need constant updates and feedback on their evolution, depending on changing environments, opportunistic combinations, stakeholder constraints, etc.

- System borders must be explicit and their modelling must be refined for operators to understand the possible direct and indirect interactions of their system with neighbours.
- The scope of system observation by the operator must be widened and enriched to include relevant elements in the environment, dependencies to and from neighbour systems, as well as potential sources of undesirable side effects.
- Conversely, system control by operators must define its exact scope – in terms of direct operator action on their infrastructure – as well as possible consequences outside of this scope – in terms of side effects on the environment and neighbours.

5. CONCLUSION AND FUTURE WORK

We have presented an initial analysis and categorisation of the latent design conditions that impact operators’ perceptions in industrial CPS settings, leading to security incidents. We have highlighted how the key characteristics of smart CPS will add further complexity to effective management of such design conditions. This poses a fundamental challenge for systems engineering but also offers interesting opportunities for novel reasoning frameworks and approaches. Our future work will focus on such approaches.

6. ACKNOWLEDGEMENT

This study is part of the Mumba project, funded by the Research Institute for Trustworthy Industrial Control Systems (RITICS).

7. REFERENCES

- [1] Anne Adams and Martina Angela Sasse. Users are not the enemy. *Commun. ACM*, 42(12):40–46, 1999.
- [2] Denis Besnard and Budi Arief. Computer security impaired by legitimate users. *Comput. Secur.*, 23, 2004.
- [3] Bloomberg. Mysterious ’08 turkey pipeline blast opened new cyberwar. <http://www.bloomberg.com/news/articles/2014-12-10/mysterious-08-turkey-pipeline-blast-opened-new-cyberwar>.
- [4] J. S. Busby and P. W. H. Chung. In what ways are designers’ and operators’ reasonable-world assumptions not reasonable assumptions? *Process Safety and Environmental Protection*, 81(2):114–120, 2003.
- [5] Bundesamt für Sicherheit in der Informationstechnik. Die lage der it-sicherheit in deutschland 2014. https://www.bsi.bund.de/DE/Publikationen/Lageberichte/lageberichte_node.html.
- [6] Tom Horlick-Jones. Meaning and contextualisation in risk assessment. *Reliability Engineering System Safety*, 59(1):79 – 89, 1998.
- [7] The Wall Street Journal. Assault on california power station raises alarm on potential for terrorism. <http://www.wsj.com/articles/SB10001424052702304851104579359141941621778>.
- [8] Raplh Langner. To kill a centrifuge. <http://www.langner.com/en/wp-content/uploads/2013/11/To-kill-a-centrifuge.pdf>.
- [9] Bev Littlewood and Lorenzo Strigini. Software reliability and dependability: a roadmap. In *22nd International Conference on Software Engineering, Future of Software Engineering Track, ICSE 2000*.
- [10] Inc Mark Zeller, Schweitzer Engineering Laboratories. Common questions and answers addressing the aurora vulnerability. <https://www.selinc.com/workarea/downloadasset.aspx?id=9487>.
- [11] National Institute of Standards and Technology. Malicious control system cyber security attack case study – maroochy water services, australia. http://csrc.nist.gov/groups/SMA/fisma/ics/documents/Maroochy-Water-Services-Case-Study_report.pdf.
- [12] Awais Rashid, Syed Asad Ali Naqvi, Rajiv Ramdhany, Matthew Edwards, Ruzanna Chitchyan, and M. Ali Babar. Discovering “unknown known” security requirements. In *Proceedings of the 38th International Conference on Software Engineering*, 2016.
- [13] James Reason. Human error: models and management. *British Medical Journal*, 2000;320(7237):768-770.
- [14] Paul Slovic. Perception of risk. *Science*, 236(4799):280–285, 1987.
- [15] J. Xu, B. Randell, A. Romanovsky, R. Stroud, A. Zorzo, E. Canver, and F. von Henke. Rigorous development of an embedded fault-tolerant system based on coordinated atomic actions. *IEEE Trans. Computers*, 51, 2002.