Chapter 11

Vulnerability and Incident Propagation in Cyber-physical Systems

By Faten Atigui, Fayçal Hamdi, Nadira Lammari and Samira Si-said Cherfi

DOI: 10.1561/9781680836875.ch11

Published in *Cyber-Physical Threat Intelligence for Critical Infrastructures Security: A Guide to Integrated Cyber-Physical Protection of Modern Critical Infrastructures* by John Soldatos, James Philpot and Gabriele Giunta (eds.). 2020. ISBN 978-1-68083-686-8. E-ISBN 978-1-68083-687-5.

Suggested citation: Faten Atigui *et al.* 2020. "Vulnerability and Incident Propagation in Cyber-physical Systems" in *Cyber-Physical Threat Intelligence for Critical Infrastructures Security: A Guide to Integrated Cyber-Physical Protection of Modern Critical Infrastructures*. Edited by John Soldatos, James Philpot and Gabriele Giunta. pp. 193–205. Now Publishers. DOI: 10.1561/9781680836875.ch11.

now
the essence of knowledge

11.1 Introduction

Hospitals are cyber-physical systems that are vulnerable by nature to a multitude of attacks that can occur at their communication, networking, and physical entry points. Such cyber-physical attacks can have detrimental effects on their operation and the safety of their patients. Thus, to properly secure these systems, it is of utmost importance to: (i) understand their underlying assets with related vulnerabilities and associated threats, (ii) quantify their effects, and (iii) prevent the potential impacts of these attacks. This implies addressing a challenging objective of understanding the tight relationships between the asset's characteristics and the propagation of attack's effects to better prevent the impacts and consequences of incidents. Such an approach needs a detailed knowledge of intrinsic and contextual assets properties. However, hospitals host a variety of medical and IT assets with very different characteristics. The next section reports on the state of the art of assets and assets interdependencies modeling as well as on incidents propagation approaches.

11.2 Related Work

This section presents existing work on impact propagation of incidents and the methods used to assess the severity of incidents and risks.

11.2.1 Characterization of Dependencies Between and Within Critical Infrastructures

Although the terms “dependency” and “inter-dependency” are commonly used interchangeably, some research work distinguish them. The consensual distinction is this of Rinaldi *et al.* (2001). The authors define a dependency as a relationship between two infrastructures in a single direction, whereas inter-dependency is bidirectional (implicitly multi-directional) with two (implicitly more) infrastructures influencing each other. This definition is also shared by Stapelberg (2008). A more precise definition of the dependency concept is given by Schmitz *et al.* (2007). The European Union Agency for Cybersecurity (ENISA) proposes to consider dependencies within critical infrastructures (CIs) and dependencies between CIs. These kinds of dependencies are qualified as upstream, internal, or downstream dependencies in Petit *et al.* (2015). An upstream dependency expresses the fact that the products or services provided to one infrastructure by another external infrastructure are necessary to support its operations and functions. Downstream dependencies are the consequences to a critical infrastructure’s consumers or recipients from the degradation of the resources provided by a critical infrastructure. Internal dependencies represent the internal links among the assets constituting a critical infrastructure. Therefore, upstream and downstream dependencies are between CIs, whereas internal ones are within CIs. Several works have focused on the characterization of dependencies between CIs. Zimmerman (2008) distinguishes spatial dependencies from functional ones. Rinaldi *et al.* (2001) and Schmitz *et al.* (2007) propose a categorization of dependencies into physical, cyber, geographic, and logical ones. Dudenhoeffer *et al.* (2006) and Clemente (2013) consider physical, informational, geo-spatial, policy/procedural, and societal dependency. For reasoning purposes, Adetoye *et al.* (2011) propose another taxonomy of dependencies. They suggest considering five types of dependencies: generic, indirect, inter, co, and redundant dependency.

11.2.2 Models Serving the Incidents’ Impact Propagation

In addition to the existing inter-dependencies between infrastructures, to deal with the impact of cascading effects that a disruption of an asset may have on the internal and external context of a critical infrastructure, one must also have, for each asset, a clear knowledge of the kinds of threats that could affect this asset, its vulnerabilities, and its relation to the other assets. These three aspects have been the subject of several research studies. From a threat perspective, to our knowledge, there is no research work that provides a high-level ontology of threats for CI. However, let us note that the European Commission reported a generic classification of threats for CI in which natural hazards are distinguished from non-malicious man-made hazards and malicious man-made hazards (Theocharidou and Giannopoulos, 2015). The HITRUST alliance have also published a threat taxonomy where at the top level logical, physical, and organizational threats are distinguished (HITRUST, 2019). Other works concentrate on specific threats. In ENISA (2016a) the most common threats affecting ICS/SCADA systems are shown. The top 10 threats affecting these systems have been published by CTED and UNOCT (2018). In the context of physical security risk assessments, Liu *et al.* (2012) propose a list of threats from terrorism. We can also find in "Common Criteria" and ANSSI portals security protection profiles for some software and physical equipment of CI where threats affecting these components are listed. In the context of the healthcare sector, ENISA (2016b) provided an overview of the cyber threats faced by smart hospitals. Taxonomies of threats for healthcare infrastructures are also proposed by Almohri *et al.* (2017) and Agrafiotis *et al.* (2018). Regarding the links between assets, we can consider research works that give much attention to the hierarchical links between assets (Silva and Jacob, 2018; Brocke *et al.*, 2014; Jakobson, 2011; Tong and Ban, 2014; Breier and Schindler, 2014). They model an infrastructure into levels to which the assets belong. The contributions differ in terms of kind and number of layers. The representation models used are also varied, ranging from simple oriented graphs to light ontologies. To define models that consider the hierarchical dependency between assets while emphasizing the links between assets within the hierarchical layers, one can rely on Enterprise Architecture (EA) modeling languages and standards or methodological guides existing in the industrial world. These tools are not specifically dedicated to critical infrastructures. As an example, we can mention ArchiMate 2.1, an open and independent EA modeling language within TOGAF Framework 9.2. We can also mention the CIM standard produced by DMTF (formerly known as the Distributed Management Task Force) that is internationally recognized by ANSI (American National Standards Institute) and ISO (International Organization for Standardization). There also exists several security risk analysis methodologies that give descriptions of assets, most of which are based on standards. These descriptions are very often informal and sometimes accompanied by catalogues. This is the case of EBIOS RM (EBIOS, 2019) and MAGERIT 3.0 (Amutio *et al.*, 2014) methodologies.

11.2.3 Incidents Propagation

Several approaches have dealt with the incidents propagation issue. We can classify these approaches into three categories: empirical, agent-based, and network-based approaches.

Empirical approaches analyze asset's interdependencies according to experts' opinions and past incidents traces. The underlying assumption is that it is difficult to identify assets' interdependencies in normal situations. Thus, analyzing the incidents helps rising intangible relationships among assets under extreme situations such as disasters, failures, or attacks. Laefer *et al.* (2006) defined accuracy, comprehensibility, timeliness, and accessibility of data as key characteristics to store, analyze, query, and visualize critical incident. This data could then be analyzed to mine records of frequent failure patterns as presented in Chou and Tseng (2010). To highlight the relationship between interdependencies and incident propagation, Mendonça and Wallace (2006) studied the 9/11 World Trade Center attacks and their impact on critical infrastructures and their services. The study showed that 20% of reported disruptions involved interdependency. Considering CI as "systems of systems" may improve response to incidents. Kotzanikolaou *et al.* (2013) combine common-cause and cascading events to assess the potential risk caused by complex situations. They considered the cumulative dependency risk of cascading chains.

Agent-based approaches consider a CIS as a complex adaptive system that could be analyzed as a complex phenomenon emerging from individual and autonomous agents. This kind of approaches captures all types of interdependencies among CIS by event simulations. It also provides scenario-based what-if analysis and the effectiveness assessment of different control strategies. Barrett *et al.* (2010) investigated cascading effects in three closely coupled systems: cellular networks, transportation networks and phone call networks. They studied the interaction between these systems and the challenges raised by their co-evolution and reaction to incidents. Gómez *et al.* (2014) proposed a method for clustering a network into agents called decision units. This method deals with the complexity by exploring relationships between agents' local decisions and their impact at the global level.

Network-based approaches represent the connected infrastructures interdependencies as a graph to show paths for incidents propagation. Shah and Babiceanu (2015) propose to evaluate the resilience of a system under attacks. The infrastructures are modeled using networks of interdependent processes. Based on this model, the authors provide simulations to predict the network behavior to face different attacks.

11.3 Impact Propagation and Decision Support Model

This part describes the impact propagation model and decision support model solution that includes the specification of the IPM ontology (SafecareOnto) and the IPM rules.

11.3.1 SAFECARE Ontology

The Safecare ontology, called SafecareOnto, describes both cyber and physical assets, their vulnerabilities, and their interdependence, as well as the risks and threats. It is the cornerstone of the knowledge graph used by the Impact Propagation and Decision Support Model module to infer the propagation of impacts over cyber and physical assets. In the following sections, we will describe the construction process of this ontology and its modular structure.

Overview of the ontology building process

For the determination of the approach to build SafecareOnto described in Figure 11.1, we have been inspired by NeOn methodological framework Suárez-Figueroa *et al.* (2012).

In the first phase, we provided information about the scope of the ontology (its purpose, the language to be used during its implementation, the target users for which it is intended, its requirements expressed under competency questions).

In the second phase, we started by studying the available resources (ontological and non-ontological) favoring the elaboration of SafecareOnto. The lack of ontological resources that perfectly meet our requirements led us to choose the option of building a first draft of the ontology from portions of non-ontological resources through an abstraction process. The objective is to identify a core of

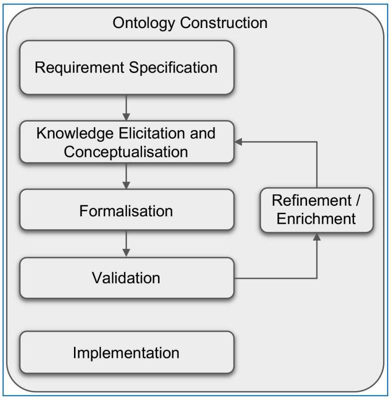

Figure 11.1. Construction process of SafecareOnto.

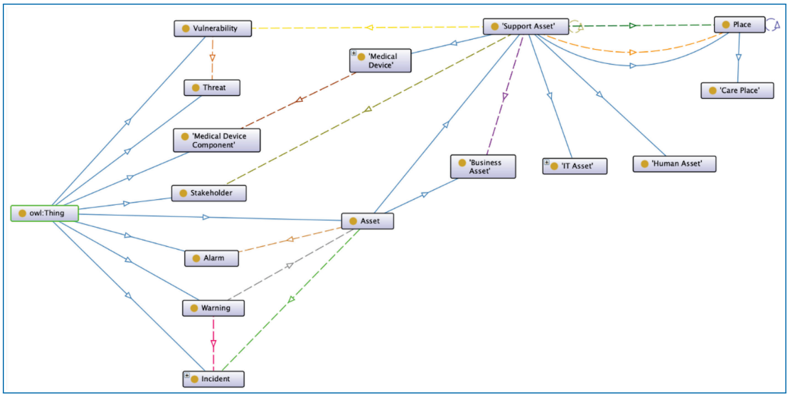

Figure 11.2. Excerpt of the SafecareOnto.

basic concepts and relationships that must be part of our ontology. As an example of non-ontological resources, we can mention the description of EBIOS RM methodology (EBIOS, 2019) and the description of medical devices of the MITRE (Connolly *et al.*, 2019). The conceptualization activity consisted of summarizing, organizing, and structuring the required knowledge into a meaningful model. In our case, for representing knowledge modeling, we opted for the UML class diagram. The benefits of such a model for ontology conceptualization have been acknowledged in several studies. One of its main advantages is that it is widely used. Furthermore, users are likely to be more familiar with a class diagram representation of the ontology (since it is a semi-formal model) than with OWL which representation is purely textual. Thus, it is more relevant for the verification of the ontology scope.

The resulting conceptual model (the first draft of SafecareOnto) has been translated, during the formalization phase, into a formal model using OWL2 that offers a highly expressive language and inference capabilities. Figure 11.2 represents an excerpt of the SafecareOnto.

The last phase consists of evaluating SafecareOnto regarding the ability of the impact propagation module to deal with the threat scenarios defined in the SAFECARE project. The validation step will lead to a refinement and enrichment of the ontology.

SafecareOnto, a modular ontology

The impact propagation and decision support model relies on both structural information about the assets and their intrinsic properties and structural relationship and on knowledge about the incidents that they suffered from. It also holds knowledge about how to infer and propagate impacts. This second knowledge evolves continuously and is more dynamic than the structural knowledge. For example,

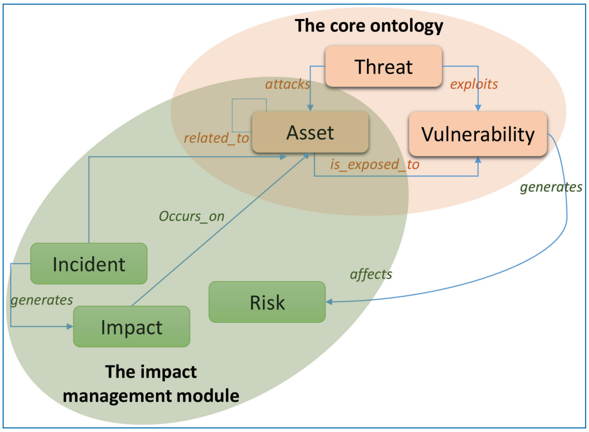

Figure 11.3. The modular structure of SafecareOnto.

the software of a medical asset could be updated to correct a known vulnerability. This kind of operation is less dynamic and more predictable that the occurrence of incidents.

To cope with the static and dynamic knowledge and to confer more stability to the IPM module, we have adopted a modular vision of the ontology. At a high level of abstraction, we could view the whole picture as depicted in Figure 11.3.

The core ontology captures essentially the static and is centered essentially on three concepts that are Asset, Vulnerability, and Threat.

An asset is any "thing" that has value. Within the SAFECARE projects assets could be business assets such as personal data about patients and personnel or the patients themselves or support assets such as medical or IT devices or medical staff. Assets are related to other assets through several kinds of relationships. A vulnerability is any weakness of an asset that could be used to generate a threat. A vulnerability assesses the protection of an asset against attacks. A threat could be accidental or malicious. As an example for "a radiology room" could have as vulnerability "likely to be subject to unauthorised access" and a "patient report" could have as vulnerability "lack of encryption."

A threat is the operationalization or a materialization of a vulnerability. An asset could be exposed to several vulnerabilities that are known or that could emerge after incident occurrence. The information about vulnerabilities is updated consequently to regular maintenance operations or after incident analysis. "Unauthorized access" or "personal data disclosure" are examples of threats. The more we know about the threats that relate to an asset, the more efficient its protection can be and the better we can react when incidents occur. These basic concepts are further refined and characterized. An excerpt is formalized in the next section. This formalization is done in such a way that it can easily be extended to meet emerging requirements.

The impact management module is an extension to the core ontology that relies on the previous concepts. It allows defining the concepts that are essential to the computation of impact propagation and provide indicators to help decide about the suitable countermeasures to face attacks consequences. It relies on concepts such as Incident, Risk, and Impact.

An incident, according to NIST (Stouffer *et al.*, 2011), is "an occurrence that actually or potentially jeopardizes the confidentiality, integrity, or availability of a system or the information the system processes, stores, or transmits or that constitutes a violation or imminent threat of violation of security policies, security procedures, or acceptable use policies." An incident could be an attack against one or several assets by exploiting vulnerabilities. In SAFECARE, we handle both physical and cyber incidents. We also have to assess the severity of an incident to better compute its propagation. An incident could be the expression of a known risk or completely unexpected. Indeed, a risk is the probability that a threat will exploit a vulnerability. When an incident occurs, it is likely to have impacts on assets. An impact needs to be qualified and/or quantified to efficiently help decide about the mitigation plans.

11.3.2 IPM Rules Specification

There are several approaches for impact propagation management such as agent-based and graph-based approaches that are mainly structure oriented. However, from our investigations, it appears that an added value that the project may produce is to combine cyber and physical incidents and to take into account the variety of interdependencies to provide a semantic oriented approach based on semantic web technologies. A first solution is consequently based on the exploitation of the ontologies' expressiveness expanded by the usage of inference rules. Indeed, the idea of the IPM module is to use axioms describing the concept and properties of SafecareOnto as well as a set of rules to deal with different threat scenarios. The creation of these rules follows the steps below (cf. Figure 11.4):

- *Knowledge elicitation*: in this phase, threat scenarios are analyzed and discussed with domain experts to identify, on the one hand, all the assets that could be impacted in each scenario, and on the other hand, the relationships between assets that lead to the propagation of impacts. Moreover, all the situations of a given scenario are analyzed to see if it is possible to generalize common parts. The objective is to avoid redundant rules.
- *Formalization*: in this phase, the concepts and properties of SafecareOnto that can be used to write rules are identified. A rule-engine (e.g., SWRL, JENA) is then used to implement these rules in the form of premises and conclusions.

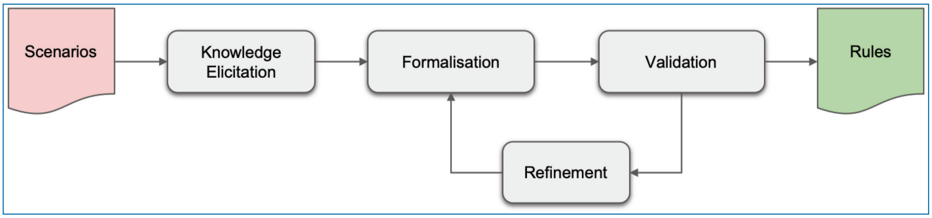

Figure 11.4. IPM rules construction process.

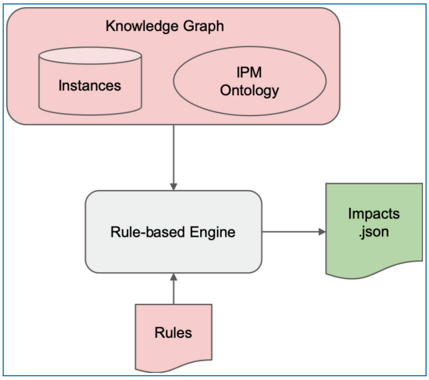

Figure 11.5. Architecture of the IPM prototype.

As existing rule-engine are often equipped with semantic reasoners, the implemented rules can be applied to automatically infer impact propagation.

- *Validation and refining*: in this phase, implemented rules are tested on different scenarios, and inferred impacts on different assets are evaluated by domain experts. At the end of the validation, IPM rules could be refined to better meet the expected results.

A first version of a prototype that simulates impacts propagation was implemented on a near-real scenario (cf. Figure 11.5). Based on the knowledge graph and on IPM rules, a reasoner is used to infer impacts propagation on assets. In this prototype, the IPM rules were expressed in terms of OWL concepts (classes, properties, individuals) using the JENA rule engine. Each rule is composed of a list of body terms (premises), a list of head terms (conclusions).

The following example presents a JENA rule that propagates warnings in case of assets located in the same places:

```
(?asset ipm:hasLocation ?place), (?warning ipm:attachedTo ?place),
(?warning ipm:hasCause ?incident), makeSkolem(?new_warning, ?warning) =>
 (?new_warning rdf:type isid:Warning), (?new_warning ipm:hasCause ?incident),
 (?new_warning ipm:attachedTo ?asset), (?asset ipm:hasWarning ?new_warning)]
```

The premise of this rule instantiates all the assets having a place, the warnings triggered in this place and the incidents causing these warnings. The conclusion attaches warnings to all assets located in the same place. An application of this rule may be a fire detection incident in a server room that could affect all the materials inside this room.

11.4 Conclusion

This chapter presents a focused view on how to handle incidents and their propagation from an assets point of view in a healthcare environment. It presents an overview of work conducted within the Safecare project. The state of the art shows that dealing with incidents and their propagation requires a detailed knowledge on assets, their context and an as precise as possible vision of the historical data about the assets, their real time state and the incidents that impacted them. From our experience within the Safecare project, it appears that collecting such data is not an easy task. It requires an additional effort from health actors, whose priority is care, although they are aware that safety is also a major issue. Consequently, we could not adopt one of the existing approaches for incidents propagation as they rely on either detailed traces, in case of empirical approaches, or a quasi-complete structure knowledge of systems as required by network-based approaches. The proposed solution is semantics based. It relies on an evolving knowledge captured by a modular ontology. The propagation is managed by rules that exploit assets states, incidents, and domain knowledge which all evolve continuously.

Acknowledgments

This work is part of SAFECARE project. This project has received funding as part of the "Secure societies – Protecting freedom and security of Europe and its citizens", challenge of the Horizon 2020 Research and Innovation programme of the European Union, under grant agreement 787002.

References

Adetoye, A. O., M. Goldsmith, and S. Creese. 2011. "Analysis of dependencies in critical infrastructures." In: *International Workshop on Critical Information Infrastructures Security*. pp. 18–29.

Agrafiotis, I., J. R. Nurse, M. Goldsmith, S. Creese, and D. Upton. 2018. "A taxonomy of cyber-harms: Defining the impacts of cyber-attacks and understanding how they propagate." *Journal of Cybersecurity*. 4(1): tyy006.

Almohri, H., L. Cheng, D. Yao, and H. Alemzadeh. 2017. "On threat modeling and mitigation of medical cyber-physical systems." In: *2017 IEEE/ACM International Conference on Connected Health: Applications, Systems and Engineering Technologies (CHASE)*. IEEE. pp. 114–119.

Amutio, M., J. Candau, and J. Mañas. 2014. "Magerit-version 3, methodology for information systems risk analysis and management, book I-the method." *Ministerio de Administraciones Públicas*.

Barrett, C., R. Beckman, K. Channakeshava, F. Huang, V. A. Kumar, A. Marathe, M. V. Marathe, and G. Pei. 2010. "Cascading failures in multiple infrastructures: From transportation to communication network." In: *2010 5th International Conference on Critical Infrastructure (CRIS)*. IEEE. 1–8.

Breier, J. and F. Schindler. 2014. "Assets dependencies model in information security risk management." In: *Information and Communication Technology-EurAsia Conference*. Springer. 405–412.

Brocke, J. vom, A. M. Braccini, C. Sonnenberg, and P. Spagnoletti. 2014. "Living IT infrastructures—an ontology-based approach to aligning IT infrastructure capacity and business needs." *International Journal of Accounting Information Systems*. 15(3): 246–274.

Chou, C.-C. and S.-M. Tseng. 2010. "Collection and analysis of critical infrastructure interdependency relationships." *Journal of Computing in Civil Engineering*. 24(6): 539–547.

Clemente, D. 2013. *Cyber Security and Global Interdependence: What is Critical?* Chatham House, Royal Institute of International Affairs.

Connolly, J., S. Christey, R. Daldos, M. Zuk, and M. Chase. 2019. "Medical Device Cybersecurity Regional Incident Preparedness and Response Playbook. October 2018." MITRE.

CTED and UNOCT. 2018. "The protection of critical infrastructures against terrorist attacks: compendium of good practices." United Nations.

Dudenhoeffer, D. D., M. R. Permann, and M. Manic. 2006. "CIMS: A framework for infrastructure interdependency modeling and analysis." In: *Proceedings of the 38th conference on Winter simulation*. Winter Simulation Conference. 478–485.

EBIOS. 2019. "EBIOS Risk Manager—The method." The French National Cybersecurity Agency (ANSSI).

ENISA. 2016a. "Communication network dependencies for ICS/SCADA Systems." European Union Agency For Network and Information Security.

ENISA. 2016b. "Cyber security and resilience for Smart Hospitals." European Union Agency For Network and Information Security.

Gómez, C., M. Sánchez-Silva, and L. Dueñas-Osorio. 2014. "An applied complex systems framework for risk-based decision-making in infrastructure engineering." *Structural Safety*. 50: 66–77.

HITRUST. 2019. "The HITRUST Threat Catalogue." Health Information Trust Alliance.

Jakobson, G. 2011. "Mission cyber security situation assessment using impact dependency graphs." In: *14th International Conference on Information Fusion*. IEEE. 1–8.

Kotzanikolaou, P., M. Theoharidou, and D. Gritzalis. 2013. "Cascading effects of common-cause failures in critical infrastructures." In: *International Conference on Critical Infrastructure Protection*. Springer. 171–182.

Laefer, D. F., A. Koss, and A. Pradhan. 2006. "The need for baseline data characteristics for GIS-based disaster management systems." *Journal of Urban Planning and Development*. 132(3): 115–119.

Liu, C., C.-K. Tan, Y.-S. Fang, and T.-S. Lok. 2012. "The security risk assessment methodology." *Procedia Engineering*. 43: 600–609.

Mendonça, D. and W. A. Wallace. 2006. "Impacts of the 2001 world trade center attack on New York City critical infrastructures." *Journal of Infrastructure Systems*. 12(4): 260–270.

Petit, F., D. Verner, D. Brannegan, W. Buehring, D. Dickinson, K. Guziel, R. Haffenden, J. Phillips, and J. Peerenboom. 2015. "Analysis of critical infrastructure dependencies and interdependencies." *Tech. Rep.* Argonne National Lab. (ANL), Argonne, IL (United States).

Rinaldi, S. M., J. P. Peerenboom, and T. K. Kelly. 2001. "Identifying, understanding, and analyzing critical infrastructure interdependencies." *IEEE Control Systems Magazine*. 21(6): 11–25.

Schmitz, W., F. Flentge, H. Dellwing, and C. Schwaegerl. 2007. "The integrated risk reduction of information-based infrastructure systems, interdependency taxonomy and interdependency approaches." *IRRIS Project*. (027568): 82.

Shah, S. S. and R. F. Babiceanu. 2015. "Resilience modeling and analysis of interdependent infrastructure systems." In: *2015 Systems and Information Engineering Design Symposium*. IEEE. 154–158.

Silva, F. and P. Jacob. 2018. "Mission-Centric Risk Assessment to Improve Cyber Situational Awareness." In: *Proceedings of the 13th International Conference on Availability, Reliability and Security*. ACM. 56.

Stapelberg, R. F. 2008. "Infrastructure systems interdependencies and risk informed decision making (RIDM): impact scenario analysis of infrastructure risks induced by natural, technological and intentional hazards." *Journal of Systemics, Cybernetics and Informatics*. 6(5): 21–27.

Stouffer, K., J. Falco, and K. Scarfone. 2011. “Guide to industrial control systems (ICS) security.” *NIST Special Publication*. 800(82): 16–16.

Suárez-Figueroa, M. C., A. Gómez-Pérez, and M. Fernández-López. 2012. “The NeOn methodology for ontology engineering.” In: *Ontology Engineering in a Networked World*. Springer. 9–34.

Theocharidou, M. and G. Giannopoulos. 2015. “Risk assessment methodologies for critical infrastructure protection. Part II: A new approach.” *Tech. Report EUR 27332 EN*.

Tong, X. and X. Ban. 2014. “A hierarchical information system risk evaluation method based on asset dependence chain.” *International Journal of Security and Its Applications*. 8(6): 81–88.

Zimmerman, R. 2008. “Understanding the implications of critical infrastructure interdependencies for water.” *Wiley Handbook of Science and Technology for Homeland Security*. 1–25.