A Formal Approach to Cyber-Physical Attacks (Extended Abstract)*

Ruggero Lanotte[1], Massimo Merro[2], Riccardo Muradore[2], and Luca Viganò[3]

[1] Dipartimento di Scienza e Alta Tecnologia
Università dell'Insubria, Como, Italy
ruggero.lanotte@uninsubria.it
[2] Dipartimento di Informatica
Università degli Studi di Verona, Italy
massimo.merro@univr.it, riccardo.muradore@univr.it
[3] Department of Informatics
King's College London, UK
luca.vigano@kcl.ac.uk

Abstract

We apply formal methods to lay and streamline theoretical foundations to reason about Cyber-Physical Systems (CPSs) and cyber-physical attacks. We focus on integrity and DoS attacks to sensors and actuators of CPSs, and on the timing aspects of these attacks. The contributions of our work are threefold: (1) we define a hybrid process calculus to model both CPSs and cyber-physical attacks; (2) we define a threat model of cyber-physical attacks and provide the means to assess attack tolerance/vulnerability with respect to a given attack; (3) we formalise how to estimate the impact of a successful attack on a CPS and investigate possible quantifications of the success chances of an attack.

1 Introduction

Cyber-Physical Systems (CPSs) are integrations of networking and distributed computing systems with physical processes that monitor and control entities in a physical environment, with feedback loops where physical processes affect computations and vice versa. For example, in real-time control systems, a hierarchy of *sensors*, *actuators* and *control processing components* are connected to control stations. Different kinds of CPSs include *supervisory control and data acquisition (SCADA)*, *programmable logic controllers (PLC)* and distributed control systems.

Historically, CPSs relied on proprietary technologies and were implemented as stand-alone networks in physically protected locations. However, in recent years the situation has changed considerably: commodity hardware, software and communication technologies are used to enhance the connectivity of these systems and improve their operation.

This evolution has dramatically increased the number of attacks to the security of cyber-physical and critical systems, e.g., manipulating sensor readings and, in general, influencing physical processes to bring the system into a state desired by the attacker. Many (in)famous examples have been so impressive to make the international news, e.g., (i) the *Stuxnet* worm, which reprogrammed PLCs of nuclear centrifuges in Iran [8], or (ii) the attack on a sewage treatment facility in Queensland, Australia, which manipulated the SCADA system to release raw sewage into local rivers and parks [29], and the (iii) *SQL Slammer worm*, which made unavailable 13000 Bank of America ATM machines, the electronic check-in kiosks of Continental Airlines, and the safety display at Davis-Besse nuclear power plant in Ohio [20].

*This is an extended abstract of the paper [18] that we presented at CSF 2017 (see also [17]).

The primary approach followed by academia and industry to face cyber-physical attacks has been to secure the communication infrastructure and hardening of control systems. There is a large body of literature on how to adapt existing IT security methods to the characteristic features of the control domain [28].

However, as pointed out by Gollmann et al. [11], attacks on CPSs usually cross the boundary between cyber-space and the physical world, possibly more than once. These attacks may manipulate sensor readings already before cryptographic security measures are applied. Attacks may try to influence physical processes to bring the system into a state desired by the attacker. Thus, the concern for consequences at the physical level puts *CPS security* apart from standard *IT security*, and demands for *ad hoc* solutions to properly address such novel research challenges.

To address the limitations of defending CPSs using only IT methods, a new line of research has focused on understanding the adversary's interactions with the physical components of cyber-physical systems. It is sometimes claimed that "once communications security is compromised the attacker can do whatever she wants". However, as explained by Krotofili and Cárdenas [14], this claim is quite imprecise. The attacker may well be able to inject any input she wants but this does not necessarily amount to being able to influence processes in the physical world at will. Both physical and logical components of CPSs have to be properly understood by the attacker in order to conduct an effective attack. In this respect, Gollmann et al. [11] fix the *stages* that a cyber-physical attacker should go through before achieving her goals: *access*, *discovery*, *control*, *damage*, and *cleanup*. In our paper [18], we only focus on the fourth stage, damage, where the attacker has already a rough idea of the control plan of the target CPS.

The works in the literature that have taken up these novel research challenges range from proposals of different notions of cyber-physical security and attacks (e.g., [3, 11, 14], to name a few) to pioneering extensions to CPS security of standard formal approaches (e.g., [3, 5, 33]). However, to the best of our knowledge, a systematic *formal approach* to cyber-physical attacks is still to be fully developed.

2 Background

The dynamic behaviour of the *physical plant* of a CPS is often represented by means of a *discrete-time state-space model*[1] consisting of two equations of the form

$$\begin{aligned} x_{k+1} &= Ax_k + Bu_k + w_k \\ y_k &= Cx_k + e_k \end{aligned}$$

where $x_k \in \mathbb{R}^n$ is the current *(physical) state*, $u_k \in \mathbb{R}^m$ is the *input* (i.e., the control actions implemented through actuators) and $y_k \in \mathbb{R}^p$ is the *output* (i.e., the measurements from the sensors). The *uncertainty* $w_k \in \mathbb{R}^n$ and the *measurement error* $e_k \in \mathbb{R}^p$ represent perturbation and sensor noise, respectively, and A, B, and C are matrices modelling the dynamics of the physical system. Here, the *next state* x_{k+1} depends on the current state x_k and the corresponding control actions u_k, at the sampling instant $k \in \mathbb{N}$. The state x_k cannot be directly observed: only its measurements y_k can be observed.

The physical plant is supported by a communication network through which the sensor measurements and actuator data are exchanged with controller(s) and supervisor(s) (e.g., IDSs), which are the *cyber* components (also called *logics*) of a CPS.

[1] See [35] for a taxonomy of the time-scale models used to represent CPSs.

3 A Formal Approach to Cyber-Physical Attacks

We focus on a formal treatment of both *integrity* and *Denial of Service (DoS)* attacks to *physical devices* (sensors and actuators) of CPSs, paying particular attention to the *timing aspects* of these attacks. The overall goal of our work is to apply formal methodologies to lay *theoretical foundations* to reason about and statically detect attacks to physical devices of CPSs. A straightforward utilisation of these methodologies is for *model-checking*, in order to be able to statically analyse security properties of CPSs before their practical implementation and deployment. In other words, we aim at providing an essential stepping stone for formal and automated analysis techniques for checking the security of CPSs (rather than for providing defence techniques).

The contributions of our work are threefold and are summarised in the following subsections; more details can be found in the full paper [18]. There, we also consider a non-trivial *running example* taken from an engineering application and use it to illustrate our definitions and cases of CPSs that tolerate certain attacks, and of CPSs that suffer from attacks that drag them towards undesired behaviours.

All our results have been formally proven. Moreover, the behaviour of our running example and of most of the cyber-physical attacks appearing in the paper have been simulated in MATLAB.

3.1 `CCPSA`: A Calculus of Cyber-Physical Systems and Attacks

The first contribution is the definition of a *hybrid process calculus*, called `CCPSA`, to formally specify both CPSs and cyber-physical attacks. In `CCPSA`, CPSs have two components:

- a *physical component* denoting the *physical plant* (also called environment) of the system, and containing information on state variables, actuators, sensors, evolution law, etc., and
- a *cyber component* that governs access to sensors and actuators, and channel-based communication with other cyber components.

Thus, channels are used for logical interactions between cyber components, whereas sensors and actuators make possible the interaction between cyber and physical components.

`CCPSA` adopts a *discrete notion of time* [12, 4, 16] and it is equipped with a *labelled transition semantics (LTS)* that allows us to observe both *physical events* (system deadlock and violations of safety conditions) and *cyber events* (channel communications). Based on our LTS, we define two trace-based system preorders: a *trace preorder*, $\sqsubseteq$, and a *timed variant*, $\sqsubseteq_{m..n}$, for $m, n \in \mathbb{N}^+ \cup \{\infty\}$, which takes into account discrepancies of execution traces within the time interval $m..n$. Intuitively, given two CPSs Sys_1 and Sys_2, we write $Sys_1 \sqsubseteq_{m..n} Sys_2$ if Sys_2 simulates the execution traces of Sys_1, except for the time interval $m..n$; if $n = \infty$ then the simulation only holds for the first $m - 1$ time slots.

3.2 Cyber-Physical Attacks

As a second contribution, we formalise a *threat model* that specifies attacks that can manipulate sensor and/or actuator signals in order to drive a CPS into an undesired state [30]. Cyber-physical attacks typically tamper with both the physical (sensors and actuators) and the cyber-layer. In our threat model, communication cannot be manipulated by the attacker, who instead may compromise (unsecured) physical devices, which is our focus. As depicted in Figure 1, our attacks may affect directly the sensor measurements or the controller commands:

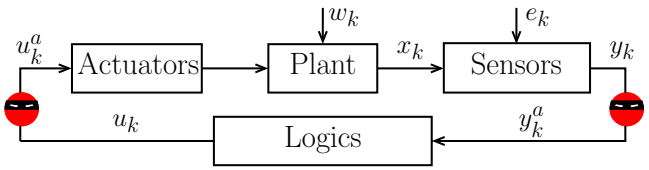

Figure 1: Our threat model for CPSs

- *Attacks on sensors* consist of reading and eventually replacing y_k (the sensor measurements) with y_k^a.
- *Attacks on actuators* consist of reading, dropping and eventually replacing the controller commands u_k with u_k^a, affecting directly the actions the actuators may execute.

We group attacks into classes, where, intuitively, a class of attacks provides information about which physical devices are accessed by the attacks of that class, how they are accessed (read and/or write), when the attack begins and when the attack ends. More specifically, a class of attacks takes into account both the *malicious activities* $\mathcal{I}$ on physical devices and the *timing parameters* m and n of the attack: begin and end of the attack. We represent a class C as a total function $C \in [\mathcal{I} \to \mathcal{P}(m..n)]$. Intuitively, for $\iota \in \mathcal{I}$, $C(\iota) \subseteq m..n$ denotes the set of time instants when an attack of class C may achieve the malicious activity ι. Thus, a class C is a total function that associates to any malicious use (read/write) of any physical device (sensor/actuator) a possibly empty set of time instants in which the attacker tampers with that specific device.

As observed in [14], timing is a critical issue in CPSs because the physical state of a system changes continuously over time and, as the system evolves in time, some states might be more vulnerable to attacks than others. For example, an attack launched when the target state variable reaches a local maximum (or minimum) may have a great impact on the whole system behaviour [15]. Furthermore, not only the timing of the attack but also the *duration of the attack* is an important parameter to be taken into consideration in order to achieve a successful attack. For example, it may take minutes for a chemical reactor to rupture [31], hours to heat a tank of water or burn out a motor, and days to destroy centrifuges [8].

In order to make security assessments on our CPSs, we adopt a well-known approach called *Generalized Non Deducibility on Composition (GNDC)* [9]. Thus, in `CCPSA`, we say that a CPS *Sys tolerates* a cyber-physical attack A if

$$Sys \parallel A \sqsubseteq Sys \ .$$

In this case, the presence of the attack A, does not affect the whole (physical and logical) observable behaviour of the system *Sys*, and the attack can be considered harmless.

On the other hand, we say that a CPS *Sys* is *vulnerable* to a cyber-physical attack A of class $C \in [\mathcal{I} \to \mathcal{P}(m..n)]$ if there is a time interval $m'..n'$ in which the attack becomes observable (physically or logically). Formally, we write:

$$Sys \parallel A \sqsubseteq_{m'..n'} Sys \ .$$

Thus, if a system *Sys* is vulnerable to an attack A of class $C \in [\mathcal{I} \to \mathcal{P}(m..n)]$, during the time interval $m'..n'$, then the attack operates during the interval $m..n$ but it influences the system under attack in the time interval $m'..n'$ (obviously, $m' \geq m$). If n' is finite we have a *temporary attack*, otherwise we have a *permanent attack*. Furthermore, if $m' - n$ is big enough

and $n - m$ is small, then we have a quick nasty attack that affects the system late enough to allow *attack camouflages* [11]. On the other hand, if m' is significantly smaller than n, then the attack affects the observable behaviour of the system well before its termination and the CPS has good chances of undertaking countermeasures to stop the attack. Finally, an attack A is called *lethal*, as it drags the system into a deadlock state. This is obviously a permanent attack.

Note that, as both notions of tolerance and vulnerability of a CPS rely on behavioural preorders, they also depend on the capability of logical components, such as *Intrusion Detection Systems* (*IDSs*), to detect and signal undesired physical behaviours. In fact, the *IDS* component might be designed to detect abnormal physical behaviours going well further than deadlocks and violations of safety conditions. Thus, we say that an attack is *stealthy* if it is able to drive the CPS under attack into an incorrect physical state (either deadlock or violation of the safety conditions) without being noticed by the *IDS* component.

Soundness criteria. We provide sufficient criteria to prove attack tolerance/vulnerability to attacks of an arbitrary class C. We define a notion of *most powerful attack* of a given class C, $Top(C)$, and prove a theorem that says that if a CPS tolerates $Top(C)$ then it tolerates all attacks A of class C (or "weaker" than C). Similarly, if a CPS is vulnerable to $Top(C)$, in the time interval $m'..n'$, then no attacks of class C can affect the system out of that time interval. This is very useful when checking for attack tolerance/vulnerability with respect to all attacks of a given class C (or "weaker" than C).

3.3 Impact of an attack

As a third contribution, we formalise how to estimate the *impact of a successful attack on a CPS* and investigate possible *quantifications of the chances* for an attack of being successful when attacking a CPS. This is important since, in industrial CPSs, before taking any countermeasure against an attack, engineers typically first try to estimate the impact of the attack on the system functioning (e.g., performance and security) and weigh it against the cost of stopping the plant. If this cost is higher than the damage caused by the attack (as is sometimes the case), then engineers might actually decide to let the system continue its activities even under attack. We thus provide a *metric* to estimate the deviation of the system under attack with respect to expected behaviour, according to its evolution law and the uncertainty of the model. Then, we prove a theorem that says that the impact of the most powerful attack $Top(C)$ represents an upper bound for the impact of any attack A of class C.

4 Related work

A number of approaches have been proposed for modelling CPSs using *hybrid process algebras* [6, 2, 27, 10, 22]. `CCPSA` shares some similarities with the ϕ-calculus [27]. However, unlike `CCPSA`, in the ϕ-calculus, given a hybrid system (E, P), the process P can dynamically change the evolution law in E. Furthermore, the ϕ-calculus does not have a representation of physical devices and measurement law, which are instead crucial for us to model cyber-physical attacks that operate in a timely fashion on sensors and actuators.

Among the 118 papers discussed in the comprehensive survey [35], 50 adopt a discrete notion of time similar to ours, 13 a continuous one, 48 a quasi-static time model, and the rest use a hybrid time model. Most of these papers investigate attacks on CPSs and their protection by relying on *simulation test systems* to validate the results.

A number of papers on CPS security have been of inspiration for us, in particular [13, 11, 14, 3]. Huang et al. were among the first to put forth an approach for developing threat models for

CPSs: in [13] they propose models for integrity and DoS attacks, and evaluate the physical and economic consequences of the attacks on a chemical reactor system.

Gollmann et al. [11] provide a clear picture of what the possible goals of a cyber-physical attacker are: *equipment damagage*, i.e., attacks aiming for physical damage of equipment or infrastructure (e.g. pipes, valves); *production damage*, when the attacker goes after the production process to spoil the product or make production more expensive; *compliance violation*, when the attacker tries to damage the safety and the environment impact of the industrial plant.

From Krotifili and Cárdenas [14] we have learned about the important role played by timing parameters on both integrity and DoS attacks. They provide an empirical analysis of the Tennessee Eastman process control challenge problem to gain insights into the behaviour of a physical process when confronted with cyber-physical attacks.

Burmester et al. [3] employed *hybrid timed automata* to give a threat framework based on the traditional Byzantine faults model for cryptographic security. However, as remarked in [30], cyber-physical attacks and faults have inherently distinct characteristics. Faults are considered as physical events that affect the system behaviour, where simultaneous events don't act in a coordinated way; cyber-attacks may be performed over a significant number of attack points and in a coordinated way.

Vigo [32], presented an attack scenario that addresses some of the peculiarities of a cyber-physical adversary, and discussed how this scenario relates to other attack models popular in the security protocol literature. Then, in [33, 34] Vigo et al. proposed an untimed calculus of broadcasting processes equipped with notions of failed and unwanted communication. These works differ quite considerably from ours, e.g., they focus on DoS attacks without taking into consideration timing aspects or impact of the attack.

Cómbita et al. [5] and Zhu and Basar [36] applied *game theory* to capture the conflict of goals between an attacker who seeks to maximise the damage inflicted to a CPS's security and a defender who aims to minimise it [21].

Finally, there are three recent papers that were developed in parallel to ours: [23, 26, 25]. Rocchetto and Tippenhaur [26] introduced a taxonomy of the diverse attacker models proposed for CPS security and outline requirements for generalised attacker models; in [25], they then proposed an extended Dolev-Yao attacker model suitable for CPSs. In their approach, physical layer interactions are modelled as abstract interactions between logical components to support reasoning on the physical-layer security of CPSs. This is done by introducing additional orthogonal channels. Time is not represented.

Nigam et al. [23] work around the notion of Timed Dolev-Yao Intruder Models for Cyber-Physical Security Protocols by bounding the number of intruders required for the automated verification of such protocols. Following a tradition in security protocol analysis, they provide an answer to the question: How many intruders are enough for verification and where should they be placed? They also extend the strand space model to CPS protocols by allowing for the symbolic representation of time, so that they can use the tool Maude [24] along with SMT support. Their notion of time is however different from ours, as they focus on the time a message needs to travel from an agent to another. The paper does not mention physical devices, such as sensors and/or actuators.

5 Conclusions and future work

Our work provides formal theoretical foundations to reason about, and statically detect, attacks to physical devices of CPSs. To that end, we have proposed a hybrid process calculus, called `CCPSA`, as a formal *specification language* to model physical and cyber components of CPSs as well

as cyber-physical attacks. Based on CCPSA and specific *timed trace semantics*, we have formalised a *threat model* for CPSs by grouping attacks in classes, according to the target physical devices and two timing parameters: begin and duration of the attacks. Then, we relied on the trace semantics of CCPSA to assess *attack tolerance/vulnerability* with respect to a given attack. Along the lines of GNDC [9], we defined a notion of *top attacker*, $Top(C)$, of a given class of attacks C, which has been used to provide sufficient criteria to prove attack tolerance/vulnerability to all attacks of class C (and weaker ones). Finally, we have provided a metric to estimate the impact of a successful attack on a CPS together with possible quantifications of the success chances of an attack. We proved that the impact of the most powerful attack $Top(C)$ represents an upper bound for the impact of any attack A of class C (and weaker ones).

While much is still to be done, we believe that our paper provides a stepping stone for the development of formal and automated tools to analyse the security of CPSs. We will consider applying, possibly after proper enhancements, existing tools and frameworks for automated security protocol analysis, resorting to the development of a dedicated tool if existing ones prove not up to the task. We will also consider further security properties and concrete examples of CPSs, as well as other kinds of cyber-physical attackers and attacks, e.g., periodic attacks. This will allow us to refine the classes of attacks we have given here (e.g., by formalising a type system amenable to static analysis), and provide a formal definition of when a CPS is more secure than another so as to be able to design, by progressive refinement, secure variants of a vulnerable CPSs.

We also aim to extend the preliminary quantitative analysis we have given here by developing a suitable behavioural theory ensuring that our trace semantics considers also the probability of a trace to actually occur. Thus, our notion of impact might be refined by taking into account quantitative aspects of an attack such as the probability of being successful when targeting a specific CPS. We expect that *n-bisimulation metrics* [7], which takes into account bounded computations of systems, will be useful to that extent [19].

Finally, for what concerns automatic approximations of the impact, while we have not yet fully investigated the problem, we believe that we can transform it into a "minimum problem". For instance, if the environment uses linear functions, then, by adapting techniques developed for linear hybrid automata (see, e.g., [1]), the set of all traces with length at most n (for a fixed n) can be characterised by a system of first degree inequalities, so the measure of the impact could be translated into a linear programming problem.

References

[1] R. Alur, C. Courcoubetis, N. Halbwachs, T. Henzinger, P.-H. Ho, X. Nicollin, A. Olivero, J. Sifakis, and S. Yovine. The algorithmic analysis of hybrid systems. *TCS*, 138(1):3–34, 1995.

[2] J. A. Bergstra and C. A. Middleburg. Process algebra for hybrid systems. *Thoretical Computer Science*, 335(2-3):215–280, 2005.

[3] M. Burmester, E. Magkos, and V. Chrissikopoulos. Modeling security in cyber-physical systems. *IJCIP*, 5(3-4):118–126, 2012.

[4] A. Cerone, M. Hennessy, and M. Merro. Modelling mac-layer communications in wireless systems. *Logical Methods in Computer Science*, 11(1), 2015.

[5] L. F. Cómbita, J. Giraldo, A. A. Cárdenas, and N. Quijano. Response and reconfiguration of cyber-physical control systems: A survey. In *CCAC*, pages 1–6. IEEE, 2015.

[6] P. Cuijpers and M. Reniers. Hybrid process algebra. *The Journal of Logic and Algebraic Programming*, 62(2):191–245, 2005.

[7] J. Desharnais, J. Gupta, R. Jagadeesan, and P. Panangaden. Metrics for Labelled Markov Processes. *TCS*, 318(3):323–354, 2004.

[8] N. Falliere, L. Murchu, and E. Chien. W32.Stuxnet Dossier, 2011.

[9] R. Focardi and F. Martinelli. A Uniform Approach for the Definition of Security Properties. In *FM*, volume 1708 of *LNCS*, pages 794–813. Springer, 1999.

[10] V. Galpin, L. Bortolussi, and J. Hillston. HYPE: Hybrid modelling by composition of flows. *Formal Asp. Comput.*, 25(4):503–541, 2013.

[11] D. Gollmann, P. Gurikov, A. Isakov, M. Krotofil, J. Larsen, and A. Winnicki. Cyber-Physical Systems Security: Experimental Analysis of a Vinyl Acetate Monomer Plant. In *ACM CCPS*, pages 1–12, 2015.

[12] M. Hennessy and T. Regan. A process algebra for timed systems. *Information and Computation*, 117(2):221–239, 1995.

[13] Y. Huang, A. A. Cárdenas, S. Amin, Z. Lin, H. Tsai, and S. Sastry. Understanding the physical and economic consequences of attacks on control systems. *IJCIP*, 2(3):73–83, 2009.

[14] M. Krotofil and A. A. Cárdenas. Resilience of Process Control Systems to Cyber-Physical Attacks. In *NordSec*, volume 8208 of *LNCS*, pages 166–182. Springer, 2013.

[15] M. Krotofil, A. A. Cárdenas, J. Larsen, and D. Gollmann. Vulnerabilities of cyber-physical systems to stale data - Determining the optimal time to launch attacks. *Int. J. Critical Infrastructure Protection*, 7(4):213–232, 2014.

[16] R. Lanotte and M. Merro. Semantic analysis of gossip protocols for wireless sensor networks. In *CONCUR 2011*, volume 6901 of *LNCS*, pages 156–170. Springer, 2011.

[17] R. Lanotte, M. Merro, R. Muradore, and L. Viganò. A Formal Approach to Cyber-Physical Attacks. *CoRR*, abs/1611.01377, 2016.

[18] R. Lanotte, M. Merro, R. Muradore, and L. Viganò. A Formal Approach to Cyber-Physical Attacks. In *30th IEEE Computer Security Foundations Symposium, CSF 2017, Santa Barbara, CA, USA, August 21-25, 2017*, pages 436–450. IEEE, 2017.

[19] R. Lanotte, M. Merro, and S. Tini. A probabilistic calculus of cyber-physical systems. *CoRR*, abs/1707.02279, 2017.

[20] E. Levy. Crossover: Online pests plaguing the offline world. *IEEE Security & Privacy*, 1(6):71–73, 2003.

[21] M. Manshaei, Q. Zhu, T. Alpcan, T. Basar, and J.-P. Hubaux. Game theory meets network security and privacy. *ACM Computer Surveys*, 45(3):25, 2013.

[22] M. Merro, J. Kleist, and U. Nestmann. Mobile objects as mobile processes. *Information and Compututation*, 177(2):195–241, 2002.

[23] V. Nigam, C. Talcott, and A. A. Urquiza. Towards the Automated Verification of Cyber-Physical Security Protocols: Bounding the Number of Timed Intruders. In *ESORICS*, volume 9879 of *LNCS*, pages 450–470. Springer, 2016.

[24] P. C. Ölveczky and J. Meseguer. Semantics and pragmatics of real-time maude. *Higher-Order and Symbolic Computation*, 20(1-2):161–196, 2007.

[25] M. Rocchetto and N. O. Tippenhauer. CPDY: Extending the Dolev-Yao Attacker with Physical-Layer Interactions. In *ICFEM*, volume 10009 of *LNCS*, pages 175–192, 2016.

[26] M. Rocchetto and N. O. Tippenhauer. On Attacker Models and Profiles for Cyber-Physical Systems. In *ESORICS*, volume 9879 of *LNCS*, pages 427–449. Springer, 2016.

[27] W. C. Rounds and H. Song. The ϕ-calculus: A language for distributed control of reconfigurable embedded systems. In *HSCC*, volume 2623 of *LNCS*, pages 435–449. Springer, 2003.

[28] S. Sicari, A. Rizzardi, L. Grieco, and A. Coen-Porisini. Security, privacy and trust in internet of things: The road ahead. *Computer Networks*, 76:146–164, 2015.

[29] J. Slay and M. Miller. Lessons Learned from the Maroochy Water Breach. In *Critical Infrastructure Protection*, volume 253 of *IFIP*, pages 73–82. Springer, 2007.

[30] A. Teixeira, I. Shames, J. Sandberg, and K. H. Johansson. A secure control framework for resource-limited adversaries. *Automatica*, 51:135–148, 2015.

[31] U.S. Chemical Safety and Hazard Investigation Board, T2 Laboratories Inc. Reactive Chemical Explosion: Final Investigation Report. Report No. 2008-3-I-FL, 2009.

[32] R. Vigo. The Cyber-Physical Attacker. In *SAFECOMP*, volume 7613 of *LNCS*, pages 347–356. Springer, 2012.

[33] R. Vigo. *Availability by Design: A Complementary Approach to Denial-of-Service*. PhD thesis, Danish Technical University, 2015.

[34] R. Vigo, F. Nielson, and H. Riis Nielson. Broadcast, denial-of-service, and secure communication. In *IFM*, volume 7940 of *LNCS*, pages 412–427. Springer, 2013.

[35] Y. Zacchia Lun, A. D'Innocenzo, I. Malavolta, and M. D. Di Benedetto. Cyber-Physical Systems Security: a Systematic Mapping Study. *CoRR*, abs/1605.09641, 2016.

[36] Q. Zhu and T. Basar. Game-theoretic methods for robustness, security, and resilience of cyber-physical control systems: games-in-games principle for optimal cross-layer resilient control systems. *IEEE Control Systems*, 35(1):46–65, 2015.