

A Formal Approach to Cyber-Physical Attacks

Ruggero Lanotte
Dipartimento di Scienza e Alta Tecnologia
Università dell'Insubria, Como, Italy

Massimo Merro and Riccardo Muradore
Dipartimento di Informatica
Università degli Studi di Verona, Italy

Luca Viganò
Department of Informatics
King's College London, 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. Our contributions 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. We illustrate definitions and results by means of a non-trivial engineering application.

Index Terms—Cyber-physical system security, formal methods, theoretical foundation, process calculus.

I. INTRODUCTION

Context and motivation: *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.

In recent years there has been a dramatic increase in 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., the Stuxnet worm, which reprogrammed PLCs of nuclear centrifuges in Iran [6] or the attack on a sewage treatment facility in Queensland, Australia, which manipulated the SCADA system to release raw sewage into local rivers and parks [25].

As stated in [10], the concern for consequences at the physical level puts *CPS security* apart from standard *information security*, and demands for *ad hoc* solutions to properly address such novel research challenges. The works that have taken up these challenges range from proposals of different notions of cyber-physical security and attacks (e.g., [3], [10], [13], to name a few) to pioneering extensions to CPS security of standard formal approaches (e.g., [3], [4], [29]). However, to the best of our knowledge, a systematic formal approach to cyber-physical attacks is still to be fully developed.

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.

Contributions: In this paper, 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 the paper 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).

Our contributions are threefold. 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.

¹See [31] for a taxonomy of the time-scale models used to represent CPSs.

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* [11] 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.

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 [26]. 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:

- *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. 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 $C \in [\mathcal{I} \rightarrow \mathcal{P}(m..n)]$. Intuitively, for $\iota \in \mathcal{I}$, $C(\iota) \subseteq m..n$ denotes the set of time instants when an attack of class C may tamper with the device 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 [13], 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 [14]. 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 [27], hours to heat a tank of water or burn out a motor, and days to destroy centrifuges [6].

In order to make security assessments on our CPSs, we adopt a well-known approach called *Generalized Non Deducibility on Composition (GNDC)* [7]. Thus, in CCPSA, we say that a




Fig. 1. Our threat model for CPSs

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 $C \in [\mathcal{I} \rightarrow \mathcal{P}(m..n)]$ if there is a time interval $m'..n'$ in which the attack becomes observable (obviously, $m' \geq m$). Formally, we write:

$$Sys \parallel A \not\sqsubseteq_{m'..n'} Sys .$$

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 that if a CPS tolerates $Top(C)$ then it tolerates all attacks A of class 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 .

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 that the impact of the most powerful attack $Top(C)$ represents an upper bound for the impact of any attack A of class C .

We introduce 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. We remark that while we have kept the example simple, it is actually far from trivial and designed to describe a wide number of attacks, as will become clear below.

All the results presented in the paper have been formally proven; due to lack of space, proofs have been sketched or omitted and full details can be found in [16]. Moreover, the

behaviour of our running example and of most of the cyber-physical attacks appearing in the paper have been simulated in MATLAB.

Organisation: In § II, we give syntax and semantics of CCPSA. In § III, we define cyber-physical attacks and provide sufficient criteria for attack tolerance/vulnerability. In § IV, we estimate the impact of attacks on CPSs, and investigate possible quantifications of the success chances of an attack. In § V, we draw conclusions and discuss related and future work.

II. THE CALCULUS

In this section, we introduce our *Calculus of Cyber-Physical Systems and Attacks*, CCPSA, which extends the *Calculus of Cyber-Physical Systems*, defined in our companion paper [15], with specific features to formalise and study attacks to physical devices. Let us start with some preliminary notations.

Notation 1. We use $x, x_k \in \mathcal{X}$ for state variables, $c, d \in \mathcal{C}$ for communication channels, $a, a_k \in \mathcal{A}$ for actuator devices, $s, s_k \in \mathcal{S}$ for sensors devices, and p, q for both sensors and actuators (generically called physical devices). Values, ranged over by $v, v' \in \mathcal{V}$, are built from basic values, such as Booleans, integers and real numbers. State variables are metavariables for physical properties like temperature, pressure, etc. Actuator names are metavariables for actuator devices like valve, light, etc. Similarly, sensor names are metavariables for sensor devices, e.g., a sensor thermometer.

Given a set of names \mathcal{N} , we write $\mathbb{R}^{\mathcal{N}}$ to denote the set of functions $[\mathcal{N} \rightarrow \mathbb{R}]$ assigning a real to each name in \mathcal{N} . For $\xi \in \mathbb{R}^{\mathcal{N}}$, $n \in \mathcal{N}$ and $v \in \mathbb{R}$, we write $\xi[n \mapsto v]$ for the function $\psi \in \mathbb{R}^{\mathcal{N}}$ such that $\psi(m) = \xi(m)$, for any $m \neq n$, and $\psi(n) = v$.

Finally, we distinguish between real intervals, such as $(m, n]$, for $m \in \mathbb{R}$ and $n \in \mathbb{R} \cup \{\infty\}$, and integer intervals, written $m..n$, for $m \in \mathbb{N}$ and $n \in \mathbb{N} \cup \{\infty\}$. As we will adopt a discrete notion of time, we will use integer intervals to denote time intervals.

Definition 1 (Cyber-physical system). In CCPSA, a cyber-physical system consists of two components:

- a physical environment E that encloses all physical aspects of a system (state variables, physical devices, evolution law, etc.) and
- a cyber component P that interacts with sensors and actuators of the system, and can communicate, via channels, with other cyber components of the same or of other CPSs.

Given a set \mathcal{S} of secured physical devices of E , we write $E \bowtie_{\mathcal{S}} P$ to denote the resulting CPS, and use M and N to range over CPSs. We write $E \bowtie P$ when $\mathcal{S} = \emptyset$.

In a CPS $E \bowtie_{\mathcal{S}} P$, the “secured” devices in \mathcal{S} are accessed in a protected way and hence they cannot be attacked.²

²The presence of battery-powered devices interconnected through wireless networks prevents the en-/decryption of all packets due to energy constraints.

Let us now define physical environments E and cyber components P in order to formalise our proposal for modelling (and reasoning about) CPSs and cyber-physical attacks.

Definition 2 (Physical environment). Let $\hat{\mathcal{X}} \subseteq \mathcal{X}$ be a set of state variables, $\hat{\mathcal{A}} \subseteq \mathcal{A}$ be a set of actuators, and $\hat{\mathcal{S}} \subseteq \mathcal{S}$ be a set of sensors. A physical environment E is an 8-tuple $\langle \xi_x, \xi_u, \xi_w, evol, \xi_e, meas, inv, safe \rangle$, where:

- $\xi_x \in \mathbb{R}^{\hat{\mathcal{X}}}$ is the state function,
- $\xi_u \in \mathbb{R}^{\hat{\mathcal{A}}}$ is the actuator function,
- $\xi_w \in \mathbb{R}^{\hat{\mathcal{X}}}$ is the uncertainty function,
- $evol : \mathbb{R}^{\hat{\mathcal{X}}} \times \mathbb{R}^{\hat{\mathcal{A}}} \times \mathbb{R}^{\hat{\mathcal{X}}} \rightarrow 2^{\mathbb{R}^{\hat{\mathcal{X}}}}$ is the evolution map,
- $\xi_e \in \mathbb{R}^{\hat{\mathcal{S}}}$ is the sensor-error function,
- $meas : \mathbb{R}^{\hat{\mathcal{X}}} \times \mathbb{R}^{\hat{\mathcal{S}}} \rightarrow 2^{\mathbb{R}^{\hat{\mathcal{S}}}}$ is the measurement map,
- $inv : \mathbb{R}^{\hat{\mathcal{X}}} \rightarrow \{\text{true}, \text{false}\}$ is the invariant function,
- $safe : \mathbb{R}^{\hat{\mathcal{X}}} \rightarrow \{\text{true}, \text{false}\}$ is the safety function.

All the functions defining an environment are total functions.

The state function ξ_x returns the current value (in \mathbb{R}) associated to each state variable of the system. The actuator function ξ_u returns the current value associated to each actuator. The uncertainty function ξ_w returns the uncertainty associated to each state variable. Thus, given a state variable $x \in \hat{\mathcal{X}}$, $\xi_w(x)$ returns the maximum distance between the real value of x and its representation in the model. Later in the paper, we will be interested in comparing the accuracy of two systems. Thus, for $\xi_w, \xi'_w \in \mathbb{R}^{\hat{\mathcal{X}}}$, we will write $\xi_w \leq \xi'_w$ if $\xi_w(x) \leq \xi'_w(x)$, for any $x \in \hat{\mathcal{X}}$. Similarly, we write $\xi_w + \xi'_w$ to denote the function $\xi''_w \in \mathbb{R}^{\hat{\mathcal{X}}}$ such that $\xi''_w(x) = \xi_w(x) + \xi'_w(x)$, for any $x \in \hat{\mathcal{X}}$.

Given a state function, an actuator function, and an uncertainty function, the evolution map $evol$ returns the set of next admissible state functions. It models the evolution law of the physical system, where changes made on actuators may reflect on state variables. Since we assume an uncertainty in our models, $evol$ does not return a single state function but a set of possible state functions. $evol$ is obviously monotone with respect to uncertainty: if $\xi_w \leq \xi'_w$ then $evol(\xi_x, \xi_u, \xi_w) \subseteq evol(\xi_x, \xi_u, \xi'_w)$.

Both the state function and the actuator function are supposed to change during the evolution of the system, whereas the uncertainty function is constant.³

The sensor-error function ξ_e returns the maximum error associated to each sensor in $\hat{\mathcal{S}}$. Again due to the presence of the sensor-error function, the measurement map $meas$, given the current state function, returns a set of admissible measurement functions rather than a single one.

The invariant function inv represents the conditions that the state variables must satisfy to allow for the evolution of the system. A CPS whose state variables don't satisfy the invariant is in deadlock.

The safety function $safe$ represents the conditions that the state variables must satisfy to consider the CPS in a safe state.

³Note that, although the uncertainty function is constant, it can be used in the evolution map in an arbitrary way (e.g., it could have a heavier weight when a state variable reaches extreme values). Another possibility is to model the uncertainty function by means of a probability distribution.

Intuitively, if a CPS gets in an unsafe state, then its functionality may get compromised.

In the following, we use a specific notation for the replacement of a single component of an environment with a new one of the same kind; for instance, for $E = \langle \xi_x, \xi_u, \xi_w, \text{evol}, \xi_e, \text{meas}, \text{inv}, \text{safe} \rangle$, we write $E[\xi_w \leftarrow \xi'_w]$ to denote $\langle \xi_x, \xi_u, \xi'_w, \text{evol}, \xi_e, \text{meas}, \text{inv}, \text{safe} \rangle$.

Let us now introduce a *running example* to illustrate our approach. We remark that while we have kept the example simple, it is actually far from trivial and designed to describe a wide number of attacks. A more complex example (say, with n sensors and m actuators) wouldn't have been more instructive but just made the paper more dense.

Example 1 (Physical environment of the CPS *Sys*). Consider a CPS *Sys* in which the temperature of an engine is maintained within a specific range by means of a cooling system. The physical environment *Env* of *Sys* is constituted by: (i) a state variable *temp* containing the current temperature of the engine, and a state variable *stress* keeping track of the level of stress of the mechanical parts of the engine due to high temperatures (exceeding 9.9 degrees); this integer variable ranges from 0, meaning no stress, to 5, for high stress; (ii) an actuator *cool* to turn on/off the cooling system; (iii) a sensor s_t (such as a thermometer or a thermocouple) measuring the temperature of the engine; (iv) an uncertainty $\delta = 0.4$ associated to the only variable *temp*; (v) the evolution law for the two state variables: the variable *temp* is increased (resp., is decreased) of one degree per time unit if the cooling system is inactive (resp., active), whereas the variable *stress* contains an integer that is increased each time the current temperature is above 9.9 degrees, and dropped to 0 otherwise; (vi) an error $\epsilon = 0.1$ associated to the only sensor s_t ; (vii) a measurement map to get the values detected by sensor s_t , up to its error ϵ ; (viii) an invariant function saying that the system gets faulty when the temperature of the engine gets out of the range $[0, 50]$; (ix) a safety function to say that the system moves to an unsafe state when the level of stress reaches the threshold 5.

Formally, $\text{Env} = \langle \xi_x, \xi_u, \xi_w, \text{evol}, \xi_e, \text{meas}, \text{inv} \rangle$ with:

- $\xi_x \in \mathbb{R}^{\{\text{temp}, \text{stress}\}}$ and $\xi_x(\text{temp}) = 0$ and $\xi_x(\text{stress}) = 0$;
- $\xi_u \in \mathbb{R}^{\{\text{cool}\}}$ and $\xi_u(\text{cool}) = \text{off}$; for the sake of simplicity, we can assume ξ_u to be a mapping $\{\text{cool}\} \rightarrow \{\text{on}, \text{off}\}$ such that $\xi_u(\text{cool}) = \text{off}$ if $\xi_u(\text{cool}) \geq 0$, and $\xi_u(\text{cool}) = \text{on}$ if $\xi_u(\text{cool}) < 0$;
- $\xi_w \in \mathbb{R}^{\{\text{temp}, \text{stress}\}}$, $\xi_w(\text{temp}) = 0.4 = \delta$ and $\xi_w(\text{stress}) = 0$;
- $\text{evol}(\xi_x^i, \xi_u^i, \xi_w^i)$ is the set of $\xi \in \mathbb{R}^{\{\text{temp}, \text{stress}\}}$ such that:
 - $\xi(\text{temp}) = \xi_x^i(\text{temp}) + \text{heat}(\xi_u^i, \text{cool}) + \gamma$, with $\gamma \in [-\delta, +\delta]$ and $\text{heat}(\xi_u^i, \text{cool}) = -1$ if $\xi_u^i(\text{cool}) = \text{on}$ (active cooling), and $\text{heat}(\xi_u^i, \text{cool}) = +1$ if $\xi_u^i(\text{cool}) = \text{off}$ (inactive cooling);
 - $\xi(\text{stress}) = \min(5, \xi_x^i(\text{stress}) + 1)$ if $\xi_x^i(\text{temp}) > 9.9$; $\xi(\text{stress}) = 0$, otherwise;
- $\xi_e \in \mathbb{R}^{\{s_t\}}$ and $\xi_e(s_t) = 0.1 = \epsilon$;
- $\text{meas}(\xi_x^i, \xi_e) = \{\xi : \xi(s_t) \in [\xi_x^i(\text{temp}) - \epsilon, \xi_x^i(\text{temp}) + \epsilon]\}$;
- $\text{inv}(\xi_x) = \text{true}$ if $0 \leq \xi_x(\text{temp}) \leq 50$; $\text{inv}(\xi_x) = \text{false}$, otherwise.

- $\text{safe}(\xi_x) = \text{true}$ if $\xi_x(\text{stress}) < 5$; $\text{safe}(\xi_x) = \text{false}$, if $\xi_x(\text{stress}) \geq 5$ (the maximum value for stress is 5).

Let us now formalise the cyber component of CPSs in CCPSA. Basically, we extend the *timed process algebra TPL* of [11] with two main ingredients:

- two different constructs to read values detected at sensors and write values on actuators, respectively;
- special constructs to represent malicious activities on physical devices.

The remaining constructs are the same as those of TPL.

Definition 3 (Processes). Processes are defined as follows:

$$\begin{aligned}
P, Q ::= & \text{nil} \mid \text{idle}.P \mid P \parallel Q \mid \text{timeout}[\pi.P]Q \mid \\
& \text{if } (b) \{P\} \text{ else } \{Q\} \mid P \setminus c \mid H\langle \tilde{w} \rangle \\
\pi ::= & \text{snd } c\langle v \rangle \mid \text{rcv } c(x) \mid \text{read } s(x) \mid \\
& \text{write } a\langle v \rangle \mid \text{read } \sharp p(x) \mid \text{write } \sharp p\langle v \rangle .
\end{aligned}$$

We write *nil* for the *terminated process*. The process *idle.P* sleeps for one time unit and then continues as *P*. We write $P \parallel Q$ to denote the *parallel composition* of concurrent *threads* *P* and *Q*. The process $\text{timeout}[\pi.P]Q$ denotes *prefixing with timeout*. Thus, $\text{timeout}[\text{snd } c\langle v \rangle].P]Q$ sends the value *v* on channel *c* and, after that, it continues as *P*; otherwise, after one time unit, it evolves into *Q*. The process $\text{timeout}[\text{rcv } c(x)].P]Q$ is the obvious counterpart for reception. The process $\text{timeout}[\text{read } s(x)].P]Q$ reads the values detected by the sensor *s*, whereas $\text{timeout}[\text{write } a\langle v \rangle].P]Q$ writes on the actuator *a*. For $\pi \in \{\text{read } \sharp p(x), \text{write } \sharp p\langle v \rangle\}$, the process $\text{timeout}[\pi.P]Q$ denotes the reading and the writing, respectively, of the physical device *p* (sensor or actuator) made by the *attack*. Thus, in CCPSA, *attack processes* have specific constructs to interact with physical devices.

The process $P \setminus c$ is the channel restriction operator of CCS. It is quantified over the set \mathcal{C} of communication channels but we sometimes write $P \setminus \{c_1, c_2, \dots, c_n\}$ to mean $P \setminus c_1 \setminus c_2 \dots \setminus c_n$. The process $\text{if } (b) \{P\} \text{ else } \{Q\}$ is the standard conditional, where *b* is a decidable guard. In processes of the form *idle.Q* and $\text{timeout}[\pi.P]Q$, the occurrence of *Q* is said to be *time-guarded*. The process $H\langle \tilde{w} \rangle$ denotes (guarded) recursion.

We assume a set of *process identifiers* ranged over by H, H_1, H_2 . We write $H\langle w_1, \dots, w_k \rangle$ to denote a recursive process *H* defined via an equation $H(x_1, \dots, x_k) = P$, where (i) the tuple x_1, \dots, x_k contains all the variables that appear free in *P*, and (ii) *P* contains only guarded occurrences of the process identifiers, such as *H* itself. We say that recursion is *time-guarded* if *P* contains only time-guarded occurrences of the process identifiers. Unless explicitly stated our recursive processes are always time-guarded.

In the two constructs $\text{timeout}[\text{rcv } c(x)].P]Q$ and $\text{timeout}[\text{read } \mu(x)].P]Q$, with $\mu \in \{p, \sharp p\}$, the variable *x* is said to be *bound*. This gives rise to the standard notions of *free/bound (process) variables* and α -conversion. A term is *closed* if it does not contain free variables, and we assume to always work with closed processes: the absence of free variables is preserved at run-time. As further notation, we

write $T\{v/x\}$ for the substitution of all occurrences of the free variable x in T with the value v .

Note that in CCPSA, a processes might use sensors and/or actuators which are not defined in the environment. To rule out ill-formed CPSs, we use the following definition.

Definition 4 (Well-formedness). *Given a process P and an environment $E = \langle \xi_x, \xi_u, \xi_w, \text{evol}, \xi_e, \text{meas}, \text{inv}, \text{safe} \rangle$, the CPS $E \bowtie P$ is well-formed if: (i) for any sensor s mentioned in P , the function ξ_e is defined on s ; (ii) for any actuator a mentioned in P , the function ξ_u is defined on a .*

Hereafter, we will always work with well-formed CPSs.

Notation 2. *To model time-persistent prefixing, we write $\pi.P$ for the process defined via the equation $Rcv = \text{timeout}[\pi.P]Rcv$, where Rcv does not occur in P . Further, we write*

- $\text{timeout}[\pi]Q$ as an abbreviation for $\text{timeout}[\pi.\text{nil}]Q$,
- $\text{timeout}[\pi.P]$ as an abbreviation for $\text{timeout}[\pi.P]\text{nil}$,
- $\text{snd } c$ and $\text{rcv } c$, when channel c is used for pure synchronisation,
- $\text{idle}^k.P$ as a shorthand for $\text{idle}.\text{idle} \dots \text{idle}.P$, where the prefix idle appears $k \geq 0$ consecutive times,
- $\text{if } (b) \{P\}$ instead of $\text{if } (b) \{P\} \text{ else } \{\text{nil}\}$.

Let $M = E \bowtie_S P$, we write $M \parallel Q$ for $E \bowtie_S (P \parallel Q)$, and $M \setminus c$ for $E \bowtie_S (P \setminus c)$.

We can now finalise our running example.

Example 2 (Cyber component of the CPS *Sys*). *Let us define the cyber component of the CPS *Sys* described in Example 1. We define two parallel processes: *Ctrl* and *IDS*. The former models the controller activity, consisting in reading the temperature sensor and in governing the cooling system via its actuator, whereas the latter models a simple intrusion detection system that attempts to detect and signal abnormal behaviours of the system. Intuitively, *Ctrl* senses the temperature of the engine at each time slot. When the sensed temperature is above 10 degrees, the controller activates the coolant. The cooling activity is maintained for 5 consecutive time units. After that time, the controller synchronises with the *IDS* component via a private channel *sync*, and then waits for instructions, via a channel *ins*. The *IDS* component checks whether the sensed temperature is still above 10. If this is the case, it sends an alarm of “high temperature”, via a specific channel, and then says to *Ctrl* to keep cooling for other 5 time units; otherwise, if the temperature is not above 10, the *IDS* component requires *Ctrl* to stop the cooling activity.*

```

Ctrl    = read s_t(x).if (x > 10) {Cooling} else {idle.Ctrl}
Cooling = write cool(on).idle5.Check
Check   = snd sync.rcv ins(y).if (y = keep_cooling)
         {idle5.Check} else {write cool(off).idle.Ctrl}
IDS     = rcv sync.read s_t(x).if (x > 10)
         {snd alarm(high_temp).snd ins(keep_cooling).
         idle.IDS} else {snd ins(stop).idle.IDS} .

```

Thus, the whole CPS is defined as:


$$\text{Sys} = \text{Env} \bowtie (\text{Ctrl} \parallel \text{IDS}) \setminus \{\text{sync}, \text{ins}\},$$


Fig. 2. Three possible evolutions of the CPS of Example 2.

where *Env* is the physical environment defined in Example 1. We remark that, for the sake of simplicity, our *IDS* component is quite basic: for instance, it does not check whether the temperature is too low. However, it is straightforward to replace it with a more sophisticated one, containing more informative tests on sensor values and/or on actuators commands.

A. Labelled transition semantics

In this subsection, we provide the dynamics of CCPSA in terms of a *labelled transition system (LTS)* in the SOS style of Plotkin. Definition 5 introduces auxiliary operators on environments.

Definition 5. Let $E = \langle \xi_x, \xi_u, \xi_w, \text{evol}, \xi_e, \text{meas}, \text{inv}, \text{safe} \rangle$.

- $\text{read_sensor}(E, s) \stackrel{\text{def}}{=} \{\xi(s) : \xi \in \text{meas}(\xi_x, \xi_e)\}$,
- $\text{update_act}(E, a, v) \stackrel{\text{def}}{=} E[\xi_u \leftarrow \xi_u[a \mapsto v]]$,
- $\text{next}(E) \stackrel{\text{def}}{=} \bigcup_{\xi \in \text{evol}(\xi_x, \xi_u, \xi_w)} \{E[\xi_x \leftarrow \xi]\}$,
- $\text{inv}(E) \stackrel{\text{def}}{=} \text{inv}(\xi_x)$,
- $\text{safe}(E) \stackrel{\text{def}}{=} \text{safe}(\xi_x)$.

The operator $\text{read_sensor}(E, s)$ returns the set of possible measurements detected by sensor s in the environment E ; due to the error $\xi_e(s)$ of sensor s , it returns a set of possible values rather than a single value. $\text{update_act}(E, a, v)$ returns the new environment in which the actuator function is updated in such a manner to associate the actuator a with the value v . $\text{next}(E)$ returns the set of the next admissible environments reachable from E , by an application of evol . $\text{inv}(E)$ checks whether the invariant is satisfied by the current values of the state variables (here, abusing notation, we overload the meaning of the function inv). $\text{safe}(E)$ checks whether the safety conditions are satisfied by the current values of the state variables.

In Table I, we provide transition rules for processes. Here, the meta-variable λ ranges over labels in the set $\{\text{idle}, \tau, \bar{c}v, cv, a!v, s?v, \ell p!v, \ell p?v, \tau:p\}$. Rules (Outp), (Inpp) and (Com) serve to model channel communication, on some channel c . Rules (Write) and (Read) denote the writing/reading of some data on the physical device p . Rule (ℓ SensWrite ℓ) models an *integrity*

TABLE I
LTS FOR PROCESSES

$\text{(Outp)} \frac{-}{\text{timeout}[\text{snd } c\langle v \rangle.P]Q \xrightarrow{\bar{c}v} P}$ $\text{(Com)} \frac{P \xrightarrow{\bar{c}v} P' \quad Q \xrightarrow{cv} Q'}{P \parallel Q \xrightarrow{\tau} P' \parallel Q'}$ $\text{(Write)} \frac{-}{\text{timeout}[\text{write } \mu\langle v \rangle.P]Q \xrightarrow{\mu!v} P}$ $\text{(\$SensWrite \$)} \frac{P \xrightarrow{\$s!v} P' \quad Q \xrightarrow{s?v} Q'}{P \parallel Q \xrightarrow{\tau:s} P' \parallel Q'}$ $\text{(Res)} \frac{P \xrightarrow{\lambda} P' \quad \lambda \notin \{cv, \bar{c}v\}}{P \setminus c \xrightarrow{\lambda} P' \setminus c}$ $\text{(Then)} \frac{[b] = \text{true} \quad P \xrightarrow{\lambda} P'}{\text{if } (b) \{P\} \text{ else } \{Q\} \xrightarrow{\lambda} P'}$ $\text{(TimeNil)} \frac{-}{\text{nil} \xrightarrow{\text{idle}} \text{nil}}$ $\text{(Timeout)} \frac{-}{\text{timeout}[\pi.P]Q \xrightarrow{\text{idle}} Q}$	$\text{(Inpp)} \frac{-}{\text{timeout}[\text{rcv } c(x).P]Q \xrightarrow{cv} P\{v/x\}}$ $\text{(Par)} \frac{P \xrightarrow{\lambda} P' \quad \lambda \neq \text{idle}}{P \parallel Q \xrightarrow{\lambda} P' \parallel Q}$ $\text{(Read)} \frac{-}{\text{timeout}[\text{read } \mu(x).P]Q \xrightarrow{\mu?v} P\{v/x\}}$ $\text{(\$ActRead \$)} \frac{P \xrightarrow{a?v} P' \quad Q \xrightarrow{\$a?v} Q'}{P \parallel Q \xrightarrow{\tau:a} P' \parallel Q'}$ $\text{(Rec)} \frac{P\{\bar{w}/\bar{x}\} \xrightarrow{\lambda} Q \quad H(\bar{x}) = P}{H(\bar{w}) \xrightarrow{\lambda} Q}$ $\text{(Else)} \frac{[b] = \text{false} \quad Q \xrightarrow{\lambda} Q'}{\text{if } (b) \{P\} \text{ else } \{Q\} \xrightarrow{\lambda} Q'}$ $\text{(Delay)} \frac{-}{\text{idle}.P \xrightarrow{\text{idle}} P}$ $\text{(TimePar)} \frac{P \xrightarrow{\text{idle}} P' \quad Q \xrightarrow{\text{idle}} Q'}{P \parallel Q \xrightarrow{\text{idle}} P' \parallel Q'}$
---	--

attack on sensor s , where the controller of s is supplied with a fake value v provided by the attack. Rule $(\text{\$ActRead \$})$ models a *DoS attack to the actuator a* , where the update request of the controller is intercepted by the attacker and it never reaches the actuator. Rule (Par) propagates untimed actions over parallel components. Rules (Res), (Rec), (Then) and (Else) are standard. The following four rules model the passage of time. The symmetric counterparts of rules (Com) and (Par) are omitted.

In Table II, we lift the transition rules from processes to systems. Except for rule (Deadlock), all rules have a common premise $\text{inv}(E)$: a system can evolve only if the invariant is satisfied. Here, actions, ranged over by α , are in the set $\{\tau, \bar{c}v, cv, \text{idle}, \text{deadlock}, \text{unsafe}\}$. These actions denote: internal activities (τ); logical activities, more precisely, channel transmission ($\bar{c}v$ and cv); the passage of time (idle); and two specific physical events: system deadlock (deadlock) and the violation of the safety conditions (unsafe). Rules (Out) and (Inp) model transmission and reception, with an external system, on a channel c . Rule (SensReadSec) models the reading of the current data detected at a *secured sensor s* , whereas rule (SensReadUnsec) models the reading of an *unsecured sensor s* . In this case, since the sensor is not secured, the presence of a malicious action $\$s!w$ prevents the reading of the sensor. We already said that rule $(\text{\$SensWrite \$})$ of Table I models integrity attacks on an unsecured sensor s , however, together with rule (SensReadUnsec), it also serves to model *DoS attacks on an unsecured sensor s* , as the controller of s cannot read its correct value if the attacker is currently supplying a fake value for it.

Rule $(\text{\$SensRead \$})$ allows the attacker to read the confidential value detected at an unsecured sensor s . Rule (ActWriteSec)

models the writing of a value v on a *secured actuator a* , whereas rule (ActWriteUnsec) models the writing on an *unsecured actuator a* . Again, if the actuator is unsecured, the presence of an attack (capable of performing an action $\$a?v$) prevents the correct access to the actuator by the controller. Rule $(\text{\$ActWrite \$})$ models an *integrity attack to an unsecured actuator a* , where the attack updates the actuator with a fake value.

Note that our operational semantics ensures a preemptive power to prefixes of the form $\text{write } \$p\langle v \rangle$ and $\text{read } \$p(x)$ on *unsecured devices p* . This because an attack process can always prevent the regular access to a unsecured physical device (sensor or actuator) by its controller.

Proposition 1 (Attack preemptiveness). *Let $M = E \bowtie_S P$.*

- *If there is Q such that $P \xrightarrow{\$s!v} Q$, with $s \notin S$, then there is no M' such that $M \xrightarrow{\tau} M'$ by an application of the rule (SensReadUnsec).*
- *If there is Q such that $P \xrightarrow{\$a?v} Q$, with $a \notin S$, then there is no M' such that $M \xrightarrow{\tau} M'$ by an application of the rule (ActWriteUnsec).*

Rule (Tau) lifts non-observable actions from processes to systems. This includes communications channels and attacks' accesses to unsecured physical devices. A similar lifting occurs in rule (Time) for timed actions, where $\text{next}(E)$ returns the set of possible environments for the next time slot. Thus, by an application of rule (Time) a CPS moves to the next physical state, in the next time slot. Rule (Deadlock) is introduced to signal the violation of the invariant. When the invariant is violated, a system deadlock occurs and then, in CCPSA, the

TABLE II
LTS FOR CPSS

$$\begin{array}{l}
\text{(Out)} \frac{P \xrightarrow{\bar{c}v} P' \quad \text{inv}(E)}{E \bowtie_S P \xrightarrow{\bar{c}v} E \bowtie_S P'} \quad \text{(Inp)} \frac{P \xrightarrow{cv} P' \quad \text{inv}(E)}{E \bowtie_S P \xrightarrow{cv} E \bowtie_S P'} \\
\text{(SensReadSec)} \frac{P \xrightarrow{s?v} P' \quad s \in S \quad \text{inv}(E) \quad v \in \text{read_sensor}(E, s)}{E \bowtie_S P \xrightarrow{\tau} E \bowtie_S P'} \\
\text{(SensReadUnsec)} \frac{P \xrightarrow{s?v} P' \quad s \notin S \quad P \not\xrightarrow{!s?v} \quad \text{inv}(E) \quad v \in \text{read_sensor}(E, s)}{E \bowtie_S P \xrightarrow{\tau} E \bowtie_S P'} \\
\text{(!SensRead !)} \frac{P \xrightarrow{!s?v} P' \quad s \notin S \quad \text{inv}(E) \quad v \in \text{read_sensor}(E, s)}{E \bowtie_S P \xrightarrow{\tau} E \bowtie_S P'} \\
\text{(ActWriteSec)} \frac{P \xrightarrow{a!v} P' \quad a \in S \quad \text{inv}(E) \quad E' = \text{update_act}(E, a, v)}{E \bowtie_S P \xrightarrow{\tau} E' \bowtie_S P'} \\
\text{(ActWriteUnsec)} \frac{P \xrightarrow{a!v} P' \quad a \notin S \quad P \not\xrightarrow{!a?v} \quad \text{inv}(E) \quad E' = \text{update_act}(E, a, v)}{E \bowtie_S P \xrightarrow{\tau} E' \bowtie_S P'} \\
\text{(!ActWrite !)} \frac{P \xrightarrow{!a!v} P' \quad a \notin S \quad \text{inv}(E) \quad E' = \text{update_act}(E, a, v)}{E \bowtie_S P \xrightarrow{\tau} E' \bowtie_S P'} \\
\text{(Tau)} \frac{(P \xrightarrow{\tau} P') \vee (P \xrightarrow{\tau:p} P' \quad p \notin S) \quad \text{inv}(E)}{E \bowtie_S P \xrightarrow{\tau} E \bowtie_S P'} \quad \text{(Deadlock)} \frac{\neg \text{inv}(E)}{E \bowtie_S P \xrightarrow{\text{deadlock}} E \bowtie_S P'} \\
\text{(Time)} \frac{P \xrightarrow{\text{idle}} P' \quad E \bowtie_S P \not\xrightarrow{\tau} \quad \text{inv}(E) \quad E' \in \text{next}(E)}{E \bowtie_S P \xrightarrow{\text{idle}} E' \bowtie_S P'} \quad \text{(Safety)} \frac{\neg \text{safe}(E) \quad \text{inv}(E)}{E \bowtie_S P \xrightarrow{\text{unsafe}} E \bowtie_S P'}
\end{array}$$

system emits a special action deadlock, forever. Similarly, rule (Safety) is introduced to detect the violation of safety conditions. In this case, the system may emit a special action unsafe and then continue its evolution.

Now, having defined the actions that can be performed by a system, we can easily concatenate these actions to define the possible execution traces of the system. Formally, given a trace $t = \alpha_1 \dots \alpha_n$, we will write \xrightarrow{t} as an abbreviation for $\xrightarrow{\alpha_1} \dots \xrightarrow{\alpha_n}$. In the following, we will use the function $\#\text{idle}(t)$ to get the number of occurrences of the action idle in t .

The notion of trace allows us to provide a formal definition of soundness for CPSs: a CPS is said to be *sound* if it never deadlocks and never violates the safety conditions.

Definition 6 (System soundness). *Let M be a well-formed CPS. We say that M is sound if whenever $M \xrightarrow{t} M'$, for some t , both actions deadlock and unsafe never occur in t .*

In our security analysis, we will focus on sound CPSs. For instance, Proposition 2 says that our running example Sys is sound and it never transmits on the channel *alarm*.

Proposition 2. *Let Sys be the CPS defined in Example 2. If $Sys \xrightarrow{t} Sys'$, for some trace $t = \alpha_1 \dots \alpha_n$, then $\alpha_i \in \{\tau, \text{idle}\}$, for any $i \in \{1, \dots, n\}$.*

Actually, we can be quite precise on the temperature reached by Sys before and after the cooling: in each of the 5 rounds

of cooling, the temperature will drop of a value laying in the real interval $[1-\delta, 1+\delta]$, where δ is the uncertainty.

Proposition 3. *Let Sys be the CPS defined in Example 2. For any execution trace of Sys , we have:*

- when Sys turns on the cooling, the value of the state variable *temp* ranges over $(9.9, 11.5]$;
- when Sys turns off the cooling, the value of the variable *temp* ranges over $(2.9, 8.5]$.

The graphic in Figure 3 collects a campaign of 100 simulations, lasting 250 time slots each, showing that the value of the state variable *temp* when the cooling system is turned on (resp., off) lays in the interval $(9.9, 11.5]$ (resp., $(2.9, 8.5]$); these bounds are represented by the dashed horizontal lines.

B. Behavioural semantics

We recall that the *observable activities* in CCPSA are: time passing, system deadlock, violation of safety conditions, and channel communication. Having defined a labelled transition semantics, we are ready to formalise our behavioural semantics, based on execution traces.

We adopt a standard notation for weak transitions: we write \Rightarrow for $(\xrightarrow{\tau})^*$, whereas $\xRightarrow{\alpha}$ means $\Rightarrow \xrightarrow{\alpha} \Rightarrow$, and finally $\xRightarrow{\hat{\alpha}}$ denotes \Rightarrow if $\alpha = \tau$ and $\xRightarrow{\alpha}$ otherwise. Given a trace $t = \alpha_1 \dots \alpha_n$, we write \xrightarrow{t} for $\xrightarrow{\alpha_1} \dots \xrightarrow{\alpha_n}$, and \xRightarrow{t} as an abbreviation for $\xRightarrow{\hat{\alpha}_1} \dots \xRightarrow{\hat{\alpha}_n}$.




Fig. 3. Simulations of the CPS of Example 2

Definition 7 (Trace preorder). We write $M \sqsubseteq N$ if whenever $M \xrightarrow{t} M'$, for some t , there is N' such that $N \xrightarrow{\hat{t}} N'$.

Remark 1. Unlike standard trace semantics, our trace preorder is able to observe deadlock thanks to the presence of the rule (Deadlock) and the special action *deadlock*: if $M \sqsubseteq N$ and M eventually deadlocks then also N must eventually deadlock.

As we are interested in examining timing aspects of attacks, such as beginning and duration, we propose a timed variant of \sqsubseteq up to (a possibly infinite) time interval. Intuitively, we write $M \sqsubseteq_{m..n} N$ if the CPS N simulates the execution traces of M , except for the time interval $m..n$.

Definition 8 (Trace preorder up to a time interval). We write $M \sqsubseteq_{m..n} N$, for $m \in \mathbb{N}^+$ and $n \in \mathbb{N}^+ \cup \{\infty\}$, with $m \leq n$, if the following conditions hold:

- m is the minimum integer for which there is a trace t , with $\#\text{idle}(t) = m - 1$, such that $M \xrightarrow{t}$ and $N \not\xrightarrow{\hat{t}}$;
- n is the infimum element of $\mathbb{N}^+ \cup \{\infty\}$, $n \geq m$, such that whenever $M \xrightarrow{t_1} M'$, with $\#\text{idle}(t_1) = n - 1$, there is t_2 , with $\#\text{idle}(t_1) = \#\text{idle}(t_2)$, such that $N \xrightarrow{t_2} N'$, for some N' , and $M' \sqsubseteq N'$.

In Definition 8, the first item says that N can simulate the traces of M for at most $m - 1$ time slots; whereas the second item says two things: (i) in time interval $m..n$ the simulation does not hold; (ii) after the time slot n the CPS N can simulate again the traces of M . Note that $\inf(\emptyset) = \infty$. Thus, if $M \sqsubseteq_{m..n} N$ then N simulates M only in the first $m - 1$ time slots.

Finally, note that we could have equipped CCPSA with a (bi)simulation-based behavioural semantics rather than a trace-based one, as done in our companion paper [15] for a core of CCPSA with no security features; however, our trace semantics is simpler than (bi)simulation and it is sensitive to deadlocks of CPSs. Thus, it is fully adequate for the purposes of this paper.

III. CYBER-PHYSICAL ATTACKS

In this section, we use CCPSA to formalise a *threat model* where attacks can manipulate sensor and/or actuator signals in order to drive a *sound* CPS into an undesired state [26]. An attack may have different levels of access to physical devices depending on the model assumed. For example, it might be able to get read access to the sensors but not write access; or it might get write-only access to the actuators but not read-access. This level of granularity is very important to model precisely how attacks can affect a CPS [4]. For simplicity, in this paper we don't represent attacks on communication channels as our focus is on attacks to physical devices.

The syntax of our cyber-physical attack is a slight restriction of that for processes: in terms of the form $\text{timeout}[\pi.P]Q$, we require $\pi \in \{\text{write } \hat{z}p(v), \text{read } \hat{z}p(x)\}$. Thus, we provide a syntactic way to distinguish attacks from genuine processes.

Definition 9 (Honest system). A CPS $E \bowtie_S P$ is honest if P is honest, where P is honest if it does not contain prefixes of the form $\text{write } \hat{z}p(v)$ or $\text{read } \hat{z}p(x)$.

We group cyber-physical attacks in classes that describe both the malicious activities and the timing aspects of the attack. 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. Thus, let \mathcal{I} be the set of all possible malicious activities on the physical devices of a system, $m \in \mathbb{N}^+$ be the time slot when an attack starts, and $n \in \mathbb{N}^+ \cup \{\infty\}$ be the time slot when the attack ends. We then say that an *attack* A is of class $C \in [\mathcal{I} \rightarrow \mathcal{P}(m..n)]$ if:

- all possible malicious activities of A coincide with those contained in \mathcal{I} ;
- the first of those actions may occur in the m -th time slot (A admits a run acting at time m) but not before;
- the last of those actions may occur in the n -th time slot (A admits a run acting at time n) but not after;
- for $\iota \in \mathcal{I}$, $C(\iota)$ returns a (possibly empty) set of time slots when A may read/tamper with the device ι ; this set is contained in $m..n$;
- C is a total function: if $C(\iota) = \emptyset$, for some $\iota \in \mathcal{I}$, then no attacks of class C can achieve the malicious activity ι .

Definition 10 (Class of attacks). Let $\mathcal{I} = \{\hat{z}p? : p \in \mathcal{S} \cup \mathcal{A}\} \cup \{\hat{z}p! : p \in \mathcal{S} \cup \mathcal{A}\}$ be the set of all possible malicious activities on physical devices. Let $m \in \mathbb{N}^+$, $n \in \mathbb{N}^+ \cup \{\infty\}$, with $m \leq n$. A class of attacks $C \in [\mathcal{I} \rightarrow \mathcal{P}(m..n)]$ is a total function such that for any attack A of class C we have:

- $C(\iota) = \{k : A \xrightarrow{t} \xrightarrow{\iota v} A' \wedge k = \#\text{idle}(t) + 1\}$, for $\iota \in \mathcal{I}$,
- $m = \inf\{k : k \in C(\iota) \wedge \iota \in \mathcal{I}\}$,
- $n = \sup\{k : k \in C(\iota) \wedge \iota \in \mathcal{I}\}$.

Along the lines of [7], [18], we can say that an attack A affects a *sound* CPS M if the execution of the composed system $M \parallel A$ differs from that of the original system M , in an observable manner. Basically, a cyber-physical attack

can influence the system under attack in at least two different ways:

- The system $M \parallel A$ might deadlock when M may not; this means that the attack A affects the *availability* of the system. We recall that in the context of CPSs, deadlock is a particular severe physical event.
- The system $M \parallel A$ might have non-genuine execution traces containing observables (violations of safety conditions or communications on channels) that cannot be reproduced by M ; here the attack affects the *integrity* of the system behaviour.

Definition 11 (Attack tolerance/vulnerability). *Let M be an honest and sound CPS. We say that M is tolerant to an attack A if $M \parallel A \sqsubseteq M$. We say that M is vulnerable to an attack A if there is a time interval $m..n$, with $m \in \mathbb{N}^+$ and $n \in \mathbb{N}^+ \cup \{\infty\}$, such that $M \parallel A \sqsubseteq_{m..n} M$.*

Thus, if a system M is vulnerable to an attack A of class $C \in [\mathcal{I} \rightarrow \mathcal{P}(m..n)]$, 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* [10]. 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, if $M \parallel A \xrightarrow{t} \xrightarrow{\text{deadlock}}$, for some trace t , the attack A is called *lethal*, as it is capable to halt (deadlock) the CPS M . This is obviously a permanent attack.

Note that, according to Definition 11, the tolerance (or vulnerability) of a CPS also depends on the capability of the *IDS* component 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.

In the following, 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.

In the rest of this section, we present a number of different attacks to the CPS Sys described in Example 2.

Example 3. *Consider the following DoS/Integrity attack on the (controller of) the actuator *cool*, of class $C \in [\mathcal{I} \rightarrow \mathcal{P}(m..m)]$ with $C(\! \not\! \text{cool}?) = C(\! \not\! \text{cool}!) = \{m\}$ and $C(\iota) = \emptyset$, for $\iota \notin \{\! \not\! \text{cool}?, \! \not\! \text{cool}!\}$; we call the attack A_m :*

$$\text{idle}^{m-1}.\text{timeout}[\text{read } \! \not\! \text{cool}(x).\text{if}(x=\text{off})\{\text{write } \! \not\! \text{cool}(\text{off})\}].$$

Here, the attack A_m operates exclusively in the m -th time slot, when it tries to steal the cooling command (on or off) coming from the controller, and fabricates a fake command to turn off the cooling system. In practice, if the controller sends a command to turn off the coolant, nothing bad will

*happen as the attack will put the same message back. When the controller sends (in the m -th time slot) a command to turn the cooling on, the attack will drop the command. We recall that the controller will turn on the cooling only if the sensed temperature is greater than 10 (and hence $\text{temp} > 9.9$); this may happen only if $m > 8$. Since the command to turn the cooling on is never re-sent by *Ctrl*, the temperature will continue to rise, and after 4 time units the system will violate the safety conditions emitting an action *unsafe*, while the *IDS* component will start sending alarms every 5 time units, until the whole system deadlocks because the temperature reaches the threshold of 50 degrees.*

Proposition 4. *Let Sys be the CPS defined in Example 2, and A_m be the attack defined in Example 3. Then,*

- $Sys \parallel A_m \sqsubseteq Sys$, for $m \leq 8$,
- $Sys \parallel A_m \sqsubseteq_{m+4..\infty} Sys$, for $m > 8$.

In this case, the *IDS* component of Sys is effective enough to detect the attack with only one time unit delay.

Example 4. *Consider the following DoS/Integrity attack to the (controller of) sensor s_t , of class $C \in [\mathcal{I} \rightarrow \mathcal{P}(2..\infty)]$ such that $C(\! \not\! s_t?) = \{2\}$, $C(\! \not\! s_t!) = 2..\infty$ and $C(\iota) = \emptyset$, for $\iota \notin \{\! \not\! s_t!, \! \not\! s_t?\}$:*

$$\begin{aligned} A &= \text{idle}.\text{timeout}[\text{read } \! \not\! s_t(x).B(x)] \\ B(y) &= \text{timeout}[\text{write } \! \not\! s_t(y).\text{idle}.B(y)]B(y) . \end{aligned}$$

*Here, the attack A does the following actions in sequence: (i) she sleeps for one time unit, (ii) in the following time slot, she reads the current temperature v at sensor s_t , and (iii) for the rest of her life, she keeps sending the same temperature v to the controller of s_t . In the presence of this attack, the process *Ctrl* never activates the *Cooling* component (and, hence, nor the *IDS* component, which is the only one which could send an alarm) as it will always detect a temperature below 10. Thus, the compound system $Sys \parallel A$ will move to an *unsafe* state until the invariant will be violated and the system will deadlock. Indeed, in the worst scenario, after $\lceil \frac{9.9}{1+\delta} \rceil = \lceil \frac{9.9}{1.4} \rceil = 8$ idle-actions (in the 9-th time slot) the value of *temp* will be above 9.9, and after further 5 idle-actions (in the 14-th time slot) the system will violate the safety conditions emitting an *unsafe* action. After $\lceil \frac{50}{1.4} \rceil = 36$ idle-actions, in the 37-th time slot, the invariant may be broken because the state variable *temp* may reach 50.4 degrees, and the system will also emit a deadlock action. Thus, $Sys \parallel A \sqsubseteq_{14..\infty} Sys$. This is a lethal attack, as it causes a shut down of the system. It is also a stealthy attack as it remains unnoticed until the end.*

*In this attack, the *IDS* component is completely ineffective as the sensor used by the component is compromised, and there is no way for the *IDS* to understand whether the sensor is under attack. A more sophisticated *IDS* might have a representation of the plant to recognise abnormal evolutions of the sensed temperature. In such case, the *IDS* might switch on a second sensor, hoping that this one has not been compromised yet. Another possibility for the designer of the CPS is to secure the sensor. Although this is not always possible,*

as encryption/decryption of all packets depends on energy constraints of the device.

Our semantics ensures that secured devices cannot be attacked, as stated by the following proposition.

Proposition 5. *Let $M = E \bowtie_S P$ be an honest and sound CPS. Let $C \in [\mathcal{I} \rightarrow \mathcal{P}(m..n)]$, with $\{p : C(\sharp p?) \cup C(\sharp p!) \neq \emptyset\} \subseteq \mathcal{S}$. Then $M \parallel A \sqsubseteq M$, for any attack A of class C .*

Now, let us examine a similar but less severe attack.

Example 5. *Consider the following DoS/Integrity attack to the controller of sensor s_t , of class $C \in [\mathcal{I} \rightarrow \mathcal{P}(1..n)]$, for $n > 0$, with $C(\sharp s_t!) = C(\sharp s_t?) = 1..n$ and $C(\iota) = \emptyset$, for $\iota \notin \{\sharp s_t!, \sharp s_t?\}$:*

$$A_n = \text{timeout}[\text{read } \sharp s_t(x).\text{timeout}[\{\text{write } \sharp s_t(x-2).\text{idle}.A_{n-1}\}A_{n-1}]A_{n-1}]$$

with $A_0 = \text{nil}$. In this attack, for n consecutive time slots, A_n sends to the controller the current sensed temperature decreased by an offset 2. The effect of this attack on the system depends on the duration n of the attack itself:

- for $n \leq 8$, the attack is harmless as the variable *temp* may not reach a (critical) temperature above 9.9;
- for $n = 9$, the variable *temp* might reach a temperature above 9.9 in the 9-th time slot, and the attack would delay the activation of the cooling system of one time slot; as a consequence, the system might get into an unsafe state in the time slots 14 and 15, but no alarm will be fired.
- for $n \geq 10$, the system may get into an unsafe state in the time slot 14 and in the following $n - 7$ time slots; this is not a stealthy attack as the IDS will fire the alarm at most two time slots later (in the 16-th time slot); this is a temporary attack which ends in the time slot $n + 7$.

Proposition 6. *Let Sys be the CPS defined in Example 2, and A_n be the attack defined in Example 5. Then:*

- $Sys \parallel A_n \sqsubseteq Sys$, for $n \leq 8$,
- $Sys \parallel A_n \sqsubseteq_{14..15} Sys$, for $n = 9$,
- $Sys \parallel A_n \sqsubseteq_{14..n+7} Sys$, for $n \geq 10$.

A. A technique for proving attack tolerance/vulnerability

In this subsection, we provide sufficient criteria to prove attack tolerance/vulnerability to attacks of an arbitrary class C . Actually, we do more than that: we provide sufficient criteria to prove attack tolerance/vulnerability to all attacks of a class C' that is somehow “weaker” than a given class C .

Definition 12. *Let $C_1 \in [\mathcal{I} \rightarrow \mathcal{P}(m_1..n_1)]$ and $C_2 \in [\mathcal{I} \rightarrow \mathcal{P}(m_2..n_2)]$ be two classes of attacks, with $m_1..n_1 \subseteq m_2..n_2$. We say that C_1 is weaker than C_2 , written $C_1 \preceq C_2$, if $C_1(\iota) \subseteq C_2(\iota)$ for any $\iota \in \mathcal{I}$.*

Intuitively, if $C_1 \preceq C_2$ then: (i) the attacks of class C_1 might achieve fewer malicious activities than any attack of class C_2 (formally, there may be $\iota \in \mathcal{I}$ such that $C_1(\iota) = \emptyset$ and $C_2(\iota) \neq \emptyset$); (ii) for those malicious activities $\iota \in \mathcal{I}$ achieved by the attacks of both classes C_1 and C_2 (i.e., $C_1(\iota) \neq \emptyset$ and

$C_2(\iota) \neq \emptyset$), if they may be perpetrated by the attacks of class C_1 at some time slot $k \in m_1..n_1$ (i.e., $k \in C_1(\iota)$) then all attacks of class C_2 may do the same activity ι at the same time k (i.e., $k \in C_2(\iota)$).

The next objective is to define a notion of *most powerful attack* (also called *top attacker*) of a given class C , such that, if a CPS M tolerates the most powerful attack of class C then it also tolerates *any* attack of class C' , with $C' \preceq C$. We will provide a similar condition for attack vulnerability: let M be a CPS vulnerable to $Top(C)$ in the time interval $m_1..n_1$; then, for any attack A of class C' , with $C' \preceq C$, if M is vulnerable to A then it is so for a smaller time interval $m_2..n_2 \subseteq m_1..n_1$.

Our notion of top attacker has two extra ingredients with respect to the cyber-physical attacks seen up to now: (i) *nondeterminism*, and (ii) *time-unguarded recursive processes*. This extra power of the top attacker is not a problem as we are looking for sufficient criteria.

For what concerns nondeterminism, we assume a generic procedure $rnd()$ that given an arbitrary set \mathcal{Z} returns an element of \mathcal{Z} chosen in a nondeterministic manner. This procedure allows us to express *nondeterministic choice*, $P \oplus Q$, as an abbreviation for the process $\text{if } (rnd(\{\text{true}, \text{false}\})) \{P\} \text{ else } \{Q\}$. Thus, let $\iota \in \{\sharp p? : p \in \mathcal{S} \cup \mathcal{A}\} \cup \{\sharp p! : p \in \mathcal{S} \cup \mathcal{A}\}$, $m \in \mathbb{N}^+$, $n \in \mathbb{N}^+ \cup \{\infty\}$, with $m \leq n$, and $\mathcal{T} \subseteq m..n$, we define the attack process $Att(\iota, k, \mathcal{T})$ as the attack which may achieve the malicious activity ι , at the time slot k , and which tries to do the same in all subsequent time slots of \mathcal{T} . Formally,

$$\begin{aligned} Att(\sharp p?, k, \mathcal{T}) &= \\ &\text{if } (k \in \mathcal{T}) \{(\text{timeout}[\text{read } \sharp p?(x).Att(\sharp p?, k, \mathcal{T})] \\ &\quad Att(\sharp p?, k+1, \mathcal{T})) \oplus \text{idle}.Att(\sharp p?, k+1, \mathcal{T})\} \text{ else} \\ &\text{if } (k < \text{sup}(\mathcal{T})) \{\text{idle}.Att(\sharp p?, k+1, \mathcal{T})\} \text{ else } \{\text{nil}\} \\ Att(\sharp p!, k, \mathcal{T}) &= \\ &\text{if } (k \in \mathcal{T}) \{(\text{timeout}[\text{write } \sharp p!(rnd(\mathbb{R})).Att(\sharp p!, k, \mathcal{T})] \\ &\quad Att(\sharp p!, k+1, \mathcal{T})) \oplus \text{idle}.Att(\sharp p!, k+1, \mathcal{T})\} \text{ else} \\ &\text{if } (k < \text{sup}(\mathcal{T})) \{\text{idle}.Att(\sharp p!, k+1, \mathcal{T})\} \text{ else } \{\text{nil}\} . \end{aligned}$$

Note that for $\mathcal{T} = \emptyset$ we assume $\text{sup}(\mathcal{T}) = -\infty$.

We can now use the definition above to formalise the notion of most powerful attack of a given class C .

Definition 13 (Top attacker). *Let $C \in [\mathcal{I} \rightarrow \mathcal{P}(m..n)]$ be a class of attacks. We define*

$$Top(C) = \prod_{\iota \in \mathcal{I}} Att(\iota, 1, C(\iota))$$

as the most powerful attack, or top attacker, of class C .

The following theorem provides soundness criteria for attack tolerance and attack vulnerability.

Theorem 1 (Soundness criteria). *Let M be an honest and sound CPS, C an arbitrary class of attacks, and A an attack of a class C' , with $C' \preceq C$.*

- If $M \parallel Top(C) \sqsubseteq M$ then $M \parallel A \sqsubseteq M$.
- If $M \parallel Top(C) \sqsubseteq_{m_1..n_1} M$ then either $M \parallel A \sqsubseteq M$ or $M \parallel A \sqsubseteq_{m_2..n_2} M$, with $m_2..n_2 \subseteq m_1..n_1$.

Proof (sketch). The top attacker $Top(C)$ can mimic any execution trace of any attack A of class C' , with $C' \preceq C$ (for

a formal proof of this statement we refer the reader to [16]. Thus, if $M \parallel A \xrightarrow{t}$, for some trace t , then $M \parallel \text{Top}(C) \xrightarrow{t}$ as well.

For any M and A , either $M \parallel A \sqsubseteq M$ or $M \parallel A \sqsubseteq_{m_2..n_2} M$, for some m_2 and n_2 ($m_2 = 1$ and $n_2 = \infty$ if the two systems are completely unrelated). Suppose by contradiction that $M \parallel A \not\sqsubseteq M$ and $M \parallel A \sqsubseteq_{m_2..n_2} M$, with $m_2..n_2 \not\sqsubseteq m_1..n_1$. There are two cases: either $n_1 = \infty$ or $n_1 \in \mathbb{N}^+$.

If $n_1 = \infty$ then $m_2 < m_1$. Since $M \parallel A \sqsubseteq_{m_2..n_2} M$, by Definition 8 there is a trace t , with $\#\text{idle}(t) = m_2 - 1$, such that $M \parallel A \xrightarrow{t}$ and $M \not\xrightarrow{t}$. Since $\text{Top}(C)$ can mimic any execution trace of any attack A of class C' , with $C' \preceq C$, this entails $M \parallel \text{Top}(C) \xrightarrow{t}$. Since $M \not\xrightarrow{t}$ and $\#\text{idle}(t) = m_2 - 1 < m_1$, this contradicts $M \parallel \text{Top}(C) \sqsubseteq_{m_1..n_1} M$.

If $n_1 \in \mathbb{N}^+$ then $m_2 < m_1$ and/or $n_1 < n_2$, and we reason as in the previous case. \square

Corollary 1. *Let M be an honest and sound CPS, and C a class of attacks. If $\text{Top}(C)$ is not lethal for M then any attack A of class C' , with $C' \preceq C$, is not a lethal attack for M . If $\text{Top}(C)$ is not a permanent attack for M then any attack A of class C' , with $C' \preceq C$, is not a permanent attack for M .*

Remark 2. *The reader may wonder about the scalability of the verification method proposed in Theorem 1. We would like to recall that in our companion paper [15] we developed a (bi)simulation (pre)congruence for a simpler version of the calculus where security features have been completely stripped off. For simplicity, in the current paper, we adopted as main behavioural semantics trace preorder instead of simulation preorder. We believe that switching to a simulation precongruence, preserved by parallel composition, would allow us to scale our verification method to bigger systems (i.e., we believe that using simulation precongruence instead of trace preorder would allow our verification method to enjoy scalability “for free”). However, this will need to be checked thoroughly in order to consider the more complex version of the calculus with security features that we focus on here.*

In the following, we provide a small example to explain how Theorem 1 could be used to infer attack tolerance/vulnerability with respect to an entire class of attacks.

Example 6. *Consider the CPS Sys of Example 2 and a class of attacks C , such that $C(\sharp s_t!) = C(\sharp s_t?) = \{13\}$, and $C(\iota) = \emptyset$, for $\iota \notin \{\sharp s_t!, \sharp s_t?\}$. Attacks of class C can read/tamper with the sensor s_t (exclusively) in the time slot 13 (i.e., in the time interval 13..13). In Proposition 7 we will show that:*

- $\text{Top}(C)$ may never lead the system Sys in deadlock,
- $\text{Top}(C)$ may drag the system Sys in an unsafe state only in the time interval 17..19 (no alarms are possible).

As a consequence, each attack of class C (or weaker than C) cannot achieve more than that. Let us provide three different attacks of classes weaker than or equal to C .

Consider the attack:

$$A_1 = \text{idle}^{12}.\text{timeout}[\text{write } \sharp s_t(11)] .$$

This attack forces the activation of the cooling system without any consequence to the observable behaviour of the whole system (see Proposition 7). The attack belongs to a class C' , with $C' \preceq C$, such that $C'(\sharp s_t!) = \{13\}$ and $C'(\iota) = \emptyset$, for $\iota \neq \sharp s_t!$.

Consider a different attack:

$$A_2 = \text{idle}^{12}.\text{timeout}[\text{write } \sharp s_t(9)] .$$

This attack is of the same class C' . However, it may delay the activation of the cooling system, and due to such delay the CPS under attack may get into an unsafe state in the time interval 17..19 (see Proposition 7).

Finally, consider the attack:

$$A_3 = \text{idle}^{12}.\text{timeout}[\text{read } \sharp s_t(x).\text{if } (x=10.5) \{ \text{timeout}[\text{write } \sharp s_t(9)] \}] .$$

This is an attack of class C which delays the activation of the cooling system as well; however, this delay may only occur when the temperature is quite low. As a consequence, the system under attack may get into an unsafe state only in the time slot 18, i.e., the time interval 18..18, (again, see Proposition 7).

The following proposition formalises the statements contained in Example 6.

Proposition 7. *Let Sys be the CPS of Example 2 and C be a class of attacks such that $C(\sharp s_t!) = C(\sharp s_t?) = \{13\}$, and $C(\iota) = \emptyset$, for $\iota \notin \{\sharp s_t!, \sharp s_t?\}$. Then,*

- $\text{Sys} \parallel \text{Top}(C) \sqsubseteq_{17..19} \text{Sys}$,
- $\text{Sys} \parallel A_1 \sqsubseteq \text{Sys}$,
- $\text{Sys} \parallel A_2 \sqsubseteq_{17..19} \text{Sys}$,
- $\text{Sys} \parallel A_3 \sqsubseteq_{18..18} \text{Sys}$.

IV. IMPACT OF AN ATTACK

In the previous section, we have grouped cyber-physical attacks by focussing on the physical devices under attack and the timing aspects of the attack (Definition 10). Then, we have provided a formalisation of when a CPS should be considered tolerant/vulnerable to an attack (Definition 11). In this section, we show that it is important not only to demonstrate the tolerance (or vulnerability) of a CPS with respect to certain attacks, but also to evaluate the disruptive impact of those attacks on the target CPS [9].

The goal of this section is thus twofold: to provide a *metric* to estimate the impact of a successful attack on a CPS, and to investigate possible quantifications of the chances for an attack of being successful when attacking a CPS.

As to the metric, we focus on the ability that an attack may have to drag a CPS out of the correct behaviour modelled by its evolution map, with the given uncertainty. Recall that *evol* is *monotone* with respect to the uncertainty. Thus, an increase of the uncertainty may translate into a widening of the range of the possible behaviours of the CPS.

In the following, for $M = E \bowtie_S P$, we write $M[\psi \leftarrow \psi']$ to mean $E[\psi \leftarrow \psi'] \bowtie_S P$.




Fig. 4. Simulation of $Sys[\delta \leftarrow \frac{29}{30}]$.

Proposition 8 (Monotonicity). *Let $M = E \bowtie_S P$ be an honest and sound CPS, and ξ_w the uncertainty of E . If $\xi_w \leq \xi'_w$ and $M \xrightarrow{t} M'$ then $M[\xi_w \leftarrow \xi_w] \xrightarrow{t} M'[\xi_w \leftarrow \xi'_w]$.*

However, a wider uncertainty in the model doesn't always correspond to a widening of the possible behaviours of the CPS. In fact, this depends on the *intrinsic tolerance* of a CPS with respect to changes in the uncertainty function.

Definition 14 (System ξ -tolerance). *An honest and sound CPS $M = E \bowtie_S P$, where ξ_w is the uncertainty of E , is ξ -tolerant, for $\xi \in \mathbb{R}^{\mathcal{X}}$ and $\xi \geq 0$, if*

$$\xi = \sup \{ \xi' : M[\xi_w \leftarrow \xi_w + \eta] \sqsubseteq M, \text{ for any } 0 \leq \eta \leq \xi' \}.$$

Intuitively, if a CPS M has been designed with a given uncertainty ξ_w , but M is actually ξ -tolerant, with $\xi > 0$, then the uncertainty ξ_w is somehow underestimated: the real uncertainty of M is given by $\xi_w + \xi$. This information is quite important when trying to estimate the impact of an attack on a CPS. In fact, if a system M has been designed with a given uncertainty ξ_w , but M is actually ξ -tolerant, with $\xi > 0$, then an attack has (at least) a “room for manoeuvre” ξ to degrade the whole CPS without being observed (and hence detected).

Let Sys be the CPS of Example 2. In the rest of the section, with an abuse of notation, we will write $Sys[\delta \leftarrow \gamma]$ to denote Sys where the uncertainty of the variable *temp* is γ .

Example 7. *The CPS Sys of Example 2 is $\frac{1}{20}$ -tolerant. This is because $\sup \{ \xi' : Sys[\delta \leftarrow \delta + \eta] \sqsubseteq Sys, \text{ for } 0 \leq \eta \leq \xi' \}$ is equal to $\frac{1}{20}$. Since $\delta + \xi = \frac{8}{20} + \frac{1}{20} = \frac{9}{20}$, the proof of this statement relies on the following proposition.*

Proposition 9. *Let Sys be the CPS of Example 2. Then:*

- $Sys[\delta \leftarrow \gamma] \sqsubseteq Sys$, for $\gamma \in (\frac{8}{20}, \frac{9}{20})$,
- $Sys[\delta \leftarrow \gamma] \not\sqsubseteq Sys$, for $\gamma > \frac{9}{20}$.

Figure 4 shows an evolution of $Sys[\delta \leftarrow \frac{29}{30}]$: the red box denotes a violation of the safety conditions because the cooling cycle wasn't sufficient to drop the (sensed) temperature below 10 (here, the controller imposes 5 further time units of cooling).

Now everything is in place to define our metric to estimate the impact of an attack.

Definition 15 (Impact). *Let $M = E \bowtie_S P$ be an honest and sound CPS, where ξ_w is the uncertainty of E . We say that an attack A has definitive impact ξ on the system M if*

$$\xi = \inf \{ \xi' : \xi' \in \mathbb{R}^{\mathcal{X}} \wedge \xi' > 0 \wedge M \parallel A \sqsubseteq M[\xi_w \leftarrow \xi_w + \xi'] \}.$$

It has pointwise impact ξ on the system M at time m if

$$\xi = \inf \{ \xi' : \xi' \in \mathbb{R}^{\mathcal{X}} \wedge M \parallel A \sqsubseteq_{m..n} M[\xi_w \leftarrow \xi_w + \xi'], n \in \mathbb{N} \cup \{\infty\} \}.$$

Intuitively, the impact of an attacker A on a system M measures the uncertainty introduced by the presence of the attacker in the compound system $M \parallel A$ with respect to the original system M .

With this definition, we can establish either the definitive (and hence maximum) impact of the attack A on the system M , or the impact at a specific time m . In the latter case, by definition of $\sqsubseteq_{m..n}$, there are two possibilities: either the impact of the attack keeps growing after time m , or in the time interval $m+1$, the system under attack deadlocks.

The impact of $Top(C)$ provides an upper bound for the impact of all attacks of class C' , with $C' \preceq C$.

Theorem 2 (Top attacker's impact). *Let M be an honest and sound CPS, and C an arbitrary class of attacks. Let A be an attack of class C' , with $C' \preceq C$.*

- *The definitive impact of $Top(C)$ on M is greater than or equal to the definitive impact of A on M .*
- *If $Top(C)$ has pointwise impact ξ on M at time m , and A has pointwise impact ξ' on M at time m' , with $m' \leq m$, then $\xi' \leq \xi$.*

Proof (sketch). Consider the case of the definitive impact. The top attacker $Top(C)$ can mimic any execution trace of any attack A of class C' , with $C' \preceq C$ (for a formal proof of this statement we refer the reader to [16]). As a consequence, $M \parallel A \xrightarrow{t}$ entails $M \parallel Top(C) \xrightarrow{t}$. This implies $M \parallel A \sqsubseteq M \parallel Top(C)$. Thus, if $M \parallel Top(C) \sqsubseteq M[\xi_w \leftarrow \xi_w + \xi]$, for $\xi \in \mathbb{R}^{\mathcal{X}}$, $\xi > 0$, then, by transitivity of \sqsubseteq , it follows that $M \parallel A \sqsubseteq M[\xi_w \leftarrow \xi_w + \xi]$.

The proof in the case of the pointwise impact is by contradiction. \square

In order to provide more intuition about what the metrics given in Definition 15 measure, we give a couple of examples.

Example 8. *Let us consider the attack A_n of Example 5, for $n = 10$. Then, A_{10} has a definitive impact of $\frac{7}{30}$ to the CPS Sys defined in Example 2. Formally,*

$$\frac{7}{30} = \inf \{ \xi' : \xi' > 0 \wedge Sys \parallel A_{10} \sqsubseteq Sys[\delta \leftarrow \delta + \xi'] \}.$$

Here, the attack may drag the system into an unsafe state during the time interval 14..17 (temporary attack). As the attack can affect the system for only 4 time slots, its impact is relatively low: the temperature may rise of at most 1.63 degrees per time slot, instead of 1.4.

Technically, since $\delta + \xi = \frac{12}{30} + \frac{7}{30} = \frac{19}{30}$, the proof of what stated in this example relies on the following proposition.

Proposition 10. Let Sys be the CPS defined in Example 2, and A_{10} be the attack defined in Example 8. Then:

- $Sys \parallel A_{10} \not\sqsubseteq Sys[\delta \leftarrow \gamma]$, for $\gamma \in (\frac{12}{30}, \frac{19}{30})$,
- $Sys \parallel A_{10} \sqsubseteq Sys[\delta \leftarrow \gamma]$, for $\gamma > \frac{19}{30}$.

The following example shows that the attack provided in Example 4 has much stronger impact on the CPS Sys .

Example 9. Let us consider the attack A of Example 4. As this attack prevents the activation of the cooling system, the temperature will keep growing until the CPS Sys gets into an unsafe state and eventually deadlocks. This is a stealthy lethal attack that has a very severe and high impact. In fact, it has a definitive impact of 8.5 on the CPS Sys . Formally,

$$8.5 = \inf \{ \xi' : \xi' > 0 \wedge Sys \parallel A \sqsubseteq Sys[\delta \leftarrow \delta + \xi'] \}.$$

Intuitively, in order to deadlock the system without being noticed, the attack must induce a rapid increase of the temperature: 9.9 degrees per time slot, instead of 1.4.

Technically, since $\delta + \xi = 0.4 + 8.5 = 8.9$, the proof of what is stated in this example relies on the following proposition.

Proposition 11. Let Sys be the CPS defined in Example 2, and A be the attack defined in Example 9. Then:

- $Sys \parallel A \not\sqsubseteq Sys[\delta \leftarrow \gamma]$, for $\gamma \in (0.4, 8.9)$,
- $Sys \parallel A \sqsubseteq Sys[\delta \leftarrow \gamma]$, for $\gamma > 8.9$.

Definition 15 provided an instrument to estimate the impact of a successful attack. However, there is at least another question that a CPS designer could ask: “Is there a way to estimate the chances that an attack will be successful during the execution of my CPS?” To paraphrase in a more operational manner: how many execution traces of my CPS are prone to be attacked by a specific attack?

For instance, consider again the simple attack A_m proposed in Example 3:

$$\text{idle}^{m-1}.\text{timeout}[\text{read } \not\text{cool}(x).\text{if } (x=\text{off}) \{ \text{write } \not\text{cool}(\text{off}) \}].$$

Here, in the m -th time slot the attack drops a command to turn on the cooling. The attack is very quick and condensed in a single time slot. The question is: what are the chances of success of such a quick attack?

Figure 5 provides a representation of an experiment in MATLAB where we launched 10000 executions of our CPS in isolation, lasting 700 time units each. From the aggregated data contained in this graphic, we note that after a transitory phase (whose length depends on several things: the uncertainty δ , the initial state of the system, the length of the cooling activity, etc.) that lasts around 300 time slots, the rate of success of the attack A_m is around 10%. The reader may wonder why exactly the 10%. This depends on the periodicity of our CPS, as in average the cooling is activated every 10 time slots.

This example shows that, as pointed out in [10], the effectiveness of a cyber-physical attack depends on the information the attack has about the functionality of the whole CPS. For instance, if the attacker were not aware of the exact periodicity of the CPS, she might try, if possible, to repeat the attack on




Fig. 5. A quantitative analysis of the attack of Example 3.

more consecutive time slots. In this case, the left graphic of Figure 6 says that the rate of success of the attack increases linearly with the length of the attack itself (data obtained by attacking the CPS after the transitory period). Thus, if the attack of Example 3 were iterated for 10 time slots, say

$$\begin{aligned} A_m^{10} &= \text{idle}^{m-1}.B_{10} \\ B_i &= \text{timeout}[\text{read } \not\text{cool}(x).\text{if } (x = \text{off}) \\ &\quad \{ \text{write } \not\text{cool}(\text{off}).\text{idle}.B_{i-1} \}]B_{i-1}, \text{ for } 1 \leq i \leq 10 \end{aligned}$$

with $B_0 = \text{nil}$, the rate of success would be almost 100%.

Finally, consider a generalisation of the attack of Example 5:

$$\begin{aligned} A_0^k &= \text{nil} \\ A_n^k &= \text{timeout}[\text{read } \not\text{s}_t(x).\text{timeout}[\text{write } \not\text{s}_t(x-k) \\ &\quad \text{idle}.A_{n-1}^k]A_{n-1}^k]A_{n-1}^k \end{aligned}$$

for $1 \leq n \leq 15$ and $2 \leq k \leq 10$. Here, the attack decreases the sensed temperature of an offset k . Now, suppose to launch this attack after, say, 300 time slots (i.e., after the transitory phase). Formally, we define the attack: $B_n^k = \text{idle}^{300}.A_n^k$. In this case, the right graphic of Figure 6 provides a graphical representation of the percentage of alarms on 5000 execution traces lasting 100 time units each. Thus, for instance, an attack lasting $n = 8$ time units with an offset $k = 5$ affects around 40% of the execution traces of the CPS.

These two examples show that 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. This would require a probabilistic extension of our approach, which we will investigate in future work.

V. CONCLUSIONS, RELATED AND FUTURE WORK

We have provided 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 its labelled transition 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,




Fig. 6. A quantitative analysis of two different attacks.

we relied on the trace semantics of *CCPSA* to assess *attack tolerance/vulnerability* with respect to a given attack. Along the lines of GNDC [7], 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). As we remarked above, we believe that we could switch to a (bi)simulation semantics, preserved by parallel composition, which would in turn allow us to scale Theorem 1 and, more generally, our approach to bigger systems, but this will require a thorough investigation. 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).

Related work: A number of approaches have been proposed for modelling CPSs using *hybrid process algebras* [5], [2], [24], [8]. *CCPSA* shares some similarities with the ϕ -calculus [24]. 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 [31], 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.

We focus on the papers that are most closely related to our work. Huang et al. [12] were among the first to propose *threat models* for CPSs. Along with [13], [14], they stressed the role played by timing parameters on integrity and DoS attacks. Alternative threat models are discussed in [9], [10], [26]. In particular, Gollmann et al. [10] discussed possible goals (*equipment damage, production damage, compliance violation*) and *stages* (*access, discovery, control, damage, cleanup*) of

cyber-physical attacks. In the current paper, we focused on the damage stage, where the attacker already has a rough idea of the plant and the control architecture of the target CPS.

A few works use formal methods for CPS security, although they apply methods, and most of the time have goals, that are quite different from ours.

Burmester et al. [3] employed *hybrid timed automata* to give a threat framework based on the traditional Byzantine faults model for crypto-security. However, as remarked in [26], 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.

In [28], Vigo 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 [29], [30] 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. [4] and Zhu and Basar [32] 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 [19].

Finally, there are three recent papers that were developed in parallel to ours: [20], [23], [22]. Rocchetto and Tippenhaur [23] introduced a taxonomy of the diverse attacker models proposed for CPS security and outline requirements for generalised attacker models; in [22], 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. [20] 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 [21] 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.

Future work: 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. We expect that *weak simulation quasimetrics* [17] will be useful to that extent.

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.

Acknowledgements: We thank the anonymous reviewers for their insightful comments.

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. *Theoretical 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] 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.
- [5] P. Cuijpers and M. Reniers. Hybrid process algebra. *The Journal of Logic and Algebraic Programming*, 62(2):191–245, 2005.
- [6] N. Falliere, L. Murchu, and E. Chien. W32.Stuxnet Dossier, 2011.
- [7] 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.
- [8] V. Galpin, L. Bortolussi, and J. Hillston. HYPE: Hybrid modelling by composition of flows. *Formal Asp. Comput.*, 25(4):503–541, 2013.
- [9] B. Genge, I. Kiss, and P. Haller. A system dynamics approach for assessing the impact of cyber attacks on critical infrastructures. *Int. J. Critical Infrastructure Protection*, 10:3–17, 2015.
- [10] 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.
- [11] M. Hennessy and T. Regan. A process algebra for timed systems. *Information and Computation*, 117(2):221–239, 1995.
- [12] 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.
- [13] 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.
- [14] 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.
- [15] R. Lanotte and M. Merro. A Calculus of Cyber-Physical Systems. In *LATA*, volume 10168 of *LNCS*, pages 115–127. Springer, 2017.
- [16] R. Lanotte, M. Merro, R. Muradore, and L. Viganò. A Formal Approach to Cyber-Physical Attacks. *CoRR*, abs/1611.01377, 2016.
- [17] R. Lanotte, M. Merro, and S. Tini. Weak simulation quasimetric in a gossip scenario. In *FORTE*, volume 10321 of *LNCS*. Springer, 2017.
- [18] D. Macedonio and M. Merro. A semantic analysis of key management protocols for wireless sensor networks. *Science of Computer Programming*, 81:53–78, 2014.
- [19] 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.
- [20] 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.
- [21] P. C. Ölveczky and J. Meseguer. Semantics and pragmatics of real-time maude. *Higher-Order and Symbolic Computation*, 20(1-2):161–196, 2007.
- [22] 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.
- [23] 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.
- [24] 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.
- [25] 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.
- [26] A. Teixeira, I. Shames, J. Sandberg, and K. H. Johansson. A secure control framework for resource-limited adversaries. *Automatica*, 51:135–148, 2015.
- [27] U.S. Chemical Safety and Hazard Investigation Board, T2 Laboratories Inc. Reactive Chemical Explosion: Final Investigation Report. Report No. 2008-3-I-FL, 2009.
- [28] R. Vigo. The Cyber-Physical Attacker. In *SAFECOMP*, volume 7613 of *LNCS*, pages 347–356. Springer, 2012.
- [29] R. Vigo. *Availability by Design: A Complementary Approach to Denial-of-Service*. PhD thesis, Danish Technical University, 2015.
- [30] 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.
- [31] Y. Zaccchia Lun, A. D’Innocenzo, I. Malavolta, and M. D. Di Benedetto. Cyber-Physical Systems Security: a Systematic Mapping Study. *CoRR*, abs/1605.09641, 2016.
- [32] Q. Zhu and T. Basar. Game-theoretic methods for robustness, security, and resilience of cyberphysical control systems: games-in-games principle for optimal cross-layer resilient control systems. *IEEE Control Systems*, 35(1):46–65, 2015.