Proceedings of the Singapore Cyber-Security Conference (SG-CRC) 2016
A. Mathur and A. Roychoudhury (Eds.)

doi:10.3233/978-1-61499-617-0-123

Simulation of Cyber-Physical Attacks on Water Distribution Systems with EPANET

Riccardo TAORMINA [a,1], Stefano GALELLI [a], Nils Ole TIPPENHAUER [b], Elad SALOMONS [c], and Avi OSTFELD [d]

[a] *ESD Pillar, Singapore University of Technology and Design*
[b] *ISTD Pillar, Singapore University of Technology and Design*
[c] *OptiWater*
[d] *Faculty of Civil and Environmental Engineering, Technion – Israel Institute of Technology*

Abstract. In this work, we discuss the use of EPANET to simulate the effects of malicious cyber-physical attacks on water distribution systems. EPANET—a standard numerical modeling environment developed by the US Environmental Protection Agency—models hydraulic and water-quality behavior of pressurized pipe networks. EPANET promises to be well suited to show the effects of direct attacks on hydraulic actuators, such as pumps, or the effects of attacks on sensors. Using the C-Town benchmark network, we show that EPANET has some limitations when modeling these attacks, and we report the workarounds needed to overcome these limitations. In particular, we describe attacks that change the control strategies of pumps, and attacks that alter the tank water level reported by sensors.

Keywords. Water Distribution Systems, Cyber-Physical System Security, EPANET

1. Introduction

Water distribution systems (WDS) are used in all urban developments to connect potable water sources (e.g., reservoirs or treatment plants) to the end users (e.g., housing and industrial estates). Typical components of WDS are pipes, pumps, valves, and storage tanks. Traditionally, these components were operated manually—or controlled locally—. In the past few decades advanced control systems were introduced by relying on SCADA systems, which allow automated collection of distributed sensor readings and centralized control of actuators. Whilst these changes promise to improve the operating performance of WDS, they are also expected to lead to novel security vulnerabilities [6]. In particular, attacks on automated controls have been recently discussed [2, 8, 9, 18]. The US Department of Homeland Security reported that 15 percent of the responses to cyber incidents were in the water sector [12]. Countermeasures against such attacks can be implemented

[1]Corresponding Author: Riccardo Taormina, Singapore University of Technology and Design, 8 Somapah Road, 487372 Singapore; E-mail: riccardo_taormina@sutd.edu.sg

through additional security measures on the sensor, network, and SCADA layers, but the fundamental problem remains: sensors readings rely on physical layer properties, which are susceptible to manipulations of the physical layer. Furthermore, these infrastructures are typically operated for extended periods of time, possibly decades. As such, there are higher chances that one or multiple components be attacked during their life cycle.

We consider cyber-physical attacks, which include both physical and cyber attacks that target the physical system (e.g. the Aurora attack [13]). The exact effect of such cyber-physical attacks effect depends on the specific physical process being altered. For WDS, both hydraulic (e.g. water spillage, tank overflow, pipe bursts, loss of pressure) and water quality processes (pollution with chemical or metals, biological matter) should be considered [1]. Complex process-based models are thus needed to assess the robustness, resilience and vulnerability of WDS against cyber-physical attacks. This is a novel field of research, and only few studies have tackled the problem of understanding the effect of cyber attacks on complex water systems. The authors of [3,4], for example, have recently studied the problem of detection and isolation of attacks on a water network comprised of cascaded canal pools. In this work, we discuss the use of EPANET for modeling the effects of malicious attacks on water distribution systems. EPANET simulates both hydraulic and water quality processes of pressurized pipe networks, so it is in principle well suited for the tasks described above. In the remainder of the paper we introduce some technical aspects of EPANET (Section 2), and then demonstrate its use for simulating cyber-physical attacks (Section 3). In particular, we introduce attacks that manipulate the control strategies of automated pumps and attacks that change the fill level reported by sensors installed in water tanks. Concluding comments are given in Section 4.

2. EPANET

EPANET is a public-domain software developed by the US Environmental Protection Agency to simulate hydraulic and water quality behavior within pressurized pipe networks [14, 15]. EPANET comes both as a self-contained application with a graphical user interface, as well as an open-source development Toolkit (written in C) that can be interfaced with other programming languages. The version used in this work, EPANET2, is de facto the standard modeling tool for both practitioners and researchers in the field of WDS. Typical applications of EPANET include long-term planning of WDS [5], design of alternative management strategies [17], as well as water quality modeling or the assessment of consumers exposure to contaminated water [10,11]. EPANET accurately reproduces WDS dynamics by employing a sophisticated network model and two engines for hydraulics and water-quality simulation.

2.1. Network Model

EPANET features *physical* components for designing the network topology, and *non-physical* components to define its operations. Physical components are divided into *nodes* and *links*, whereas non-physical components include *curves*, *time patterns* and *controls*.

Table 1. Properties and output variables of EPANET physical components

Component	Principal input properties	Output variables
Junctions	Coordinates and elevation Water demand Initial water quality	Hydraulic head Pressure Water quality
Reservoirs	Coordinates, Elevation Initial water quality	None (network boundary)
Tanks	Coordinates and bottom elevation Diameter (or shape) Initial, minimum and maximum level Initial water quality and mixing model	Hydraulic head Pressure Water quality
Pipes	Start and end nodes Diameter, Length, Roughness Status Bulk/Wall reaction coefficient	Flow rate, Velocity, headloss Friction factor Average reaction rate Average water quality
Pumps	Start and end nodes Pump curve Status, Speed	Flow rate Head gain Energy consumed
Valves	Start and end nodes Diameter, Setting, Status	Flow rate headloss

2.1.1. Nodes

Nodes can be either *reservoirs*, *junctions* or *tanks*, and represent source, sink and storage points of the WDS. Reservoirs are infinite external sources or sinks of water, e.g., lakes, rivers, groundwater aquifers, and tie-ins to other systems. They are treated as boundary points by the simulation engines. Junctions are points in the network where links join together, and where water enters or leaves the network. They are usually employed to represent the demand nodes in the network, but they can also model water-quality sources. Tanks are modelled as nodes with a given maximum capacity, and whose storage varies in time with the simulation.

EPANET allows the specification of several properties for each physical component, so that the modelled network closely resembles the real one. Table 1 summarizes the customizable input properties and available output variables for each component.

2.1.2. Links

The network nodes are connected by links, which are either *pipes* conveying water from one point to another or actionable components of the WDS, i.e., *pumps* and *valves*. In EPANET, all pipes are supposed to be full at all simulation time steps. The water in the pipes flows from nodes with higher energy—i.e., higher hydraulic head—to those with lower energy. When flowing, water experiences headloss due to friction with the pipe walls. The hydraulic engine of EPANET computes the headloss between the start and end node of a pipe using different empirical formulas, all expressing the headloss as a function of flow rate, length of the pipe and its roughness coefficient. The software also accounts for local headlosses due to turbulence occurring at pipes bends and fittings.

Pumps raise the hydraulic head of the water by imparting energy to it. They are modeled using a pump curve defining all possible combinations of head and flow that

Table 2. Examples of simple controls in EPANET

Control statement	Meaning
VALVE-2 CLOSED IF TANK-1 ABOVE 20	Close VALVE-2 if the level in TANK-1 exceeds 20 ft.
VALVE-2 OPEN IF NODE 130 BELOW 30	Open VALVE-2 if the pressure at Node 130 $<$ 30 psi
PUMP-3 1.5 AT TIME 16	Set relative speed of PUMP-3 to 1.5 after 16 hours
PUMP-3 CLOSED AT CLOCKTIME 10 AM	Shuts down PUMP-3 at 10 AM

a pump can produce, or as a device that supplies a constant amount of energy to the fluid flowing through it. EPANET also features pumps with variable speed—the relative velocity with respect to the original pump curve can be changed during the simulation. A change of speed affects the position and shape of the pump curve.

Valves are links limiting the pressure or flow at a specific point in the network. EPANET supports several types of valves, such as pressure reducing valves (PRV) and flow control valves (FCV). General purpose valves (GPV) can also be employed to let the user define particular components or conditions. Shutoff and check valves (CV), which completely open or close pipes, are also featured in EPANET but as a property of the pipe in which they are placed.

2.1.3. Controls

Controls are statements that determine how the network is operated over time. They determine the functioning of selected links as a function of time or the value of some nodal variables. There are two categories of controls in EPANET. *Simple controls* change the status or setting of a link as a function of either 1) the water level in a tank, 2) the pressure at a junction, 3) the time into the simulation, or 4) the time of the day. Some examples are given in Table 2. *Rule-based controls* extend simple controls with logical operators, thus allowing to command a link based on a combination of conditions.

2.2. EPANET Simulation Models

EPANET's hydraulic engine is demand-driven—consumer demands are always satisfied regardless of the pressures throughout the system. This entails that if nodal demands exceed the maximum flow, the software extrapolates the pump curves to meet the required flow, even if this produces a negative head. Such cases are regarded as unrealistic and the software returns a warning message. EPANET hydraulically balances the network for heads and flows at a particular point in time by solving simultaneously the conservation of flow equation for each junction and the headloss relationship across the links. This process is carried out using the Gradient Algorithm [16].

3. Modeling Cyber-Physical Attacks with EPANET

3.1. Attacker Model

The goal of the attacker is the disruption of the network's hydraulic behavior. The attacker can achieve this through *direct* and *indirect* attacks. In a direct attack, the attacker takes direct control of an actionable component, e.g., by physical manipulation, remote compromise of the controller, or manipulation of the control messages. In an indirect

attack, the attacker manipulates the reported reading of a nodal sensor (by either physical manipulation, remote compromise of the sensor, or manipulation of the network traffic). The manipulated sensor reading then leads to incorrect control action to be taken. Both attacks are altering the normal operations of pumps and valves in the network.

3.2. Attack Impact Metric

To better characterize an attack, appropriate criteria should be defined to assess the level of service disruption in the network. For instance, pressure-related criteria can be employed to check whether the pressure at the demand nodes is within an optimal functioning range. These measures may be evaluated on the whole network or on specific zones affected by the attack. In case storage tanks are targeted, the amount of overflow witnessed during the attack can be used to estimate its impact. Other measures may consider economic costs, such as those deriving from damage to pumps and valves.

3.3. EPANET-based Implementation of Cyber-Physical Attacks

Since EPANET was not developed to study WDS behavior from a cyber-physical perspective, some workarounds are necessary to model attack scenarios. The degree of flexibility needed for such implementation is achieved through the programming Toolkit. In this study, threats are simulated by modifying the control statements commanding the targeted components in the network. The attacks strike at a chosen time during the simulation and last for a given period, during which the control statements are altered consistently. This procedure is carried out by running the hydraulic simulation in a step-by-step fashion and interposing extra conditional constructs between consecutive steps of the simulation.

3.3.1. The C-TOWN Network

The implementation of cyber-physical attacks in EPANET is illustrated with some examples for the C-Town benchmark network [7], which is depicted in Figure 1. C-Town is based on a real-world medium-sized network with two main storage tanks (T1 and T2), and a main pumping station (S1) consisting of three pumps (PU1, PU2 and PU3). The operations of PU1 and PU2 are regulated by the water levels in T1, while PU3 is a redundant pump—it is kept off during standard operating conditions. The remaining five tanks are refilled by four secondary boosting stations that pump water from T1 and T2. While T1 is directly connected to S1, the branch pertaining T2 is connected to the source through a FCV (V1) controlled by the water level in T2.

The attack scenarios described in the following sections were generated for one week-long simulations (168 hours), and performed with an hydraulic time step of 15 minutes. The demand was generated by 5 different week-long time patterns with a time step of 1 hour. The attacks start 48 hours into the simulation and last for 2 whole days.

3.3.2. Example 1: Overflow of T1 via Attack on PU1 and PU2

The first example is an indirect attack on S1 aimed at producing an overflow of T1. This scenario assumes that there is no additional overflow sensor in the tank or that it has been deactivated by the attacker. In normal conditions, the pumps PU1 and PU2 activate and deactivate according to the following controls written in the input file:

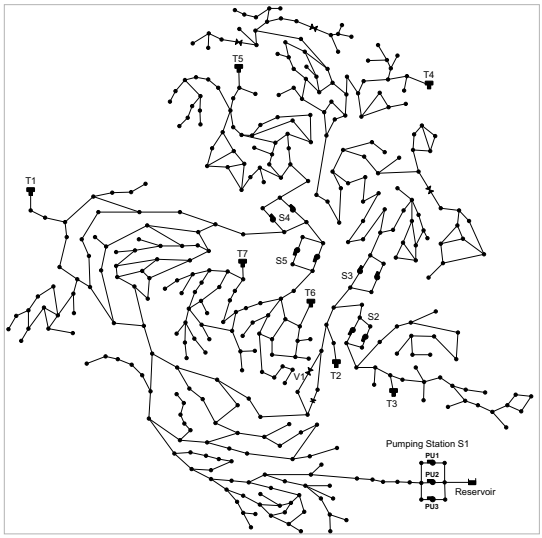

Figure 1. Subset of the C-TOWN network. We focused on tanks T1 and T2, the pumping station S1, and the main valve V1.

```
LINK  PU1  OPEN    IF  NODE  T1  BELOW  4
LINK  PU1  CLOSED  IF  NODE  T1  ABOVE  6.3

LINK  PU2  OPEN    IF  NODE  T1  BELOW  1
LINK  PU2  CLOSED  IF  NODE  T1  ABOVE  4.5
```

The attack consists in tampering the readings of the water level in T1 so that the value transmitted to the controller is always below 4.5 meters. In this way, once started, PU1 and PU2 keep pumping water to the system causing T1 to overflow when the level rises above the maximum capacity of 6.5 meters. This attack is simulated by temporarily changing the settings at the end of the second and fourth control statements to a value above 6.5 meters, so that the pumps never switch to the CLOSED status when the attack is in place. Since EPANET cannot directly model tank overflows, the original network map is modified to amend for such shortcoming. This is done by 1) duplicating the original pipe connecting the tank to the network, 2) connecting it to a dummy storage tank, and 3) including controls that keep the link closed unless the level in the original tank reaches the maximum capacity. `NEWLINK` represents this additional pipe, and the additional control statements for modeling the overflow of T1 are written as:

```
LINK  NEWLINK OPEN    IF  NODE  T2  ABOVE  6.499999
LINK  NEWLINK CLOSED  IF  NODE  T2  BELOW  6.499999
```

The amount of overflow due to the attack is therefore equal to the amount of water stored in the dummy tank at the end of the attack, or to the volume of water that flows through `NEWLINK` during the attack. The results are illustrated in Figure 2, where the attack scenario is compared against the time series expected for the same initial conditions under normal operations. Due to the system inertia and prior state, it takes around 10 hours for the attack to drive the system outside its normal regime, while almost the entire duration of the attack is needed to cause T1 to overflow.

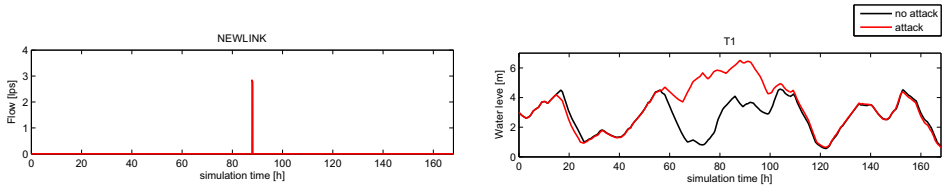

Figure 2. Trajectory of water flow in NEWLINK (left) and storage in T1 (right) during normal operations (black line) and under attack (red line). Results show that, when the attack strikes, a large volume of water overflows from T1.

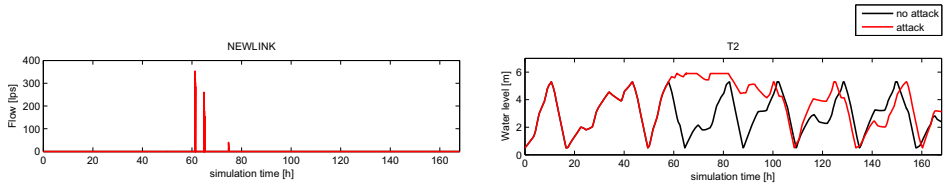

Figure 3. Trajectory of water flow in NEWLINK (left) and storage in T2 (right) during normal operations (black line) and under attack (red line). Results show that the attack leads to large overflow at T2.

3.3.3. Example 2: Overflow of T2 via Attack on V1

Overflow in T2 occurs if the valve connecting it to the main line of the network is forced to stay open for a sufficient amount of time. In this case, the attacker intercepts the communication to the valve controller and tampers the real readings of the water level in T2. Under normal conditions, the controller operates V1 as follows:

```
LINK  V1  OPEN    IF  NODE  T2  BELOW  0.5
LINK  V1  CLOSED  IF  NODE  T2  ABOVE  5.3
```

In this example, the valve is forced to stay in the OPEN status by substituting the setting at the end of the second statement to a value above the maximum level of T2 (for example, 5.9 meters). The effects of this attack are shown in Figure 3—under the same assumptions as for the first example, with dummy tank connected to T2.

4. Conclusion

In this work we discuss the opportunities and challenges of simulating cyber-physical attacks to WDS using EPANET—a numerical modeling environment commonly adopted to simulate WDS behavior under standard operating conditions. We present two examples, related to the overflow of storage tanks due to the direct attack to pumps controls and manipulation of valves settings. The examples show that some workarounds are necessary to implement these attacks. A limit of EPANET that surely deserve further research is the assumption that the consumers' demand is always satisfied (regardless of the pressure at the demand nodes). This implies that EPANET cannot simulate sudden shortfalls of water supply as well as other pressure-related issues, such as pipe bursts. Adopting pressure-driven numerical simulation environments seems to be a viable option; this is subject of ongoing research.

Acknowledgments

This work was supported by the National Research Foundation (NRF), Singapore, under its National Cybersecurity R&D Programme (Award No. NRF2014NCR-NCR001-40), by the Chief Scientist (OCS) Ministry of Science, Technology and Space (MOST), and by the Germany Federal Ministry of Education and Research (BMBF), under project no. 02WA1298.

References

[1] M. Abrams and J. Weiss. Malicious control system cyber security attack case study–maroochy water services, australia.

[2] S. Amin, X. Litrico, S. Sastry, and A. Bayen. Cyber security of water SCADA systems; Part I: Analysis and experimentation of stealthy deception attacks. *Control Systems Technology, IEEE Transactions on*, 21(5):1963–1970, 2013.

[3] S. Amin, X. Litrico, S. Sastry, and A. M. Bayen. Cyber security of water scada systemspart i: analysis and experimentation of stealthy deception attacks. *Control Systems Technology, IEEE Transactions on*, 21(5):1963–1970, 2013.

[4] S. Amin, X. Litrico, S. S. Sastry, and A. M. Bayen. Cyber security of water scada systemspart ii: Attack detection using enhanced hydrodynamic models. *Control Systems Technology, IEEE Transactions on*, 21(5):1679–1693, 2013.

[5] G. Fu, Z. Kapelan, J. R. Kasprzyk, and P. Reed. Optimal design of water distribution systems using many-objective visual analytics. *Journal of Water Resources Planning and Management*, 2012.

[6] W. Gao. *Cyberthreats, attacks and intrusion detection in supervisory control and data acquisition networks*. PhD thesis, Mississippi State University, December 2013.

[7] O. Giustolisi, L. Berardi, D. Laucelli, D. Savic, and Z. Kapelan. Operational and tactical management of water and energy resources in pressurized systems: Competition at WDSA 2014. *Journal of Water Resources Planning and Management*, page C4015002, 2015.

[8] O. Kosut, L. Jia, R. Thomas, and L. Tong. Malicious data attacks on smart grid state estimation: Attack strategies and countermeasures. In *Proc. of the IEEE Conference on Smart Grid Communications (SmartGridComm)*, pages 220–225, Oct 2010.

[9] Y. Liu, P. Ning, and M. K. Reiter. False data injection attacks against state estimation in electric power grids. *ACM Transactions on Information and System Security (TISSEC)*, 14(1):13, 2011.

[10] A. Ostfeld and E. Salomons. Optimal layout of early warning detection stations for water distribution systems security. *Journal of Water Resources Planning and Management*, 130(5):377–385, 2004.

[11] A. Ostfeld, J. G. Uber, E. Salomons, J. W. Berry, W. E. Hart, C. A. Phillips, J.-P. Watson, G. Dorini, P. Jonkergouw, Z. Kapelan, et al. The battle of the water sensor networks (bwsn): A design challenge for engineers and algorithms. *Journal of Water Resources Planning and Management*, 134(6):556–568, 2008.

[12] Industrial Control Systems Cyber Emergency Response Team. Malware infections in the control environment. *ICS-CERT Monitor*, pages 1–15, 2012.

[13] North American Electric Reliability Corporation. *NERC Issues AURORA Alert to Industry*. 2010.

[14] L. A. Rossman. The EPANET programmers toolkit for analysis of water distribution systems. In *Proceedings of Conference on Water Resources Planning and Management*, pages 39–48, 1999.

[15] L. A. Rossman. EPANET 2: users manual, 2000.

[16] E. Todlini and S. Pilati. A gradient method for the analysis of pipe networks. In *Proceedings of the conference on computer applications for water supply and distribution*, 1987.

[17] J. E. Van Zyl, D. A. Savic, and G. A. Walters. Operational optimization of water distribution systems using a hybrid genetic algorithm. *Journal of water resources planning and management*, 130(2):160–170, 2004.

[18] L. Xie, Y. Mo, and B. Sinopoli. False data injection attacks in electricity markets. In *Proc. of the IEEE Conference on Smart Grid Communications (SmartGridComm)*, pages 226–231, Oct 2010.