

Synergistic Security for the Industrial Internet of Things: Integrating Redundancy, Diversity, and Hardening

Aron Laszka Waseem Abbas Yevgeniy Vorobeychik Xenofon Koutsoukos
 University of Houston Information Technology University Washington University in Saint Louis Vanderbilt University
 Houston, TX, USA Lahore, Pakistan St. Louis, MO, USA Nashville, TN, USA

Abstract—As the Industrial Internet of Things (IIoT) becomes more prevalent in critical application domains, ensuring security and resilience in the face of cyber-attacks is becoming an issue of paramount importance. Cyber-attacks against critical infrastructures, for example, against smart water-distribution and transportation systems, pose serious threats to public health and safety. Owing to the severity of these threats, a variety of security techniques are available. However, no single technique can address the whole spectrum of cyber-attacks that may be launched by a determined and resourceful attacker. In light of this, we consider a multi-pronged approach for designing secure and resilient IIoT systems, which integrates redundancy, diversity, and hardening techniques. We introduce a framework for quantifying cyber-security risks and optimizing IIoT design by determining security investments in redundancy, diversity, and hardening. To demonstrate the applicability of our framework, we present a case study in water-distribution systems. Our numerical evaluation shows that integrating redundancy, diversity, and hardening can lead to reduced security risk at the same cost.

I. INTRODUCTION

Emerging industrial platforms such as the Industrial Internet (II) in the US and Industrie 4.0 in Europe are creating novel systems that include the devices, systems, networks, and controls used to operate and/or automate Industrial Internet of Things (IIoT) systems. IIoT systems abound in modern society, and it is not surprising that many of these systems are targets for attacks. Critical infrastructure such as water management and transportation systems, in particular, have been growing more connected following recent advances in co-engineered interacting networks of physical and computational components. Due to the tightly coupled nature between the cyber and physical domains, new attack vectors are emerging. Attacks can include physical destruction, network spoofing, malware, data corruption, malicious insiders, and others. Further, the impacts of attacks propagate because of tight interactions. As IIoT systems become more ubiquitous, the risks posed by cyber-attacks becomes severe. The steady increase in the number of reported cyber-incidents evidences how difficult it is in practice to secure such systems against determined attackers.

A variety of techniques have been proposed for providing resilience against cyber-attacks, ranging from hardening techniques (e.g., address-space layout randomization) to increasing

system diversity (e.g., [1]). However, defending complex and large-scale IIoT systems is particularly challenging. These systems often face a variety of threats, have large attack surfaces, and may contain a number of undiscovered vulnerabilities. In light of these factors, it is clear that there is no “silver bullet” technique that could protect a complex system against every kind of attack. Instead of relying on a single technique, defenders must employ multi-pronged solutions, which combine multiple techniques for improving the security and resilience of IIoT. We can divide many of existing techniques into three canonical approaches:

- *Redundancy* for deploying additional redundant components in a system, so that even if some components are compromised or impaired, the system may retain normal (or at least adequate) functionality;
- *Diversity* for implementing components using a diverse set of component types, so that vulnerabilities which are present in only a single type have limited impact on the system; and
- *Hardening* for reinforcing individual components or component types (e.g., tamper-resistant hardware and firewalls), so that they are harder to compromise or impair.

While it is possible to combine these approaches easily by designing and implementing them independently, security and resilience of IIoT systems can be significantly improved by designing and implementing them in an integrated manner. However, a sound framework and methodology for combining techniques from different approaches is lacking. In lieu of a unified framework or methodology, defenders must follow best practices and intuition when integrating techniques, which can result in the deployment of ineffective—or even vulnerable—combinations.

In this paper, we propose a framework for integrating redundancy, diversity, and hardening techniques for designing secure and resilient IIoT systems. The objective is to develop a systematic framework for prioritizing investments for reducing security risk. The contributions of the paper are as follows:

- Establishing a system model that can capture (1) a wide variety of components that are found in IIoT as well as the interactions between them, (2) a security investment model for redundancy, diversity, and hardening, and (3) a security risk model which quantifies the impact of attacks and defense mechanisms (Section II).




Fig. 1. Example cyber-physical system. Arrows represent flows of sensor data and control signals.

- Formulating the resilient IIoT design problem as an optimization problem for prioritizing security investments and showing that the problem is NP-hard (Section III).
- Developing an efficient meta-heuristic design algorithm for finding near-optimal designs in practice (Section III).
- Evaluating the applicability of the approach using two case studies in canonical IIoT domains of water distribution and transportation systems (Sections IV and V and [2]).

We give an overview of related work in Section VI.

II. MODEL

An IIoT system is comprised of a variety of components: sensors, controllers, actuators, and human-machine interfaces for interacting with users as shown in Figure 1. Our first step introduces a general system model for evaluating security risk. First, we present a high-level model of IIoT systems. Then, we introduce a model of security investments in redundancy, diversity, and hardening, and we quantify risks posed by cyber-attacks, considering both probability and impact. Based on this model, we formulate the problem of optimal system design. For a list of symbols used in this paper, see [2].

A. System Model

We model the cyber part of the system as a directed graph $G = (C, E)$. The set of nodes C represents the components of the system, while the set of directed edges E represents connections between the components, which are used to send data and control signals. For each component $c \in C$, we let $O_c \subseteq C$ denote the set of origin components of the incoming edges of component c . Further, we let T_c denote the type of component c , which is one of the following:

- *sensor*: components that measure the state of physical processes (e.g., pressure sensors);
- *actuator*: components that directly affect physical processes (e.g., valves);
- *processing*: components that process and store data and control signals (e.g., PLCs);
- *interface*: components that interact with human users (e.g., HMI workstations).

The implementation of each component is chosen from a set of implementation types. We let I_c denote the set of types that may be used to implement component c , and we let I

denote the set of all implementation types that may be used in the system (i.e., $I = \cup_{c \in C} I_c$).

B. Security Investment Model

1) *Redundancy*: We model redundancy as deploying multiple instances of the same component. For simplicity, we assume that for each component, at most one instance of each suitable implementation type is deployed.¹ We make this assumption because our goal is to address security risks posed by deliberate attacks, and if a security vulnerability exists in an implementation type, then attackers can typically compromise all instances of that type.

We let $r_c \subseteq I_c$ denote the set of implementation types that are deployed for component $c \in C$. To quantify the cost of redundancy, we let R_i denote the cost of deploying an instance of type $i \in I_c$. Then, the total cost of redundancy is

$$\text{cost of redundancy} = \sum_{c \in C} \sum_{i \in r_c} R_i. \quad (1)$$

2) *Diversity*: We model diversity as deploying a diverse set of implementation types. In other words, diversity is modeled as selecting different implementations r_c to be deployed for each component $c \in C$ (or at least attempting to use as many distinct sets as possible).

To quantify the cost of diversity, we let D_i denote the cost of using an implementation type $i \in I$ in any non-zero number of components (i.e., D_i is the cost incurred when the first instance of type i is deployed). Then, the cost of diversity is

$$\text{cost of diversity} = \sum_{i \in \cup_{c \in C} r_c} D_i. \quad (2)$$

3) *Hardening*: We model the hardening of an implementation type as decreasing the probability that a zero-day security vulnerability is discovered by an attacker. We assume that hardening is applied in steps (e.g., performing a code review), resulting in a discrete set of hardening levels.

We let L_i denote the set of hardening levels available for implementation type $i \in I$, and we let l_i denote the chosen level. To model the amount of security provided by hardening level $l \in L_i$, we let S_l denote the probability that the implementation type will be secure (i.e., no zero-day vulnerability is discovered) if level l is chosen. To quantify the cost of hardening, we let H_l denote the cost of attaining level $l \in L_i$. Then, the total cost of hardening is

$$\text{cost of hardening} = \sum_{i \in I} H_{l_i}. \quad (3)$$

C. Security Risk Model

Next, we quantify the risks faced by a system with given redundancy, diversity, and hardening design. In principle, risk can be quantified as

$$\text{Risk} = \sum_{\text{outcome}} \Pr[\text{outcome}] \cdot \text{Impact}(\text{outcome}). \quad (4)$$

¹Note that relaxing this assumption would be straightforward; however, such a generalization would provide little further insight into security.

In our model, an outcome can be represented as a set of components that have been compromised by an attacker:

$$Risk(\mathbf{r}, \mathbf{l}) = \sum_{\hat{C} \subseteq C} \Pr[\hat{C} \text{ is compromised}] \cdot Impact(\hat{C}), \quad (5)$$

where $Impact(\hat{C})$ is the amount of loss inflicted on the system by an attacker who has compromised components \hat{C} . In the remainder of this subsection, we discuss how to measure $\Pr[\hat{C} \text{ is compromised}]$ and $Impact(\hat{C})$.

1) *Probability*: We quantify the probability that an attacker compromises a set of components $\hat{C} \subseteq C$ implicitly by describing a probabilistic process that models how an attacker can take control of the components of a system one-by-one. We consider two alternative attack models in our framework: non-stealthy attacks and stealthy attacks. The two attack models are summarized in Table I.

TABLE I
COMPONENT COMPROMISE RULES

Attack Type	Component Type			
	sensor	actuator	processing	interface
stealthy attack	if all instances are compromised	if all instances are compromised or all input components are compromised		
non-stealthy attack	if majority of instances are compromised	if majority of instances are compromised or majority of input components are compromised		

a) *Non-Stealthy Attacks*: First, an attacker attempts to find exploitable vulnerabilities in the implementation types that are deployed in the system. Based on our hardening model, the attacker discovers a zero-day vulnerability in each implementation type $i \in I$ with probability $1 - S_i$ (independently of the other types). We then consider all instances of the vulnerable implementation types to be compromised, and let \hat{I} denote the set of vulnerable implementations.

Next, we determine the set of compromised components \hat{C} . We start with $\hat{C} = \emptyset$, and then extend the set \hat{C} in iterations based on the following rules:

- a *sensor* component c is considered to be compromised if the majority of its instances r_c are vulnerable (i.e., if $|r_c \cap \hat{I}| \geq |r_c|/2$),
- an *actuator*, *processing*, or *interface* component c is considered to be compromised if the majority of its instances r_c are vulnerable or if the majority of its inputs are compromised (i.e., if $|O_c \cap \hat{C}| \geq |O_c|/2$).

We repeat the above steps until the set of compromised components \hat{C} cannot be extended any further.

b) *Stealthy Attacks*: For stealthy attacks, the process is the same except that “majority” is replaced in both rules with “all” (i.e., $|r_c \cap \hat{I}| = |r_c|$ and $|O_c \cap \hat{C}| = |O_c|$).

2) *Impact*: We let $Impact(\hat{C})$ denote the financial and physical loss resulting from an attack that compromises and maliciously controls components in \hat{C} . The exact formulation of $Impact(\hat{C})$ depends on the system and the characteristics of its physical processes. We present examples from two

domains, water-distribution (Section IV) and transportation systems [2].

Finally, we formulate the problem of finding an optimal design as follows.

Definition 1 (Optimal Design Problem). Given redundancy, diversity, and hardening investments R , D , and H , an *optimal design* (\mathbf{r}, \mathbf{l}) is

$$\operatorname{argmin}_{\mathbf{r}, \mathbf{l}} Risk(\mathbf{r}, \mathbf{l}) \quad (6)$$

subject to

$$\begin{aligned} & \forall c \in C : r_c \subseteq I_c; \quad \forall l \in I : l_i \in L_i \\ & \sum_{c \in C} \sum_{i \in r_c} R_i \leq R; \quad \sum_{i \in \cup_{c \in C} r_c} D_i \leq D; \quad \sum_{i \in I} H_i \leq H. \end{aligned}$$

III. COMPUTATIONAL ANALYSIS AND META-HEURISTIC ALGORITHMS

Since the number of feasible designs to choose from may be very large even for small systems, finding an optimal design using exhaustive search is computationally infeasible. In light of this, a key question for the practical application of the proposed framework is whether there exist efficient algorithms for finding optimal or near-optimal designs. We first show that finding an optimal design is computationally hard. Then, we introduce an efficient meta-heuristic algorithm that can find a near-optimal solution in polynomial time.

A. Computational Complexity

The objective of the design problem depends on the impact function, which could be any function, even one that is hard to compute. To show that the design problem is inherently hard (not only due to the potential complexity of computing the impact function), we assume a simplistic impact function, whose value is simply the number of compromised components. Formally, we consider $Impact(\hat{C}) = |\hat{C}|$.

Theorem 1. *The Optimal Design Problem is NP-hard.*

The proof of Theorem 1 can be found in [2].

B. Meta-Heuristic Design Algorithm

We propose an efficient meta-heuristic algorithm for finding near-optimal designs in practice. Our algorithm is based on simulated annealing, which requires randomly generating feasible solutions that are “neighbors” of (i.e., similar to) a given solution. Unfortunately, in our solution space (i.e., in the set of designs that satisfy the budget constraints), the feasible neighbors of a solution are not naturally defined. Hence, we first introduce an alternative representation of feasible designs, which we call design plans.

Definition 2 (Design Plan). A *design plan* is a pair $(\mathbf{r}\mathbf{o}, \mathbf{l}\mathbf{o})$, where

- $\mathbf{r}\mathbf{o}$ is a list of component-implementation pairs $(c, i) \in C \times I$ such that $i \in I_c$ holds for every pair $(c, i) \in \mathbf{r}\mathbf{o}$, and each possible pair (c, i) appears exactly once in $\mathbf{r}\mathbf{o}$;

- \mathbf{lo} is an ordered multiset of implementation types such that each implementation type $i \in I$ appears exactly $|L_i| - 1$ times in \mathbf{lo} .

ALGORITHM 1: *MapToDesign*(\mathbf{ro}, \mathbf{lo})

Data: optimal design problem, list \mathbf{ro} , ordered multiset \mathbf{lo}

Result: design (\mathbf{r}, \mathbf{l})

$\forall c \in C: r_c \leftarrow \emptyset; \forall i \in I: l_i \leftarrow \operatorname{argmin}_{l \in L_i} H_l$

for $(c, i) \in \mathbf{ro}$ **do**

$\mathbf{r}' \leftarrow \mathbf{r}; \mathbf{r}'_c \leftarrow r_c \cup \{i\}$

if $(\mathbf{r}', \mathbf{l})$ is feasible **then**

$\mathbf{r} \leftarrow \mathbf{r}'$

end

end

for $i \in \mathbf{lo}$ **do**

$\mathbf{l}' \leftarrow \mathbf{l}; \mathbf{l}'_i \leftarrow \operatorname{argmin}_{l \in L_i: H_l > H_{l_i}} H_l$

if $(\mathbf{r}, \mathbf{l}')$ is feasible **then**

$\mathbf{l} \leftarrow \mathbf{l}'$

end

end

output (\mathbf{r}, \mathbf{l})

Next, we show how to translate a design plan $(\mathbf{ro}, \mathbf{lo})$ into a feasible design. The translation is presented formally in Algorithm 1, and it is described in detail in the extended version of our paper [2]. Note this mapping is surjective.

ALGORITHM 2: Meta-Heuristic Design Algorithm

Data: optimal design problem, number of iterations k_{\max} , initial temperature T_0 , cooling parameter β

Result: design (\mathbf{r}, \mathbf{l})

choose $(\mathbf{ro}, \mathbf{lo})$ at random

$\rho \leftarrow \operatorname{Risk}(\operatorname{MapToDesign}(\mathbf{ro}, \mathbf{lo}))$

for $k = 1, \dots, k_{\max}$ **do**

$(\mathbf{ro}', \mathbf{lo}') \leftarrow \operatorname{Perturb}(\mathbf{ro}, \mathbf{lo})$

$\rho' \leftarrow \operatorname{Risk}(\operatorname{MapToDesign}(\mathbf{ro}', \mathbf{lo}'))$

$T \leftarrow T_0 \cdot e^{-\beta k}; pr \leftarrow e^{(\rho' - \rho)/T}$

if $(\rho' < \rho) \vee (\operatorname{rand}(0, 1) \leq pr)$ **then**

$\mathbf{ro} \leftarrow \mathbf{ro}', \mathbf{lo} \leftarrow \mathbf{lo}'$

end

end

output $\operatorname{MapToDesign}(\mathbf{ro}, \mathbf{lo})$

Finally, we present our meta-heuristic design algorithm (see Algorithm 2), which can find a near-optimal design in polynomial time. The algorithm starts by choosing a random design plan $(\mathbf{ro}, \mathbf{lo})$. In practice, we can implement this simply as choosing a random permutation of the list of component-implementation pairs and a random permutation of the multiset of implementation types. The algorithm is described in detail in the extended version of our paper [2].

IV. EVALUATION

To demonstrate the applicability of our framework, we present a case study from a canonical IIoT domain, water distribution. We present an additional case study from the transportation domain in the extended version of our paper [2].

IIoT systems have a particularly significant and wide application in water distribution systems. Examples include

monitoring water quality and detecting leaks. On the one hand, IIoT offers significant advantages, such as improved service and better maintenance at a low cost, but on the other hand, potential challenges include cost of the cyber infrastructure, reliability of communications, and of course, cyber-security.

As evidenced by the recent water crisis in Flint, MI [3], ensuring the quality of drinking water is of critical importance. Compromising systems that control the treatment and distribution of drinking water may allow adversaries to suppress warnings about contaminations or to decrease the quality of water [4]. Cyber-attacks can also have a devastating environmental impact. For example, in 2000, a disgruntled ex-employee launched a series of attacks against the SCADA system controlling sewage equipment in Maroochy Shire, Australia [5], [6]. As a result of these attacks, approximately 800,000 liters of raw sewage spilt out into local parks and rivers, killing marine life.

Here, we apply our framework to model cyber-physical contamination attacks against water-distribution systems. The system is modeled as a graph, in which links represent pipes, and nodes represent junctions of pipes, residential consumers, reservoirs, pumps, etc. IIoT components include:

- *Sensors*: water-quality sensors, which are located at certain nodes of the water-distribution network;
- *Processing*: components that collect, process, and forward water-quality data;
- *Interfaces*: components with human-machine interfaces, which can alert operators about contaminations.

We consider a malicious adversary who tries to cause harm by contaminating the water network with harmful chemicals. We assume that the adversary can introduce contaminants at certain nodes, such as unprotected reservoirs or tanks, which will then spread in the network, eventually reaching the residential consumers. We measure the impact of this physical attack as the amount of contaminants consumed by residential consumers before the detection of the attack.

To detect contaminations, each sensor continuously monitors the water flowing through the node at which it is deployed, and raises an alarm when the concentration of a contaminant reaches a threshold level. The alert generated by a *sensor* node is sent to a *processing* node, which forwards the alert to an *interface* node that can notify operators. Once operators are alerted, they respond immediately by warning residents not to consume water from the network.

We measure the impact of a physical attack as the amount of contaminants consumed by residential consumers before they are warned. This amount depends on the time between the physical attack and its detection, the contaminant concentration levels at the consumer nodes in this time interval, and the amount of water consumed in this interval. To increase the impact of the physical attack, the adversary launches a cyber-attack, which compromises and disables some of the components \hat{C} . Since the adversary's goal is to suppress warnings, this attack can be modeled as a *stealthy attack* (Section II-C1b). We assume that the adversary first compromises a set of components \hat{C} , and then decides where to introduce

the contaminant, maximizing the impact $Impact(\hat{C})$.

V. NUMERICAL RESULTS

We evaluate our approach numerically using a case study of a real-world water distribution network. We present additional numerical results in the extended version of our paper [2].

We use a real-world water-distribution network from Kentucky, which we obtained from the Water Distribution System Research Database[7]. The topology of this network, which is called KY3 in the database, is shown by Figure 2. In addition to topology, the database also contains hourly water-demand values for each network node.




Fig. 2. Topology of the water-distribution network. Colors show the spread of the contaminant from the first reservoir two hours after its introduction.

We assume that the adversary can introduce a contaminant at one of six given nodes in the network, which model three tanks and three reservoirs. Once the contaminant is introduced, we simulate its spread throughout the network using EPANET. From the simulation, we obtain the contaminant concentration values at the various nodes as functions of time. For a given set of compromised components \hat{C} , we then use these values to compute the time of detection and the resulting impact $Impact(\hat{C})$ (i.e., amount of contaminant consumed by the time of detection). Finally, we use the following numerical parameter values: $I = \{i_1, i_2, i_3, i_4, i_5\}$; for every $c \in C$, $I_c = I$; $R_{i_1} = R_{i_2} = R_{i_3} = 0$ and $R_{i_4} = R_{i_5} = 1$; $D_{i_1} = 0$ and $D_i = 1$ for every $i \in \{i_2, i_3, i_4, i_5\}$; for every $i \in I$, $L_i = \{1, 2, 3, \dots, 10\}$; for every $l \in L_i$, $S_l = 1 - 0.5^{0.5 \cdot l + 1}$ and $H_l = 4 \cdot l^2$.

Figure 3 shows the security risk in the water-distribution network for various budget values invested into the canonical approaches (i.e., redundancy, diversity, or hardening) and their optimal combination. Again, we note the logarithmic scaling on the vertical axis. We see that investing in a combination of redundancy, diversity, and hardening results in significantly lower risks than investing in only one of these approaches, thus demonstrating the efficacy and superior performance of a synergistic approach.

Figure 4 shows the optimal combination of redundancy, diversity, and hardening investments in the water-distribution network for various budget values. In this example, the optimal design is primarily a combination of diversity and hardening.




Fig. 3. Security risk in the water-distribution network when investing only in redundancy, only in diversity, only in hardening, or in their combination.




Fig. 4. Optimal combination of redundancy, diversity, and hardening investments in the water-distribution network.

However, with higher budget values, designers also need to invest in redundancy.

To illustrate the performance of the proposed design algorithm, we use the water-distribution network with $R = 10$ and $D = H = 100$. We find that the meta-heuristic algorithm (Algorithm 2) is very efficient: a single iteration takes less than 6.4×10^{-4} seconds (more than 1,500 iterations per second) on an average laptop computer. To determine the number of iterations that are necessary to find a good solution in practice, we focus on the solution quality (i.e., security risk) as a function of the number of iterations.




Fig. 5. Security risk in each iteration of one execution of the the meta-heuristic algorithm (Algorithm 2).

Figure 5 shows the security risk in each iteration of one particular execution of the meta-heuristic algorithm (Algorithm 2) with the current solution (solid red line) and with

the best solution found so far (dashed blue line). Please note the logarithmic scaling on the vertical axis. We have executed the algorithm a number of times, but since the results are qualitatively the same, we plot only one particular execution for illustration. The figure shows that risk decreases rapidly in the first few hundred iterations, but after around 400 iterations, the decrease becomes much slower. At around one thousand iterations, the risk reached its lowest value, so we omit the remaining iterations from the plot. In light of this, it is clear that the running time of the meta-heuristic algorithm is very low since it settles in a matter of seconds.

VI. RELATED WORK

IIoT and cyber-physical systems (CPS) have significantly improved the overall functionality, reliability, observability, and operational efficiency of industrial control systems and critical infrastructure networks [8], [9]. The integration and connectivity between various system components allow data exchange and information processing to fine tune system processes, but this integration and connectivity also opens new threat channels in the form of cyber- and cyber-physical attacks, against which these systems need to be secured [10], [11]. Conventional cybersecurity mechanisms are inadequate and thus need to be expanded to incorporate the complexity and physical aspects of such systems [10], [11], [12]. A detailed overview of the security issues in industrial automation systems that are based on open communication systems is provided in [13]. Similarly, security issues associated with various documented standards in SCADA systems are highlighted in [14], [15], and it is concluded that such issues cannot be resolved by employing only IT security mechanisms. There are various other studies that mainly highlight the security threats and associated risk assessment in the domain of industrial IoT, for instance [16], [17], [18]. All of these studies discuss and point towards a holistic security framework to address the security issues in industrial IoT. In this paper, we provide a framework for synergistic security that combines various security mechanisms to effectively secure such systems.

The water-supply industrial sector can benefit significantly from applying the ideas and technology of industrial Internet [19]. The adoption of new technologies (such as IoT, CPS) and networking devices enhances the monitoring capability, service reliability, and operational efficiency of water distribution systems, but also exposes them to malicious intrusions in the form of cyber- and cyber-physical attacks [4], [20]. A number of attack scenarios against water distributions systems are specified and demonstrated through simulations in [4]. Recently, in [21], several attacks on simulated and a real water distribution testbed (WADI [22]) are demonstrated through cyber-physical botnets capable of performing adversarial control strategies under CPS constraints. The security breach in the SCADA system of Maroochy Water Services, Australia [6] is a famous incident, which also highlights the need for effective security mechanisms. To effectively address the security challenge in such complex, interconnected, and spatially expanded systems,

we need to employ a combination of security mechanisms to protect them against cyber-physical attacks.

REFERENCES

- [1] A. J. O'Donnell and H. Sethu, "On achieving software diversity for improved network security using distributed coloring algorithms," in *Proceedings of the 11th ACM Conference on Computer and Communications Security (CCS)*. ACM, 2004, pp. 121–131.
- [2] A. Laszka, W. Abbas, Y. Vorobeychik, and X. Koutsoukos, "Synergistic security for the Industrial Internet of Things: Integrating redundancy, diversity, and hardening," *arXiv preprint arXiv:1808.09090* <https://arxiv.org/pdf/1808.09090>, 2018.
- [3] M. Kennedy, "Lead-laced water in Flint: A step-by-step look at the makings of a crisis," NPR, April 2016.
- [4] R. Taormina, S. Galelli, N. O. Tippenhauer, E. Salomons, and A. Ostfeld, "Characterizing cyber-physical attacks on water distribution systems," *Journal of Water Resources Planning and Management*, vol. 143, 2017.
- [5] M. Abrams and J. Weiss, "Malicious control system cyber security attack case study – Maroochy Water Services, Australia," http://csrc.nist.gov/groups/SMA/fisma/ics/documents/Maroochy-Water-Services-Case-Study_report.pdf, July 2008.
- [6] J. Slay and M. Miller, "Lessons learned from the Maroochy water breach," in *Critical Infrastructure Protection*, E. Goetz and S. Shenoi, Eds. Springer, 2008, pp. 73–82.
- [7] M. D. Jolly, A. D. Lothes, L. Sebastian Bryson, and L. Ormsbee, "Research database of water distribution system models," *Journal of Water Resources Planning and Management*, vol. 140, no. 4, 2014.
- [8] J. R. Moyne and D. M. Tilbury, "The emergence of industrial control networks for manufacturing control, diagnostics, and safety data," *Proceedings of the IEEE*, vol. 95, no. 1, pp. 29–47, 2007.
- [9] A. W. Colombo, S. Karnouskos, O. Kaynak, Y. Shi, and S. Yin, "Industrial cyberphysical systems: A backbone of the fourth industrial revolution," *IEEE Industrial Electronics Magazine*, vol. 11, no. 1, 2017.
- [10] A.-R. Sadeghi, C. Wachsmann, and M. Waidner, "Security and privacy challenges in industrial Internet of Things," in *Proceedings of the 52nd Annual Design Automation Conference*. ACM, 2015, p. 54.
- [11] X. Koutsoukos, G. Karsai, A. Laszka, H. Neema, B. Potteiger, P. Volgyesi, Y. Vorobeychik, and J. Sztipanovits, "SURE: A modeling and simulation integration platform for evaluation of secure and resilient cyber-physical systems," *Proceedings of the IEEE*, vol. 106, no. 1, pp. 93–112, 2018.
- [12] M. Cheminod, L. Durante, and A. Valenzano, "Review of security issues in industrial networks," *IEEE Transactions on Industrial Informatics*, vol. 9, no. 1, pp. 277–293, 2013.
- [13] D. Dzung, M. Naedele, T. P. Von Hoff, and M. Crevatin, "Security for industrial communication systems," *Proceedings of the IEEE*, vol. 93, no. 6, pp. 1152–1177, 2005.
- [14] J. Gao, J. Liu, B. Rajan, R. Nori, B. Fu, Y. Xiao, W. Liang, and C. Philip Chen, "SCADA communication and security issues," *Security and Communication Networks*, vol. 7, no. 1, pp. 175–194, 2014.
- [15] Y. Cherdantseva, P. Burnap, A. Blyth, P. Eden, K. Jones, H. Soulsby, and K. Stoddart, "A review of cyber security risk assessment methods for SCADA systems," *Computers & Security*, vol. 56, pp. 1–27, 2016.
- [16] L. Da Xu, W. He, and S. Li, "Internet of Things in industries: A survey," *IEEE Transactions on Industrial Informatics*, vol. 10, no. 4, 2014.
- [17] D. Meltzer, "Securing the industrial Internet of Things," *Information Systems Security Association Journal*, pp. 24–30, 2015.
- [18] J. Lee, B. Bagheri, and H.-A. Kao, "A cyber-physical systems architecture for industry 4.0-based manufacturing systems," *Manufacturing Letters*, vol. 3, pp. 18–23, 2015.
- [19] S. Kartakis, "Next generation cyber-physical water distribution systems," Ph.D. dissertation, Imperial College London, 2016.
- [20] L. Perelman and S. Amin, "A network interdiction model for analyzing the vulnerability of water distribution systems," in *Proceedings of the 3rd International Conference on High Confidence Networked Systems*. ACM, 2014, pp. 135–144.
- [21] D. Antonioli, G. Bernieri, and N. O. Tippenhauer, "Taking control: Design and implementation of botnets for cyber-physical attacks with cpsbot," *arXiv preprint arXiv:1802.00152*, 2018.
- [22] C. M. Ahmed, V. R. Palleti, and A. P. Mathur, "WADI: A water distribution testbed for research in the design of secure cyber physical systems," in *Proceedings of the 3rd International Workshop on Cyber-Physical Systems for Smart Water Networks*. ACM, 2017, pp. 25–28.