When Cyber Got Real: Challenges In Securing Cyber-Physical Systems

(Invited Paper)

Stefano Zanero
Dipartimento di Elettronica, Informazione e Bioingegneria
Politecnico di Milano
Email: stefano.zanero@polimi.it

***Abstract*—The security challenges arising at the layer between their digital and physical components of Cyber-Physical Systems (CPSs) are, at the same time, pressing, important and complex. In this paper, we discuss how the unique nature of CPSs influences and drives security analysis through two case studies. We also discuss how sensors and their networks, being a key element of the aforementioned interaction surface, are critical to this effort.**

I. Introduction

In many real-world systems, computational and physical resources are strictly interconnected: embedded computers receive inputs from the external, physical world through an array of sensors, and in turn govern physical actuators, creating a smart, flexible control loop, capable of adaptation, autonomy and improved efficiency. These systems are commonly and broadly defined cyber-physical systems (CPS) [1].

Common examples include industrial control systems (now particularly interesting, in light of the "Industry 4.0" paradigm), computerized controls of vehicles, wireless sensor networks, and almost all of the devices usually encompassed by the broad term "Internet of Things".

We can identify at least three reasons why security, safety and reliability of CPS are research areas of foremost importance.

First, such systems are often a key component of modern critical infrastructure: They are vital to our societys viability, to its economical and social stability. Typical examples include the smart grid, control systems used by other utilities, transportation systems, but also industrial control systems more in general. Many medical devices, in particular implantable ones, are also becoming CPSs [2], [3], often remotely monitored.

Second, critical CPS have a very challenging and wide threat landscape, ranging from intentional targeted attacks to terrorist acts, from unplanned disruptions to market disturbances, up to and including state-level threats connected to international policy shifts or even cyberwarfare [4]. In [5], half of the interviewed critical infrastructure operators mention what they perceived as politically-motivated cyberattacks, claiming an average of 10 of those over 5 years, causing an average cost of $850,000 per company.

Third, many modern CPS are increasingly autonomous, lacking a human in the control loop. A significant example, which is the subject of a lot of attention at the moment [6], are self-driving cars: highly complex CPSs, with a large array of sensors and actuators, a complex world to navigate in, and external connectivity to the road infrastructure and, often, to the Internet. But many CPSs are increasingly autonomous, and this exacerbates the potential security hazards if a malicious attacker compromises them.

Due to the reasons above, many researchers (both in academia and throughout the industry) developed potential attacks, analyzed vulnerabilities and pointed out flaws. However, such academic results have triggered in some cases limited response throughout the industry. One of the challenges is the lack of a perceived threat level that can justify the investment of significant resources [2], [7]. Conversely, some works in the area may appear over-hyped, or may even unwillingly overstate the potentially dire consequences of an attack. Since security investment is a risk-based decision, in order to bake security in the CPS development phases, as opposed to trying to deliver it as an afterthought, a *risk-based* design process suitable for each environment must be adopted. We summarily describe one, in an automotive use-case, in Section II.

Another very significant aspect is the unique set of challenges arising from the combination of a computational nature (which is, by definition, discrete and adherent to rigid specifications) and of a continuous physical system which is often not easily modeled and not completely understood. In parallel, safety-critical physical systems are commonly designed with safety, not security, in mind, contrarily to the digital systems they are coupled with. The interconnection between such systems creates new attack surfaces that are neither purely physical, nor purely digital. These surfaces cannot be identified if such systems are studied separately. In other words, besides finding digital vulnerabilities (e.g., a buffer overflow bug), we must assess and validate to what extent such vulnerabilities can be exploited in practice to facilitate or enable *new* physical attacks (e.g., instability); and vice-versa. An example of such a technique, applied to the security analysis of industrial robots, is summarized in Section III.

As we will show throughout the remainder of this paper, sensors and their networks, being a key element of the interaction surface between the digital and the physical world, are also on the front line of these new concepts in security vulnerabilities and secure design.

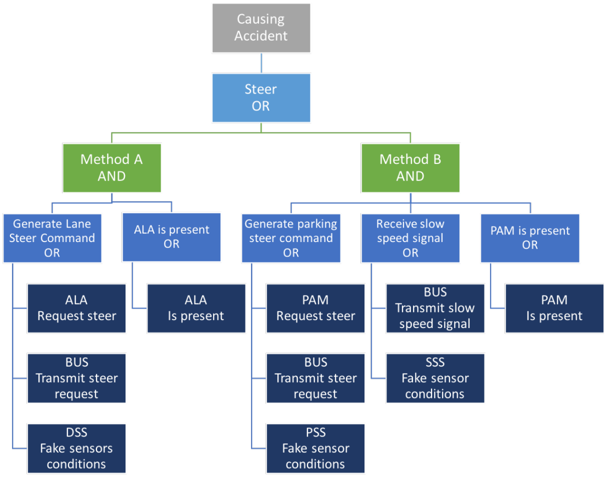
Causing Accident
Steer OR
Method A AND
Method B AND
Generate Lane Steer Command OR
ALA is present OR
Generate parking steer command OR
Receive slow speed signal OR
PAM is present OR
ALA Request steer
ALA Is present
PAM Request steer
BUS Transmit slow speed signal
PAM Is present
BUS Transmit steer request
BUS Transmit steer request
SSS Fake sensor conditions
DSS Fake sensors conditions
PSS Fake sensor conditions

Fig. 1. An example of Attack Tree, as shown in [7]

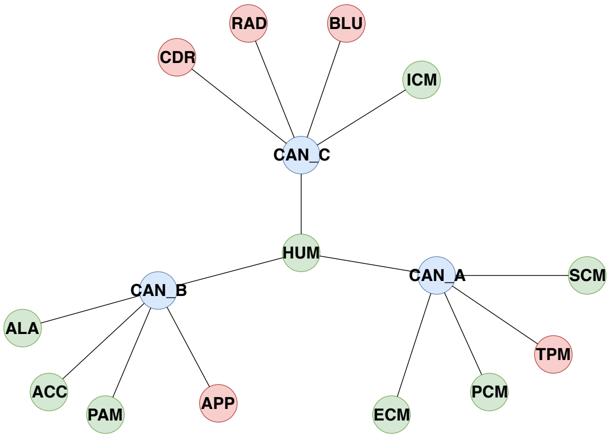

Fig. 2. Topology example: three buses (colored in blue) and multiple components connected to them, among which the ones in Figure 1.

II. RISK-BASED DESIGN OF AUTOMOTIVE NETWORKS

Cybersecurity of vehicular networks is a relatively recent concern: The very first comprehensive cybersecurity assessment of such networks is less than 10 years old [8], [9]. The first vehicle cybersecurity standard by SAE very correctly mandates risk-based design of automotive networks [10].

In [7] we proposed and automated a concept of analysis based on attack trees (inspired by [11]). Our proposal integrates a formal risk analysis methodology based on the development of very general attack trees that can be easily maintained and updated along with the evolution of domain knowledge.

We can see an example of such a tree in Figure 1, we can model an attacker who has the goal of causing an accident. The root of the tree is the abstract attacker *goal*. In this simplistic example, the only way to reach that goal is to steer the car regardless of driver input, which in turn can be obtained through either of two different methods. One of them, for instance, consists of generating a "Lane Steer Command", but it requires an ALA (Active Lane Assist) component to be present. The lane steering command can be generated in turn through three different actions.

Attack trees allow the model to be strongly modular, making it easy to add new methods by just combining steps, or by adding new ones. If new information arises (or new methods are identified) the model can be updated easily and updates can be propagated throughout the tree chain and throughout the other steps of our methodology.

The leaf nodes of these threes can be mapped on the topology of the network under design, as shown in Figure 2. This is actually the core novelty in our proposed approach, and it allows to (i) consider only attacks that are actually achievable given the topology under assessment and (ii) considering the impact of the topology on the feasibility of those attacks, with the final aim of proposing structural countermeasures.

By using a modified version of the risk functions proposed in [11] we can obtain a qualitative but precise assessment of risk for each attack and attack path.

For each network, an automated algorithm can propose a set of architectural changes to the layout, and the insertion of control points such as security gateways, to improve the overall security of the system (we do not currently consider the insertion of intrusion detection systems, e.g. [12]). The current version of the algorithm considers several types of design constraints before proposing mutations, but it can definitely be refined more in future versions of the work. Also, currently we do not take into account potential attacks to the availability of the CAN bus itself, as shown in [13].

III. CYBER-PHYSICAL ATTACKS ON AN INDUSTRIAL ROBOT

Our seminal work on the security of industrial robots [14] was motivated by a few key observations: First, the increased connectivity of robot systems is (and will be) exposing robots to cyber attacks. Indeed, nowadays, industrial robots—originally conceived to be isolated—are exposed to corporate networks and to the Internet.

Second, whereas robots were traditionally designed to operate in a *protected* workspace, physically separated from humans, vendors are introducing several models of *collaborative robots* able to work nearby humans. This, along with the gradual shift of safety devices' implementations from hardwired logic to more flexible software-based implementations, increases the potential safety impact of cybersecurity issues.

Third, a survey of robotics users and integrators revealed scarce awareness of security risks.

Motivated by these observations, we defined an attacker model, introduced industrial-robot-specific attacks based on the properties that a robot must possess, and experimentally verified their feasibility on a standard robot architecture.

In order to understand what is a robot-specific attack, we must first observe that robots are expected to follow three "laws" or requirements. First, *accuracy*: they should read precise values from sensors, and issue correct and accurate commands to the actuators, so that the movements are performed

TABLE I
ATTACK SUMMARY WITH VIOLATED REQUIREMENTS

Attack	**Safety**	**Integrity**	**Accuracy**
Control Loop Alteration	✓	✓	✓
User-perceived Robot State Alteration	✓	✗	✗
Robot State Alteration	✓	✓	✓
Production Logic Tampering	✓	✓	✓
Calibration Parameters Tampering	✓	✓	✓

within acceptable error margins. The second requirement is *safety*: Robots must expose sufficient and correct information so that operators can take safe and informed decisions; allow operators to engage emergency procedures; execute emergency procedures quickly and safely. Finally, robots should ensure their *integrity*, minimizing the risk that badly written control logic may result in damage to their physical parts.

We consider a *robot-specific cyberattack* any violation of these requirements that can be initiated through a digital vector.

We model attackers according to their *goals*, their level of *access* to the system, and their *capabilities*. For the complete taxonomy, we refer the interested reader to [14]. It is however relevant for the scope of the present paper to briefly summarize our results related to attacker goals. While the goal of an attacker who targets a digital system is usually either to gain *unauthorized access* to data and software capabilities, or to *disable* the system itself (Denial-of-Service), in a robot other goals are also possible: the attacker can aim to alter the *production outcome*, injecting faults and micro-defects (which can cause financial losses or even safety concerns); they can also aim to *physically damage* machinery or personnel. This extended threat model naturally led us to develop several new types of robot-specific cyberattacks. Table I summarizes the attacks and the corresponding requirements they violate.

The first type of attack leverages the fact that in robots kinematics and configuration parameters are stored in configuration files that attackers can modify (and users, often, cannot even edit or see). The most interesting parameters to modify are the ones affecting the robot movements, which can lead to several consequences. By detuning closed- or open-loop control parameters, the robot's trajectories can be made unstable or modified (violating the accuracy requirement, but depending on the type of modification safety and integrity can also be affected: for instance by making the robot apply forces beyond specifications).

A second type of attack deals with the information provided to the user by the robot's UI. This information can be relevant for user safety, and is remarkably easy to alter for an attacker. In a variant of this attack, motor state, and even emergency stop signals (in some architectures) are also prone to modification by an attacker.

Of course, attackers can also tamper with the production logic (i.e., with the programming of the robot itself), obtaining results that range from obvious disruption to the insidious insertion of micro-defects that are difficult or impossible to spot.

Finally, attackers can modify calibration data. After any configuration change, the sensing equipment of robots must be calibrated to compensate for known measurement errors when triggering servo motors. This data, initially stored in the sensing equipment, is transmitted to the controller during system boot. Then, the controller uses its local copy of the data, which is vulnerable to manipulation by an attacker. Such manipulation can force a servo motor to move erratically or unexpectedly, because the true error in the measured signal (e.g., joint position) is different from the error that the controller knows. This can lead to abnormal movements, and also to excessive speed or force being applied.

It is important to note how all of these attacks are made possible by the unique digital/physical surface of CPSs, and happen precisely at that layer. The role of sensing equipment (in particular for the last attack we mentioned) is very evident.

IV. CONCLUSIONS

In this paper, we briefly discussed the security challenges arising at the layer between the digital and the physical components of CPSs. We showed how critical the problem is, because of the vital role such systems play, because of the existence of threat actors, and because of the increasing degree of autonomy of CPSs.

We analyzed two examples of how the unique nature of CPSs influences and drives security analysis. In particular, we showed how the physical topology of a CPS network can influence risk assessment and proposed risk-based design approaches for a specific type of systems (in the automotive field). We also showed how the physical impact of attacks can drive the identification of new threats and threat vectors (with an example drawn from robotics).

Sensors and their networks, being a key element of the interaction surface between the digital and the physical world, are heavily impacted by these new types of attacks and their corresponding novel defense strategies.

As CPSs evolve in complex interconnections of autonomous systems, interacting both digitally and physically, sensing the world and reacting smartly to the external conditions, becoming a cloud-like, transparent infrastructure around our lives and our society, it becomes imperative to ensure that their emerging safety and security properties are accurately assessed. More research is needed to make this a systematic engineering approach.

ACKNOWLEDGMENTS

The author would like to thank his current and past PhD students and collaborators (in particular Eric Evenchick, Stefano Longari, Federico Maggi, Andrea Palanca, Marcello Pogliani, Mario Polino, Davide Quarta) for contributing to the research that inspired this paper. The research leading to these results has been partially supported by BVTech SpA under project grant UCSA; by the "FilieraSicura" project, funded by Cisco Systems Inc. and Leonardo SpA; and by the European Union's Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement nr. 690972.

References

[1] S. Zanero, "Cyber-physical systems," *IEEE Computer*, vol. 50, no. 4, pp. 14–16, 2017. [Online]. Available: https://doi.org/10.1109/MC.2017.105

[2] S. Zanero and E. Evenchick, "Up close and personal: Cybersecurity in medical iot devices," in *38th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC 2016)*, USA, 2016.

[3] G. Zheng, R. Shankaran, M. A. Orgun, L. Qiao, and K. Saleem, "Ideas and challenges for securing wireless implantable medical devices: A review," *IEEE Sensors Journal*, vol. 17, no. 3, pp. 562–576, Feb 2017.

[4] J. Carr, *Inside Cyber Warfare: Mapping the Cyber Underworld*. O'Reilly Media, Inc., 2009.

[5] , "Symantec 2010 critical infrastructure protection study," Symantec Corporation, Tech. Rep., October 2010.

[6] C. Miller and C. Valasek, "Securing self-driving cars (one company at a time)," August 2018. [Online]. Available: http://illmatics.com/securing_self_driving_cars.pdf

[7] A. Cannizzo, S. Longari, and S. Zanero, "A cybersecurity-by-design methodology and tool for vehicular networks," July 2018, submitted for publication.

[8] K. Koscher, A. Czeskis, F. Roesner, S. Patel, T. Kohno, S. Checkoway, D. McCoy, B. Kantor, D. Anderson, H. Shacham, and S. Savage, "Experimental security analysis of a modern automobile," in *2010 IEEE Symposium on Security and Privacy*, May 2010, pp. 447–462.

[9] S. Checkoway, D. McCoy, B. Kantor, D. Anderson, H. Shacham, S. Savage, K. Koscher, A. Czeskis, F. Roesner, and T. Kohno, "Comprehensive experimental analyses of automotive attack surfaces," in *Proceedings of the 20th USENIX Conference on Security*, ser. SEC'11. Berkeley, CA, USA: USENIX Association, 2011, pp. 6–6. [Online]. Available: http://dl.acm.org/citation.cfm?id=2028067.2028073

[10] SAE International, "Cybersecurity guidebook for cyber-physical vehicle systems," Standard J3061, 2017. [Online]. Available: https://www.sae.org/standards/content/j3061/

[11] A. Ruddle, D. Ward, B. Weyl, S. Idrees, Y. Roudier, M. Friedewald, T. Leimbach, A. Fuchs, S. Gürgens, O. Henniger *et al.*, "Security requirements for automotive on-board networks based on dark-side scenarios," EVITA Deliverable D.2.3, 2009.

[12] H. al Khateeb, G. Epiphaniou, A. Reviczky, P. Karadimas, and H. Heidari, "Proactive threat detection for connected cars using recursive bayesian estimation," *IEEE Sensors Journal*, vol. 18, no. 12, pp. 4822–4831, June 2018.

[13] A. Palanca, E. Evenchick, F. Maggi, and S. Zanero, "A stealth, selective, link-layer denial-of-service attack against automotive networks," in *Detection of Intrusions and Malware, and Vulnerability Assessment - 14th International Conference, DIMVA 2017, Bonn, Germany, July 6-7, 2017, Proceedings*, ser. Lecture Notes in Computer Science, M. Polychronakis and M. Meier, Eds., vol. 10327. Springer, 2017, pp. 185–206. [Online]. Available: https://doi.org/10.1007/978-3-319-60876-1_9

[14] D. Quarta, M. Pogliani, M. Polino, F. Maggi, A. M. Zanchettin, and S. Zanero, "An experimental security analysis of an industrial robot controller," in *2017 IEEE Symposium on Security and Privacy, SP 2017, San Jose, CA, USA, May 22-26, 2017*. IEEE Computer Society, 2017, pp. 268–286. [Online]. Available: https://doi.org/10.1109/SP.2017.20