Battle of the Attack Detection Algorithms: Disclosing Cyber Attacks on Water Distribution Networks

Riccardo Taormina[1]; Stefano Galelli, M.ASCE[2]; Nils Ole Tippenhauer[3]; Elad Salomons[4]; Avi Ostfeld, F.ASCE[5]; Demetrios G. Eliades[6]; Mohsen Aghashahi, S.M.ASCE[7]; Raanju Sundararajan[8]; Mohsen Pourahmadi[9]; M. Katherine Banks, F.ASCE[10]; B. M. Brentan[11]; Enrique Campbell[12]; G. Lima[13]; D. Manzi[14]; D. Ayala-Cabrera[15]; M. Herrera[16]; I. Montalvo[17]; J. Izquierdo[18]; E. Luvizotto Jr.[19]; Sarin E. Chandy[20]; Amin Rasekh, M.ASCE[21]; Zachary A. Barker[22]; Bruce Campbell[23]; M. Ehsan Shafiee[24]; Marcio Giacomoni[25]; Nikolaos Gatsis[26]; Ahmad Taha[27]; Ahmed A. Abokifa, S.M.ASCE[28]; Kelsey Haddad[29]; Cynthia S. Lo[30]; Pratim Biswas[31]; M. Fayzul K. Pasha[32]; Bijay Kc[33]; Saravanakumar Lakshmanan Somasundaram[34]; Mashor Housh[35]; and Ziv Ohar[36]

Abstract: The BATtle of the Attack Detection ALgorithms (BATADAL) is the most recent competition on planning and management of water networks undertaken within the Water Distribution Systems Analysis Symposium. The goal of the battle was to compare the performance of algorithms for the detection of cyber-physical attacks, whose frequency has increased in the last few years along with the adoption of smart water technologies. The design challenge was set for the C-Town network, a real-world, medium-sized water distribution system operated through programmable logic controllers and a supervisory control and data acquisition (SCADA) system. Participants were provided with data sets containing (simulated) SCADA observations, and challenged to design an attack detection algorithm. The effectiveness of all submitted algorithms was evaluated in terms of time-to-detection and classification accuracy. Seven teams participated in the battle and proposed a variety of successful approaches leveraging data analysis, model-based detection mechanisms, and rule checking. Results were presented at the Water Distribution Systems Analysis Symposium (World Environmental and Water Resources Congress) in Sacramento, California on May 21–25, 2017. This paper summarizes the BATADAL problem, proposed algorithms, results, and future research directions. **DOI: 10.1061/(ASCE)WR.1943-5452.0000969.**

Author keywords: Water distribution systems; Cyber-physical attacks; Cyber security; EPANET; Smart water networks; Attack detection.

[1]Postdoctoral Research Fellow, iTrust Centre for Research in Cyber Security, Singapore Univ. of Technology and Design, 8 Somapah Rd., Singapore 487372, Singapore.

[2]Assistant Professor, Pillar of Engineering Systems and Design, Singapore Univ. of Technology and Design, 8 Somapah Rd., Singapore 487372, Singapore (corresponding author). Email: stefano_galelli@sutd.edu.sg

[3]Assistant Professor, Pillar of Information Systems Technology and Design, Singapore Univ. of Technology and Design, 8 Somapah Rd., Singapore 487372, Singapore.

[4]Independent Water Resources Consultant, OptiWater, 6 Amikam Israel St., Haifa 3438561, Israel.

[5]Professor, Faculty of Civil and Environmental Engineering, Technion–Israel Institute of Technology, Haifa 32000, Israel.

[6]Research Associate, KIOS Research and Innovation Center of Excellence, Univ. of Cyprus, 75 Kallipoleos Ave., CY-1678, Nicosia, Cyprus.

[7]Ph.D. Student, Zachry Dept. of Civil Engineering, Texas A&M Univ., College Station, TX 77840.

[8]Ph.D. Student, Dept. of Statistics, Texas A&M Univ., College Station, TX 77843.

[9]Professor, Dept. of Statistics, Texas A&M Univ., College Station, TX 77843.

[10]Professor, Zachry Dept. of Civil Engineering, Texas A&M Univ., College Station, TX 77840.

[11]Ph.D. Student, Centre de Recherche en Automatique de Nancy, Universitet de Lorraine, BP 70239, 54506 Nancy, France.

Note. This manuscript was submitted on November 5, 2017; approved on February 28, 2018; published online on June 9, 2018. Discussion period open until November 9, 2018; separate discussions must be submitted for individual papers. This paper is part of the ***Journal of Water Resources Planning and Management***, © ASCE, ISSN 0733-9496.

[12]Engineer, Fluing-Instituto de Matemática Multidisciplinar, Universitat Politècnica de València, 46022 Valencia, Spain; Berliner Wasserbetriebe, Berlin, Germany.

[13]Research Assistant, Computational Hydraulic Laboratory, Universidade Estadual de Campinas, CEP 13083-889, Campinas, Brazil.

[14]Research Assistant, Computational Hydraulic Laboratory, Universidade Estadual de Campinas, CEP 13083-889, Campinas, Brazil.

[15]Research Fellow, Dept. of Water, Irstea, BP 5095, 34196 Cestas, France.

[16]Research Associate, Dept. of Architecture and Civil Engineering, Univ. of Bath, Bath BA2 7AY, UK.

[17]Research Engineer, Ingeniousware GmbH, 76137 Karlsruhe, Germany.

[18]Professor, Fluing-Instituto de Matemática Multidisciplinar, Universitat Politècnica de València, 46022 Valencia, Spain.

[19]Associate Professor, Computational Hydraulic Laboratory, Universidade Estadual de Campinas, CEP 46022, Campinas, Brazil.

[20]Data Scientist, Sensus, Inc., 8601 Six Forks Rd., Suite 700, Raleigh, NC 27615.

[21]Lead R&D Engineer, Sensus, Inc., 8601 Six Forks Rd., Suite 700, Raleigh, NC 27615.

[22]Water Resources Engineer, Sensus, Inc., 8601 Six Forks Rd., Suite 700, Raleigh, NC 27615.

[23]Chief Data Scientist, Sensus, Inc., 8601 Six Forks Rd., Suite 700, Raleigh, NC 27615.

[24]Lead R&D Engineer, Sensus, Inc., 8601 Six Forks Rd., Suite 700, Raleigh, NC 27615.

[25]Assistant Professor, Dept. of Civil and Environmental Engineering, Univ. of Texas at San Antonio, San Antonio, TX 78249.

[26]Assistant Professor, Dept. of Electrical and Computer Engineering, Univ. of Texas at San Antonio, San Antonio, TX 78249.

Introduction

Recent decades have witnessed the transition of water distribution systems from traditional physical infrastructures to cyber-physical systems that combine physical processes with computation and networking: physical assets—such as pipes, pumps, and valves—work in unison with networked devices that monitor and coordinate the operations of the entire system. These devices include programmable logic controllers (PLCs), supervisory control and data acquisition (SCADA) systems, remote terminal units (RTUs), static and mobile sensor networks, and smart meters (Hill et al. 2014; Gong et al. 2016; Sønderlund et al. 2016). The adoption of such smart water technologies plays a pivotal role in enhancing the a utomation and reliability of water distribution systems, but simultaneously exposes them to cyber-physical attacks (Rasekh et al. 2016)—namely the deliberate exploitation of computer systems aimed at accessing sensitive information or compromising the operations of the underlying physical system. Water (and wastewater) systems represent one of the 16 critical infrastructure sectors identified by the US Department of Homeland Security (2017), according to which the number of reported attacks on water infrastructures has been growing steadily (ICS-CERT 2014, 2015, 2016), making them the third most targeted sector after critical manufacturing and energy (ICS-CERT 2016). To take remedial actions, several countries are establishing research centers and international collaborations, such as the Israel–New York collaboration to defend water systems from infrastructure terrorists (The Times of Israel 2018).

Protecting water distribution systems from cyber attacks requires (as with other cyber-physical systems) a combination of proactive and reactive mechanisms (Cardenas et al. 2008). Proactive mechanisms comprise all tools that reduce the chances to penetrate the system, such as appropriate measures for traffic authentication and confidentiality protection, access control, and device hardening (Graham et al. 2016; Adepu et al. 2017). Because it is not possible to rule out all attacks, cyber-physical systems should also be equipped with intrusion detection schemes that assist with the recovery phase (Anderson 2010). Disclosing cyber attacks—without issuing false alarms—is thus crucial. Unfortunately, this does not come without some system-specific challenges. First, the definition of anomalous behaviors should not only be related to point, or content, anomalies—i.e., data points lying beyond some specific thresholds—because cyber-physical attacks can tamper with one or multiple network components while keeping the performance characteristics within historical bounds (Abokifa et al. 2017). This implies that detection schemes should be capable of disclosing both content and contextual anomalies, namely, data points that are considered abnormal when viewed against meta-information associated with the data points (Hayes and Capretz 2015). For example, unaccounted high volumes of water leaving tanks during the night, when demand is generally low, may be seen as a contextual anomaly revealed by looking at the flow data in the context of time. Second, the same hydraulic response of a water network (e.g., low water levels in a tank) can be obtained through different attacks (Taormina et al. 2017). Therefore, detection schemes should also identify the cyber components that have been attacked; this is a nonnegligible challenge in large water networks. Third, all networked devices, including SCADA systems, represent potential targets. This means that the information provided by SCADA systems may not be fully reliable.

As the field of intrusion detection continues to grow, so too does the need for an objective comparison of attack detection algorithms for water distribution systems. The BATtle of the Attack Detection ALgorithms (BATADAL) was organized for this purpose. Participants were provided with data sets containing (simulated) SCADA data for a water distribution system that was the target of cyber attacks, and were tasked with the design of an attack detection mechanism. The design goals of a detection algorithm were to: (1) disclose the presence of an ongoing attack in the minimum amount of time possible, (2) avoid issuing false alarms, and (3) identify which components of the system have been compromised (optional). Seven teams, from both academia and industry, contributed novel solutions, which were evaluated using specific evaluation criteria—i.e., time-to-detection and classification accuracy. The BATADAL results were presented at a special session of the Water Distribution Systems Analysis Symposium (World Environmental and Water resources Congress), in Sacramento, California on May 21–25, 2017.

This paper summarizes the main solutions and outcomes of the BATADAL and proposes future research directions for event detection in the realm of cyber-physical security. The remainder of the paper (1) describes the BATADAL problem, data, and evaluation criteria; (2) presents a synopsis of the proposed attack detection algorithms; (3) analyzes the results; and (4) presents conclusions and future research directions.

Problem Description

The operators of the C-Town water distribution system have observed anomalous behaviors in some hydraulic components, e.g., tank overflows, reduction in pump speed, and anomalous activation/deactivation of pumps. They suspect that the anomalies are attributable to cyber-physical attacks that interfered with the system operations and tampered with the readings recorded by the SCADA system. The aim of the participants was to develop an attack detection mechanism that detects the presence of attacks—in the shortest amount of time—from the available hourly SCADA data. In particular, attack detection algorithms must classify the system state as either safe or under attack. The rest of this section summarizes C-Town and presents the development data and evaluation criteria. BATADAL rules, problem details, and data are available in the Supplemental Data.

C-Town Network

The C-Town water distribution system is based on a real-world, medium-sized network first introduced for the Battle of the Water

[27]Assistant Professor, Dept. of Electrical and Computer Engineering, Univ. of Texas at San Antonio, San Antonio, TX 78249.
[28]Ph.D. Student, Dept. of Energy, Environmental, and Chemical Engineering, Washington Univ., St. Louis, MO 63130-4899.
[29]Ph.D. Student, Dept. of Energy, Environmental, and Chemical Engineering, Washington Univ., St. Louis, MO 63130-4899.
[30]Assistant Professor, Dept. of Energy, Environmental, and Chemical Engineering, Washington Univ., St. Louis, MO 63130-4899.
[31]Professor, Dept. of Energy, Environmental, and Chemical Engineering, Washington Univ., St. Louis, MO 63130-4899.
[32]Associate Professor, Dept. of Civil and Geomatics Engineering, California State Univ., Fresno, CA 93740.
[33]Research Assistant, Dept. of Civil and Geomatics Engineering, California State Univ., Fresno, CA 93740.
[34]Research Assistant, Dept. of Civil and Geomatics Engineering, California State Univ., Fresno, CA 93740.
[35]Assistant Professor, Faculty of Management, Dept. of Natural Resources and Environmental Management, Univ. of Haifa, Haifa 3498838, Israel.
[36]Research Assistant, Faculty of Management, Dept. of Natural Resources and Environmental Management, Univ. of Haifa, Haifa 3498838, Israel.

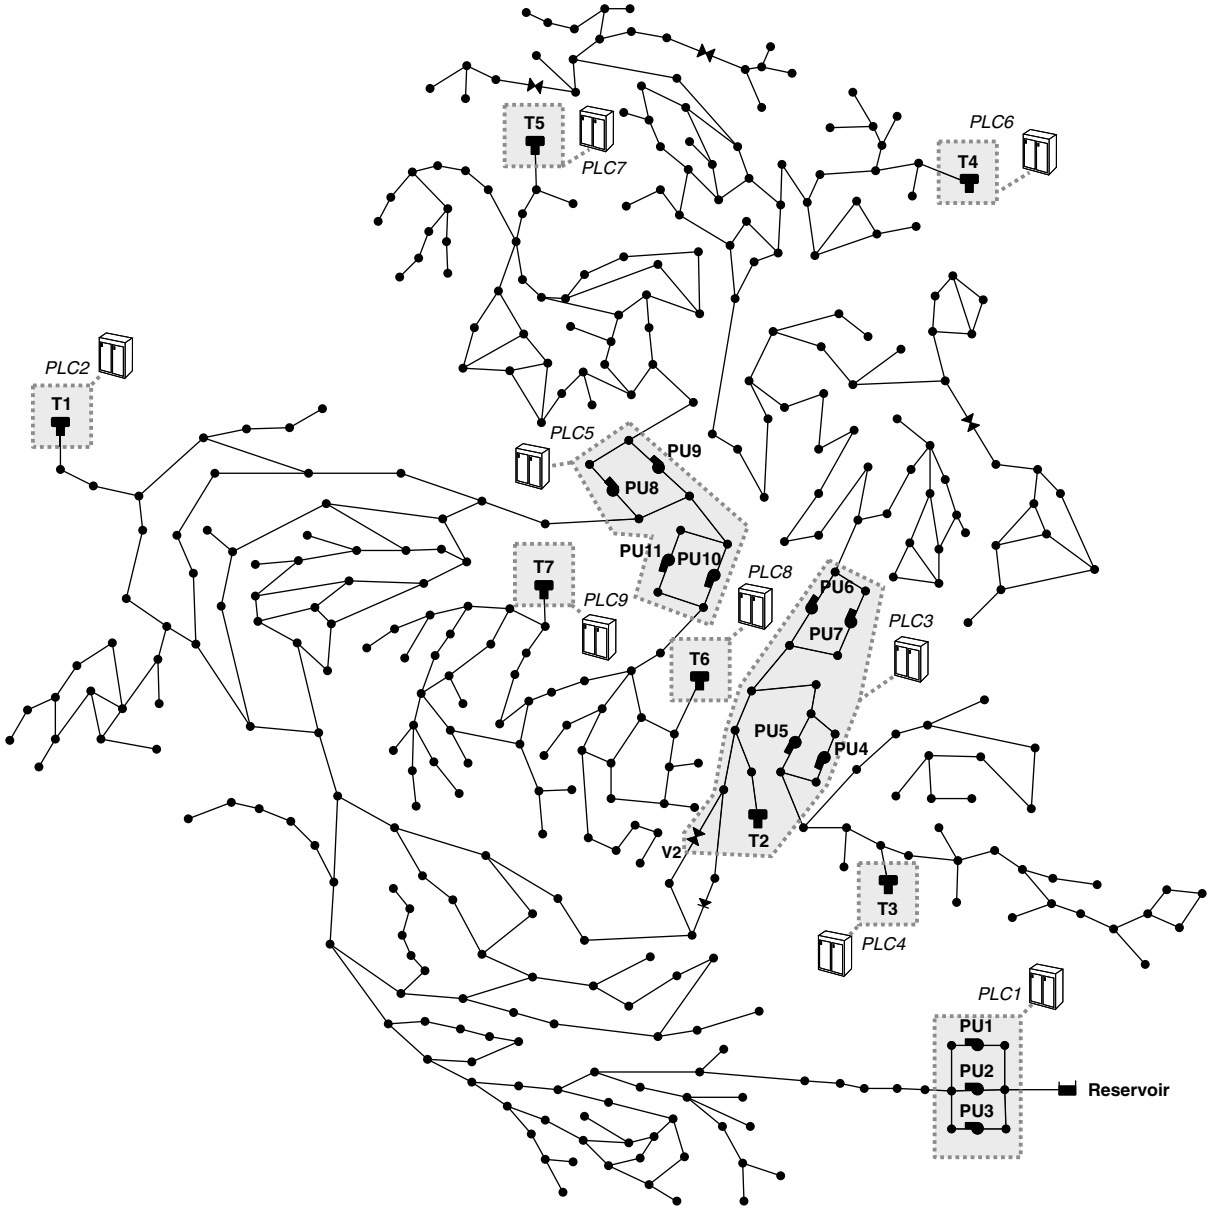

Fig. 1. C-Town water distribution system. (Adapted from Taormina et al. 2017.)

Calibration Network (Ostfeld et al. 2012). The network consists of 429 pipes, 388 junctions, 7 storage tanks, 11 pumps (distributed across 5 pumping stations), 5 valves, and a single reservoir (Fig. 1). Water consumption is fairly regular throughout the year. These physical assets are augmented with a network of 9 PLCs, which are located in proximity to pumps, storage tanks, and valves. Most of the PLCs controlling the pumps receive the information needed by the control logic from other PLCs—for instance, PLC1 controls Pumps PU1 and PU2 on the basis of Tank T1 water level, which is monitored by PLC2 (Table 1). PLCs controlling pumps and valves record information on the device status (on/off or open/closed), the flow passing through it, and the inlet and outlet pressure of pumping stations. The cyber network includes a SCADA system, which coordinates the operations and stores the readings provided by the PLCs. All information regarding the distribution system were incorporated into the EPANET2 (Rossman 2000) input file C-Town.inp which was provided to the participants. Water demand in all nodes of C-Town was not shared, meaning that participants could not run the model for the same period and then compare the results with the provided SCADA data.

Development Data

Participants were provided with three data sets containing SCADA readings for 43 system variables, i.e., tank water levels (7 variables, denoted L_<tank id>), inlet and outlet pressure for one actuated valve and all pumping stations (12 variables, denoted P_<junction id>) and their flow and status (24 variables, denoted F_<actuator id> and S_<actuator id>, respectively). All variables were continuous, with the exception of the status of valve and pumps, represented by binary variables. The data sets were generated via simulation with epanetCPA, a MATLAB toolbox that allows the design of a variety of cyber attacks and the simulation, with EPANET2 (version 2.0.12), of the hydraulic response of a water distribution network (Taormina et al. 2017). The toolbox is available on GitHub (Taormina 2018). The hydraulic time step was set

Table 1. Sensors and actuators (pumps and valves) monitored/controlled by PLCs

PLC	Sensor	Actuators (controlling sensor)
PLC1	—	PU1(T1), PU2(T1)
PLC2	T1	—
PLC3	T2	V2(T2), PU4(T3), PU5(T3), PU6(T4), PU7(T4)
PLC4	T3	—
PLC5	—	PU8(T5), PU9(-), PU10(T7), PU11(T7)
PLC6	T4	—
PLC7	T5	—
PLC8	T6	—
PLC9	T7	—

Note: For each PLC, the corresponding controlling sensor provides the information needed to operate the actuators. A PLC-to-PLC connection is established whenever an actuator and the corresponding control sensor are connected to two different PLCs.

to 15 min, although the SCADA data reported to the participants were sampled with fixed hourly intervals. The first two data sets, hereafter named Training Data Set 1 and Training Data Set 2, were provided at the beginning of the competition, whereas the third (Test Data Set) was subsequently used to evaluate and rank the attack detection algorithms.

- Training Data Set 1 was generated with a simulation horizon of 365 days. A key aspect of the data set was the absence of cyber attacks, which made it suitable for studying the operations of the water distribution system under normal operating conditions.
- Training Data Set 2 contained 7 attacks, spanning 492 hourly time steps. One attack was entirely revealed to the participants (by appropriately labelling the corresponding time steps), whereas the remaining attacks were either partially revealed or hidden (Table 2). This corresponds to a postattack scenario in which forensics experts carry out an investigation to determine whether, when, and where the water distribution system has been affected.
- Test Data Set contained 7 additional attacks, spanning over 407 hourly time steps (Table 3). Naturally, no information regarding the attacks was revealed. Participants were required to run the detection algorithms on Test Data Set and submit a detection report containing the following information: number of attacks detected, start and end time of each attack (in DD-MM-YYYY HH format), and the label of the attacked device(s) (optional).

The operations of the water system were altered through malicious activation of hydraulic actuators, change of actuator settings, and deception attacks—among the most common attacks on cyber-physical systems (Cardenas et al. 2009). The latter were aimed at manipulating the information sent or received by sensors and PLCs, with the ultimate goal of affecting the operations of an actuator (Urbina et al. 2016). Deception attacks were also used to alter the information received by the SCADA system, therefore concealing the real, physical outcomes of the attacks. SCADA concealment was performed by either adding an offset to the transmitted sensor readings or by replacing actual traffic information between the PLCs and the SCADA system with previously recorded data, a type of manipulation known as a replay attack (Urbina et al. 2016). The replay attacks featured in the BATADAL consisted of replacing data for a given hour of the day with those recorded during the same hour 1 or 2 days previous. Fig. 2 illustrates Attack #3 (Training Data Set 2), in which both pump operations and SCADA data were compromised. In this case, a deception attack manipulated Tank T1 water level readings sent by PLC2–PLC1. PLC1 received a reading of 0.5 m, which was below the low-level thresholds that activate Pumps PU1 and PU2 (4 and 1 m, respectively). This resulted in both pumps working for the entire period of the attack, which lasted 60 h. Consequently, the water level in Tank T1 reached the full tank level (6.5 m), with the excess water being spilled. The adversary tried to conceal the surge in T1 water level with a second deception attack that altered the signal sent by PLC2 to the SCADA system with a time-varying offset.

Evaluation Criteria

The attack detection algorithms were evaluated by comparing the detection report submitted by each team against the provided Test Data Set. The assessment was based on two scores that accounted for (1) the time taken to detect an attack and (2) the classification

Table 2. Attacks featured in Training Data Set 2

Identifier	Starting time (DD/MM/YYYY HH)	Ending time (DD/MM/YYYY HH)	Duration (h)	Attack description	SCADA concealment	Label (h)
1	13/09/2016 23	16/09/2016 00	50	Attacker alters SCADA transmission to PLC9 and changes L_T7 thresholds determining when pumps PU10/PU11 are switched on/off. Low levels in T7.	Replay attack on L_T7.	42
2	26/09/2016 11	27/09/2016 10	24	Like Attack #1.	Like Attack #1 but replay attack extended on PU10/PU11 flow and status.	0
3	09/10/2016 09	11/10/2016 20	60	Attack alters L_T1 readings sent by PLC2 to PLC1, which reads a constant low level and keeps pumps PU1/PU2 on. Overflow in T1.	Polyline to offset L_T1 increase.	60
4	29/10/2016 19	02/11/2016 16	94	Like Attack #3.	Replay attack on L_T1, PU1/PU2 flow and status, as well as on pressure at pumps outlet (P_J269).	37
5	26/11/2016 17	29/11/2016 04	60	Working speed of PU7 reduced to 0.9 of nominal speed. Lower water levels in T4.		7
6	06/12/2016 07	10/12/2016 04	94	Like Attack #5, but speed reduced to 0.7.	Replay attack on L_T4.	73
7	14/12/2016 15	19/12/2016 04	110	Like Attack #6.	Replay attack on L_T4, as well as on PU6/PU7 flow and status.	0

Table 3. Attacks featured in Test Data Set

Identifier	Starting time (DD/MM/YYYY HH)	Ending time (DD/MM/YYYY HH)	Duration (h)	Attack description	SCADA concealment
8	16/01/2017 09	19/01/2017 06	70	Attacker gains control of PLC3 and changes L_T3 thresholds determining when pumps PU4/PU5 are switched on/off. Low levels in T3.	Replay attack on L_T3, as well as on PU4/PU5 flow and status.
9	30/01/2017 08	02/02/2017 00	65	Attack alters L_T2 readings arriving to PLC3, which reads a low level and keeps valve V2 OPEN. Attack leads T2 to overflow.	Polyline to offset L_T2 increase.
10	09/02/2017 03	10/02/2017 09	31	Malicious activation of pump PU3	—
11	12/02/2017 01	13/02/2017 07	31	Similar to Attack #10	—
12	24/02/2017 05	28/02/2017 08	100	Similar to Attack #9	Replay attack on L_T2, V2 flow and status, as well as on V2 inlet and outlet pressure readings (P_J14, P_J422)
13	10/03/2017 14	13/03/2017 21	80	Attacker gains control of PLC5 and changes the L_T7 thresholds determining when pumps PU10/PU11 are switched on/off. The pumps are forced to switch on/off continuously during attack.	Replay attack on L_T7, PU10/PU11 flow and status, as well as on pumps inlet and outlet pressure readings (P_J14, P_J422). Inlet pressure concealment terminates before that of other variables.
14	25/03/2017 20	27/03/2017 01	30	Alteration of T4 signal arriving to PLC6. Overflow in T6.	—

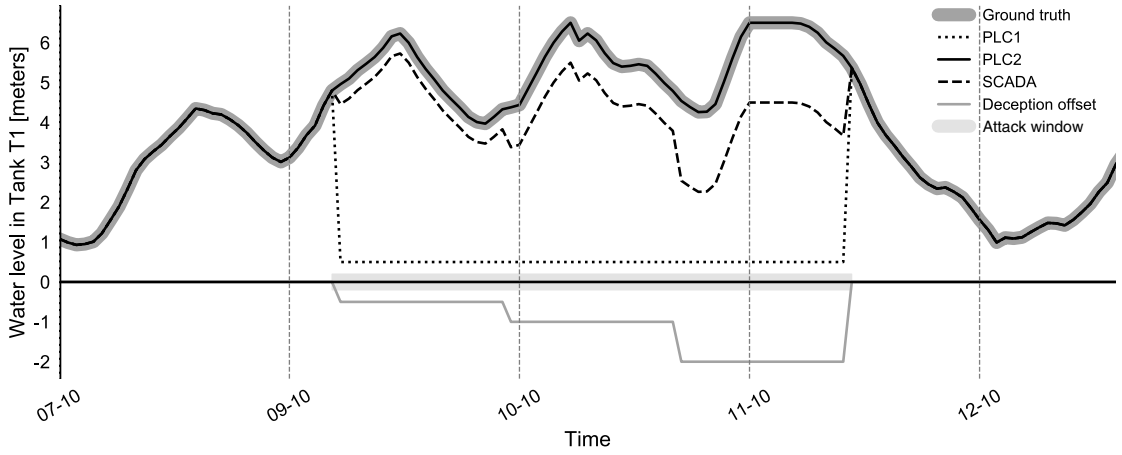

Fig. 2. Attack #3 (from Training Data Set 2). The attacker alters Tank T1 water level readings (solid line) sent by PLC2 to PLC1, which reads a constant low level (dotted line) and keeps Pumps PU1/PU2 on. This causes an overflow in Tank T1 (thick darkly shaded line). To conceal the action, the attacker alters the signal sent by PLC2 to the SCADA system (dashed line) by adding a time-varying offset (thin shaded line). The duration of the entire attack is highlighted by the thick lightly shaded line on the horizontal axis.

accuracy. The two scores were eventually combined into an overall ranking score.

Time-to-Detection

The time-to-detection (TTD) is the time needed by an algorithm to disclose a threat. It is defined as the difference between the time t_d at which the attack is detected and the time t_0 at which the attack starts

$$\text{TTD} = t_d - t_0 \tag{1}$$

The value of t_d is inferred from the detection report, and it corresponds to the first time stamp flagged as under attack while the attack is ongoing. The lower the value of TTD, the better the algorithm performs. If an attack is detected

$$0 \leq \text{TTD} \leq \Delta t \tag{2}$$

where Δt = total duration of the attack. If the attack is not detected while it is ongoing (or at all), $\text{TTD} = \Delta t$. To facilitate the comparison of all algorithms under different attack scenarios, the following performance score (S_{TTD}) was computed:

$$S_{\text{TTD}} = 1 - \frac{1}{n_a}\sum_{i}^{n_a}\frac{\text{TTD}_i}{\Delta t_i} \tag{3}$$

where n_a = number of attacks contained in a data set; TTD_i = time-to-detection of the ith attack; and Δt_i = corresponding duration. The value of S_{TTD} varies between 0 and 1, with $S_{\text{TTD}} = 1$ being the ideal case in which all attacks are immediately detected and $S_{\text{TTD}} = 0$ being the case in which none of the attacks are detected.

Classification Performance

The accuracy of an algorithm was determined as its ability to disclose threats without raising false alarms. In the context of binary classification problems such as the BATADAL, the ability

to identify threats is generally assessed with the true positive rate (TPR), also known as recall or sensitivity, which is defined as

$$\text{TPR} = \frac{\text{TP}}{\text{TP} + \text{FN}} \tag{4}$$

where TP and FN = number of true positives and false negatives, respectively. In other words, the true positive rate is the ratio between the number of time steps correctly classified as under attack and the total number of time steps during which the system is under attack.

The ability to avoid false alarms is measured with the true negative rate (TNR), or specificity, defined as

$$\text{TNR} = \frac{\text{TN}}{\text{FP} + \text{TN}} \tag{5}$$

where FP and TN = number of false positives and true negatives, respectively. The true negative rate is thus the ratio between the number of time steps correctly classified as safe conditions and the total number of time steps during which the system is in safe conditions.

To ease the comparison across all algorithms, the true positive and true negative rates were combined into a single classification performance score (S_{CLF}), defined as the mean of TPR and TNR

$$S_{\text{CLF}} = \frac{\text{TPR} + \text{TNR}}{2} \tag{6}$$

This score accounts for both correct detection and false alarms, so it is suited for binary classification problems in which the sample distribution is biased toward one of the two classes—i.e., safe conditions, in the BATADAL. The value of S_{CLF} varies between 0 and 1, with 1 representing a perfect classification.

Ranking Score

The time-to-detection and accuracy scores were finally merged into an overall ranking score (S), defined as

$$S = \gamma \cdot S_{\text{TTD}} + (1 - \gamma) \cdot S_{\text{CLF}} \tag{7}$$

where $\gamma (0 \le \gamma \le 1)$ determines the relative importance of the two evaluation scores. The coefficient γ was set to 0.5 for the analysis reported subsequently, so early detection and accurate classification were equally weighed. A naïve detection mechanism that predicts the system to be always in safe conditions receives a score S equal to 0.25 ($S_{\text{TTD}} = 0$, $S_{\text{CLF}} = 0.5$). On the other hand, flagging the system as always under attack yields a value of S equal to 0.75 ($S_{\text{TTD}} = 1$, $S_{\text{CLF}} = 0.5$). This reflects the fact that S is intrinsically biased toward attack identification, because the consequences of failing to disclose an attack are deemed more costly than issuing false alarms. These naïve detection methods have the same value of S_{CLF} (0.5), yet TPR and TNR are equal to 0 and 1 in the first case, and to 1 and 0 in the second case. This highlights the contrasting nature of the two components of S_{CLF} and suggests how increased sensitivity may come at the cost of issuing more false alarms (and vice versa). Similarly, a potential conflict seems to exist between ensuring a timely detection of the attacks (high S_{TTD}) and issuing few false alarms, as pointed out by Housh and Ohar (2017c).

Attack Detection Algorithms

Seven teams participated in the BATADAL. This section briefly describes each team's attack detection algorithm.

- Aghashahi et al. (2017) adopted a two-stage method that first extracts a four-dimensional feature vector from the observed (multidimensional) time series data and then constructs a classifier to detect attacks. In the first stage, the periods of attack/no attack are used to extract four features that capture information about the covariance and mean structure. For every time instance, a local neighborhood is used to construct estimates of mean and covariance. In the second stage, a supervised classification technique, random forests (Breiman 2001) is used to classify the system state as safe or under attack.
- Brentan et al. (2017) reduced the dimensionality of the problem by exploiting the division of the C-Town network in district metered areas (DMAs). For each DMA, they used data on normal operating conditions to create recurrent neural networks that forecast tank water levels as a function of pump flow, upstream pressure (of the corresponding pump station), and hour of the day (Díaz et al. 2016). A statistical control process identifies abrupt changes in the neural network error time series when the latter are applied to data containing cyber attacks (Guralnik and Srivastava 1999). The rationale behind this approach is that it is plausible to expect an increase in the error time series when the system is under attack, because all neural networks are trained with data pertaining to normal operations.
- Chandy et al. (2017) developed two detection models running sequentially. The first uses features of the SCADA data (e.g., combined flow of pump stations or volume pumped and stored) to check whether physical and/or operating rules have been violated (e.g., tank levels within the bounds or hydraulic relationships between nodes hold). The outcome of this model is a set of flagged events, which are confirmed by the second model. The latter is a convolutional variational autoencoder—belonging to the family of deep learning methods (Kingma and Welling 2013; Doersch 2016)—that calculates the reconstruction probability of the data: the lower the probability, the higher the chance of the data being anomalous.
- Giacomoni et al. (2017) proposed two detection methods. The first verifies the integrity of the actuator rules and SCADA data by (1) checking whether the SCADA readings are consistent with the actuator rules defined for the water distribution system and (2) comparing the data for all variables to identify values falling below or above thresholds created by analyzing data corresponding to normal operating conditions. The second method builds on unveiling low-dimensionality components in the available data as well as the sparse nature of anomalies, thereby facilitating the separation of anomalies from the overall data. The separation of data into normal and anomalous components can be performed using prinicipal component analysis (PCA) (Lakhina et al. 2004) or a convex optimization routine (Mardani et al. 2013). The results reported subsequently for Giacomoni et al. (2017) correspond to the second detection method based on PCA.
- Abokifa et al. (2017) introduced a three-stage detection method, with each stage targeting a specific class of anomalies. The first step features outlier detection techniques to find statistical outliers in the data, thereby focusing on local anomalies that affect each sensor individually. The second stage employs an artificial neural network—in the form of a multilayer perceptron—to detect contextual anomalies that do not conform to normal operating conditions. The third stage targets global anomalies that simultaneously affect multiple sensors. To disclose these anomalies, the layer uses principal component analysis to decompose the high-dimensional data sets of sensor measurements into two subspaces representing normal and anomalous conditions (Lee et al. 2013).
- Pasha et al. (2017) presented an algorithm consisting of three main interconnected modules working on control rules and

consistency checks, pattern recognition, and hydraulic and system relationships. The first module checks the consistency of the data against the set of control rules characterizing the water system, whereas the second uses statistical analysis to identify patterns for single hydraulic parameters and combinations thereof. The idea is that patterns under cyber attacks may not follow the original patterns. The anomalous behaviors detected by the first two modules are finally confirmed by the third module, which develops relationships for some physical quantities (e.g., tank levels or flows) and compares their estimates against those reported by the first two modules.

- Housh and Ohar (2017b) proposed a model-based approach that uses EPANET to simulate the hydraulic processes of the water distribution systems, and then uses the error between EPANET-simulated values and the available SCADA readings to detect anomalous behaviors. The approach consists of three main steps: first, available SCADA readings are used in a mixed-integer linear program to estimate the water demand in all nodes of C-Town; second, EPANET is used to generate reference values for the SCADA readings which are used to produce simulation errors when compared with actual readings; and third, a multilevel classification approach is implemented to classify the obtained simulation errors into event and normal conditions. A similar approach was successfully developed by Housh and Ohar (2017a) to detect contamination events in water distribution systems.

Results

Algorithm Performance

Table 4 reports the values of the ranking, time-to-detection, and classification scores (S, S_{TTD}, and S_{CLF}) obtained by the competing algorithms on the test data set. The table also reports the number of attacks detected, the values of TPR and TNR yielding the classification score, and the elements of the confusion matrix (i.e., TP, FP, TN, and FN). The scatter plot of Fig. 3 visually compares S, S_{TTD}, and S_{CLF}.

Fig. 3 highlights a cluster of four high-performing algorithms, all achieving a ranking score S greater than or close to 0.90. The group is led by the algorithm proposed by Housh and Ohar (2017b), which showed the best overall performance ($S = 0.970$). This algorithm was the top scorer in terms of both time-to-detection, S_{TTD}, and classification score, S_{CLF}. Indeed, the detection trajectory in Fig. 4(a) shows that all attacks were immediately detected, with the exception of the last one, which was disclosed a few hours after its starting time. The algorithm of Abokifa et al. (2017) was a close second, with $S = 0.949$. This method was almost as quick as that of Housh and Ohar (2017b) in identifying the attacks, but it was more prone to false alarms. Abokifa et al.'s (2017) algorithm disclosed Attacks #10 and #11 as a single continuous episode, erroneously flagging the system as under attack for the period between them [Fig. 4(b)]. The algorithm proposed by Giacomoni et al. (2017) had the same TNR as that of Housh and Ohar (2017b), meaning that both algorithms were the most successful in avoiding false alarms. However, Giacomoni et al.'s (2017) algorithm was less sensitive, resulting in lower TPR and minor timing errors [Fig. 4(c)] that led to a score of 0.927. With $S = 0.896$, the algorithm proposed by Brentan et al. (2017) can also be regarded as a strong performer. This algorithm consistently and accurately detected most of the attacks, but it failed to identify the last one [Fig. 4(d)].

Although outdistanced by the leading group, the contributions of Chandy et al. (2017) and Pasha et al. (2017) were still sensibly better than the naïve detection mechanisms described in section "Problem Description." Their scores were 0.802 and 0.773, respectively. Figs. 4(e and f) show that these two detection algorithms appear to suffer from opposite problems. The algorithm of Chandy et al. (2017) was oversensitive—meaning that it was able to identify most of the attack instances, but at the cost of issuing numerous false alarms. This is reflected in a relatively high value of the TPR, which, however, coincided with the lowest overall value of the TNR. On the other hand, the algorithm of Pasha et al. (2017) issued just a few false alarms, but it lacked sensitivity, thus failing to flag the system as under attack for the entire duration of events. This resulted in a very high value of the TNR and the overall lowest TPR. Finally, the contribution of Aghashahi et al. (2017) detected only three attacks, leading to a score of 0.534.

General Observations

The main insights from the results are summarized as follows:

- All algorithms but one achieved a ranking score S greater than 0.75, meaning that they performed better than naïve detection mechanisms. However, their performance varied widely.
- Both time-to-detection and classification score are important aspects of performance. Logically, the algorithms that performed consistently well for both metrics achieved a higher ranking score. There appears to be a strong correlation between these two metrics for most of the proposed algorithms (Fig. 3).
- Interestingly, the BATADAL was won by the only model-based approach. The idea of estimating the water demands to simulate system dynamics with EPANET and then measuring the errors with respect to the SCADA readings proved successful. In this regard, it is important to note that the BATADAL demand patterns were fairly regular and consistent across the three data sets. Similarly, the participants were given the same computational model of the C-Town network that was used to generate the SCADA data (i.e., the input file C-Town.inp). Therefore, successful application of this approach in real-world settings might

Table 4. Attack detection algorithms ranked by overall ranking score (S), assessed in terms of number of attacks detected; time-to-detection (S_{TTD}); accuracy (S_{CLF}); true positive ratio (TPR); true negative ratio (TNR); and number of true positives (TP), false positives (FP), true negatives (TN), and false negatives (FN)

Rank	Team	Number of attacks detected	S	S_{TTD}	S_{CLF}	TPR	TNR	TP	FP	TN	FN
1	Housh and Ohar	7	0.970	0.965	0.975	0.953	0.997	388	5	1,677	19
2	Abokifa et al.	7	0.949	0.958	0.940	0.921	0.959	375	69	1,613	32
3	Giacomoni et al.	7	0.927	0.936	0.917	0.838	0.997	341	5	1,677	66
4	Brentan et al.	6	0.894	0.857	0.931	0.889	0.973	362	45	1,637	45
5	Chandy et al.	7	0.802	0.835	0.768	0.857	0.678	349	541	1,141	58
6	Pasha et al.	7	0.773	0.885	0.660	0.329	0.992	134	14	1,668	273
7	Aghashahi et al.	3	0.534	0.429	0.640	0.396	0.884	161	195	1,487	246

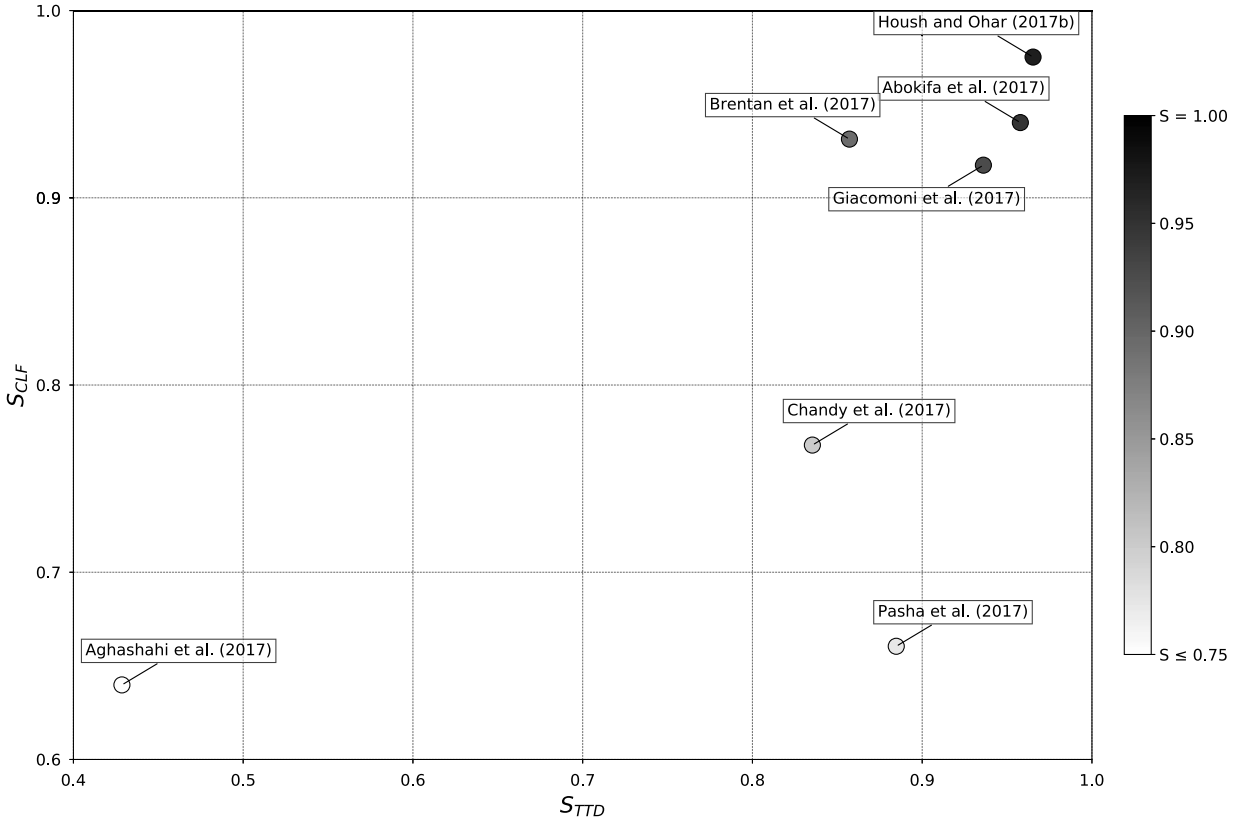

Fig. 3. Algorithm performance, measured in terms of time-to-detection (S_{TTD}, horizontal axis), classification performance (S_{CLF}, vertical axis), and overall ranking score (S, gradient bar).

be hindered by various factors, such as the intrinsic variability of demand patterns, key uncertainties in the hydraulic model (e.g., actual status of each component, pipe roughness, or pump performance curves), or the unavailability of a reliable system model.

- Three data-driven algorithms belonged to the cluster of high-performing detection mechanisms. This indicates that both model-based and data-driven approaches may be suitable for attack detection problems, although their performance would probably vary with the modeling context at hand.
- Only a few algorithms provided information on the attacked devices. Among these, the algorithms proposed by Brentan et al. (2017) and Giacomoni et al. (2017) were the most accurate.
- Most teams presented multistage detection methods. Comparing and confirming the detection issued by different modules can help decrease classification errors.
- Detection algorithms adopting a multivariate approach may be better suited than algorithms analyzing a single time series per time. The inherent interdependence of the elements in the water network should theoretically allow for the detection of anomalies even when adversaries try to conceal their actions by altering the SCADA readings of one or a few deployed sensors. Such interdependence generally presents a nonlinear nature, which can be well described by nonlinear models such as those belonging to the class of artificial neural networks.
- The adoption of supervised classification algorithms that learn how to classify the system state as either safe or under attack may not be ideal, because the number of attacks in the available data is generally limited. Supervised classification algorithms should always be combined with cross-validation schemes.
- It appears that consistency checks and the analysis of control rules should lead to the identification of the simplest attacks.

The results described previously were obtained using three specific data sets, which represent only a small portion of the entire set of cyber attacks that could threaten a water distribution system. Hence, the generation of different attacks is likely to produce different results—a limitation observed in other battles (e.g., Ostfeld et al. 2008).

Another factor that influenced the BATADAL results relates to the evaluation criteria. First, the time-to-detection score S_{TTD} was based on the ratio between the time taken to detect an attack and the attack duration; this implies that a 2-h attack detected within 1 h would have the same score as a 10-h attack detected in Hour 5. Some operators may prefer to define scores that account explicitly for the absolute value of the attack duration or its corresponding damage. Second, the classification performance score S_{CLF} is based on TPR and TNR, which are common metrics for classification problems. However, other metrics may be adopted, such as the $F1$ score (Sokolova and Lapalme 2009). Third, time-to-detection and classification performance scores were given the same importance [the coefficient $\gamma = 0.50$ in Eq. (7)]. Depending on the problem at hand, it may be desirable to overweight the time-to-detection or the classification accuracy.

Future Research Directions

The BATADAL highlighted the following gaps that may need additional research efforts:

- Robustness analysis: As mentioned previously, the performance of an attack detection algorithm may depend to a certain extent on the data used during the calibration and validation process. To limit the impact of data when evaluating the robustness of an algorithm, it is thus advisable to generate stochastic simulation

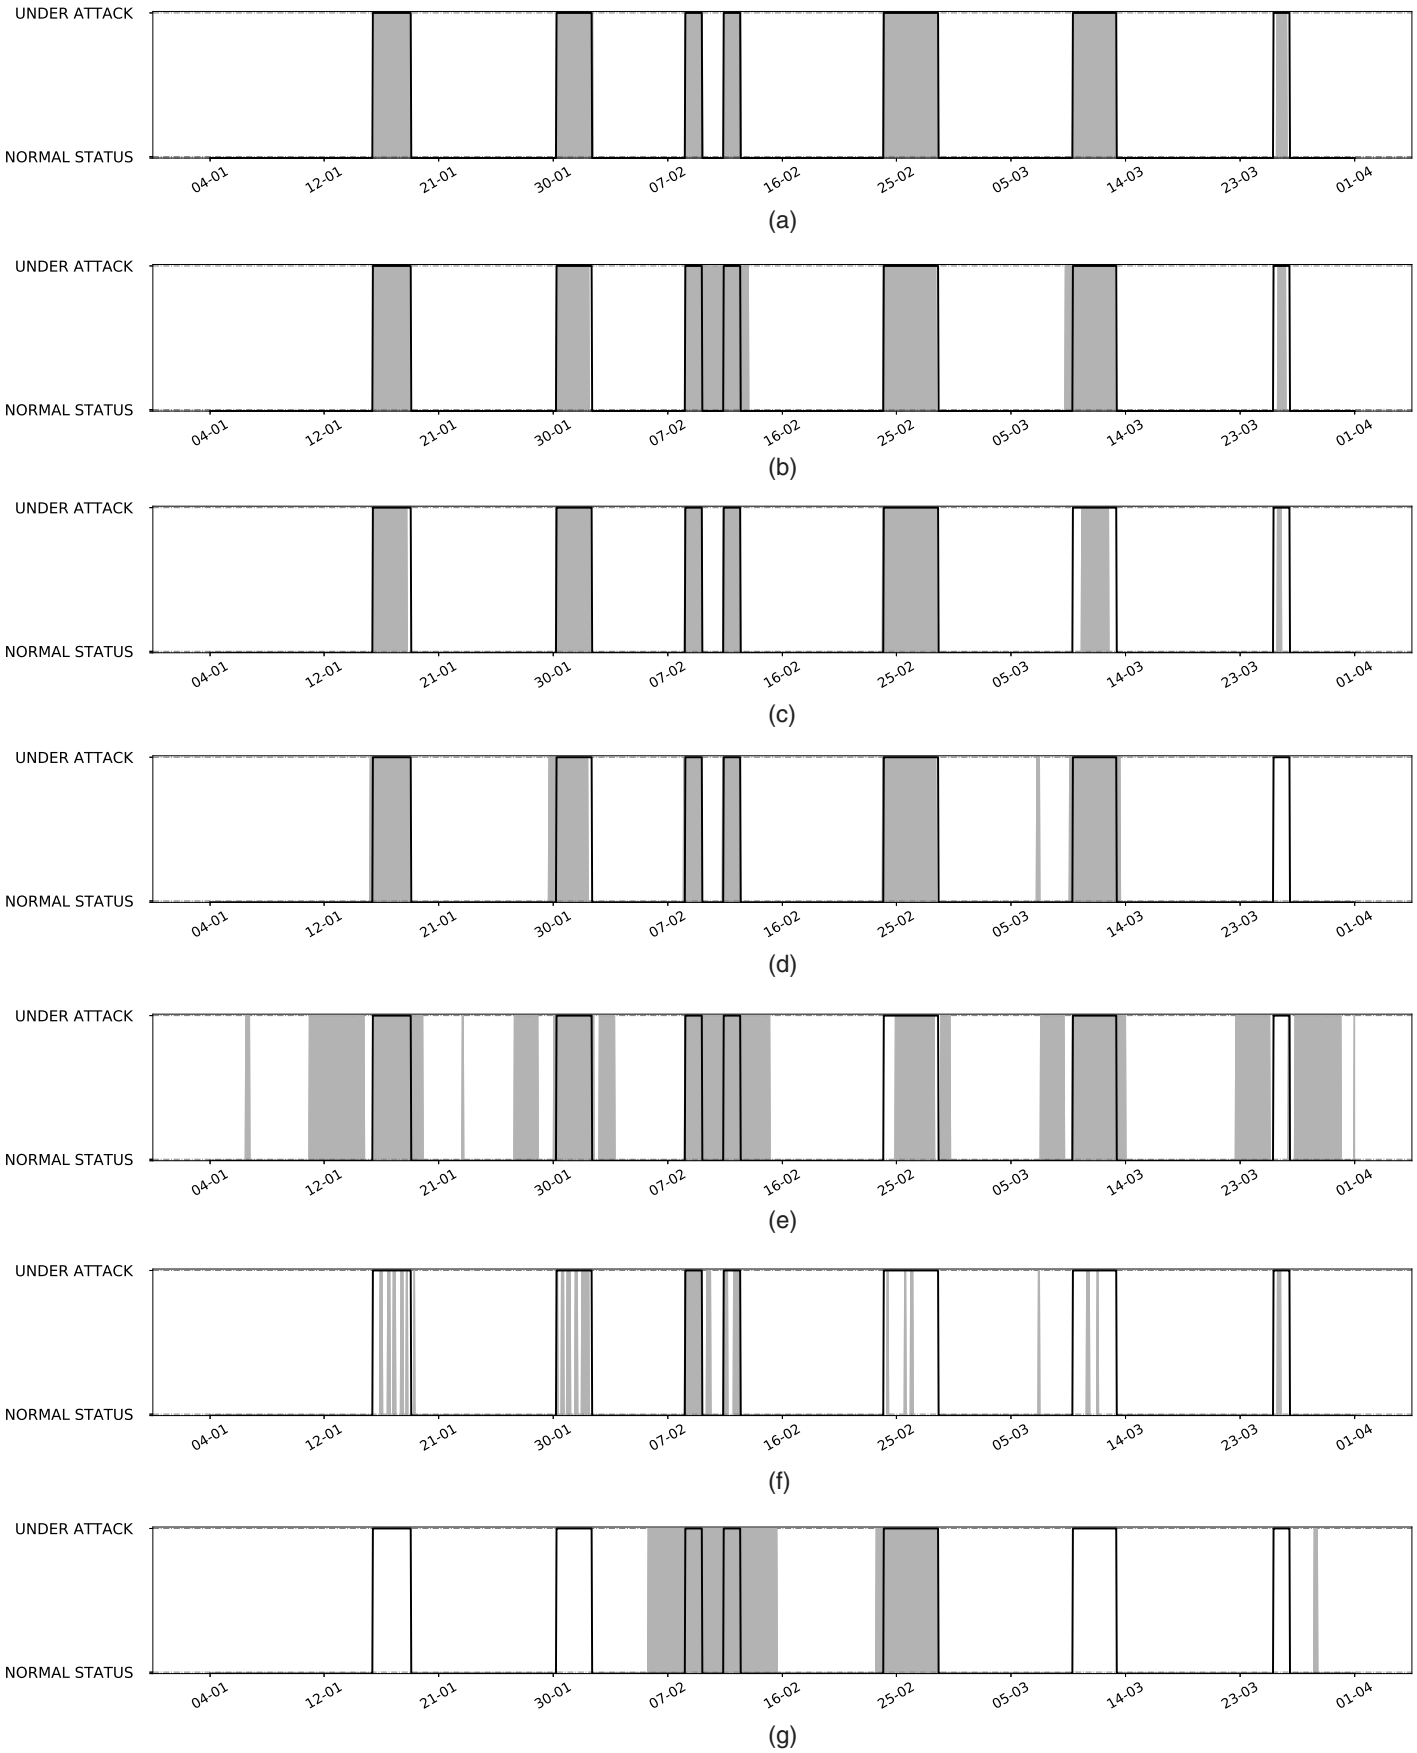

Fig. 4. Comparison of actual and detected attacks (shaded area and solid line, respectively) for the Test Data Set. Each panel corresponds to a different attack detection algorithm: (a) Housh and Ohar (2017b); (b) Abokifa et al. (2017); (c) Giacomoni et al. (2017); (d) Brentan et al. (2017); (e) Chandy et al. (2017); (f) Pasha et al. (2017); and (g) Aghashahi et al. (2017).

scenarios comprising varying hydraulic conditions (i.e., water demand or initial tank levels) and multiple attack sequences.

- Use of real SCADA data: A major limitation of the current research in cybersecurity is the absence of detailed information about cyber attacks on water utilities (e.g., timing, compromised devices, or hydraulic response of the system). Access to such information and to the corresponding SCADA data—perhaps in some anonymized forms—would drastically enhance the understanding of the skills and limitations of detection algorithms. Another challenge with SCADA data is that they often contain noise and measurement errors, so attack detection algorithms should be coupled with data preprocessing techniques.
- Pressure deficient conditions and water quality problems: A limitation of this battle is its reliance on data generated with a demand-driven engine (Taormina et al. 2017). The range of attacks should be thus extended to include pressure-deficient conditions, water quality problems, and adversial attempts aimed at threatening emergency responses, such as firefighting operations. In the absence of real SCADA data, simulated data could be generated by combining epanetCPA with more-sophisticated hydraulic engines (e.g., Sayyed et al. 2015) or water quality models, e.g., EPANET-MSX (Shang et al. 2007).
- Sensitivity analysis: The definition of the cutoff criteria defining outliers regulates the trade-off between TPR and TNR for most of the algorithms, so there is a need to adopt or develop sensitivity analysis tools that draw the appropriate line between normal and anomalous data (Abokifa et al. 2017). This step should always precede the application of an algorithm to new data sets—or its deployment in a SCADA system.
- Computational requirements and scalability to large networks: The algorithms presented in this paper were applied to a medium-sized water distribution system comprising one SCADA system and nine PLCs. Because attack detection algorithm are meant to run in real-time, it is necessary to evaluate their computational requirements as well as their scalability to larger networks.
- Attack localization: To facilitate and hasten incident resolution, an ideal detection mechanism should be able to identify which components of the network are being attacked. This is a rather challenging task due to the intrinsic correlation among the hydraulic variables. For data-driven detection mechanisms, the task may be solved with variable (or feature) selection algorithms (Galelli et al. 2014; Karakaya et al. 2016), which identify the variables that are strongly related to the detected anomalies.
- Integration with other fault detection mechanisms: Because attack detection mechanisms aim to disclose outliers and contextual anomalies in the system behavior, they may accidentally disclose anomalous behaviors that are not necessarily caused by cyber attacks (e.g., a water level sensor reporting wrong readings or a malfunctioning pump). Hence, there is a need to disclose the nature of each problem being identified—for example, by combining the attack detection algorithms with fault detection mechanisms that monitor the operations of PLCs.
- Cost effectiveness of attack detection: In the BATADAL, the different algorithms were evaluated based on their responsiveness and classification performance. Although these metrics provide some insight into the potential benefits of deploying an attack detection mechanism, a more comprehensive evaluation is needed. For example, the damage or cost associated with each cyber-physical attack could be estimated and the corresponding cost savings guaranteed by a detection algorithm.

Closure

The BATADAL was the first battle competition dealing with the emerging topic of cyber-physical security of water distribution systems. This battle provided an opportunity to develop, test, and compare attack detection algorithms for SCADA data. The solutions provided by seven teams suggest that timely and accurate detection can be obtained by both model-based and data-driven approaches, usually made of multiple sequential stages. Although the data and algorithms presented here provide a first step toward an objective comparison of attack detection algorithms for water distribution systems, they do not represent the entire spectrum of modeling contexts that practitioners and researchers would encounter. Hence, the authors hope that the availability of a dedicated website (BATADAL 2017) will help share more data sets and case studies.

Acknowledgments

Riccardo Taormina, Stefano Galelli, and Nils Ole Tippenhauer are supported by the National Research Foundation (NRF), Singapore, under its National Cybersecurity R&D Programme (Award No. NRF2014NCR-NCR001-40). Demetrios Eliades is supported by the European Union's Horizon 2020 research and innovation program under grant agreement No. 739551 (KIOS CoE). Mohsen Aghashahi and M. Katherine Banks are supported by Qatar National Research Fund (QNRF) under the grant NPRP8-1292-2-548. B. M. Brentan, Enrique Campbell, G. Lima, D. Manzi, D. Ayala-Cabrera, M. Herrera, I. Montalvo, J. Izquierdo, and E. Luvizotto Jr. are supported CAPES and CNPq founding agencies. The work of Marcio Giacomoni, Nikolaos Gatsis, and Ahmad Taha is supported by the US National Science Foundation under Grant No. 1728629. Ahmed Abokifa, Kelsey Haddad, Cynthia Lo, and Pratim Biswas' work was carried out with the partial support from the Lucy and Stanley Lopata Endowment at Washington University in St. Louis.

Supplemental Data

Files containing BATADAL rules, an EPANET input file, problem details, and additional data are available online in the ASCE Library (www.ascelibrary.org).

References

Abokifa, A. A., K. Haddad, C. S. Lo, and P. Biswas. 2017. "Detection of cyber physical attacks on water distribution systems via principal component analysis and artificial neural networks." In *Proc., World Environmental and Water Resources Congress 2017*, 676–691. Reston, VA: ASCE.

Adepu, S., G. Mishra, and A. Mathur. 2017. "Access control in water distribution networks: A case study." In *Proc., IEEE Int. Conf. on Software Quality, Reliability and Security (QRS), 2017*, 184–191. Piscataway, NJ: IEEE.

Aghashahi, M., R. Sundararajan, M. Pourahmadi, and M. K. Banks. 2017. "Water distribution systems analysis symposium—Battle of the attack detection algorithms (BATADAL)." In *Proc., World Environmental and Water Resources Congress 2017*, 101–108. Reston, VA: ASCE.

Anderson, R. J. 2010. *Security engineering: A guide to building dependable distributed systems*. New York, NY: Wiley.

BATADAL. 2017. "The BATtle of the Attack Detection ALgorithms." Accessed May 22, 2018. http://www.batadal.net.

Breiman, L. 2001. "Random forests." *Mach. Learn.* 45 (1): 5–32. https://doi.org/10.1023/A:1010933404324.

Brentan, B. M., E. Campbell, G. Lima, D. Manzi, D. Ayala-Cabrera, M. Herrera, I. Montalvo, J. Izquierdo, and E. Luvizotto. 2017. "On-line cyber attack detection in water networks through state forecasting and control by pattern recognition." In *Proc., World Environmental and Water Resources Congress 2017*, 583–592. Reston, VA: ASCE.

Cardenas, A., S. Amin, B. Sinopoli, A. Giani, A. Perrig, and S. Sastry. 2009. "Challenges for securing cyber physical systems." In Vol. 5 of *Proc., Workshop on Future Directions in Cyber-Physical Systems Security*. Washington, DC: Department of Homeland Security.

Cardenas, A. A., S. Amin, and S. Sastry. 2008. "Secure control: Towards survivable cyber-physical systems." In *Proc., Conf. on Distributed Computing Systems Workshops (ICDCS)*, 495–500. Piscataway, NJ: IEEE.

Chandy, S. E., A. Rasekh, Z. A. Barker, B. Campbell, and M. E. Shafiee. 2017. "Detection of cyber-attacks to water systems through machine-learning-based anomaly detection in SCADA data." In *Proc., World Environmental and Water Resources Congress 2017*, 611–616. Reston, VA: ASCE.

Díaz, S., J. González, and R. Mínguez. 2016. "Uncertainty evaluation for constrained state estimation in water distribution systems." *J. Water Resour. Plann. Manage.* 142 (12): 06016004. https://doi.org/10.1061/(ASCE)WR.1943-5452.0000718.

Doersch, C. 2016. *Tutorial on variational autoencoders*. Ithaca, NY: Cornell Univ. Library.

Galelli, S., G. B. Humphrey, H. R. Maier, A. Castelletti, G. C. Dandy, and M. S. Gibbs. 2014. "An evaluation framework for input variable selection algorithms for environmental data-driven models." *Environ. Modell. Software* 62: 33–51. https://doi.org/10.1016/j.envsoft.2014.08.015.

Giacomoni, M., N. Gatsis, and A. Taha. 2017. "Identification of cyber attacks on water distribution systems by unveiling low-dimensionality in the sensory data." In *Proc., World Environmental and Water Resources Congress 2017*, 660–675. Reston, VA: ASCE.

Gong, W., M. A. Suresh, L. Smith, A. Ostfeld, R. Stoleru, A. Rasekh, and M. K. Banks. 2016. "Mobile sensor networks for optimal leak and backflow detection and localization in municipal water networks." *Environ. Modell. Software* 80: 306–321. https://doi.org/10.1016/j.envsoft.2016.02.001.

Graham, J., R. Olson, and R. Howard. 2016. *Cyber security essentials*. Boca Raton, FL: CRC.

Guralnik, V., and J. Srivastava. 1999. "Event detection from time series data." In *Proc., 5th ACM SIGKDD Int. Conf. on Knowledge Discovery and Data Mining*, 33–42. New York: ACM.

Hayes, M. A., and M. A. Capretz. 2015. "Contextual anomaly detection framework for big sensor data." *J. Big Data* 2 (1): 2. https://doi.org/10.1186/s40537-014-0011-y.

Hill, D., B. Kerkez, A. Rasekh, A. Ostfeld, B. Minsker, and M. K. Banks. 2014. "Sensing and cyberinfrastructure for smarter water management: The promise and challenge of ubiquity." *J. Water Resour. Plann. Manage.* 140 (7): 01814002. https://doi.org/10.1061/(ASCE)WR.1943-5452.0000449.

Housh, M., and Z. Ohar. 2017a. "Integrating physically based simulators with event detection systems: Multi-site detection approach." *Water Res.* 110: 180–191. https://doi.org/10.1016/j.watres.2016.12.003.

Housh, M., and Z. Ohar. 2017b. "Model based approach for cyber-physical attacks detection in water distribution systems." In *Proc., World Environmental and Water Resources Congress 2017*, 727–736. Reston, VA: ASCE.

Housh, M., and Z. Ohar. 2017c. "Multiobjective calibration of event-detection systems." *J. Water Resour. Plann. Manage.* 143 (8): 06017004. https://doi.org/10.1061/(ASCE)WR.1943-5452.0000808.

ICS-CERT (Industrial Control Systems-Cyber Emergency Response Team). 2014. *NCCIC/ICS-CERT year in review: FY 2013*. Rep. No. 13-50369. Washington, DC: US Dept. of Homeland Security.

ICS-CERT (Industrial Control Systems-Cyber Emergency Response Team). 2015. *NCCIC/ICS-CERT year in review: FY 2014*. Rep. No. 14-50426. Washington, DC: US Dept. of Homeland Security.

ICS-CERT (Industrial Control Systems-Cyber Emergency Response Team). 2016. *NCCIC/ICS-CERT year in review: FY 2015*. Rep. No. 15-50569. Washington, DC: US Dept. of Homeland Security.

Karakaya, G., S. Galelli, S. D. Ahipaşaoğlu, and R. Taormina. 2016. "Identifying (quasi) equally informative subsets in feature selection problems for classification: A max-relevance min-redundancy approach." *IEEE Trans. Cybern.* 46 (6): 1424–1437. https://doi.org/10.1109/TCYB.2015.2444435.

Kingma, D. P., and M. Welling. 2013. *Auto-encoding variational bayes*. Ithaca, NY: Cornell Univ. Library.

Lakhina, A., M. Crovella, and C. Diot. 2004. "Diagnosing network-wide traffic anomalies." In *Proc., 2004 Conf. on Applications, Technologies, Architectures, and Protocols for Computer Communications, SIGCOMM '04*, 219–230. New York: Association for Computing Machinery's Special Interest Group on Data Communications.

Lee, Y.-J., Y.-R. Yeh, and Y.-C. F. Wang. 2013. "Anomaly detection via online oversampling principal component analysis." *IEEE Trans. Knowl. Data Eng.* 25 (7): 1460–1470. https://doi.org/10.1109/TKDE.2012.99.

Mardani, M., G. Mateos, and G. B. Giannakis. 2013. "Recovery of low-rank plus compressed sparse matrices with application to unveiling traffic anomalies." *IEEE Trans. Inf. Theory* 59 (8): 5186–5205. https://doi.org/10.1109/TIT.2013.2257913.

Ostfeld, A., et al. 2008. "The battle of the water sensor networks (BWSN): A design challenge for engineers and algorithms." *J. Water Resour. Plann. Manage.* 134 (6): 556–568. https://doi.org/10.1061/(ASCE)0733-9496(2008)134:6(556).

Ostfeld, A., et al. 2012. "Battle of the water calibration networks." *J. Water Resour. Plann. Manage.* 138 (5): 523–532. https://doi.org/10.1061/(ASCE)WR.1943-5452.0000191.

Pasha, M. F. K., B. Kc, and S. L. Somasundaram. 2017. "An approach to detect the cyber-physical attack on water distribution system." In *Proc., World Environmental and Water Resources Congress 2017*, 703–711. Reston, VA: ASCE.

Rasekh, A., A. Hassanzadeh, S. Mulchandani, S. Modi, and M. K. Banks. 2016. "Smart water networks and cyber security." *J. Water Resour. Plann. Manage.* 142 (7): 01816004. https://doi.org/10.1061/(ASCE)WR.1943-5452.0000646.

Rossman, L. A. 2000. *EPANET 2 users manual*. Washington, DC: US Environmental Protection Agency.

Sayyed, M. A. H. A., R. Gupta, and T. T. Tanyimboh. 2015. "Noniterative application of EPANET for pressure dependent modelling of water distribution systems." *Water Resour. Manage.* 29 (9): 3227–3242. https://doi.org/10.1007/s11269-015-0992-0.

Shang, F., J. G. Uber, and L. A. Rossman. 2007. "Modeling reaction and transport of multiple species in water distribution systems." *Environ. Sci. Technol.* 42 (3): 808–814. https://doi.org/10.1021/es072011z.

Sokolova, M., and G. Lapalme. 2009. "A systematic analysis of performance measures for classification tasks." *Inf. Process. Manage.* 45 (4): 427–437. https://doi.org/10.1016/j.ipm.2009.03.002.

Sønderlund, A. L., J. R. Smith, C. J. Hutton, Z. Kapelan, and D. Savic. 2016. "Effectiveness of smart meter-based consumption feedback in curbing household water use: Knowns and unknowns." *J. Water Resour. Plann. Manage.* 142 (12): 04016060. https://doi.org/10.1061/(ASCE)WR.1943-5452.0000703.

Taormina, R. 2018. "epanetCPA: A MATLAB toolbox for assessing the impacts of cyber-physical attacks on water distribution systems." Accessed May 22, 2018. https://github.com/rtaormina/epanetCPA.

Taormina, R., S. Galelli, N. O. Tippenhauer, E. Salomons, and A. Ostfeld. 2017. "Characterizing cyber-physical attacks on water distribution systems." *J. Water Resour. Plann. Manage.* 143 (5): 04017009. https://doi.org/10.1061/(ASCE)WR.1943-5452.0000749.

The Times of Israel. 2018. "Israel tech to protect NY water systems from cyberattacks." Accessed September 24, 2017. https://www.timesofisrael.com/israel-tech-to-protect-ny-water-systems-from-attack/.

Urbina, D., J. Giraldo, N. O. Tippenhauer, and A. Cárdenas. 2016. "Attacking fieldbus communications in ICS: Applications to the SWaT testbed." In *Proc., Singapore Cyber Security Conf. (SG-CRC)*. Amsterdam, Netherlands: IOS Press.

US Department of Homeland Security. 2017. "Critical infrastructure sectors." Accessed September 24, 2017. https://www.dhs.gov/critical-infrastructure-sectors.