Exploring the quantitative resilience analysis of cyber-physical systems

Romain Dagnas, Michel Barbeau, Maxime Boutin, Joaquin Garcia-alfaro, Reda Yaich

Exploring the Quantitative Resilience Analysis of Cyber-Physical Systems

Romain Dagnas* , Michel Barbeau† , Maxime Boutin* , Joaquin Garcia-Alfaro‡ , Reda Yaich*
*Institut de Recherche Techologique SystemX, Palaiseau, France
†School of Computer Science, Carleton University, Ottawa, Canada
‡SAMOVAR, Télécom SudParis, Institut Polytechnique de Paris, Palaiseau, France

***Abstract*—With technological advances, Cyber-Physical Systems (CPSs), specifically critical infrastructures, have become strongly connected. Their exposure to cyber adversaries is higher than ever. The number of cyber-attacks perpetrated against critical infrastructure is growing in number and sophistication. The protection of such complex systems is of paramount importance. Resilience applied to critical infrastructures aims at protecting these vital systems from cyber-attacks and making them continue to deliver a certain level of performance, even when attacks occur. In this work, we explore new advances related to cyber-resilience applied to CPSs. We also explore using a metric to quantify the resilience of critical infrastructures. As a use case, we consider a water treatment system.**

***Index Terms*—Cyber-resilience, critical infrastructures, cyber-physical system, resilience, spectral radius.**

I. Introduction

Nowadays, the increase in competitiveness leads to a race to digitization, making the Cyber-Physical Systems (CPS) domain, and specifically, critical infrastructures, more and more connected and, thus, more exposed to cyber-attacks. The term CPS was first coined by Gill [11]. They comprise hardware and software components communicating with the real and the cyber worlds. We live in a digitization era. The race for competitiveness in every field (health, robotics, avionics, railways, etc.) implies that critical systems have become increasingly connected to the cyber-space. Thus, they become more and more exposed to cyber-attacks. Critical infrastructures are considered as complex CPS. Every day, the news proves how essential systems protection is. Indeed, critical infrastructures are considered vital for industries, organizations, and countries. An attack perpetrated against complex systems could have disastrous consequences and damages. As an example of such devastating consequences, we can cite the cyber-attack perpetrated against the Florida water treatment station [18]. The cyber-adversary responsible for this attack has delivered a quantity of sodium hydroxide 100 times higher than the usual amount. A significant quantity of such a product added to drinking water can harm citizens' health and, in extreme cases, may cause death. Infrastructures such as water treatment systems, nuclear plants, electrical distribution stations, petroleum stations, and railways are called *critical* because they directly impact people's lives. Resilience applied to critical infrastructures aims at protecting these systems from the adverse effects of cyber adversaries.

ISBN 978-3-903176-57-7

We explore recent works published in the field of quantification and assessment of cyber resilience. We seek to use metrics for quantifying resilience by applying them to a critical infrastructure example, namely a water treatment system whose architecture is inspired by the Secure Water Treatment System (SWaT). We have deployed the STPA method (Systems-Theoretic Process Analysis), a hazard analysis based on STAMP (Systems-Theoretic Accident Model and Processes), on the water treatment system.

This paper is organized as follows. Section II goes over commonly recognized definitions of the notion of cyber-resilience. Section III reviews related work. In Section IV, we apply a resilience assessment strategy to a water treatment system and compare the results obtained with two metrics previously introduced in the literature. We discuss axes for further research in Section V. Finally, Section VI concludes this work.

II. Cyber-resilience, a consensus for a definition

The resilience term, initially applied in ecology by Holling [15] to quantify a population's ability to recover from changes, became an important research axis in the cybersecurity community. Resilience is used in many fields [16] such as ecology (resilience facing natural events), psychology (resilience facing a trauma), economy (resilience facing market changes), etc. In the field of computer sciences, the notion of resilience is called *cyber-resilience* when it is related to cyber-physical systems and, in a more general way, to cybersecurity. However, in many cybersecurity-related works, the authors do not usually use the *cyber-resilience* designation and use the *resilience* appellation. Various definitions of resilience can be found in the literature. Many of these definitions have the notion of *performance* as a common basis. Indeed, a well-known way to quantify the resilience of a system is to consider its performance degradation facing a disruptive event, for example, a cyber-attack. The performance of a system is highly related to *mission delivery*, and most of the definitions of resilience applied in the cybersecurity field consider this aspect. A well-known definition from [23] defines cyber-resilience as *The system's ability to recover or regenerate its performance after a cyber-attack produces a degradation to its performances*. Another similar definition describes resilience as *The ability to anticipate, withstand, recover from, and adapt to adverse conditions, stresses, attacks, or compromises on*

systems that use or are enabled by cyber resources. [31]. Other works such as [8] provide a more exhaustive classification of resilience definitions found in the literature, depending on systems properties, service delivery, and events handling. According to [9], it is possible to highlight the main principles of resilience as: (i) Anticipate, i.e., maintain a state of informed preparedness to forestall compromises of mission functions from adversarial activities; (ii) Withstand, i.e., continue essential mission/business functions despite successful execution of an attack by an adversary; (iii) Recover, i.e., restore mission/business functions to the maximum extent possible after the successful execution of an attack by an adversary; and (iv) Adapt, i.e., change the mission functions or the supporting cyber capabilities to minimize adverse impacts from actual or predicted attacks. Without being able to define and measure cyber resilience adequately, it will be difficult for organizations to self-monitor and share critical infrastructure among other organizations and policymakers to observe the posture of systems and enforce policies [33], and as noted by Linkov and Kott [24], to improve the cyber resilience of a system, you have to measure it. That's why in the past years, there have been massive efforts from academics and organizations to develop new metrics, models, and frameworks to identify cyber resilience goals, objectives, practices, and costs [6]. According to Linkov and Kott, there are two different approaches to measuring resilience [23]. The first approach considers the use of metrics based, seen as technical measurements of individual properties of system components or functions, to assess overall system performance. The second approach relies on modeling, using system configuration modeling and scenario analysis to predict system evolution.

We focus on metrics to measure CPS cyber resiliency. Metrics are defined as measurable properties of the system that quantify the degree to which the objectives of the system are achieved. They can be either qualitative (connected with what something is like rather than how much of it there is) or quantitative (connected with the amount or number rather than with how good it is). The development of the Cyber Resiliency Metrics Catalog by MITRE [7] alongside a new NIST standard on Developing Cyber-Resilient Systems [31] and a standard Cyber Resiliency Engineering Framework (CERF) by the MITRE [6] show that this domain is taken seriously by stakeholders. The next section presents works related to the vulnerability quantification of systems, formalization of loss scenarios, attack impact quantification techniques, functional diversity and variability for enhancing resilience, and resilience quantification techniques.

III. Related Work

A. Vulnerability quantification

The use of prediction techniques to estimate the evolution of adversarial activities can combine the use of attack surface metrics [14], combined with tree structures and graph theoretic-solutions [19], [22], [32]. The goal is to analyze and quantify the likelihood of being affected by cyber-attacks due to existing vulnerabilities in the system. Extended solutions can also lead to proactive automation of countermeasures right after mapping threats to handle vulnerability exploitation [29] properly. Early solutions based on attack graphs can be constructed using the results of network scanners [30], i.e., building a mapping between network topology and existing vulnerabilities in each node. This type of graph offers a well-structured strategy to add weights to network resources based on their likelihood of being affected by cyber-attacks. This way, attack evolution can be formally represented as how adversaries keep accessing new resources by exploiting existing vulnerabilities. Another strategy relies on Bayesian-based reasoning, in which graph edges and vertices are assigned probabilities based on their likelihood of being affected by the exploitation of underlying vulnerabilities. However, the nature of adversaries is driven by the achievement of detrimental goals against the system rather than probabilities or topological weights [3].

B. The importance of critical sequences

STAMP-STPA is a hazard analysis method used to accurately find loss scenarios from different families of complex system incidents and control structure models [25]. The outputs are the following ones: (1) loss scenarios leading to unsafe control actions and (2) scenarios where control actions are not or are improperly executed. These loss scenarios describe the causal factors that can lead to unsafe control actions and hazards and, as such, can be viewed as critical sequences that cybersecurity stakeholders must avoid to occur. STAMP-STPA considers not only cyber-based threats but also errors or control actions improperly executed by human staff interacting with the systems. Indeed, it is well-known that the human factor could be the cause of many failures or security breaches, which is not necessarily a correct assumption as operator behavior are a product of their environment. To reduce operator *error*, we must change the environment in which they work [25].

C. Attack impact quantification

Any cyber-security enhancement strategy must take into consideration attack impact quantification metrics. Kotenko et al. [21] present such quantification methods by considering the resilience assessment of computer networks and their stochastic network representations. Several previous works deal with this notion of attack impact quantification. Kotenko and Chechulin present a tool used for modeling and analysis of attacks using attack graph and service dependency models combined with several cybersecurity metrics such as the Common Vulnerability Scoring Systems (CVSS) scores, and several lists of known information security vulnerabilities such as Common Vulnerabilities and Exposures (CVE) or the National Vulnerability Database (NVD) to evaluate the security and quantify the possible impact of an attack [20]; Dudorov et al. introduced an approach for the generation of stochastic models of attacks [10]; Abraham and Nair have presented a framework that allows to take into account vulnerabilities and to calculate quantitative values for security metrics [2].

D. Functional diversity and variability for resilience purposes

Resilience strategies applied to complex systems are very diversified. Functional diversity and redundancy are significant aspects of resilience strategies. Redundancy is more related to safety and robustness. Functional diversity is generally applied in resilience enhancement strategies. Indeed, it is known that redundancy does not consist in a sufficient challenge for cyber-adversaries than functional diversity in resilience improvement of CPSs. However, functional diversity includes components that measure or act at the level of different physical phenomena, which is a higher challenge for an adversary who attempts to remain invisible from an attack detection strategy [4], [5].

E. Resilience quantification

1) Spectral radius: Lewis highlighted the importance of quantifying system resilience with metrics [26]. Design alternatives can be objectively evaluated and compared. It can be done using either empirical data and/or analyzing the structure of a representation of a system. In his article, two empirical metrics are introduced, exceedance probability and probable loss, and one structural metric, spectral radius.

Spectral radius is a metric calculated over a CPS modeled as a graph. The network nodes correspond to the components of the CPS while the links represent relationships between components. Let n be the number of components. A CPS can be represented as a $n \times n$ connection matrix C with elements in $\{0, 1\}$. For $i, j \in 1, \ldots, n$, $c_{i,j} = 1$ when component i is related to component j. Otherwise, is it equal to 0. A relationship may be a physical or logical connection. Let $\lambda_1, \ldots \lambda_n$ be the eigenvalues of C. The spectral radius of the connection matrix C, denoted as $\rho(C)$, is defined as

$$\rho(C) = \underset{i=1,\ldots,n}{\arg\max} |\lambda_i|. \quad (1)$$

It is the largest absolute value of the eigenvalues. Lewis observed that the spectral radius increases with the density of links, i.e., the number of links per node, or the number of connections that network hubs may have [26]. The high density of links and presence of hubs increase the risk of cascading failures, i.e., propagation of failures from neighbor component to neighbor component. However, this is not the only way to interpret the spectral radius. Indeed, an increase in the value of the spectral radius versus different architectures also reflects a lower risk of getting unreachable nodes or states in the case of a node deletion, i.e., a component failure. Thus, this interpretation implies that a higher spectral radius reflects a higher resiliency due to more connections between the graph nodes.

2) (k, ℓ)-resilience: The works [4], [5] have introduced the (k, ℓ)-resiliency metric and how to compute it. We present the necessary material to understand the (k, ℓ)-resiliency property, applied to CPSs modeled by state-space representations with A, B, C and D be respectively the state, input, output and direct transmission matrices of a system based on a differential equations representation; and x, u and y be respectively the state, input and output vectors. The (k, ℓ)-resiliency is based on inter-variable dependencies and on dependency graphs. The notion of inter-variable dependencies is based on the Pearson correlation coefficient. This correlation coefficient is unitless between -1 and 1. If $\rho(A, B) = 1$ their is a perfect positive correlation between A and B [34].

Now we consider u, x, and y, which are respectively p-element, m-element, and n-element column vectors representing the input, state, and output variables of a CPS modeled as a state-space representation. Correlation coefficient matrices are used to capture in structures the relationships between state variables and input or output variables. The following notions have been presented in the works [4], [5].

An $m \times p$ [$m \times n$] input [output] correlation coefficient matrix Q [R] is equal to $(q_{i,j})$ [$(r_{i,j})$] with $i = 1, \ldots, m$, $j = 1, \ldots, p$ [$j = 1, \ldots, n$]. An entry of this matrix $q_{i,j}$ [$r_{i,j}$] is the correlation coefficient $\rho(x_i, u_j)$ [$\rho(x_i, y_j)$] between x_i and u_j [y_j], i.e., between the state variable and the input [output] variable.

An input [output] dependency graph consists of a bipartite graph $G_U = (X, U, E)$ [$G_Y = (X, Y, E)$] where the two sets of vertices are $X = \{x_1, \ldots, x_m\}$ and $U = \{u_1, \ldots, u_p\}$ [$Y = \{y_1, \ldots, y_n\}$] the state and input [output] variables, respectively. There is an edge (x_i, u_i) [(x_i, y_i)] in E if-and-only-if the absolute value of the correlation between x_i and u_i [y_i] is greater than or equal to a threshold T, chosen to be close to one.

Following these notions, by considering a dependency graph G_U [G_Y] and a vertex x in X, the expression $\deg(x)$ represents its input [output] degree, i.e., the number of adjacent vertices in U [Y]. The ℓ-monitorability degree has been built to reflect the availability of at least ℓ sensor output signals for monitoring any state variable. With G_Y the output dependency graph of a CPS, let ℓ be equal to $\min_{x \in X} \deg(x)$. A CPS has ℓ-monitorability degree (ℓ-monitorability for short). The notion of k-steerability has also been introduced, indicating the availability of at least k actuator input signals for acting on every plant state variable. With G_U the input dependency graph of a CPS, let k be equal to $\min_{x \in X} \deg(x)$. The CPS is said to have k-steerability degree (k-steerability for short). A CPS which has k-steerability and ℓ-monitorability (with $k \leq \ell$) is said to be (k, ℓ)-resilient.

The (k, ℓ)-resiliency metric calculates a CPS's resilience according to the ability to steer and monitor the plant and make it return to its original state when an attack occurs. Furthermore, it has been admitted that the relationship $k \leq \ell$ is desirable. Indeed, a higher steerability does not increase the resiliency potential without sufficient monitorability capacities.

The next section explores a quantitative assessment method of the SWaT's resilience.

IV. Quantitative resilience assessment of SWaT

As a use-case, we choose the SWaT CPS [13], [27]. It is a testbed built by the Singapore University of Technology and Design (SUTD). This system reproduces on a small scale the behavior of a real treatment station in Singapore. The

criticality and need to protect such a system from cyber adversaries are topical.

Figure 1 presents the architectures of the water treatment system that we consider, inspired by SWaT [17]. It comprises six phases. The first phase pumps water from a reservoir to the second phase. This second phase includes a chemical dosing station and a mixer to purify the water from bacteria and contaminants. In the third phase, an Ultrafiltration (UF) membrane removes micro-particles from the water. Then, the water is sent at high pressure to an ultraviolet system in the dichlorination fourth phase. A reverse osmosis system eliminates residual viruses and impurities during the fifth phase. The water arriving in the sixth phase is stored in two different tanks. One of them is intended for water distribution. The second one is connected to a pump that sends a backwash flow of water to clean the UF membrane of the third phase. Indeed, a cake layer formation occurs at the UF membrane during the water treatment process. Thus, the system alternates between filtration and backwash cycles. Depending on the concentration of particles in the water, a filtration cycle could last between thirty minutes to one hour. A backwash cycle generally lasts for less than one minute. A manual-cleaning step of the UF membrane is required every twelve hours. A water treatment system such as SWaT is controlled by a Human Machine Interface-Supervisory Control and Data Acquisition (HMI-Scada) system. There is one controller for each phase. The controllers communicate with each other for global management purposes.

A. Spectral radius evaluation

The spectral radius calculation approach considers a whole system as a graph structure. The nodes of the graph represent the different components. Links between the nodes represent physical connections, e.g., a level sensor attached to a tank, and logical relationships corresponding to communication links between components, such as the flow rate readings from a sensor to a controller. We present several architectural variations of SWaT. Each of them incorporates n different components contributing to the steerability and monitorability capabilities of the system. We compute the spectral radius for every architecture by mapping it to a $n \times n$ connection matrix C. Element $c_{i,j}$ equals k when there is a logical or physical link between the components i and j, where k is called the functional diversity, i.e., the number of links between i and j. We notice that the numerical values presented at the level of the steerability and monitorability components of the architectures shown in Fig 1 highlight a functional diversity that increases the resiliency potential of the system. Thus, we have identified a family of steerability components composed of controllers, pumps, and valves. Monitorability components are the sensors.

Architecture A_0, presented in Fig. 1 is similar to the original SWaT. In A_1, two controllers are working in tandem in each phase. When one of them fails, the water treatment process is not impacted. The redundant controllers increase the steerability of the system [4], [5]. Architecture A_2 is based on A_0. It integrates new components such as sensors and redundant pumps to achieve a higher degree of monitorability [4], [5]. Thus, from an illustrative point of view, A_2 does not differ from A_0. However, it has higher numerical degrees on its links, i.e., greater functional diversity. Architecture A_{12} is based on A_0. It also comprises the new elements added into the architectures A_1 and A_2. The calculation of the spectral radius for each architecture has been obtained by modeling each architecture as square-adjacent matrices with degrees representing the functional diversity. Then, we built a MATLAB script to read the matrices from an Excel file and compute the spectral radius of each architecture. These results are presented in Table I, and the associated files can be found at [1]. The numerical results are consistent with an expected increase in the degree of resilience, $\rho(A_1)$ and $\rho(A_2)$ both greater than $\rho(A_0)$, and $\rho(A_{12})$ greater than both $\rho(A_1)$ and $\rho(A_2)$.

B. (k, ℓ)-resilience evaluation

The k-steerability means that at least distinct k actuators are available for affecting every single plant state variable [4], [5]. Besides, the ℓ-monitorability signifies the readings from at least ℓ different sensors observe any state variable. The water levels in the tanks are the state variables of the SWaT-derived architectures. Nine tanks (water and chemical) are distributed in the six phases. A maximum of two pumps are connected to each tank, and only one is in Phase 6. There is a total of fifty-one sensors of all kinds in the system. Thus, a state-space representation of SWaT is as follows: A state matrix A of size (9×9); an input matrix B of size (9×25); an output matrix C of size (51×9); and a direct transition matrix D of size (51×25). We know that the parameters k and ℓ are the minimum degrees of adjacent vertices in the input and output dependency graphs, which are bipartite. Thus, in the A_0 architecture, we have (i) 1-steerability capacity on all the tanks (since there is only one pump on one of the tanks in Phase 6); (ii) at least 5-monitorability capacity on all the tanks (the sensors of SWaT have a high correlation with the water levels of the tanks since the physical phenomena they measure are related to water and the dosage of all kinds of chemicals depends on a specific quantity of water). Hence, A_0 is $(1, 5)$-resilient. By applying the same reasoning to the other architectures, we get the results provided in the third column of Table I. A_1 has redundant controllers, which doubled the k value. A_2 has a k and ℓ parameters greater than A_1 because new sensors have been incorporated in the whole system (which increments the ℓ).

TABLE I
SPECTRAL RADIUS AND (k, ℓ)-RESILIENCE EVALUATION.

Architecture (A)	***Spectral radius*** $(\rho(A))$	***(k,ℓ)-resilience***
A_0	10.91	$(1, 5)$
A_1	15.95	$(2, 5)$
A_2	14.85	$(2, 9)$
A_{12}	21.41	$(4, 9)$

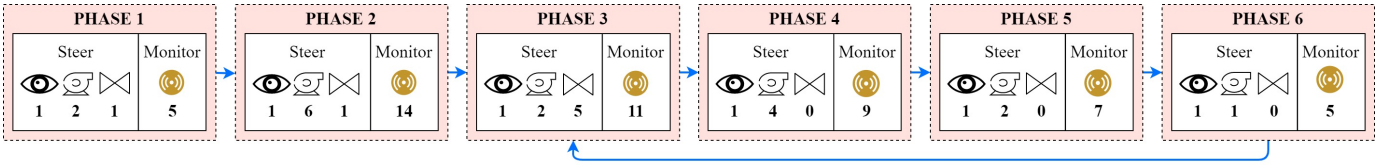

(a) Architecture A_0

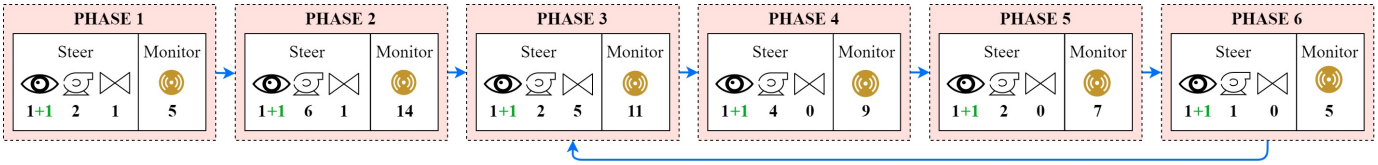

(b) Architecture A_1

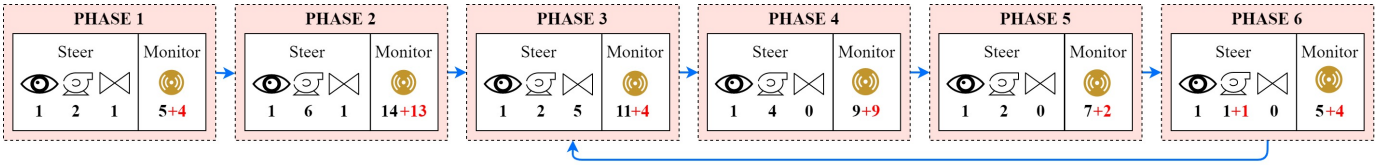

(c) Architecture A_2

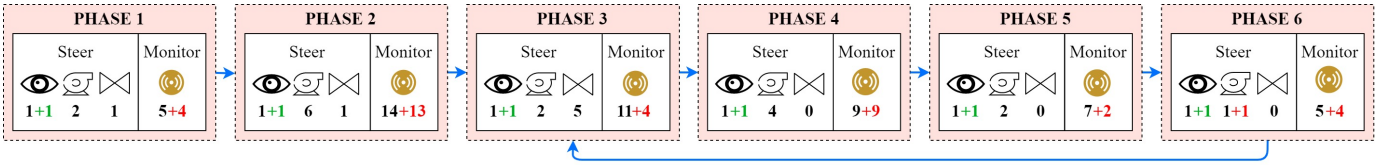

(d) Architecture A_{12}

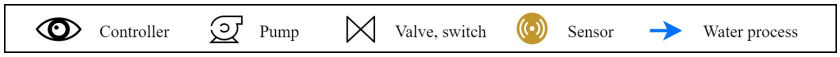

(e) Legends

Fig. 1. A_0 presents the original architecture of SWaT [17]. A_1 is based on A_0 and it incorporates redundant controllers green-marked. A_2 is similar to A_0 and incorporates new sensors and one red-marked redundant pump. A_{12} incorporates the additional elements of A_1 and A_2.

Also, a redundant pump has been added in the sixth phase (which increments the k parameter). Table I provides two different interpretations of the architectures' resilience. The spectral radius emphasizes a resilience estimation based on the functional structure of a CPS. In contrast, the (k, ℓ)-resilience metric focuses on hardware and software components that make a system resilient by design.

V. Discussion

We are searching for methodologies for the quantitative assessment of the resilience of a CPS. As a case study, we use architectural variants of the SWaT CPS. While small, it comprises several aspects of a full-scale CPS, making it an ideal candidate for exploratory studies. For every architecture, we calculated the spectral radius. The results are consistent with an expected increase in resilience. We also evaluated the (k, ℓ)-resilience, with somewhat consistent results for the architectures A_0, A_1, A_2, and A_{12}.

We have identified several research axes that must be investigated in the quantitative resilience analysis of CPSs. Firstly, it is essential to notice that the resilience capabilities of a specific CPS' architecture are limited. In the case of resilience-by-design strategies, adding new components such as controllers or sensors can bring more steerability and monitorability, which are resilience enhancement as shown in the works [4], [5]. However, it has also been proven that adding components to a CPS increases the attack surface. Thus, a delicate balance between risk mitigation and enhancing resiliency capacities must be achieved. Another approach to consider in resilience enhancement is in relation to the different operating layers of a CPS. A CPS can be viewed as a stack of layers. Each layer communicates with the others via different mechanisms. The physical layer represents the plant itself. The mission layer consists of a functional view of the system's mission. The business layer can be viewed as an interconnection between the stakeholders involved in the system's operating cycle. Resilience enhancement strategies must be exercised at each layer, but with no detrimental effects to the other ones. Extending a physical plant with a digital twin is also a topic to explore for resilience enhancement purposes.

One may also find inspiration in automated synthesis approaches that use symbolic control refinement to guarantee

that desired properties are obtained, such as safety [12]. Top-down formal refinement of CPS configurations could also model Boolean satisfiability of loss scenarios. From the high-level description of CPS configurations (e.g., those abstracting the details of the concrete system), the automated refinement could interface low-level descriptions satisfying the absence of loss scenarios, assuring that both high- and low-level descriptions behave identically (low-level configurations refined automatically from high-level verifiable descriptions with good resilience properties). The approach could combine property-proving techniques, at the high layers, with model-checking techniques, at the low layers [28].

VI. Conclusion

Resilience assessment of critical infrastructures is challenging because of their complexity. Many modeling strategies reason on virtual representations, making obtaining a complete overview of the elements used in resiliency-quantification strategies difficult. For example, the (k, ℓ)-resilience and spectral radius metrics return different interpretations of the architectures of SWaT. This is due to the elements the metrics are considering. Indeed, the (k, ℓ)-resilience metric focuses on design configurations. In contrast, the spectral radius metric builds upon a graphical representation of the system capturing node and state distribution. The system itself is not the only element to consider in resilience quantification and enhancement strategies. Its environment and related interactions must also be examined.

References

[1] Exploring SWaT's resilience. https://github.com/IRT-SystemX/exploring_resilience_SWaT. GitHub repository, created: 2023-03-21.

[2] S. Abraham and S. Nair. A predictive framework for cyber security analytics using attack graphs. *arXiv preprint arXiv:1502.01240*, 2015.

[3] F.-X. Aguessy, L. Gaspard, O. Bettan, and V. Conan. Remediating Logical Attack Paths Using Information System Simulated Topologies. In *C&ESAR 2014*, pages 187–205, Rennes, France, 2014.

[4] M. Barbeau, F. Cuppens, N. Cuppens, R. Dagnas, and J. Garcia-Alfaro. Metrics to enhance the resilience of cyber-physical systems. In *19th IEEEInternational Conference on Trust, Security and Privacy in Computing and Communications*, pages 1167–1172. IEEE, 2020.

[5] M. Barbeau, F. Cuppens, N. Cuppens, R. Dagnas, and J. Garcia-Alfaro. Resilience estimation of cyber-physical systems via quantitative metrics. *IEEE Access*, 9:46462–46475, 2021.

[6] D. J. Bodeau, R. Graubart, J. Picciotto, and R. McQuaid. Cyber resiliency engineering framework. Technical report, MITRE CORP BEDFORD MA, 2011.

[7] D. J. Bodeau, R. D. Graubart, R. M. McQuaid, and J. Woodill. Cyber resiliency metrics catalog. Technical report, MITRE CORP BEDFORD MA, 2018.

[8] T. Clédel, N. Cuppens, F. Cuppens, and R. Dagnas. Resilience properties and metrics: how far have we gone? *Journal of Surveillance, Security and Safety*, 1(2):119–139, 2020.

[9] C. Disasters, C. Policy, and N. Academies. *Disaster Resilience: A National Imperative*. 12 2012.

[10] D. Dudorov, D. Stupples, and M. Newby. Probability analysis of cyber attack paths against business and commercial enterprise systems. In *2013 European Intelligence and Security Informatics Conference*, pages 38–44. IEEE, 2013.

[11] H. Gill. Nsf perspective and status on cyber-physical systems: National workshop on cyber-physical systems austin, tx october 16-17. 2006.

[12] A. Girard and A. Eqtami. Least-violating symbolic controller synthesis for safety, reachability and attractivity specifications. *Automatica*, 127:109543, 2021.

[13] J. Goh, S. Adepu, K. N. Junejo, and A. Mathur. A dataset to support research in the design of secure water treatment systems. In *Critical Information Infrastructures Security: 11th International Conference, CRITIS 2016, Paris, France, October 10–12, 2016, Revised Selected Papers 11*, pages 88–99. Springer, 2017.

[14] G. Gonzalez-Granadillo, S. Dubus, A. Motzek, J. Garcia-Alfaro, E. Alvarez, M. Merialdo, S. Papillon, and H. Debar. Dynamic risk management response system to handle cyber threats. *Future Generation Computer Systems*, 83:535–552, 2018.

[15] C. S. Holling. Resilience and stability of ecological systems. *Annual Review of Ecology, Evolution, and Systematics*, 4:1–23, 1973.

[16] S. Hosseini, K. Barker, and J. E. Ramirez-Marquez. A review of definitions and measures of system resilience. *Reliability Engineering & System Safety*, 145:47–61, 2016.

[17] iTrust (Center for Research in Cyber Security). Secure Water Treatment (SWaT Testbed). Technical report, SUTD (Singapore University of Technology and Design), July 2021. Version 4.4.

[18] S. Kardon. Florida water treatment plant hit with cyber attack. https://www.industrialdefender.com/blog/florida-water-treatment-plant-cyber-attack, Feb 2021.

[19] B. Kordy, S. Mauw, S. Radomirović, and P. Schweitzer. Foundations of attack–defense trees. In *Formal Aspects of Security and Trust: 7th International Workshop, FAST 2010, Pisa, Italy, September 16-17, 2010. Revised Selected Papers 7*, pages 80–95. Springer, 2011.

[20] I. Kotenko and A. Chechulin. A cyber attack modeling and impact assessment framework. In *2013 5th International Conference on Cyber Conflict (CYCON 2013)*, pages 1–24. IEEE, 2013.

[21] I. Kotenko, I. Saenko, and O. Lauta. Modeling the impact of cyber attacks. In A. Kott and I. Linkov, editors, *Cyber Resilience of Systems and Networks*, chapter 7, pages 135–169. Springer, Switzerland, 2018.

[22] I. Kotenko and M. Stepashkin. Attack graph based evaluation of network security. *Lecture Notes in Computer Science*, 4237:216–227, 2006.

[23] A. Kott and I. Linkov. *Cyber Resilience of Systems and Networks*. Springer Publishing Company, Incorporated, 1st edition, 2018.

[24] A. Kott and I. Linkov. To improve cyber resilience, measure it. *Computer*, 54(2):80–85, 2021.

[25] N. G. Leveson. *Engineering a Safer World: Systems Thinking Applied to Safety*. The MIT Press, 01 2012.

[26] T. G. Lewis. The many faces of resilience. *Communications of the ACM*, 66(1):56–61, 2023.

[27] A. P. Mathur and N. O. Tippenhauer. Swat: a water treatment testbed for research and training on ics security. In *2016 International Workshop on Cyber-physical Systems for Smart Water Networks (CySWater)*, pages 31–36, 2016.

[28] O. A. Mohamed et al. LCF-style Platform based on Multiway Decision Graphs. *Electronic Notes in Theoretical Computer Science*, 246:3–26, 2009.

[29] A. Motzek, G. Gonzalez-Granadillo, H. Debar, J. Garcia-Alfaro, and R. Möller. Selection of pareto-efficient response plans based on financial and operational assessments. *EURASIP Journal on Information Security*, 2017(1):1–22, 2017.

[30] N. Poolsappasit, R. Dewri, and I. Ray. Dynamic security risk management using bayesian attack graphs. *IEEE Transactions on Dependable and Secure Computing*, 9(1):61–74, 2012.

[31] R. S. Ross, V. Y. Pillitteri, R. Graubart, D. Bodeau, and R. McQuaid. *Developing Cyber-Resilient Systems: A Systems Security Engineering Approach*. Special Publication (NIST SP), National Institute of Standards and Technology, Gaithersburg, MD, 12 2021.

[32] B. Schneier. Modelling security threats. *Dr. Dobb's Journal*, 1999.

[33] R. Singh, S. Hutton, M. Donahoo, and D. Sicker. Toward grading cybersecurity & resilience posture for cyber physical systems. *SSRN Electronic Journal*, 01 2021.

[34] S. Tutorials. Pearson Correlation. Retrieved on April 2020. https://libguides.library.kent.edu/SPSS/PearsonCorr, 2014.