

Synergic Security for Smart Water Networks: Redundancy, Diversity, and Hardening

Aron Laszka
Vanderbilt University
Nashville, Tennessee 37212
aron.laszka@vanderbilt.edu

Yevgeniy Vorobeychik
Vanderbilt University
Nashville, Tennessee 37212
yevgeniy.vorobeychik@vanderbilt.edu

Waseem Abbas
Vanderbilt University
Nashville, Tennessee 37212
waseem.abbas@vanderbilt.edu

Xenofon Koutsoukos
Vanderbilt University
Nashville, Tennessee 37212
xenofon.koutsoukos@vanderbilt.edu

ABSTRACT

Smart water networks can provide great benefits to our society in terms of efficiency and sustainability. However, smart capabilities and connectivity also expose these systems to a wide range of cyber attacks, which enable cyber-terrorists and hostile nation states to mount cyber-physical attacks. Cyber-physical attacks against critical infrastructure, such as water treatment and distribution systems, pose a serious threat to public safety and health. Consequently, it is imperative that we improve the resilience of smart water networks. We consider three approaches for improving resilience: redundancy, diversity, and hardening. Even though each one of these “canonical” approaches has been thoroughly studied in prior work, a unified theory on how to combine them in the most efficient way has not yet been established. In this paper, we address this problem by studying the synergy of these approaches in the context of protecting smart water networks from cyber-physical contamination attacks.

CCS CONCEPTS

•**Security and privacy** → *Formal security models; Economics of security and privacy*; •**Networks** → *Sensor networks*; •**Computer systems organization** → *Embedded and cyber-physical systems*;

KEYWORDS

Smart water networks, cybersecurity, cyber-physical attacks, redundancy, diversity, hardening, cyber-physical systems

ACM Reference format:

Aron Laszka, Waseem Abbas, Yevgeniy Vorobeychik, and Xenofon Koutsoukos. 2017. Synergic Security for Smart Water Networks: Redundancy, Diversity, and Hardening. In *Proceedings of the 3rd International Workshop on Cyber-Physical Systems for Smart Water Networks, Pittsburgh, PA USA, April 2017 (CySWATER'17)*, 4 pages.
DOI: 10.1145/3055366.3055376

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
CySWATER'17, Pittsburgh, PA USA

© 2017 Copyright held by the owner/author(s). Publication rights licensed to ACM.
978-1-4503-4975-8/17/04...\$15.00
DOI: 10.1145/3055366.3055376

1 INTRODUCTION

Smart water networks promise to provide great benefits to our society in terms of efficiency and sustainability. For instance, smart water-distribution and waste-water systems may facilitate conserving water, thereby reducing consumer costs and environmental impact at the same time. In a smart water network, physical processes, sensor devices, controllers, and actuators form a connected cyber-physical system. Unfortunately, enhanced capabilities and connectivity also have a downside: previously secluded infrastructure is now susceptible to cyber-attacks.

Cyber-attacks against cyber-physical systems can pose a severe threat to public safety and health. For instance, 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. As evidenced by the recent water crisis in Flint, MI [6], ensuring the quality of drinking water is of critical importance. 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 [1, 12]. As a result of these attacks, approximately 800,000 liters of raw sewage spilt out into local parks and rivers, killing marine life.

Considering the importance of the issue, there has been an increasing concern to develop tools and approaches for assessing vulnerabilities in water networks [10]. In this direction, the emphasis is on identifying components in water networks that could be exposed to cyber-physical attacks, as well as the types of attacks that could be carried out [2, 14]. Recently, a simulation-based approach is presented in [13] for realistically assessing the risks associated with cyber-physical attacks on water distribution networks.

In this paper, we consider three canonical approaches for improving the resilience of a smart water network against cyber-attacks: redundancy, diversity, and hardening.

Redundancy means adding extra components to a system, which are not strictly necessary for achieving desired system functionality. Similar to hardening, redundancy can increase the cost of an attack, or reduce its success probability. In a cyber-physical systems, redundancy can be implemented by, e.g., deploying multiple sensors for monitoring the same physical processes. Further, redundancy can be implemented not only for components providing functionality, but also for security mechanisms. For example, multi-factor authentication methods grant a user access to a system only after

the user's identity has been successfully verified by multiple authentication methods. The rationale behind redundancy is that an adversary needs to disable or circumvent multiple components to compromise a system, which can significantly decrease the success probability of an attack.

Components that are based on the same hardware or software implementation typically suffer from the same vulnerabilities. Consequently, if an adversary can automate the exploitation of vulnerabilities, it may compromise a multitude of components with relatively little effort. *Diversity* can prevent the adversary from compromising a large number of system components using the same vulnerability. Diversity means employing multiple software or hardware implementations for components that perform the same tasks. In practice, different implementations are typically susceptible to different vulnerabilities, which limits the number of components that the adversary may compromise using a single vulnerability. In prior work, diversity has been explored using both static approaches (e.g., the effect of software diversity on cyber-risks [7, 11]) and dynamic approaches (e.g., moving target defense [9]).

Hardening means eliminating potential vulnerabilities from a component of the system. In a deterministic model, hardening increases the effort that an adversary needs to spend in order to find an exploitable vulnerability, while in a non-deterministic model, hardening decreases the probability of finding an exploitable vulnerability. A component can be hardened at multiple levels, ranging from hardware protection to software techniques. On the hardware level, employing tamper-resistant devices can prevent adversaries from mounting simple attacks based on physical access. On the software level, hardening approaches range from following secure-coding principles to setting up firewalls. Operators can also find and eliminate vulnerabilities by hiring security experts for penetration testing or by outsourcing vulnerability discovery through bug-bounty programs [8, 15]. Optimal security investments have been thoroughly studied in the economics of security literature [3, 4].

In this paper, we provide theoretical foundations for finding optimal combinations of these approaches in a smart water network. First, we introduce a model of cyber-physical contamination attacks and security investments into redundancy, diversity, and hardening (Section 2). Based on this model, we perform a case study of a real-world water network using simulated contaminations (Section 3). In our case study, we evaluate various combinations of the three approaches and compare them with each other. Finally, we present our concluding remarks (Section 4).

2 MODEL

In this section, we introduce our framework for studying the security of smart water networks. We first establish our system model, which captures the physical network and the sensor devices monitoring water quality. Then, we discuss our security-investment model, which includes redundancy, diversity, and hardening. Finally, we introduce our model of cyber-physical attacks.

2.1 System Model

We model the water network as a graph $G = (V, E)$, where the set of links E models pipes, and the set of nodes V models reservoirs, tanks, consumers, junctions of pipes, etc. Every consumer node

$v \in V$ has a demand value U_v , which quantifies the amount of water consumed at node v . For notational simplicity, we assign a demand value of zero to sources, such as water tanks.

To detect harmful contaminants in the water, the network is monitored by a set of sensor devices S . We assume that the sensors are deployed at the nodes of the network, and we let the location of sensor $s \in S$ be denoted by $l_s \in V$. A sensor continuously monitors the water flowing through its node, and raises an alarm when the concentration of a contaminant reaches a threshold level τ .

2.2 Security-Investment Model

To increase the impact of its physical attack, an adversary may compromise and disable sensor devices. Here, we discuss three approaches that defenders can implement to thwart cyber-attacks.

Redundancy. Firstly, a defender can increase resilience by deploying additional sensor devices. We let the minimum number of sensors that can adequately monitor the water network – without cyber-attacks – be denoted by S_{\min} . Then, we let the *level of redundancy* R be the number of sensor devices deployed above the bare minimum S_{\min} , that is, the level of redundancy is $R = |S| - S_{\min}$. Assuming that the cost of deploying and operating an additional sensor is C_R , the total cost of implementing redundancy is $C_R \cdot R$.

Diversity. Secondly, a defender can increase resilience by employing a diverse set of hardware and software to implement the sensor devices. We let the set of implementation types employed by the defender be denoted by T , and let the type of sensor s be denoted by $t_s \in T$. An implementation type $t \in T$ defines the choice of both hardware and software (e.g., hardware architecture and operating system). We let the *level of diversity* D be the number of different implementation types employed minus one, that is, the level of diversity is $D = |T| - 1$. Assuming that the cost of employing an additional sensor type is C_D , the total cost of diversity is $C_D \cdot D$.

Hardening. Thirdly, a defender can increase resilience by investing in hardening the implementation types (e.g., performing thorough testing for software vulnerabilities), as well as the individual devices (e.g., using tamper-resistant hardware). We let the defender's investment in hardening implementation type $t \in T$ be h_t , and the investment in hardening sensor device $s \in S$ be denoted by h_s . Then, we let the *level of hardening* H be the sum of all of these investments, that is, the level of hardening is $H = \sum_{t \in T} h_t + \sum_{s \in S} h_s$.

2.3 Cyber-Physical Attack Model

Finally, we introduce our model of cyber-physical attacks. We first discuss physical attacks, and then extend them with cyber-attacks.

2.3.1 Physical Attacks. 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 of the network, such as unprotected reservoirs or tanks. We let the possible introduction points for contamination be denoted by $P \subseteq V$.

Following its introduction at a node $p \in P$, the contaminant spreads in the network according to a function $C_p : \mathbb{N} \times V \mapsto \mathbb{R}_{\geq 0}$. For a given number of time steps $n \in \mathbb{N}$ and node $v \in V$, the value $C_p(n, v)$ is the concentration of the contaminant at node v , n time

steps after its introduction. Since a sensor detects the contaminant when its concentration reaches a threshold level τ , the number of time steps L_p until the contamination is detected is

$$L_p(S) = \min \left\{ n \in \mathbb{N} \mid \exists s \in S : C_p(n, l_s) \geq \tau \right\}. \quad (1)$$

We measure the impact of a physical attack as the amount of contaminants consumed with the water, which is proportional to both the demand values and the concentration of the contaminant. More formally, we quantify the impact of an undetected physical attack p in time step n as

$$\sum_{v \in V} U_v \cdot C_p(n, v). \quad (2)$$

Further, we assume that once the contamination is detected, operators can take instant countermeasures, such as warning customers not to consume water from the network. Consequently, we quantify the total impact I_p of a physical attack p as

$$I_p(S) = \sum_{n=1}^{L_p(S)} \sum_{v \in V} U_v \cdot C_p(n, v). \quad (3)$$

In other words, the cumulative impact of the attack up to the time step in which it is first detected.

2.3.2 Cyber-Physical Attacks. To increase the impact of the physical attack, an adversary can launch a cyber-attack, which compromises and disables some of the sensors. Here, we introduce our probabilistic model of cyber-attacks against sensors.

First, for each implementation type $t \in T$, the adversary finds a common vulnerability (e.g., a software bug) with probability

$$\Pr[\text{finding a vulnerability in type } t] = V_t \cdot e^{-h_t/C_H^T}, \quad (4)$$

where V_t is the vulnerability probability of type t without hardening, and C_H^T is the unit cost of hardening an implementation type. If the adversary finds a common vulnerability in type t , it compromises and disables all sensors of type t (i.e., removes every sensor $s \in S$ such that $t_s = t$).

Second, for each remaining sensor device s , the adversary compromises and disables the device with probability

$$\Pr[\text{compromising sensor } s] = V_s \cdot e^{-h_s/C_H^S}, \quad (5)$$

where V_s is the vulnerability probability of sensor s without hardening, and C_H^S is the unit cost of hardening a sensor.

Due to the cyber-attack, only a subset S_A of the sensors remains active and monitors the network for contaminants. Consequently, the *expected total impact* of a cyber-physical attack p is

$$\mathbb{E}_{S_A} [I_p(S_A)], \quad (6)$$

where $I_p(S_A)$ is the total impact of physical attack p given that only sensors S_A are active.

2.4 Problem Statement

We assume that the malicious adversary launches a worst-case attack against the system. Formally, the adversary mounts

$$\operatorname{argmax}_{p \in P} \mathbb{E}_{S_A} [I_p(S_A)]. \quad (7)$$

For given levels of redundancy R , diversity D , and hardening H , an optimal defense minimizes the expected impact of cyber-physical attacks, assuming that the adversary will launch a worst-case attack:

$$\min_{\substack{S, T, \langle l_s, t_s, h_s \rangle_{s \in S}, \langle h_t \rangle_{t \in T}: \\ |S| - S_{\min} \leq R, |T| - 1 \leq D, \sum_{t \in T} h_t + \sum_{s \in S} h_s \leq H}} \max_{p \in P} \mathbb{E}_{S_A} [I_p(S_A)]. \quad (8)$$

More generally, for a given amount of security investment C , the optimal combination of redundancy, diversity, and hardening is

$$\min_{\substack{S, T, \langle l_s, t_s, h_s \rangle_{s \in S}, \langle h_t \rangle_{t \in T}: \\ |S| - S_{\min} \leq R, |T| - 1 \leq D, \sum_{t \in T} h_t + \sum_{s \in S} h_s \leq H, \\ C_R \cdot R + C_D \cdot D + H \leq C}} \max_{p \in P} \mathbb{E}_{S_A} [I_p(S_A)]. \quad (9)$$

Since these problems are very challenging computationally, we use a simple but effective greedy heuristic to find optimal placements, type assignments, and distributions of hardening expenditure.

3 NUMERICAL ILLUSTRATIONS

For our numerical illustrations, we used a real-world water-distribution network from Kentucky, which we obtained from the Water Distribution System Research Database [5]¹. The topology of this network, which is called KY3 in the database, is shown by Figure 1. In addition to the topology, the database also contains hourly water-demand values for each node of the network.

We assumed that the adversary may introduce contaminants into the network at one of six nodes, which model three tanks and three reservoirs. We simulated every one of these six physical contamination attacks using EPANET², and recorded the resulting water-quality values at each node. Figure 1 shows the spread of the contaminant from the first reservoir two hours after its introduction.




Figure 1: Topology of the water-distribution network. Colors show the spread of the chemical contaminant from the first reservoir two hours after its introduction.

To detect contamination attacks, a defender can deploy water-quality sensors at any node of the network. We assumed the level of redundancy R to be equal to the number of sensors minus one (i.e., $S_{\min} = 1$), the level of diversity D to be equal to the number of different sensor types (e.g., different architectures or operating systems) minus one, and the cost of hardening to be $C_H^D = 1$ and $C_H^T = 100$ for sensor devices and sensor types, respectively. For each combination of redundancy, diversity, and hardening levels, we simulated 1 million cyber-attack scenarios to estimate the

¹<http://www.uky.edu/WDST/database.html>

²<https://www.epa.gov/water-research/epanet>

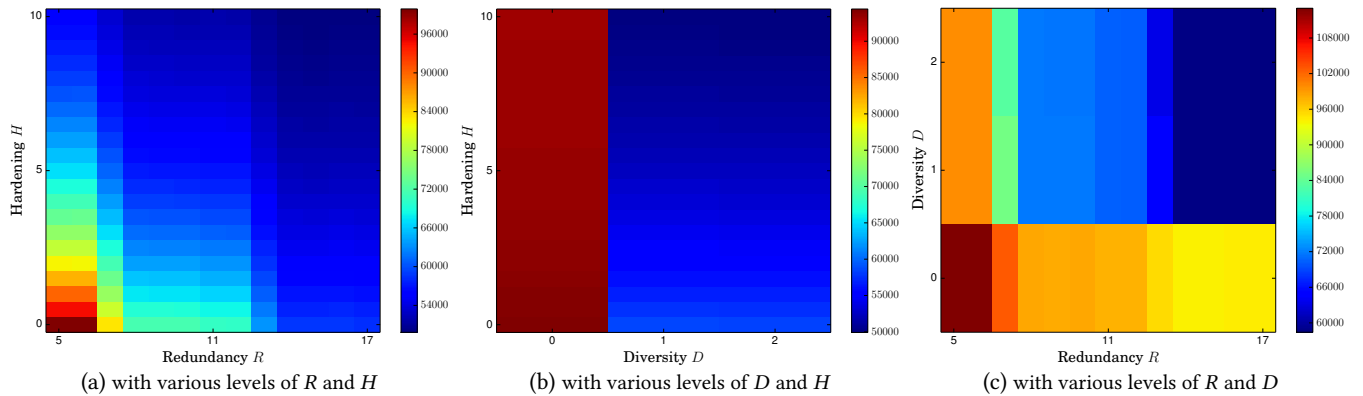


Figure 2: Expected impact of cyber-physical attacks with various levels or redundancy R , diversity D , and hardening H .

expected impact of each cyber-physical attack, and find the worst-case attack (Equation (7)). Finally, we used our greedy heuristic to find optimal placements, type assignments, and distributions of hardening expenditure (Equation (8)).

Figure 2(a) shows the impact of cyber-physical attacks with various levels of redundancy and hardening. In this figure, the level of diversity is fixed at $D = 2$. We observe that both redundancy and hardening are effective, but focusing only one approach may leave the network vulnerable.

Figure 2(b) shows the impact of cyber-physical attacks with various levels of diversity and hardening. In this figure, the level of redundancy is fixed at $R = 17$. We see that investing in diversity can significantly improve security; however, increasing the level of diversity above 1 leads to negligible improvement. On the other hand, investing in hardening provides a more modest but smooth improvement in security.

Figure 2(c) shows the impact of cyber-physical attacks with various levels of redundancy and diversity. In this figure, the level of hardening is fixed at $H = 0$ in this figure. We observe that neither redundancy nor diversity can solve security problems alone; however, their combination can significantly reduce the vulnerability of the network.

4 CONCLUSION

Cyber-physical attacks against smart water networks pose a severe threat to public health and safety. To protect a system from cyber-physical attacks, defenders may invest in multiple approaches: redundancy, diversity, and hardening. In this paper, we provided theoretical foundations for finding an optimal combination of these approaches in a smart water network. We first introduced a model of cyber-physical contamination attacks and security investments into redundancy, diversity, and hardening. Based on this model, we performed a case study of a real-world water network using simulated contaminations. We found that the three approaches are much more effective when combined, but finding optimal an combination can be challenging since expected impact is not a smooth function of the security investment levels. In future work, we will investigate the computational complexity of optimally combining the three approaches, and provide efficient algorithms for improving the resilience of a system in practice.

ACKNOWLEDGMENTS

This work was supported in part by the National Science Foundation (CNS-1238959), by the Air Force Research Laboratory (FA 8750-14-2-0180), and by the National Institute of Standards and Technology (70NANB15H263).

REFERENCES

- [1] Marshall Abrams and Joe Weiss. 2008. 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).
- [2] Saurabh Amin, Xavier Litrico, Shankar Sastry, and Alexandre M Bayen. 2013. Cyber security of water SCADA systems – Part I: Analysis and experimentation of stealthy deception attacks. *IEEE Transactions on Control Systems Technology* 21, 5 (2013), 1963–1970.
- [3] Ross Anderson and Tyler Moore. 2006. The economics of information security. *Science* 314, 5799 (2006), 610–613.
- [4] Lawrence A Gordon and Martin P Loeb. 2002. The economics of information security investment. *ACM Transactions on Information and System Security (TISSEC)* 5, 4 (2002), 438–457.
- [5] Matthew D Jolly, Amanda D Lothes, L Sebastian Bryson, and Lindell Ormsbee. 2014. Research database of water distribution system models. *Journal of Water Resources Planning and Management* 140, 4 (2014), 410–416.
- [6] Merrit Kennedy. 2016. Lead-Laced Water In Flint: A Step-By-Step Look At The Makings Of A Crisis. *NPR*, <http://www.npr.org/sections/thetwo-way/2016/04/20/465545378/>. (April 2016).
- [7] Aron Laszka and Jens Grossklags. 2015. Should Cyber-Insurance Providers Invest in Software Security?. In *Proc. of the 20th European Symposium on Research in Computer Security (ESORICS)*, 483–502.
- [8] Aron Laszka, Mingyi Zhao, and Jens Grossklags. 2016. Banishing Misaligned Incentives for Validating Reports in Bug-Bounty Platforms. In *Proc. of the 21st European Symposium on Research in Computer Security (ESORICS)*, 161–178.
- [9] Hamed Okhravi, Thomas Hobson, David Bigelow, and William Streilein. 2014. Finding focus in the blur of moving-target techniques. *IEEE Security & Privacy* 12, 2 (2014), 16–26.
- [10] Lina Perelman and Saurabh Amin. 2014. A network interdiction model for analyzing the vulnerability of water distribution systems. In *Proc. of the 3rd International Conference on High Confidence Networked Systems*. ACM, 135–144.
- [11] Fred B. Schneider and Kenneth P. Birman. 2009. The Monoculture Risk Put into Context. *IEEE Security & Privacy* 7, 1 (2009), 14–17.
- [12] Jill Slay and Michael Miller. 2008. Lessons learned from the Maroochy water breach. In *Critical Infrastructure Protection*. Springer, 73–82.
- [13] Riccardo Taormina, Stefano Galelli, Nils Ole Tippenhauer, Elad Salomons, and Avi Ostfeld. 2017. Characterizing Cyber-Physical Attacks on Water Distribution Systems. *Journal of Water Resources Planning and Management* (2017), 04017009.
- [14] Gurudeo Anand Tularam and Mark Properjohn. 2011. An investigation into modern water distribution network security: Risk and implications. *Security Journal* 24, 4 (2011), 283–301.
- [15] Mingyi Zhao, Jens Grossklags, and Peng Liu. 2015. An empirical study of web vulnerability discovery ecosystems. In *Proc. of the 22nd ACM SIGSAC Conference on Computer and Communications Security (CCS)*. ACM, 1105–1117.