Established in collaboration with MIT

Pressure-Driven Hydraulic Modelling of Cyber-Physical Attacks on Water Distribution Systems

Submitted by

Hunter C. Douglas

Thesis Advisor

Dr. Stefano Galelli

Pillar of Engineering Systems and Design

A thesis submitted to the Singapore University of Technology and Design in fulfillment of the requirement for the degree of Master of Engineering by Research, Pillar of Engineering Systems and Design

June 28, 2017

Declaration of Authorship

I, Hunter C. Douglas, declare that this thesis titled, "Pressure-Driven Hydraulic Modelling of Cyber-Physical Attacks on Water Distribution Systems" and the work presented in it are my own. I confirm that:

- This work was done wholly or mainly while in candidature for a research degree at this University.
- Where any part of this thesis has previously been submitted for a degree or any other qualification at this University or any other institution, this has been clearly stated.
- Where I have consulted the published work of others, this is always clearly attributed.
- Where I have quoted from the work of others, the source is always given. With the exception of such quotations, this thesis is entirely my own work.
- I have acknowledged all main sources of help.
- Where the thesis is based on work done by myself jointly with others, I have made clear exactly what was done by others and what I have contributed myself.

Signed: 

Date: 28/06/2017

"*Since all models are wrong the scientist cannot obtain a 'correct' one by excessive elaboration. On the contrary following William of Occam he should seek an economical description of natural phenomena. Just as the ability to devise simple but evocative models is the signature of the great scientist so overelaboration and overparameterization is often the mark of mediocrity.*"

George E. P. Box

SINGAPORE UNIVERSITY OF TECHNOLOGY AND DESIGN

Abstract

Pillar of Engineering Systems and Design

Master of Engineering by Research

Pressure-Driven Hydraulic Modelling of Cyber-Physical Attacks on Water Distribution Systems

by Hunter C. Douglas

Water distribution systems the world over are being augmented with sensors and logic controllers to make them run more automatically, efficiently, and reliably. However, these interconnected devices also expose such systems to greater risk of getting attacked, which can result in unexpected behaviour. As the world's water distribution systems become increasingly threatened by cyber-physical attacks, the ability to realistically simulate the hydraulic effects of these attacks has never been more important. A standard approach in hydraulic modelling is to use demand-driven analysis, where the consumers' water requirements are assumed to be met at all times. This approach does not allow for atypical scenarios, such as cyber-physical attacks, to be accurately simulated because the pressure may be insufficient to actually supply the required demand. In order to rectify this, an existing hydraulic modelling and attack toolkit, *epanetCPA*, was modified to add pressure-driven analysis capabilities. The toolkit was tested, calibrated, and verified in experiments using a real-world testbed network. It was then used to simulate a range of attack scenarios on a town-scale benchmark network model. This work showed that the toolkit can be used to realistically replicate and predict the performance of real-world networks subject to cyber-physical attacks. The findings of these simulations also have important implications for the design and operation of water distribution systems. The toolkit can thus help system operators continue to provide water to consumers despite the threat of attacks.

Acknowledgements

First and foremost, I want to thank my advisor, Dr. Stefano Galelli, and Dr. Riccardo Taormina, both of whom were absolutely crucial to this project. Not only did they produce the research and software upon which this work was based, but they were supportive guides and mentors for me all through my graduate studies. I am hugely grateful for the trust that they have placed in me and the time that they've spent to help me learn. It has been a real pleasure working as a part of the Resilient Water Systems Group and becoming friends with the other members: Debasish, Tian, Jia Yi, Hung, Gulten, Lauren, Rizwan, Christoph, and Sean.

I am indebted to Lauren Goh for carrying out the experiments using the WADI testbed and performing the initial calibration of the network model. I would also like to thank the rest of the team at the iTrust Centre under director Dr. Aditya P. Mathur for the use of their facilities. I hope that this project can help extend the impact of the Centre's research. I am grateful to Dr. Nils Ole Tippenhauer for generously taking the time to review this thesis.

A big thanks to Dr. Ami Preis and his team at Visenti, Pte. Ltd. for taking me onboard as a research intern. I really enjoyed investigating the applications of my research to real-world problems and learnt a lot through your personal guidance. Thank you for providing me with this fruitful opportunity; I hope that my work can be of service to your own projects.

Of course, none of this would have been possible without support from the MIT-SUTD Dual Masters' Programme. I would like to thank all the people at MIT and SUTD who run this programme that afforded me the incredible opportunity to study at two amazing universities. It has been a life-changing two years. A special thank you to Ai Ling Ong and her colleagues at the Office of Graduate Studies for helping to make my time in Singapore so enjoyable.

Finally, I couldn't have done any of this without the love and support of my family in New Zealand.

Contents

List of Figures

List of Tables

List of Abbreviations

DDA	**D**emand-**D**riven **A**nalysis
PDA	**P**ressure-**D**riven **A**nalysis
HFR	**H**ead-**F**low **R**elationship
DSR	**D**emand **S**atisfaction **R**atio
CMH	**C**ubic **M**etres per **H**our
LPS	**L**itres **P**er **S**econd
RMSE	**R**oot **M**ean **S**quare **E**rror

Chapter 1

Introduction

This chapter provides a background to the work in this thesis, describing the problems that the work seeks to address. It also provides an outline, a brief description of the work contained in each chapter.

1.1 Motivation

In 2008, the U.S. National Academy of Engineering released its list of 14 Engineering Grand Challenges: 14 problems facing world in the 21^{st} Century that engineering can help to solve. Number seven on that list? Restore and improve urban infrastructure. For years, engineers have been using the latest developments in sensor and communications technologies to improve our infrastructure, but in doing so we may have actually made things worse for two of the other challenges: number eight, secure cyberspace, and number nine, provide access to clean water. By connecting the infrastructure that provides clean water to communications networks, this critical lifeblood is suddenly subject to the same threats as our email accounts, credit card numbers, and photos in the cloud. Already, the U.S.-based Industrial Control Systems Cyber Emergency Response team has reported a significant number of cybersecurity breaches in the water sector, responding to 13 suspected incidents in 2013, 14 in 2014, 25 in 2015, and 18 in 2016 (ICS-CERT, 2014; ICS-CERT, 2015; ICS-CERT, 2016; U.S. Department of Homeland Security, 2017).

The potential impact of attacks on water infrastructure is severe. Attacks range from simply stealing data to damaging equipment, cutting off water supply, or even compromising water quality (Slay and Miller, 2008). Such attacks can lead to economic losses, environmental impacts, human health impacts, and even loss of life. These outcomes make water distribution systems prime targets for hostile foreign powers and terrorist organisations (Rasekh et al., 2016). That's not to mention less malevolent cyber vandals or bots who may not appreciate the possible severity of their actions. Accordingly, it is critical for the welfare of society to prepare for such attacks and prevent them from ever occurring.

There are many techniques employed by information technology professionals in order to improve cybersecurity, but there is also a role for hydraulic engineers to play. Perhaps the best way for engineers and utilities operators to prepare for cyber-physical attacks on water distribution systems is to simulate what these attacks could look like in real life. In doing so, they can better understand how to identify an attack when it happens, and how to best mitigate any damage that an attack may cause. Currently, there is a gap in

the range of tools available for performing these simulations; there is no pressure-driven model that is designed to simulate a range of attack scenarios on a range of networks. Having a pressure-driven model, as explained in Section 1.2.1 below, is crucial for accurately simulating what happens in the network when things go wrong. The purpose of this work was to develop this type of model. It is the author's hope that the tools developed over the course of this work can gain use in the water industry. Also, the resulting, more accurate estimates of the impacts of cyber-physical attacks should aid in increasing funding for and public awareness of this key issue.

1.2 Background

1.2.1 Pressure-Driven Hydraulic Modelling

Water distribution systems are the networks of pipes, tanks, reservoirs, pumps, valves, and other components that deliver drinking water to end users. These systems are not only complex but also very expensive to build and maintain. In order to design and improve water distribution systems, engineers rely on computer-based hydraulic models. Hydraulic models can simulate water flow through networks by representing the network as a series of nodes and links. The nodes can supply water to the network (as in reservoirs and emptying tanks) or "demand" water from the network (as in end users and filling tanks). The links, pipes or valves, allow flow from one node to another. Pumps that add pressure to the network can also be represented as links.

Most hydraulic models use what is known as Demand-Driven Analysis (DDA), an approach that assumes that all demands are met. In DDA, a required demand (either constant or time-varying) is set for each consumer, and that amount of water is always withdrawn from the system. This is an appropriate approach for designing a system to cope with expected demands, or for modelling the typical operations of a water distribution system, but it can fall apart when attempting to model unusual scenarios such as pipe bursts, pump failures, or withdrawals for fire-fighting (Germanopoulos, 1985; Gupta and Bhave, 1996). Such scenarios can cause the model to report unrealistically low pressures, cause tanks to run dry, or cause parts of the network to become disconnected. Depending on the type of model used, the software may crash or abort because these scenarios violate the systems of continuity equations used to model the system (e.g. Rossman, 2000). In order to better simulate these kinds of atypical scenarios, it is necessary to use a class of models that use Pressure-Driven Analysis (PDA).

In PDA, the model also has a demand pattern set for each node, but the actual amount of water delivered to each node is determined by the available head. Head can be thought of as the height to which water will flow freely at any given point in a network. It is dependent on the water pressure, elevation, and velocity, though the last term is usually small enough to be ignored in the model (Walski et al., 2003). As such, if the head at a given node is insufficient to satisfy the full demand, only a fraction of the water demanded will actually be delivered at that node. By allowing nodal outflow to drop to zero, PDA models prevent unrealistic pressures from being simulated. When properly calibrated, PDA models can better capture real-world pressure-deficient scenarios (Todini, 2003).

1.2.2 Cyber-Physical Attacks

As sensor and computing technology is improving, water distribution systems are becoming more and more automated. All sorts of interconnected devices are being added to these systems, including sensors that can measure the level of water in a tank or the pressure in a pipe, and programmable logic controllers, PLCs, which directly control actuators in pumps, valves, or other components. There are also Supervisory Control and Data Acquisition (SCADA) systems, which monitor, record, and control all of the devices across a network. In addition to making these complex systems easier to manage, these innovations can also improve service reliability, efficiency, and even water quality. However, they also expose the system to potential attacks. Any component that provides real-time readings or control of a system must be connected to a central computer, such as a SCADA system. Communications across the network provide an "attack surface" which can allow attackers access to the data sent between components and/or the components themselves (Rasekh et al., 2016). While measures such as air gaps (physically isolating the system from unsecured networks such as the Internet) can decrease the risk of attack, no system is invulnerable.

There are a variety of classes of attacks that can be employed by attackers (Taormina et al., 2017). *Eavesdropping attacks* compromise confidentiality by simply reading the data sent across a network. This data can then be used to develop more sophisticated attacks or be exploited by the attacker if it is intrinsically valuable. *Denial of service attacks* render the system unusable by preventing sensors from sending or receiving data, preventing actuators from being (de)activated, and/or preventing controllers from issuing commands. They can be achieved by, for example, overloading the communications between devices with unanticipated, large amounts of traffic. *Deception attacks* are more sophisticated; the attacker manipulates or replaces the data sent across a network. Such attacks can be used to issue unwanted commands to network components or even mask the effects of an attack by reporting "business as usual" sensor readings to the PLCs and SCADA system if the attacker has previously eavesdropped on the data stream. Deception attacks can thus be very difficult to detect, either with human operators or detection algorithms.

One relatively simple goal of a cyber-physical attack is to intentionally lower or cut off water supply to an area, depriving the population of drinking water. This could be achieved by remotely closing valves, turning off pumps, or diverting water to other outflow points. If water quality is instead targeted, the utility may have to respond to the attack by diverting or isolating contaminated water, potentially cutting off or reducing supply to part of the network (Rasekh and Brumbelow, 2014). Thus, attacks can either directly or indirectly result in pressure-deficient scenarios that require PDA in order to be accurately simulated. (Section 3.3 discusses in detail the specific benefits of using a PDA model to simulate attacks on water distribution systems.) Through using models equipped with PDA, we can simulate cyber-physical attacks and measure their hydraulic impacts more comprehensively than with traditional DDA models. This will allow operators to better prepare for any potential attacks, helping to keep populations safe.

1.3 Outline

- Chapter 2, *Literature Review*, discusses in detail much of the previous work done in: developing pressure-driven hydraulic models (Section 2.1), evaluating the performance of water distribution networks (Section 2.2), and simulating cyber-physical attacks on water distribution systems (Section 2.3).

- Chapter 3, *Description of Pressure-Driven Model*, describes the computer model that was developed over the course of this project. It begins with an explanation of the model that served as a basis for the present work, before outlining the modifications and additional features that were added in order to incorporate pressure-driven modelling. The resulting benefits are also described in detail.

- Chapter 4, *Model Testing, Calibration, and Validation*, presents the results of an experiment that applied attacks to a small-scale, physical testbed to test, calibrate, and validate the computer model. The testbed, named WADI, is described in detail, as are the experimental design and results. As is discussed in Section 4.5, very little work has previously been done to experimentally verify pressure-driven hydraulic models of water distribution systems. These experiments thus present an advancement of knowledge in the field.

- Chapter 5, *Hydraulic Effects of Cyber Attacks on a Town-Scale Network*, then explores the effects of cyber-physical attacks on a medium-sized water distribution system. A benchmark network, C-Town, (which was designed to be a realistic substitute for real-world systems) was chosen. The added capabilities of the computer model allowed for new metrics of performance to be tested, and so this work also presents an advancement of knowledge in the field.

- Chapter 6, *Conclusions*, discusses the implications of the work as a whole, including recommended directions for future research to take so that this work can be built upon and improved.

1.4 Research Output

The following products have been or will be completed as a result of the work done to fulfil requirements for the degree of Master of Engineering by Research, Pillar of Engineering Systems and Design. Items 1 and 2 are the required products, while the other items are natural extensions of the work.

1. The thesis contained herein.
2. An accompanying poster, which summarises the main findings of this thesis.
3. The pressure-driven hydraulic modelling toolkit extension that I developed. This will be made publicly available as open-access software.
4. A journal article discussing the effects of cyber-physical attacks on water distribution systems, using metrics that require pressure-driven modelling to be calculated, based primarily off of the content in Chapter 5. This is currently undergoing peer review.
5. A journal article describing the use of WADI to calibrate a pressure-driven hydraulic model, based on the work done in Chapter 4. This is currently being written.
6. An oral presentation at the 2017 EWRI World Environmental & Water Resources Congress, delivered by Dr. Riccardo Taormina, described the model and the experiments conducted in this thesis to test it.
7. Additionally, I have prepared a report for Visenti, PTE, LTD. summarising volunteer work that I have done for them. This work involved using the toolkit to aid in calibration of hydraulic models and drew on the findings of Chapter 4.

Chapter 2

Literature Review

This chapter goes into detail to discuss the state of the art of research relevant to this thesis. It is arranged into three sections that summarise previous work done by other researchers to address specific problems. The sections discuss: developing pressure-driven hydraulic models (Section 2.1), evaluating the performance of water distribution networks (Section 2.2), and simulating cyber-physical attacks on water distribution systems (Section 2.3).

2.1 Approaches to Pressure-Driven Analysis

2.1.1 Incorporating a Head-Flow Relationship

The earliest computer models to simulate water distribution networks were developed in the 1960s (Walski et al., 2003). It wasn't for a couple of decades that computational power allowed for the added complexity of PDA to be considered. Some of the earliest work in this field was done by Bhave, 1981. He proposed a simple relationship where the full demand was delivered if the pressure was above a set threshold and zero water was delivered if the pressure was below this threshold (Figure 2.1a). (We will refer to such relationships as Head-Flow Relationships (HFR), though they are also known as pressure-outflow relationships.) Germanopoulos, 1985 later proposed a more complex HFR (Figure 2.1b) that allowed for intermediate flow rates:

$$C_i = C_i^*(1 - a_i e^{-b_i P_i / P_i^*}) \tag{2.1}$$

where:
C_i = consumer outflow at node i
C_i^* = nominal consumer demand
a_i, b_i = constants for the particular node
P_i = current pressure at node i
P_i^* = pressure at which a known outflow is provided

Wagner, Shamir, and Marks, 1988 introduced PDA as part of extended-period simulations, where the operations of water distribution networks over a 200-year time period were simulated. The authors modelled pipe and pump failures and repairs as stochastic events, and recorded the response of the system to these failure events. The reliability of the network was captured by a number of metrics developed by the authors, including the annual shortfall (amount of demanded water that was not delivered), and the percentage of time spent in failure mode. By using PDA, the authors were able to quantify

the amount of water that was not delivered to customers, which would not have been possible with DDA. The simulations were carried out using SDP8, a commercial software of the time.

Unlike typical extended period analyses conducted today, where demands vary in a daily or seasonal cycle, Wagner, Shamir, and Marks, 1988 set nodal demands at constant values. This was because they were concerned solely with the long-term reliability of networks. Utilities operators, on the other hand, are also concerned with the reliability of water distribution networks in real-time, and so will employ models where demand varies on an hourly basis or even more frequently.

Wagner, Shamir, and Marks, 1988 made outflows pressure-driven by using an HFR with two thresholds: a minimum head, below which no flow would occur, and a "service" or desired head, above which the supplied flow was assumed to be constant (Figure 2.1c). That is, once the head was sufficient to supply the users' demands at that node, any additional head would not result in greater flow. Between these thresholds, the flow was assumed to be proportional to the square root of the available head. The terms in this HFR relate more closely to actual physical parameters than do those in the HFR proposed by Germanopoulos. The HFR can be represented by the following set of equations:

$$\left.\begin{aligned} q_j = q_j^{req} \end{aligned}\right\} H_j \geq H_j^{des} \tag{2.2}$$

$$\left.\begin{aligned} q_j = q_j^{req} \left(\frac{H_j - H_j^{min}}{H_j^{des} - H_j^{min}} \right)^{\frac{1}{n_j}} \end{aligned}\right\} H_j^{min} < H_j < H_j^{des} \tag{2.3}$$

$$\left.\begin{aligned} q_j = 0 \end{aligned}\right\} H_j \leq H_j^{min} \tag{2.4}$$

where:
q_j & H_j = the available flowrate and head, respectively, at node j
H_j^{des} = the desired head at node j, above which the flowrate is constant
q_j^{req} = the maximum flowrate, achieved at the desired head
H_j^{min} = the minimum head at node j, below which the flowrate is zero
n_j = a constant that defines the shape of the HFR curve, set to 2 to mimic the theoretical orifice equation (see Equation 2.8).

1988 also saw another major contribution to the development of water distribution network models by Todini and Pilati, 1988, who introduced a gradient algorithm for solving a model's system of equations. This algorithm was proposed for conducting DDA, but was later modified for conducting PDA (Todini, 2003). Water distribution network models can be thought of as a set of linear and non-linear equations that are based on the principles of conservation of mass and conservation of energy, as well as fluid dynamics equations that describe the nature of flow through pipes. The amounts of water and energy coming into and out of a junction must balance, and the values at each junction depend on the values at all of their neighbours. To complicate matters, sections of the network are often connected in loops. A model describes all of these interdependent equations and attempts to solve them in an iterative process. The algorithm introduced

by Todini & Pilati (herein referred to as the Global Gradient Algorithm) presented a reliable and efficient way to do this. The algorithm was later incorporated into EPANET, an industry-standard software package for modelling water distribution networks (Rossman, 2000).

Different HFRs were compared by Gupta and Bhave, 1996 in order to see which performed best for simulating pressure-deficient conditions. A simple five-node network with a tank and four demand nodes in series (Figure 2.2) was modelled using the HFRs proposed by Bhave, 1981, Germanopoulos, 1985, and Wagner, Shamir, and Marks, 1988, as well as two other simple relationships. The network was first modelled as a primary network and then reduced to a single node assumed to be part of a larger network. This approximates the process known as skeletonisation, where the demands of many nodes are grouped together in order to reduce the size and complexity of a water distribution network model. The authors found that the HFR proposed by Wagner, Shamir, and Marks, 1988, when properly calibrated, best matched the behaviour of the primary network.

Another HFR was introduced by Fujiwara and Ganesharajah, 1993 in their study of the reliability of water distribution networks under pressure-deficient conditions. The HFR was defined by minimum and desired head thresholds, similar to that of Wagner, Shamir, and Marks, 1988, but the curve between these heads was smooth and differentiable (Figure 2.1d). This leads to fewer convergence issues when using the optimisation algorithm in the model (Fujiwara and Li, 1998; Siew and Tanyimboh, 2010). The HFR is defined as:

$$\left.\begin{aligned} \rho(H_j) &= 1 \end{aligned}\right\} H_j \geq H_j^{des} \tag{2.5}$$

$$\left.\begin{aligned} \rho(H_j) &= \frac{(H_j - H_j^{min})^2(3H_j^{des} - 2H_j - H_j^{min})}{(H_j^{des} - H_j^{min})^3} \end{aligned}\right\} H_j^{min} < H_j < H_j^{des} \tag{2.6}$$

$$\left.\begin{aligned} \rho(H_j) &= 0 \end{aligned}\right\} H_j \leq H_j^{min} \tag{2.7}$$

where:
$\rho(H_j)$ = the fraction of demand available at node j
$H_j, H_j^{des}, H_j^{min}$ = as in Equations 2.2-2.4 above

The aforementioned studies all followed the same general approach to conducting PDA – incorporating an HFR into the energy and mass balance equations, which are then solved by the algorithm in the model. A variety of different solver algorithms can be used, such as Sequential Quadratic Programming (Ackley et al., 2001), or the Newton-Raphson Method (Kalungi and Tanyimboh, 2003). The Global Gradient Algorithm proposed by Todini and Pilati, 1988 is perhaps the most widely used (Cheung, Van Zyl, and Reis, 2005; Wu and Walski, 2006; Giustolisi, Savic, and Kapelan, 2008; Formiga and Chaudhry, 2008; Siew and Tanyimboh, 2010; Muranho et al., 2014). The method for adding an HFR to the system of equations was summarised by Todini, 2003, and was adopted with little modification in the aforementioned papers. The robustness of this

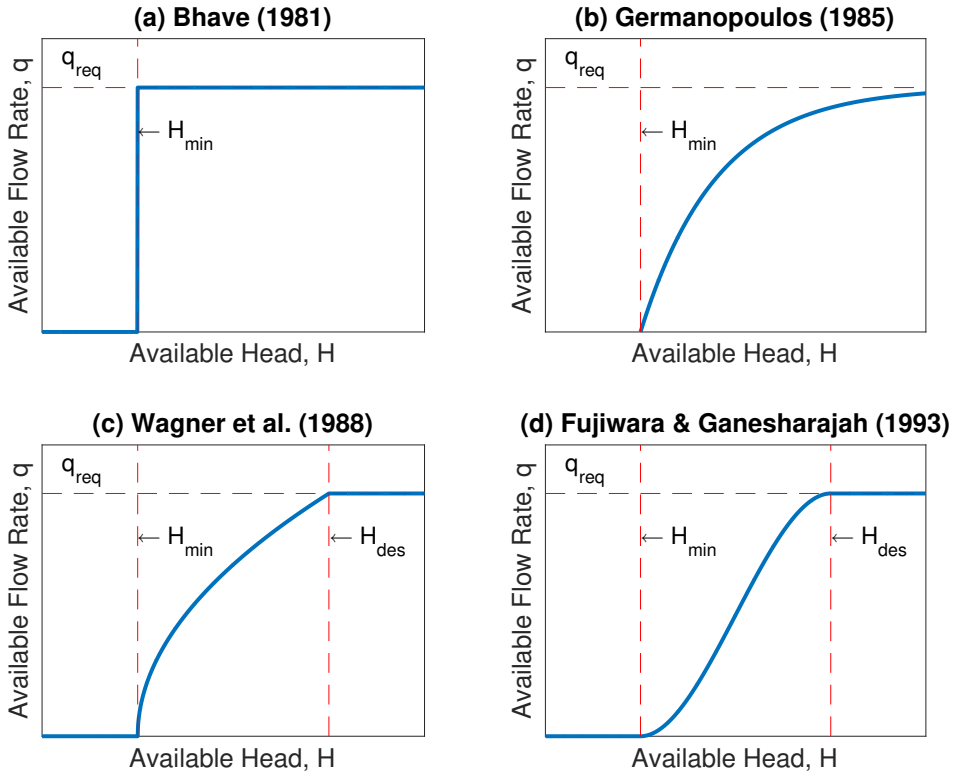

FIGURE 2.1: Illustrative head-flow relationship curves

approach depends on two factors: how realistically the HFR can reproduce real-world conditions, and how consistently the algorithm will converge to a solution. Some researchers have adopted techniques from computer science to tackle the second of these factors (Elhay et al., 2016), while comparatively little work has been done to address the first.

The challenge of verifying an HFR with real-world results has been investigated more thoroughly when considering leakage in pipes. Theoretically, flow through an open orifice is described by the following equation:

$$q = C_d A \sqrt{2gh} \tag{2.8}$$

where:
q = flowrate
C_d = discharge coefficient
A = orifice area
g = acceleration due to gravity
h = pressure head

This corresponds to the HFR used by Wagner, Shamir, and Marks, 1988, with an exponent value of 0.5. However, experimental studies and field tests have shown that for flow through leaks in pipes, the actual exponent value typically varies between 0.5 and 1.5, and can reach as high as 2.5 (Thornton and Lambert, 2005; Fu et al., 2013;

Fontana, Giugni, and Marini, 2016). The results are highly dependent on pipe material and leak geometry. When it comes to pressure-dependent demand, one must also take into account human behaviour, such as opening a tap less fully when the pressure is high. Accordingly, the exponent may be even lower than 0.5 (van Zyl and Clayton, 2007). It is likely, then, that the variables that define the HFR will vary between different networks (and nodes within a network), and that calibration of these variables is required in order to obtain accurate simulation results.

2.1.2 Adding Artificial Components

After demonstrating how the global gradient algorithm could be modified to include an HFR, Todini, 2003 also proposed a fundamentally different approach to conducting PDA, one where artificial components are added to the network. This paper described a three-step process where: 1) a traditional DDA is first conducted for the network, 2) artificial reservoirs are added to those nodes where negative pressures are encountered, and another DDA run is conducted, and then 3) the nodal demands are updated based on the results of the second model run, and a third DDA run is conducted. By adding artificial reservoirs in this way, one can avoid the complications introduced by assuming an HFR. However, because the process is iterative, it can add to the computation time required to run a simulation (Abdy Sayyed, Gupta, and Tanyimboh, 2015).

Other authors have also developed approaches to PDA based on the idea of adding artificial reservoirs. Ang and Jowitt, 2006 described another iterative approach that instead begins with setting all demands to zero and adding artificial reservoirs at all nodes where available head exceeds the minimum threshold. Artificial reservoirs are added, removed, or replaced with demand nodes based on the outcomes of repeated DDA model runs until a set of conditions are met. The authors termed this approach the pressure-dependent network algorithm (PDNA). The results obtained using PDNA are equivalent to those using the HFR proposed by Bhave (Figure 2.1a), where the desired head threshold is equal to the minimum head threshold (Morley and Tricarico, 2008; Elhay et al., 2016). The PDNA approach has subsequently been modified and extended (e.g. Sharoonizadeh, Mamizadeh, and Sarvarian, 2016).

In the artificial reservoir approach, as available head increases, the nodal outflow will tend to increase at only one node at a time. Of the nodes with unmet demand, the one with the lowest elevation (i.e. lowest minimum head) will see outflow increase as available head increases, while the other nodes will continue to exhibit the same (possibly zero) outflow (Morley and Tricarico, 2008). This node-by-node behaviour is equivalent to adopting the HFR proposed by Bhave, 1981 (Figure 2.1a), where the model will attempt to deliver the maximum demand to the node as soon as there is available pressure. This does not necessarily mean that maximum flow is delivered as soon as the minimum head at a node is achieved; only the demand which can be provided by the current system pressure is delivered. In contrast, a pressure-driven analysis that assumes a more complex HFR will tend to exhibit outflow increases at multiple nodes simultaneously.

As part of this thesis, a small experiment was carried out to demonstrate this behaviour. The experiment used a simple, gravity-driven, 4-node benchmark network model first

proposed by Gupta and Bhave, 1996 (Figure 2.2). Using the same parameters proposed by Ang and Jowitt, 2006 (artificial components approach) and Cheung, Van Zyl, and Reis, 2005 (Wagner HFR approach), the available head at the source was varied from 84 m to 130 m. Flow at the nodes (at elevations of 90, 88, 90, and 85 m, respectively) was found to increase with greater available head, as expected, but the rate of increase was different for the two different approaches (Figure 2.3). It is clear to see that ouflow increases at one node at a time in the artificial reservoir approach and at multiple nodes simultaneously using the HFR approach.

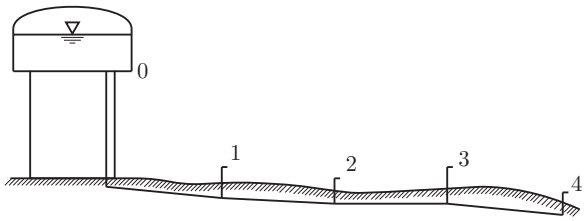

FIGURE 2.2: Benchmark network with 4 nodes in series (based on Gupta and Bhave, 1996)

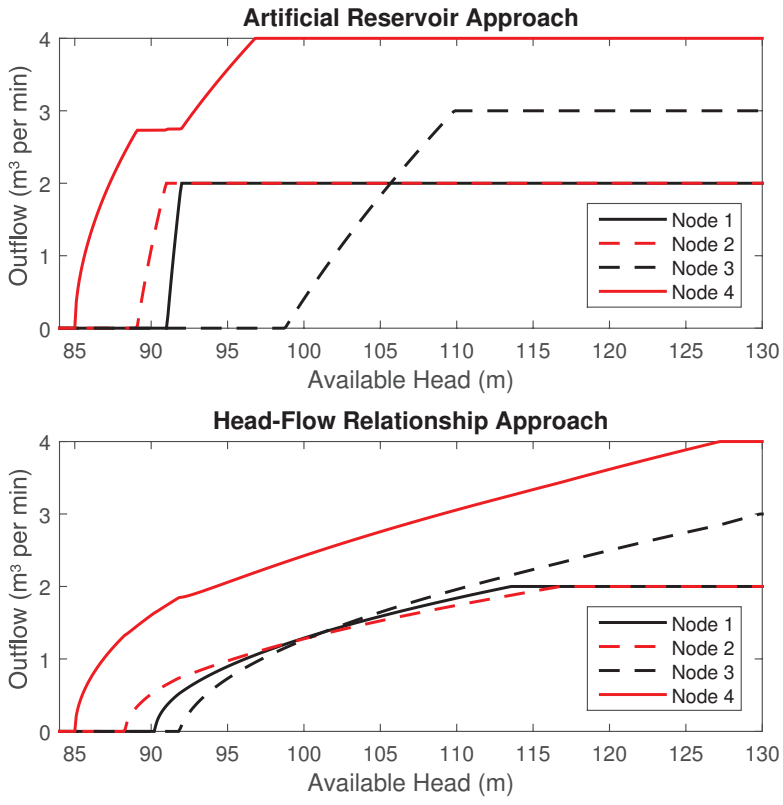

FIGURE 2.3: Nodal outflow vs. available head, benchmark network

While the artificial reservoir approach gives results that can be seen as unrealistic (Ackley et al., 2001), little work has been done to experimentally verify which approach (artificial

reservoirs or HFRs) best matches real-world conditions. The results from analysing pressure-dependent leakage suggest that the characteristics of the flow will differ from network to network (Thornton and Lambert, 2005). At any rate, adopting an HFR is the more versatile approach because the parameter values can be calibrated within a range that allows for results equivalent to the artificial reservoir approach.

2.1.3 Hybrid Approach

In addition to these two approaches to PDA (incorporating an HFR into the system of equations and adding artificial components), a "hybrid" approach can be taken, where artificial components that include an HFR are added. EPANET includes built-in functionality to model pressure-dependent flow at "emitter" nodes. Emitter nodes can be thought of as openings in pipes, where the outflow depends directly on the pressure, not on how much water is demanded by the user. They are commonly used to model sprinklers and leaks in pipes. Flow at these nodes is governed by the equation:

$$q = C(p)^{\gamma} \tag{2.9}$$

where:
q = outflow from the node
p = available pressure at the node
C = emitter coefficient
γ = emitter exponent
(Rossman, 2000)

There are significant drawbacks to using emitters for PDA. There are no upper bounds on the nodal outflow; outflow will continue to increase as pressure increases, even if the outflow exceeds consumer demand at the node. Also, if the network is in a pressure-deficient situation and negative pressures are simulated at the node, "outflow" at the emitter will become negative, and water will be supplied to the network from the emitter. Obviously, unless the emitter is located underwater, this is an unrealistic scenario. Nevertheless, some researchers have built on the emitter functionality to develop approaches to PDA (e.g. Morley and Tricarico, 2008; Abdy Sayyed, Gupta, and Tanyimboh, 2015).

Morley and Tricarico, 2008 developed a pressure-driven extension for EPANET by modifying the equation governing emitters. They named this extension EPANETpdd. It worked by introducing two pressure thresholds, P_{min} and $P_{critical}$, which correspond to the minimum and desired heads described by Wagner, Shamir, and Marks, 1988. These thresholds were incorporated into equations describing an HFR that is applied to emitter nodes, which are used in place of regular demand nodes. Instead of the standard emitter equation (Equation 2.9), the flow was governed by the HFR proposed by Wagner, Shamir, and Marks, 1988 (Equations 2.2-2.4). The authors additionally allowed for alternative HFRs to be used in place of Equation 2.3; the user could select the HFR proposed by Fujiwara and Li, 1998 (Equation 2.6) or a sinusoidal HFR proposed by Tucciarelli, Criminisi, and Termini, 1999, which closely matches the Fujiwara HFR. Morley and Tricarico, 2008 went on to compare the results of using EPANETpdd against prior PDA simulations. They achieved excellent correlation with the artificial reservoir approach

of Ang and Jowitt, 2006 and the Global Gradient Algorithm modification approach of Cheung, Van Zyl, and Reis, 2005 when the emitter equations were calibrated in ways equivalent to these prior works.

Another approach to PDA that uses emitters in EPANET was proposed by Abdy Sayyed, Gupta, and Tanyimboh, 2015. The authors added an artificial string of a check valve, a flow control valve, and an emitter to each demand node, all set at the elevation of the demand node (Figure 2.4). The check valve prevents backflow from the emitter and the flow control valve is set to limit outflow to be equal to or less than the nodal demand at each timestep. The nodal demands are all set to zero to prevent double-counting. By using built-in EPANET components, this approach avoids the drawbacks to using emitters for PDA, without having to make significant modifications to the EPANET source code. (The approach of Morley and Tricarico, 2008 and all techniques that add HFRs directly to the system of equations require source code modifications.) It also can be run in a single execution of EPANET, unlike the iterative approaches that add artificial reservoirs, thus saving computation time (Wu et al., 2009).

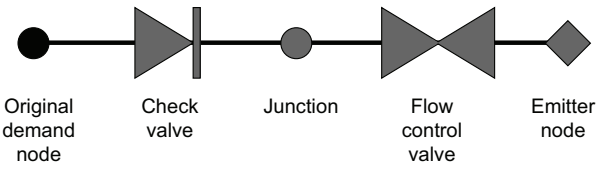

FIGURE 2.4: String of artificial components added to each demand node by Abdy Sayyed, Gupta, and Tanyimboh, 2015

Abdy Sayyed, Gupta, and Tanyimboh, 2015 determined that the emitter equation (Equation 2.9, above) is equivalent to the HFR proposed by Wagner, Shamir, and Marks, 1988 (Equation 2.3, above) if the following are observed:

$$C = \frac{q_j^{req}}{(H_j^{des} - H_j^{min})^{\frac{1}{n_j}}} \tag{2.10}$$

$$\gamma = \frac{1}{n_j} \tag{2.11}$$

$$p = H_j - H_j^{min} \tag{2.12}$$

where:
C = emitter coefficient
q_j^{req} = the maximum flowrate, i.e. the demand
H_j^{des} = the desired head at node j, above which the flowrate is constant
H_j^{min} = the minimum head at node j, below which the flowrate is zero
n_j = a constant that defines the shape of the HFR curve, set to 2 by Wagner, Shamir, and Marks, 1988
γ = emitter exponent
p = available pressure at the node
H_j = the available head at node j

Because all of the parameters can vary from node to node, the emitter coefficient and exponent may be node-specific. For the applications in their paper, Abdy Sayyed, Gupta, and Tanyimboh, 2015 set $(H_j^{des} - H_j^{min})$ and n_j to constant values across the entire network, in line with prior PDA studies. Other studies have adopted a similar approach of adding a series of artificial components that incorporate an HFR to the consumer nodes. Namely, a combination of a valve and an emitter (Bertola and Nicolini, 2007; Mahmoud, Savić, and Kapelan, 2017) or a valve and a reservoir (Jinesh Babu and Mohan, 2012; Gorev and Kodzhespirova, 2013; Pacchin, Alvisi, and Franchini, 2017).

2.2 Performance Evaluation

Whichever approach is taken to modelling a water distribution system, the end goal is always to have an accurate way to test the performance of the system. But this introduces a question: what exactly is meant by "performance"? When it comes to assessing the response of water distribution systems to cyber-physical attacks, this means measuring how reliable the system is when subject to adverse conditions. How much clean water is actually delivered to customers when they need it? Various ways to do this assessment have been proposed.

Perhaps the simplest metric for assessing network performance is the demand satisfaction ratio, where a score of 1 means that all water demanded is supplied, and 0 means that no water is supplied (Siew and Tanyimboh, 2010). This can be computed at the level of an individual node, across an entire network, or any level in between. Network operators may wish to additionally incorporate considerations of equity, for example by rating more highly a network that supplies 50% of demand to two nodes than a network that supplies 100% of demand at one node and 0% at the other. Designing a network and its operating rules to ensure equitable distribution during pressure-deficient scenarios can be formulated as a multi-objective optimisation problem (e.g. Fujiwara and Li, 1998).

Another consideration that can be added to network performance assessment is efficiency. Efficiency takes into account water losses due to leakage. The amount of water supplied to the system and the amount of water that reaches customers are measured and compared. Systems that lose a larger proportion of water are deemed less efficient. In a given network, it may be that efficiency and demand satisfaction are negatively correlated, and so determining overall network performance requires appropriate weighting of the two factors (Creaco, Franchini, and Todini, 2016). Additional performance factors, such as the duration of service interruption or the risk of backflow due to low pressures, can also be considered.

A water distribution system's reliability can be thought of as how likely it is to successfully deliver water over a period of time, taking into account possible failure mechanisms. There is no universal approach to defining or measuring this reliability (Huang, McBean, and James, 2005). The types of failure considered can be mechanical (e.g. pipe breaks), hydraulic (e.g. insufficient pressure), and/or water quality-related (e.g. water age becoming too old) in nature. Many of these measures involve defining statistical

models of the failure modes and then running extended period simulations with failures occurring as stochastic events, so-called simulation-based approaches (Gheisi, Forsyth, and Naser, 2016). An alternative is to use a heuristic-based approach, where the network topology (the characteristics, layout, and connectivity of the pipes) is used to directly assess reliability. Such approaches allow engineers to quickly compare the reliability of different designs, changing things like pipe diameters and the number of redundant pipes in a system, until they achieve a design that balances cost and reliability.

One heuristic approach that is widely used is the combined resilience-failure index first proposed by Todini, 2000. This index measures the hydraulic power across network; a resilient network has more power than is required to satisfy all demands, while a network experiencing failure has insufficient power to satisfy all demands. Power, in this sense, is defined as $QH\gamma$: the product of flow, head, and the specific weight of water. The index was recently updated to account for pressure-driven modelling (Creaco, Franchini, and Todini, 2016). It is defined as:

$$I_r = \frac{max(\mathbf{q}_{user}^T\mathbf{H} - \mathbf{d}^T\mathbf{H}_{des}, 0)}{\mathbf{Q}_0^T\mathbf{H}_0 + \mathbf{Q}_p^T\mathbf{H}_p - \mathbf{d}^T\mathbf{H}_{des}} \tag{2.13}$$

$$I_f = \frac{min(\mathbf{q}_{user}^T\mathbf{H} - \mathbf{d}^T\mathbf{H}_{des}, 0)}{\mathbf{d}^T\mathbf{H}_{des}} \tag{2.14}$$

where:
I_r = resilience index
I_f = failure index
$\mathbf{q}_{user}$ = vector of outflows at consumer nodes
$\mathbf{H}$ = vector of heads at consumer nodes
$\mathbf{d}$ = vector of demands at consumer nodes
$\mathbf{H}_{des}$ = vector of the desired heads at consumer nodes, as in previous equations
$\mathbf{Q}_0, \mathbf{Q}_p$ = vector of flows from source nodes and pumps, respectively
$\mathbf{H}_0, \mathbf{H}_p$ = vector of heads at source nodes and pumps, respectively
and ^{T} denotes the transpose of a vector.

The combined index is given by $I_R + I_f$. A network where no water is delivered has an index of -1, a network where exactly enough power is provided to satisfy all demands has an index of 0, and an "ideal" network with no headloss or leakage has an index of 1. The specific weight of water is assumed to be constant and so cancels out of the equation. The index will vary over time in an extended period simulation as the demands and flows across the network change.

Deciding which performance metric to use depends on the intended use of the model. For assessing the impacts of different attacks on the same network, a simple metric such as the demand satisfaction ratio may work best because what matters is how the consumers are affected. For comparing the ability of different network layouts to cope with a given attack, a heuristic approach may be better because it allows for comparisons to be made without running extended period simulations.

2.3 Modelling Attacks on Water Distribution Systems

The need for tools to quantitatively assess the impacts of attacks on water distribution systems has been identified for years (Haimes et al., 1998). In order to simulate cyber-physical attacks, both the hydraulic behaviour of the network and the actions of the attackers must be modelled. Early work has been done in combining these, but the approaches developed so far were only applied at a small scale, were capable of simulating a limited range of attacks, and/or employed a demand-driven hydraulic model. In order to be useful to water distribution system operators, a cyber-physical attack modelling approach should be scalable and should produce realistic results for a range of attack scenarios.

Bespoke approaches that integrate the system of hydraulic equations and attacker actions into a single model have been developed by Do, 2015 and Perelman and Amin, 2014. Do, 2015 developed a model for testing attack detection algorithms. While this model allowed for the simulation of a range of many different attack scenarios, it was only demonstrated on small example networks with two consumer nodes. The model linearised the mass and energy balance equations, an assumption that is often inappropriate for real-world networks (Walski et al., 2003). The hydraulic model was also demand-driven; this limits the number and type of failure scenarios that can be accurately modelled.

Perelman and Amin, 2014 developed a bespoke network interdiction model for simulating attacks on water distribution systems, where the system of mass- and energy-balance equations were formulated as a convex optimisation problem. They also included pressure-dependent flows by using an artificial reservoir technique, and simulated the behaviour of both attackers and operators. The attackers attempted to minimise flow provided to customers by removing one pipe at a time from the network, while the operators attempted to re-route water to maximise flow to customers. Only this one type of attack scenario was incorporated into the model. The simulations were carried out on a widely used benchmark network with six consumer nodes.

Taormina et al., 2017 developed an approach which allowed for the simulation of a wide range of attack scenarios. The authors created a toolkit, *epanetCPA*, which uses the hydraulic engine of EPANET and runs simulations in MATLAB, a widely used engineering computation software package. The toolkit is explained in more detail in Section 3.1. The authors also developed an attacker model to represent the actions of a range of cyber-physical attacks. They then simulated six different attack scenarios that targeted different components of a medium-sized benchmark network. These attacks resulted in tanks overflowing or emptying down to very low levels. However, because the hydraulic engine was demand-driven, the tanks could not be allowed to be fully emptied and the consumers' full demands were assumed to always be satisfied. One important finding was that different attacks could have very similar impacts on the system, meaning that simply observing a failure in a network isn't necessarily sufficient to identify the cause. Another finding was that the impact of the attack depended strongly on the initial conditions, meaning that an informed attacker could time an attack to have maximum impact. The authors also showed that eavesdropping attacks can lead to more sophisticated deception attacks, and so have ramifications even more serious than compromising data

confidentiality.

Ahmed, Murguia, and Ruths, 2017 similarly used EPANET as the (demand-driven) hydraulic engine for a modelling approach that was used to test attack detection methods. The inputs to and outputs from EPANET were mapped onto a system of linear equations and a Kalman filter was used to estimate the state of the system. Three different attack scenarios were simulated: two where sensor readings were manipulated in the system of linear equations and one where a modified consumer demand pattern was applied directly to the EPANET model. These attack scenarios were applied to a small network with four consumer nodes. While useful for detailed testing of attack detection algorithms, this approach is limited in its application due to its small network size, linearised hydraulic equations, and reliance on a demand-driven hydraulic engine.

Chapter 3

Description of Pressure-Driven Model

The primary aim of this research project was to adapt an existing software package used for modelling cyber-physical attacks on water distribution systems by adding the ability to conduct pressure-driven analyses, and then to use this adapted software to explore a variety of attack scenarios. This chapter describes the computer model that was developed over the course of this project. It begins with an explanation of the model that served as a basis for the present work, before outlining the modifications and additional features that were added in order to incorporate pressure-driven modelling. The resulting benefits are also described in detail.

3.1 Existing Hydraulic Toolkit

epanetCPA is a toolkit developed by Taormina et al., 2017 that allows users to simulate cyber-physical attacks on water distribution systems. The toolkit uses the hydraulic engine of EPANET and runs simulations in MATLAB. *epanetCPA* introduces a cyber layer that contains the digital components, e.g. tank level sensors, PLCs, and a centralised SCADA system. The PLCs collect and store readings from the sensors, transmit those data to the SCADA system, and take actions based on the readings. The toolkit separately reports the actual physical status of the system (pressures, flows, etc.) and the cyber layer status, meaning that these can diverge in the case of a simulated attack. The SCADA system stores the (potentially erroneous) data from all PLCs in the network and can send new settings back to the PLCs. Typically, hydraulic components in an EPANET model will be activated according to simple rules, such as a pump turning on when a tank level drops below a certain threshold or a valve opening if the pressure at a node drops below a certain threshold. In *epanetCPA*, these rules are assigned to PLCs in the added cyber layer.

The attacker model allows for customisation of cyber-physical attacks by setting the attack start and end times, the target components, and the override settings. In this way, components can be forced to remain on or off even if activation thresholds are crossed, or the thresholds themselves can even be modified. Examples include keeping valves in an open or closed position and forcing pumps to remain on or off. The user can also specify more sophisticated attacks where the true status of a component is concealed by replacing readings with fabricated data. These "deception attacks" interfere with the PLCs' control operations; the PLCs command actuators as normal, but based

on erroneous readings. At the end of each simulation run, the toolkit can plot the status of the affected components over time, allowing for assessment of the impact of attacks.

Until this current work, *epanetCPA* has relied on the standard EPANET hydraulic engine, which is demand-driven. This limits the type of failure scenarios that can be simulated (e.g. tanks cannot necessarily be fully emptied, sections of the network cannot be hydraulically disconnected, and unmet demand cannot be computed). Hence, there is much to be gained by adding PDA capabilities to the toolkit. The code for *epanetCPA* is open-source, and the modifications made during the course of this project are intended to be incorporated into future releases.

3.2 Incorporating PDA

As discussed in Section 2.1, there are several different approaches that have been used for pressure-driven analyses of water distribution systems. The approach of Abdy Sayyed, Gupta, and Tanyimboh, 2015 was selected over other pressure-dependent EPANET modifications for its simple implementation and accurate results. It was found by Pacchin, Alvisi, and Franchini, 2017 to produce reliable results and has also been used successfully in extended period simulations (Mahmoud, Savić, and Kapelan, 2017). This approach adds a string of artificial components to each demand node for which the settings are updated at each time step in the simulation. It also utilises a generic head-flow relationship that can be easily modified, and so calibrated for a range of networks.

New modifications to the *epanetCPA* code allow for this approach to be automatically implemented. This takes place in two phases:

Before beginning the hydraulic simulation:

1. Store the user-specified P_{des} & P_{min} values.
2. Store the user-specified exponent of the emitter equation (γ in Equation 2.9).
3. Store the user's choice of HFR equation (Wagner, Bhave, or Fujiwara).
4. Read and store the base demands for all nodes.
5. Set the base demands to zero to prevent double-taking of water.
6. Add a junction and an emitter node near each original demand node in the EPANET input file.
7. Join the original demand node to the junction with a check valve of negligible resistance and join the junction to the emitter node with a flow control valve of negligible resistance.
8. Set the elevations of the junction and the emitter node equal to the elevation of the original demand node, plus the value of P_{min}. This modification, also proposed by Mahmoud, Savić, and Kapelan, 2017, prevents flow when the pressure is below P_{min}.

During the hydraulic simulation:

1. Before each step in the simulation, calculate the settings of the flow control valves and emitters:

   - For the flow control valve, the setting is the actual demand at the original demand node at that particular timestep. This is calculated by multiplying the stored base demand value for that node by the corresponding pattern multiplier value. Patterns are components of the EPANET model. They are user-specified vectors of values that determine how demand varies over time, according to:

   $$D_i(t) = BD_i \times M_j(t) \tag{3.1}$$

   where:
   $D_i(t)$ = demand at node i at time t
   BD_i = base demand at node i
   $M_j(t)$ = multiplier value of pattern j at time t.
   The number of patterns can vary from one for the whole network to one for each node. Typically, network models will have one pattern for each of a few distinct regions.

   - For the emitter, the setting is the coefficient C in Equation 2.9. It is determined by the specified HFR.
   For Wagner, C is as in Equation 2.10.
   For Bhave, C is set to an arbitrarily large value.
   For Fujiwara,

   $$C = \frac{d}{P_i^{des2}} \quad \Big\} P_i \geq P_i^{des} \tag{3.2}$$

   $$C = d\left[\frac{3P_i^{des} - 2P_i - P_i^{min}}{(P_i^{des} - P_i^{min})^3}\left(1 + \frac{P_i^{min2} - P_i P_i^{min}}{P_i^2}\right)\right] \quad \Big\} P_i < P_i^{des} \tag{3.3}$$

   where:
   d = the demand at node i at time t
   P_i = pressure at node i at time *t-1*
   P_i^{min}, P_i^{des} = minimum and desired pressure thresholds for node i

   Note that when using the Fujiwara HFR, because the pressure at the node at time t has not yet been calculated at this stage in the simulation, the value of C is approximated by taking the pressure at the previous timestep. This is a constraint of using the emitter approach of Abdy Sayyed, Gupta, and Tanyimboh, 2015, and one that may introduce errors. Additionally, the emitter exponent, γ, must be set to 2 when using the Fujiwara HFR. The code performs a check for this and alerts the user if it is not.

2. Update the flow control valve and emitter settings with the calculated values, then perform the next step of the hydraulic simulation.

Beyond these steps, the *epanetCPA* toolkit performs exactly as it did prior to modification. The toolkit first adds a cyber layer to the network before carrying out hydraulic simulations on both layers, either with or without cyber attacks.

3.3 Benefits of the PDA Model

These modifications to the *epanetCPA* toolkit allow for the simple comparison of simulation results using pressure-driven analysis and traditional demand-driven analysis. Executions of the toolkit can be set up using the same network and demand pattern inputs, and then the model can be run using PDA and DDA in sequence, both with and without attacks. The results can then be directly compared to assess the differences between the approaches.

3.3.1 Identifying the Need for PDA

The modified toolkit can be used to determine whether or not the network might be experiencing pressure-deficient conditions. Normally, when running a DDA simulation, pressure-deficient conditions are only detected when the model calculates negative pressures in pipes and produces a warning message. However, the conditions in real life may be such that the head is too low to deliver full demand despite the model not calculating negative pressures. In such a case, it may be unclear that pressure deficient conditions are being exhibited.

Using the modified *epanetCPA* toolkit, if there are negligible differences between the PDA and DDA simulation runs, then the network likely has sufficient head for the specified demands. However, if there are differences between the runs, such as lower outflow at demand nodes, then pressure-deficient conditions likely exist. In such a case, the results from the DDA run are likely to be unrealistic and so inadequate for making predictions about real-world network behaviour. An operator can then know that a PDA simulation is required for better results.

3.3.2 Quantifying Network Failure

Perhaps the most powerful aspect of a pressure-driven model is that it allows the operator to calculate shortfalls in water supplied to customers. While a DDA model assumes that all demands are met, a PDA model will reflect lower outflow at consumer nodes if there is insufficient head in the network. The difference in outflow between the DDA and PDA results can thus be thought of as "undelivered demand" – the amount of water that was not able to be provided to customers. This is a useful metric for determining the impacts of attacks or other network failures.

A PDA model also allows for tanks to empty completely because outflow drops to zero as the available water is used up. Tanks cannot run dry in DDA models (unless there is an alternative source of water available) because the demand does not diminish. Attempting to consume more water than is available results in a violation of the continuity equations that describe the system, and the model then crashes or aborts. It is important for the tanks in the model to be able to empty completely when simulating atypical operating

conditions. Operators need to be able to simulate this physically possible scenario in order to best understand how to prevent it from happening during events such as attacks, pipe bursts, or fire fighting.

Chapter 4

Model Testing, Calibration, and Validation

This chapter describes an experiment that used a small-scale, physical testbed to test, calibrate, and validate the computer model described in Chapter 3. An attack that resulted in pressure-deficient conditions was carried out on the testbed. Three different head-flow relationship (HFR) equations were tested, and all produced results with less error than the traditional DDA approach. The testbed, named WADI, is described in detail, as are the experimental design and results. The chapter begins with a discussion of why it's difficult to calibrate PDA models, and thus why the work in this chapter presents a valuable addition of knowledge to the field of hydraulic modelling.

4.1 Goals and Background

Because water distribution systems are large, expensive networks that provide a critical resource to the populace, it is difficult (and indeed unethical) to carry out experiments on real systems to calibrate and verify hydraulic models. Instead, model developers have typically relied on calibrating their models with measured field data (Thornton and Lambert, 2005). This approach is useful for assessing "business-as-usual" scenarios, but little publicly available data exist for systems experiencing pressure-deficient conditions. Part of the challenge for modelling PDA is that metrics like demand satisfaction ratio are quantities that cannot truly be measured, only estimated.

The goal of the work done in this chapter was to experimentally determine what values of the parameters governing the HFR equations best replicate real-world behaviour. Essentially, what inputs produce a simulation with the most realistic results? The process was repeated using each of the three HFRs used in the model to determine which HFR best simulates real-world flows. Because the Bhave HFR produces results equivalent to an artificial reservoir approach, this experiment should also reveal which approach to PDA (HFRs vs. artificial components) produces more realistic results.

Determining the parameters of the HFR equation is not a trivial task. As far as this author can tell, the closest that any prior work has come to experimentally determining these parameters is in testing flow through leaks (e.g. Fontana, Giugni, and Marini, 2016, van Zyl and Clayton, 2007). Such experiments, while useful, aren't directly applicable to pressure-driven flow through existing, fixed orifices because they (rightfully in the case of leaks) assume that parameters such as orifice size depend on pressure. For

Paper	γ	P_{min} (m)	P_{des} (m)
Wagner, Shamir, and Marks, 1988	0.5	14.06	28.12
Tanyimboh, Tahar, and Templeman, 2003	0.5	0	30
Cheung, Van Zyl, and Reis, 2005	0.5	0	20
Morley and Tricarico, 2008	0.5	0	20
Wu et al., 2009	0.5	0	14.06-70.31
Abdy Sayyed, Gupta, and Tanyimboh, 2015	$\frac{2}{3}$	0	15

TABLE 4.1: Wagner HFR parameter values in the literature

those researchers who have used the Wagner HFR in network simulations, the parameter values have typically fallen within a similar range. Table 4.1 summarises some of these values. Here, $P_{min/des} = H_{min/des} -$ the node's elevation, and γ, the emitter exponent, is as in Equations 2.9 and 2.11. The values of P_{min} and P_{des} will largely be dependent on the network. Wu et al., 2009 recommended either basing the value of P_{des} off of water industry guidelines, or making the assumption that simulated nodal pressures in a pressure-sufficient scenario represent the desired pressures. Setting $\gamma = 0.5$ is consistent with the theoretical orifice equation (Equation 2.8), but as discussed at the end of Section 2.1.1, flows in real-world situations may not strictly follow this. Experimentally determining and verifying the HFR parameters will lead to more realistic models with results in which users can place greater confidence.

4.2 The WADI Testbed

WADI, short for Water Distribution, is a physical testbed installed in the iTrust Centre for Research in Cyber Security at SUTD (Figure 4.1). The testbed is designed to allow researchers to carry out experiments on an analogue of an extended water distribution system; it includes both the physical and software aspects. WADI is comprised of inlets leading to two storage tanks, two elevated tanks (herein referred to as reservoirs), six consumer tanks, pumps to power water through the PVC pipes, valves to direct flow, and pressure sensors and flowmeters throughout. The storage tanks and elevated reservoirs are 0.5m in depth with capacities of 2.5m^3 and 1.25m^3, respectively. Typical network flows are less than 2.5m^3/hr. There is a net elevation drop of only 10cm from the elevated reservoir outlets to the consumer tank inlets, but with the booster pumps activated, pressures in the network can exceed 2 bars (∼20m of water column). Operators have the option of directing flow either through the booster pumps or pipes that rely on gravity only. Demand patterns can be set individually for each of the 6 consumer tanks, and flows into these tanks are controlled by variable valves. The six consumer tanks are effectively identical; they're at the same elevation and fed by hydraulically connected pipes. During runs of the testbed, data at each of the pressure sensors and flowmeters are recorded at 1-second intervals. A schematic representation of WADI as an input network in EPANET is shown in Figure 4.1. (Not shown in this schematic is a tank that stores water emptied out from the consumer tanks and a return line for recirculating the water.) A computer-generated rendering of the system is shown in Figure 4.2.

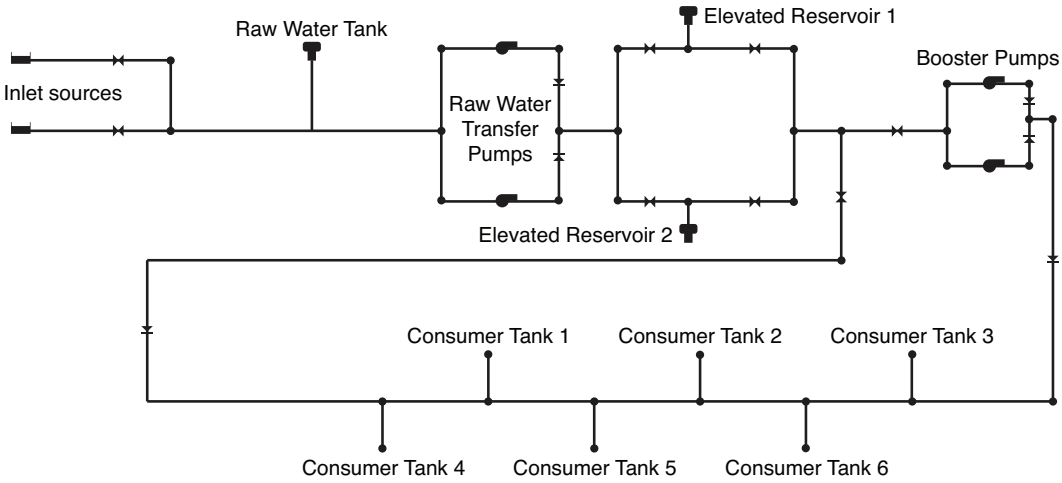

FIGURE 4.1: Schematic of the WADI testbed EPANET network file

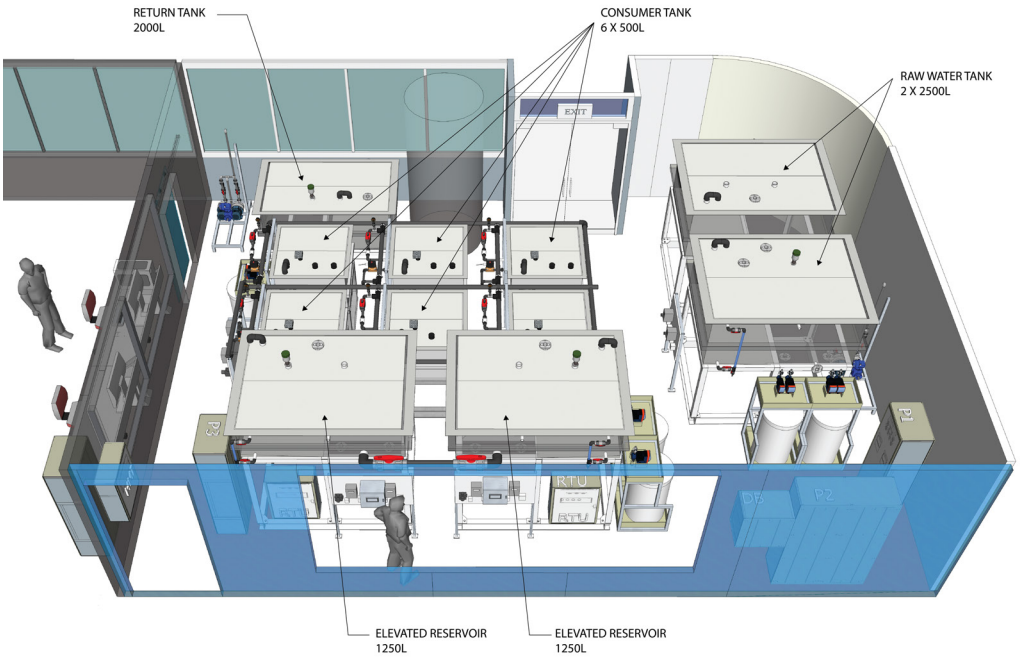

FIGURE 4.2: Computer-generated rendering of the WADI testbed

4.3 Method

There were two main steps to this experiment: calibration of the parameter values and validation of the values found. Real-world data from the WADI testbed were used for both steps. Calibration of the parameters was carried out algorithmically using optimisation.

4.3.1 Calibration

Obtaining the Real-World Data

The HFR parameters were calibrated using data from three different runs of the WADI network. Each run had the following characteristics:

- Runs were three hours in length.
- The elevated reservoirs were initially over 75% full.
- Demand patterns at each of the 6 consumer tanks varied every 30 minutes. The demand patterns were equal to the modelled DDA flows shown in Appendix A.
- To simulate an attack, the raw water transfer pumps and booster pumps were switched off from $t = 0$ to $t = 2$ hrs, meaning that the elevated reservoir tanks would empty through the consumer nodes by gravity flow only.
- For 1 hour of the attack, total demand from the consumers was less than the maximum flow rate from the elevated reservoirs under gravity flow only (pressure-sufficient conditions). For the other hour of the attack, total demand from the consumers was greater than this maximum flow rate (pressure-deficient conditions).
- The demand patterns for each consumer tank were different from run to run to compensate for any systematic variation in the behaviour of the tanks.

Crucially, the HFR parameters determined should be able to replicate flows in the network during both pressure-sufficient and pressure-deficient conditions. That is why the runs included periods with both conditions during gravity flow and periods with the pumps switched on. The maximum flow rate under gravity flow (around $1.0\text{m}^3/\text{hr}$) was determined in an initial run where consumer demand was set to a high value and the pumps were switched off.

For the purposes of calibration, the data from WADI were post-processed to remove high frequency variations in the flow rate. The WADI system works on a volumetric basis: in a given time period (the time step), the system will attempt to deliver the demanded volume of water to each of the consumer tanks, varying the setting of the consumer tank supply valves. While the system attempts to adjust the valves to ensure a steady flow rate, the flow rate can sometimes vary within each time step; for example if the demanded volume is delivered in the first half of the time step, then zero water will be delivered during the second half of the time step. To get around this, a short time step of 5 minutes was chosen. Despite this, some high frequency variations in the flow rate were experienced (such as when the pumps switched back on). In order to remove these high frequency variations (which could skew the calibration), a new time series was generated where the flow rate was set at a constant value during each 5-minute time step: the mean of the reported flow rate during the same time step. This is illustrated in Figure 4.3, below. This processing of the data ensured that the WADI results used for calibration matched the format of the EPANET results, where flow rates are constant during a given time step.

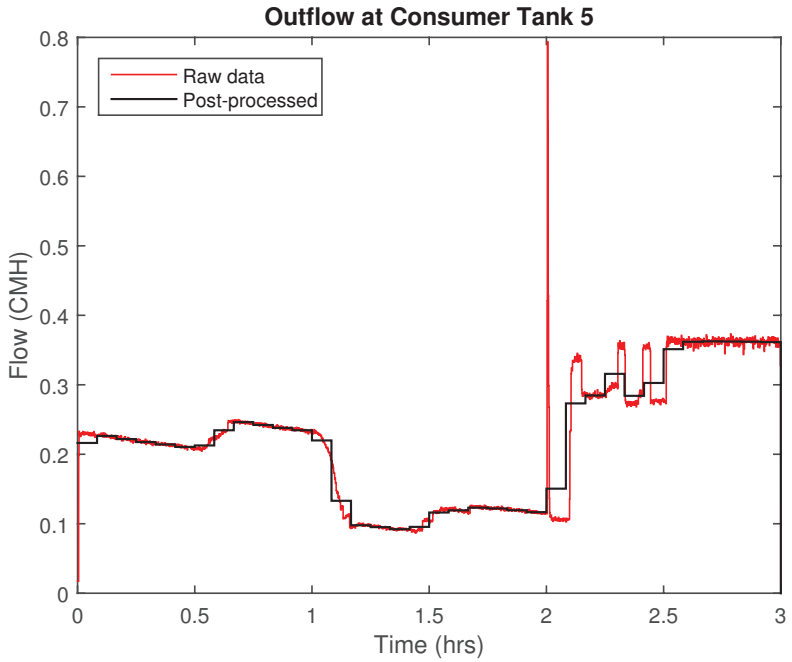

FIGURE 4.3: An example of raw data from WADI and the corresponding post-processed data used for calibration

Optimisation Scheme

After the runs were completed, the results were replicated in EPANET using the modified *epanetCPA* toolkit. A model of the network that had previously been calibrated for demand-driven analyses was used as a basis for the calibration. This model was modified for each of the experimental runs to replicate the corresponding input conditions (i.e. the initial levels in the tanks, the demand patterns, and the timing of the attacks).

An optimisation routine was then carried out to determine the optimal values of the HFR parameters: γ, P_{min}, and P_{des}. Simulation-based optimisation works by taking different sets of inputs to a model (so-called decision variables) and comparing the different outputs that they produce (the "objective"). Large numbers of simulations of the runs (using the modified *epanetCPA* toolkit) were carried out automatically using a differential evolution solver (Buehren, 2014). A differential evolution solver is a type of genetic algorithm that mimics the biological concept of adaptation through natural selection:

1. An initial "generation" of individuals is created, up to a certain population size. Each individual has a different combination of input parameter values, i.e. a different solution to the objective function. The parameters are usually bounded in some way, so the values will only vary within a given range.

2. The simulation is carried out using the parameter values from each of the individuals, and then each individual is scored according to some measure (such as cost or closeness of fit to real-world data). The function that determines the score is called the objective function.

3. The individuals with the best scores "survive" into the next generation, and new individuals with parameters closer to the best individuals are generated to fill out the population. Additional techniques can be used in this step, such as "mating" the best individuals to produce "children" with parameters similar to their "parents", or introducing wildly different individuals (with "mutations"), to prevent the optimisation routine from converging to a local (instead of global) solution.

4. This process is repeated for multiple generations until the parameters converge on a single best individual or some other stopping criteria is met.

In this case, the score was determined by calculating the root mean square error (RMSE) for flow at each consumer tank, compared to the data from the real-world WADI runs. RMSE is a commonly used measure of goodness-of-fit; the lower the RMSE value, the more closely the simulated data matches the real data. The simulation data was interpolated to ensure that the time steps aligned with the one-second resolution data from the analogue run. Only the flows at the consumer nodes were considered. The first 10 minutes, when the testbed was initialising flows, were not counted. The data were normalised by *feature scaling*, so that RMSE values from separate runs could be combined without bias. That is, the flow data from each analogue run were transformed to vary from 0 to 1, and the same transformation was applied to the simulation data:

$$q_{n,t}^{anlg'} = \frac{q_{n,t}^{anlg} - min(q^{anlg})}{max(q^{anlg}) - min(q^{anlg})} \tag{4.1}$$

$$q_{n,t}^{sim'} = \frac{q_{n,t}^{sim} - min(q^{anlg})}{max(q^{anlg}) - min(q^{anlg})} \tag{4.2}$$

where:
q^{anlg} = The outflow at all consumer nodes over all time in the analogue WADI run
$q_{n,t}^{anlg}$ = The outflow at consumer node n at time step t in the analogue WADI run
$q_{n,t}^{sim}$ = The outflow at consumer node n at time step t in the simulation
$'$ indicates the normalised data

The RMSE values were then calculated and the score was set equal to the mean value of all RMSE values as follows:

$$RMSE_n = \left[\frac{\sum_{t=1}^{T}(q_{n,t}^{sim'} - q_{n,t}^{anlg'})^2}{T}\right]^{0.5} \tag{4.3}$$

$$score = \frac{\sum_{n=1}^{N} RMSE_n}{N} \tag{4.4}$$

where:
N = The number of consumer nodes (n) in the network
T = The number of time steps (t) in the simulation/analogue run

Here, a lower score means a better fit. Accordingly, the objective function of the differential evolution solver was to minimise the score. Because WADI was run three times for calibration, each time with different demand patterns, a score was calculated for each

run. The mean of these three scores constituted the overall score for each set of parameter values, and the set of values with the lowest overall score was deemed the best. This means that the values had to work well generally, not just for one specific run of WADI.

The settings for optimisation using differential evolution were as follows:

Parameter	Lower Bound	Upper Bound	Seed Values
γ (unitless)	0	2.5	0.1, 0.5
P_{min} (m)	0	P_{des}	0.1, 0.5
P_{des} (m)	P_{min}	100	1.5, 3.0

TABLE 4.2: Optimisation routine settings

Recall that P_{min} is the pressure at which flow will begin at the node, P_{des} is the pressure above which flow will be constant and equal to demand, and the emitter exponent, γ, determines the shape of the HFR curve for intermediate pressures. The bounds on γ were determined by the range of values reported in the literature (see the end of Section 2.1.1). The upper bound on the pressure thresholds was determined by considering the maximum pressures typically encountered in municipal water distribution systems. Mckenzie and Wegelin, 2009 give this as 100 m for a South African system, while Chan, 1983 reports 90 m for a system in Hong Kong.

For the Bhave HFR, the value of P_{des} has no effect on the results. This is because the emitter coefficient was set at an arbitrarily high value in the toolkit (1×10^9 was used in this case). γ was held constant at 0.5 because if the pressure was less than 1m and γ was more than 1.0, then the value of p^γ in Equation 2.9 could cancel out the arbitrarily high emitter coefficient and result in lower than expected flow values. For the Fujiwara HFR, the value of γ must be equal to 2, and so only the pressure threshold values were varied in this case.

Each HFR was calibrated with the differential evolution solver at least twice, using different sets of seed values (initial guesses for the parameters) each time. The population size was set at $10(x + 1)$, where x was the number of parameters being varied. The suggested population size for this solver was $10x$ (Buehren, 2014) but a larger number was chosen because at times only one parameter was being varied. The maximum number of generations was initially set at 100, and later lowered to 50 when convergence was observed to occur after no more than 45 generations. After initial calibration runs showed that the optimal P_{des} value was much smaller than 100 m, the upper bound was lowered to 5 m for subsequent runs to improve calibration time and accuracy. When the different trials converged to the same result (to a precision of 0.01), these were taken to be the optimal values for the respective HFR.

4.3.2 Validation

As mentioned above, the model calibration would only be successful if the resulting parameter values provide good fits for data from WADI under a wide range of conditions. In order to test this, a fourth run of the testbed was designed with characteristics similar

to the calibration runs, but with different demand patterns and slightly different initial tank levels. In this run, the demand in the first 2 hours started low, went high, then low again, as opposed to the earlier runs, which only changed once. Details are shown in Appendix A. The parameter values found during the calibration phase were used as input values, and then the goodness of fit was determined by computing the RMSE for consumer tank outflows. As an additional check, the levels in the elevated reservoirs during the simulation were compared to the corresponding data from the WADI runs. This was not a parameter that was optimised for, but because the levels in the elevated reservoirs depend directly on the outflows to the consumer tanks, a good fit in the latter should lead to a good fit in the former. If not, then there may be a problem with the model.

4.4 Results

Plots of the flows from all consumer tanks for all three runs along with the simulated flows using the three HFRs are included in Appendix A. Also shown are the total network flow and the tank levels in one of the elevated reservoirs.

	Optimal Values			RMSE			Overall	**RMSE**
HFR	γ	P_{min}	P_{des}	Run 1	Run 2	Run 3	Score	**Validation**
Wagner	1.60	0.00	0.45	0.1047	0.0642	0.0730	0.0807	0.1524
Fujiwara	2*	0.02	0.52	0.1047	0.0647	0.0733	0.0809	0.1524
Bhave	0.5*	0.23	20*	0.1298	0.0724	0.0963	0.0995	0.2319
DDA	-	-	-	0.1474	0.1514	0.1645	0.1544	0.2443

Table 4.3: HFR calibration results (*indicates a set value)

The optimal HFR parameters found for the Wagner and Fujiwara HFRs yielded results with the least error. Results for these HFRs, which allow for intermediate flow rates when pressure is between P_{min} and P_{des}, were close to identical in all runs. The Bhave HFR produced worse results, followed by the DDA model. Each HFR performed worse in the validation run, but still achieved reasonable results (with less error than the DDA model). Figures 4.4, 4.5, and 4.6 show the total network flow (the sum of the outflows at the six consumer nodes over time) for the validation run using the Wagner, Fujiwara, and Bhave HFRs, respectively.

All three HFR equations over-predicted how much the flow rate depends on the tank pressure; the flow rate changed more rapidly in the simulations than in the analogue data. In order to try to improve this, parameter values that resulted in flow depending less on pressure were manually selected. However, this resulted in overall underestimates of flow rates. i.e. the slopes of the lines were shallower, but they were translated further down. The optimal values represent something of a tradeoff between these two inaccuracies. All of the HFRs displayed overestimates of flow in some consumer tanks and underestimates in others, though this behaviour is masked when only viewing overall network flow. It is possible that better results could be obtained by finding optimal parameter values at

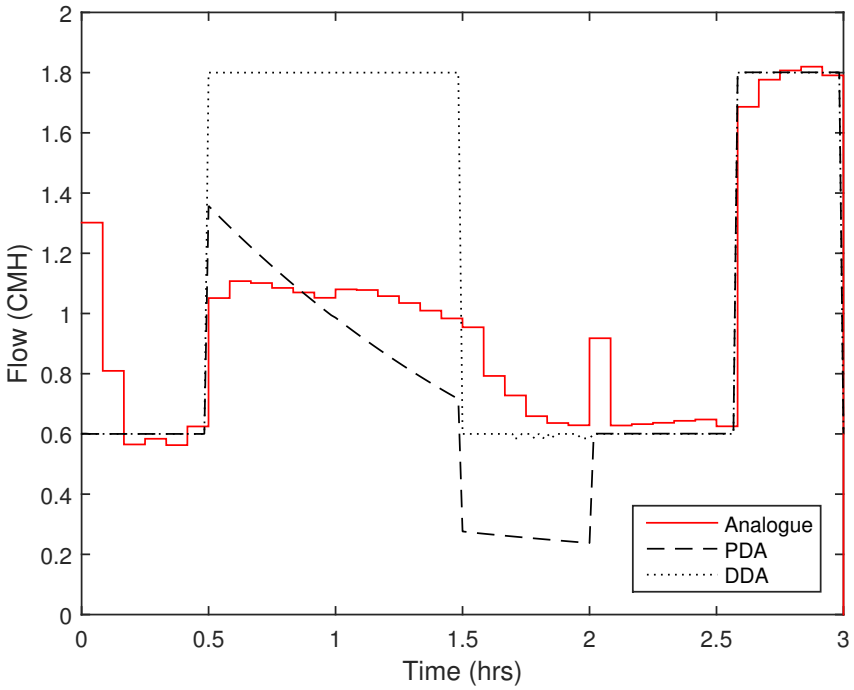

FIGURE 4.4: Validation run – Total network flow (Wagner HFR)

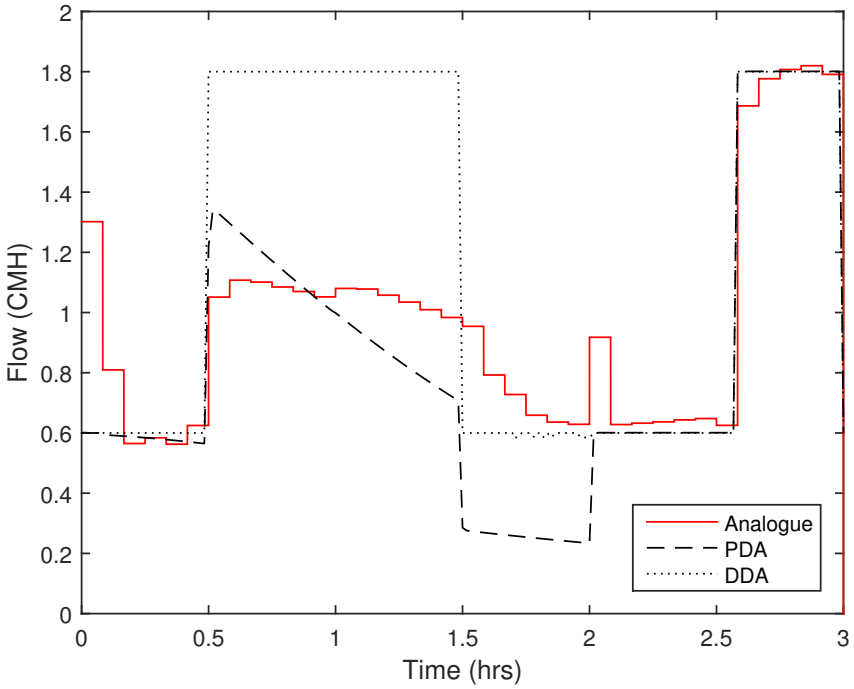

FIGURE 4.5: Validation run – Total network flow (Fujiwara HFR)

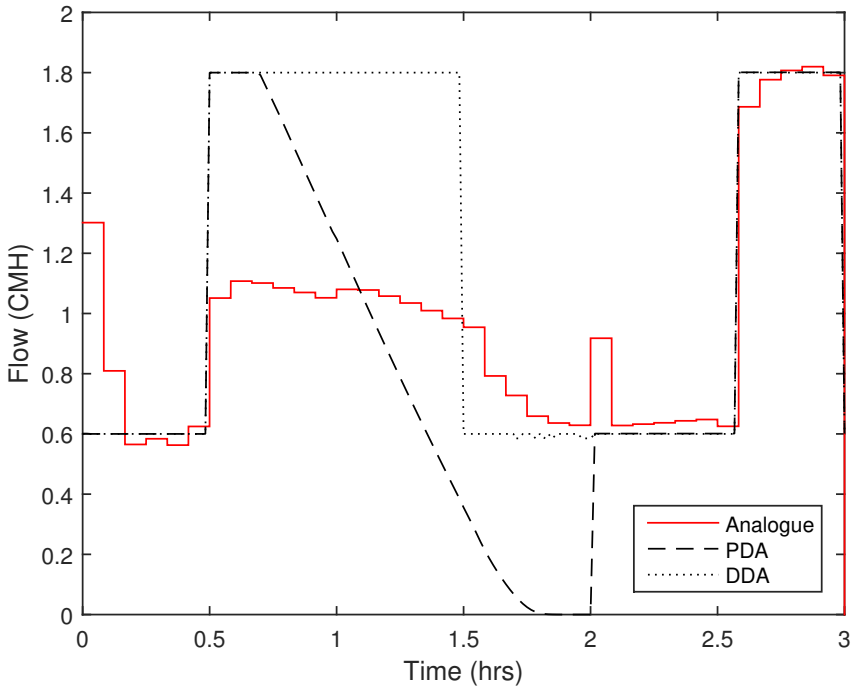

FIGURE 4.6: Validation run – Total network flow (Bhave HFR)

each individual node. The Fujiwara HFR had a constraint that the other two did not: the equation depends on the pressure from the previous time step in order to fit it to the emitter equation (see the discussion following Equation 3.3). However, the results did not appear to be adversely affected by this.

An additional way to check goodness of fit is to compare the actual and simulated tank levels in the elevated reservoirs. These are shown in Figures A.1–A.12. Because the elevated reservoir tanks in WADI are directly connected, the levels in each should be equivalent. In this case, Elevated Reservoir 2 was selected because of data recording issues with the level sensor in Elevated Reservoir 1. The tank level data support the finding from comparing flows at the consumer tanks: the Wagner and Fujiwara HFRs produce the best results, followed by Bhave, then the traditional DDA approach. In the WADI system, the elevated reservoirs are located 0.1 m above the consumer tanks and so a full tank (0.5 m) represents a head difference of 0.6 m between the tank and the consumers. This is not much above the optimal values of P_{des} found using the Wagner and Fujiwara HFRs: 0.45 m and 0.52 m, respectively. Indeed the full demand tended to be met when the elevated reservoirs were full, though this was dependent on the demand patterns.

One notable feature in the analogue data from WADI is that there appears to be a lag in adjusting flow when shifting from a higher to a lower flow rate. It tended to take around 10 minutes for the flow rate to decrease down to the demand. This may be an idiosyncrasy of the WADI system caused by the flow control valves adjusting. Because

the *epanetCPA* toolkit does not allow flow rates greater than the demand, this led to systematic underestimates of flow. Because this affected all HFRs (and the DDA model) equally, it was considered to not have a detrimental impact on calibration.

4.5 Discussion

As has been noted earlier, it is likely that the best values of HFR parameters will differ from network to network. The optimal values found for WADI will not necessarily be optimal values for other networks. This experiment nonetheless reveals important findings than can be applied to any system. Unsurprisingly, it confirmed that using any kind of PDA model (with reasonable parameters) to simulate pressure-deficient conditions yields results that are more realistic than those from a DDA model. Notably, all HFRs precisely matched the DDA model when the pumps were switched on. This experiment also offers compelling evidence that an HFR approach that allows for intermediate flow rates more closely matches real-world flows than an artificial reservoir approach (equivalent to the Bhave HFR results).

For situations where modellers cannot calibrate the HFR parameters, they will have to make educated guesses. For γ, the optimal value for the Wagner HFR was 1.60. This might seem larger than expected, as Thornton and Lambert, 2005 show in studies of pipe leakage that larger values of γ are measured when the orifice size increases with pressure. In WADI, the flow control valves should open up more when there is less pressure, the opposite effect. The nature of the HFR is such that when a node is experiencing pressure-deficient conditions (i.e. $P < (P_{des} - P_{min})$), then (all else being equal), a larger value of γ means that the flow rate actually varies *less* with pressure. Accordingly, a value of γ larger than 1.0 makes sense for the WADI system. For modelling of pressure-deficient scenarios, values of γ greater than 1.0 should be used for systems where it can be expected that consumers will attempt to compensate for low pressures by, for instance, opening taps more or opening more taps at once.

For P_{min} and P_{des}, the optimal values will depend largely on the network topology. For WADI, P_{min} was close to or equal to 0 because the consumer tank elevations were known accurately and input as the node elevations in the EPANET model. The optimal value of P_{des} was around 75-80% of the largest pressure difference experienced from tank to node. It may be that this is a rule of thumb that scales up to larger networks. As mentioned earlier, a better calibration may be achieved by allowing different P_{min} and P_{des} values at each consumer node. However, because the consumer tanks are close to identical, there is little justification in the network hydraulics for doing so. The optimal values may be influenced more by the demand patterns, which are subject to change. The emitter exponent is set globally in EPANET, so this could not be varied from node to node.

Because the Wagner and Fujiwara HFRs produced nearly identical results, the Wagner HFR may be overly complex; it has three degrees of freedom while the Fujiwara HFR has only two. This means that calibration can be achieved more quickly using the Fujiwara HFR. However, it is uncertain whether or not having the equation depend

on the pressure from the previous time step would lead to significant errors in other simulations. Adapting the toolkit to avoid this constraint could improve confidence in results using the Fujiwara HFR. The other benefit of the Fujiwara HFR over the other two is that it is smooth and differentiable, which has been shown to lead to fewer solver convergence issues (Fujiwara and Li, 1998; Siew and Tanyimboh, 2010; Elhay et al., 2016). However because the (non-differentiable) emitter equation remains unmodified in the toolkit, this benefit is nullified.

Chapter 5

Hydraulic Effects of Cyber Attacks on a Town-Scale Network

This chapter uses the modified *epanetCPA* toolkit to explore the effects of cyber-physical attacks on full-scale water distribution systems. A medium-sized benchmark network, C-Town, (which was designed to be a realistic substitute for real-world systems) was chosen. The attacks revealed a range of behaviours that could be useful for improving system vulnerability.

5.1 Goals

The goals of testing on the C-Town network were twofold: 1) conduct a sensitivity analysis to see how much the HFR parameters affect the model results for a network much larger than WADI, and 2) implement a range of illustrative attack scenarios to evaluate the impacts that cyber-physical attacks can have on full-scale water distribution systems.

5.2 The C-Town Network

C-Town is an EPANET network model that was introduced for the Battle of the Water Calibration Networks in late 2009 (Ostfeld et al., 2011). The network, shown in Figure 5.1 and summarised in Table 5.1, includes a feature common to medium-to-large municipal water distribution networks: distinct areas with their own pumping stations and storage tanks. These areas, known as district meter areas (DMAs), will herein be referred to as districts. How big is the C-Town network? By comparing the average consumption of the consumer nodes in the network (∼0.5 litres per second, LPS) to an estimate of the OECD average per capita household water consumption (∼100 kL per person per year – Grafton et al., 2011), we can estimate that each node serves around 160 people, and that the population of C-Town is about 54,000. The model thus represents a water distribution system for a small-to-medium-sized town. (Note that there are 334 junctions with a nonzero base demand, fewer than the 388 junctions.)

C-Town was chosen primarily because it is the same network used by Taormina et al., 2017 to demonstrate the capabilities of the *epanetCPA* toolkit, and so the results of these experiments could be directly compared to earlier, published results. Those authors initially selected C-Town because it is structured similarly to a real network and (having

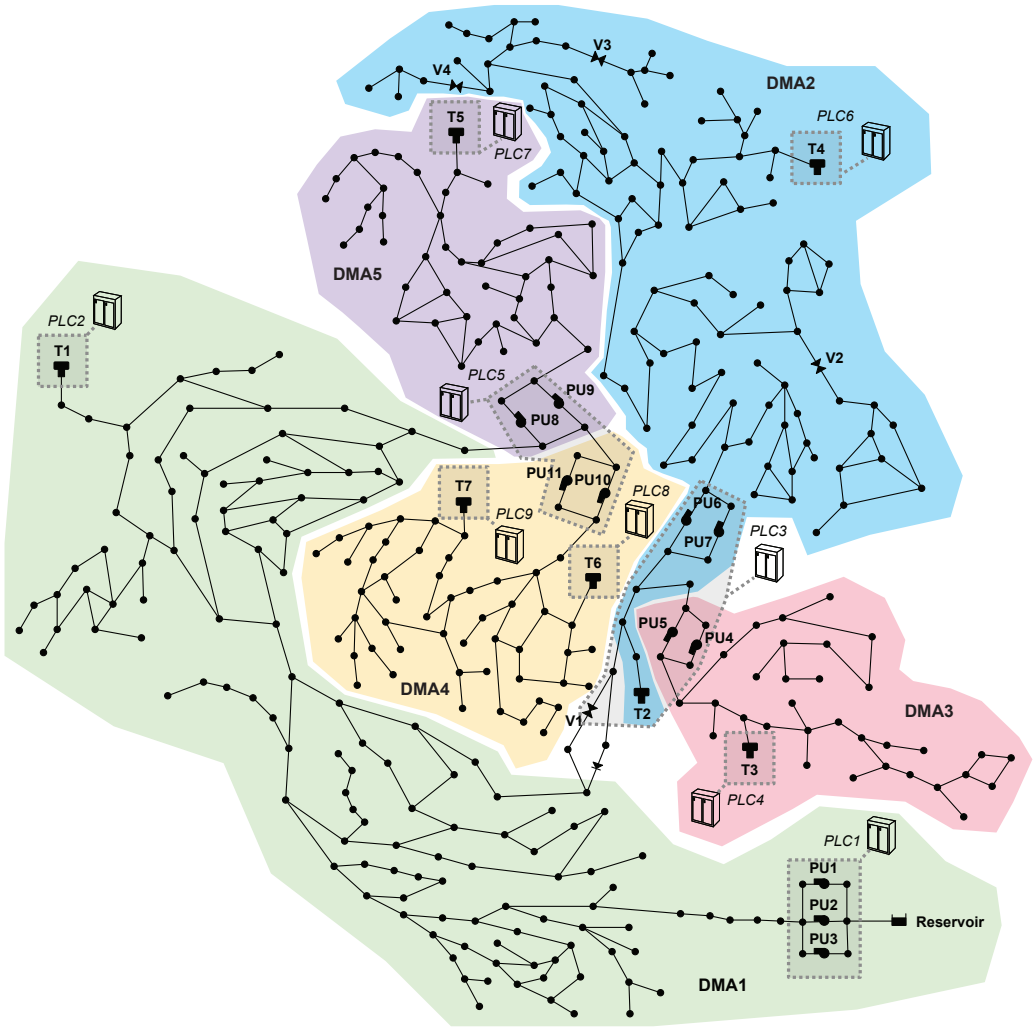

FIGURE 5.1: Schematic of the C-Town network model. Adapted from Ostfeld et al., 2011 and Taormina et al., 2017.

	District 1	District 2	District 3	District 4	District 5	**Total**
Sources	1	-	-	-	-	**1**
Junctions	137	111	36	54	47	**388**
Tanks	1	2	1	2	1	**7**
Pumps	3	2	2	2	2	**11**
Pipes	157	123	37	59	55	**429**
PLCs	2	2*	2*	2	2	**9**
SCADA system	-	-	-	-	-	**1**
Total Average Demand (LPS)	67.9	42.2	14.8	25.4	21.2	**171.6**

TABLE 5.1: Physical and cyber components and attributes of the C-Town Network; (LPS = Litres per Second); *PLC3 controls components in both District 2 and District 3

been used for studies of calibration, leakage reduction, and optimal design and operation) its behaviour is well understood.

5.3 Sensitivity Analysis

The initial step in testing the toolkit on the C-Town network was to perform a sensitivity analysis on the HFR parameters. Knowing how the network responds to changes in these parameters can reveal how critical it is to get accurate values for them. As has been previously mentioned, the complexity of a model should depend on its intended use. If the desired performance metrics are very sensitive to changes in the input parameters, then it is essential to ensure that accurate values are used. Conversely, if the performance metrics don't change much, then it should be acceptable to use "good enough" estimates of the parameter values.

5.3.1 Method

A simulation was designed with the following characteristics:

- The simulation was 24 hours in length.
- To simulate an attack, pumps PU1, PU2, and PU3 were switched off from $t = 10$ hrs to $t = 20$ hrs. This attack cut off total water supply to the network, meaning that the only sources were storage tanks, which could run dry.
- The performance metric of interest was the network-wide average Demand Satisfaction Ratio, DSR, during the attack. (Two DSR values were calculated: one during the hours when the attack was active, and one during all other hours.) The combined resilience-failure index (Equations 2.13 and 2.14) was not used because, as discussed in Section 5.4.2, setting pump flow to zero leads to an overestimation of network resilience.
- The Wagner HFR was chosen so as to investigate the effects of all three HFR parameters. Default values were set at: $\gamma = 0.5$, $P_{min} = 0$ m, and $P_{des} = 20$ m. These values have previously been used by other authors for similar simulations (see Table 4.1).

This simulation was repeated while varying one parameter at a time, so as to test the influence of each parameter individually. For each value of the given parameter, 10 simulations were conducted with randomly assigned initial tank levels and demand patterns (resulting in 260 total simulations for γ and 210 each for P_{min} and P_{des}). Varying these conditions results in a more comprehensive understanding of how the HFR parameters affect the simulation across a range of conditions. This is important because, as demonstrated by Taormina et al., 2017, the performance of the network is sensitive to such boundary conditions.

The same method used in the aforementioned paper was used for setting the tank levels and demand patterns. The initial tank levels were drawn from a matrix of simulated tank levels for the 7 tanks in C-Town during 5740 hours of operation. The toolkit selected one of these hours at random and set the initial tank levels equal to the corresponding

tank levels at that hour. For the demand patterns, the pattern multiplier values for each district at each one hour time step during a 24-hour day were determined according to a normal distribution. The mean and standard deviation of the normal distribution were set for each hour, ensuring that the resulting patterns, while somewhat random, followed a consistent diurnal pattern. Both of these methods of initialising the simulation ensured that different, though uniformly realistic results were created.

5.3.2 Results

Figures 5.2 a–c show the results of varying the values of γ, P_{min}, and P_{des}, respectively. The metric considered was the average DSR aggregated across the entire network during the 10 hours for which the attack was active. Note that the highest value of P_{min} was set at 19.99 m and the lowest values of P_{des} and γ were set at 0.01, in order to avoid division by zero in calculating the emitter coefficient (see Equation 2.10). All other values are as they appear in the figure. Values of γ below 0.3 and above 2.0 caused the EPANET hydraulic solver to fail to converge in almost all simulations. This also occurred for the extreme values, $P_{min} = 19.99$ m and $P_{des} = 0.01$ m. These results were excluded from the analysis.

γ and P_{min} do not appear to have a clear effect on the average DSR during the attacks. The median average DSR is relatively constant across all successfully tested values of these parameters. For P_{des}, the average DSR during the attacks does not appear to have a clear trend until around $P_{des} = 55$ m. For greater values of P_{des}, a higher pressure threshold tended to result in lower average DSR. For all sets of simulations, the randomly-determined boundary conditions (tank initial levels) and simulation setup (demand patterns) had a much larger effect on average DSR than the HFR parameter values. The range of average DSR values across simulations with any one parameter value is generally greater than the difference in average DSR between the smallest and largest parameter values.

5.3.3 Discussion

None of the HFR parameters appear to dramatically affect the results for assessing the effects of cyber-physical attacks at a network level, unless extreme values are used. This is likely because the pressures in the network were sufficient to meet demand while the tanks had water in them, so intermediate nodal outflow rates weren't often experienced. That is, the nodal outflow tended to be either full or zero. It is thus reasonable to conclude that educated estimates for the HFR parameters should be sufficient for this type of analysis. This further suggests that HFR parameters may not need to be assigned on a node-by-node basis; a single set of parameters for each district, or even the entire network, will likely produce adequate results. Similar work on the pressure dependence of leakage rates suggests that assuming a single exponent value for the flow equation across the whole network is sufficient for large-scale modelling (Thornton and Lambert, 2005).

The initial levels in the tanks and the demand patterns had a more significant effect on DSR than the HFR parameter values. Most modern water distribution systems have

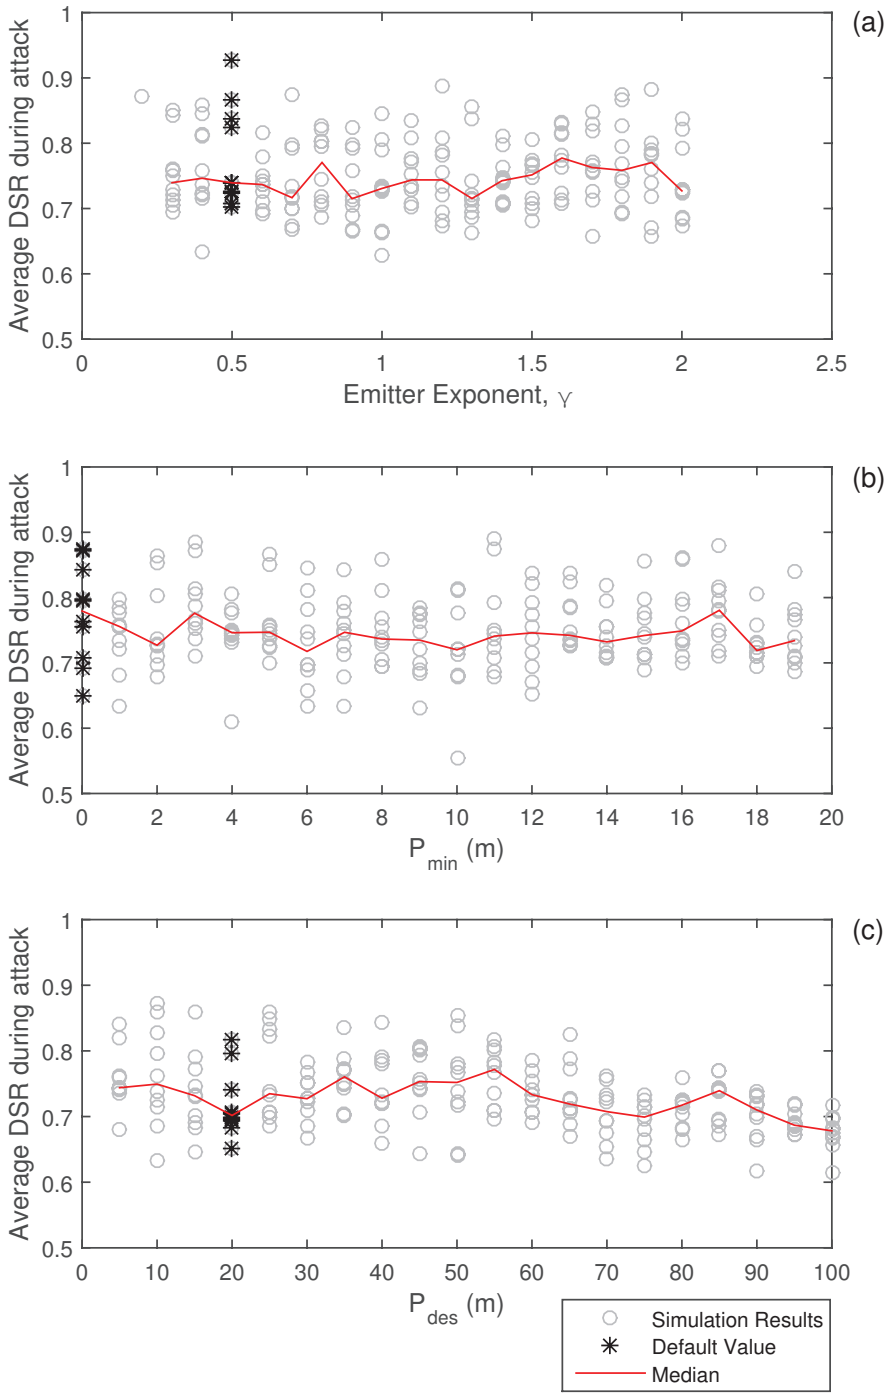

FIGURE 5.2: Sensitivity analysis: Average DSR during attack for varying values of the HFR parameters: (a) γ, the emitter exponent; (b) P_{min}, the minimum pressure at all nodes, below which the flow rate is zero; and (c) P_{des}, the desired pressure at all nodes, above which the flow rate is constant.

good, real-time data on tank levels; the more challenging factor is demand patterns. While an operator may have a good idea of diurnal trends in demand at a network or district level, finer resolution data (such as from smart meters) can give a more accurate prediction of future demand (Cominola et al., 2015). Accurate demand patterns form just as important a factor in pressure-driven modelling of cyber attacks as the HFR parameters, if not more so.

The finding that the EPANET solver would occasionally fail to converge is in line with prior studies that demonstrated how non-differentiable HFR equations (such as the emitter equation used in this study) can lead to more instances of non-convergence (Fujiwara and Li, 1998; Siew and Tanyimboh, 2010; Elhay et al., 2016). The HFR equation proposed by Wagner, Shamir, and Marks, 1988 and used in this work is not differentiable at the threshold values. However, the solver only failed to converge when extreme values of the HFR parameters were used. Selecting parameter values within a reasonable range should prevent convergence issues from occurring when using this modelling approach.

5.4 Illustrative Attack Scenarios

5.4.1 Method

A range of attack scenarios were designed to test different aspects of the model and produce different results in the C-Town system. Attack scenarios 1 and 2 were designed to allow for comparison with the earlier *epanetCPA* toolkit. Attack scenarios 3, 4, and 5 were designed to produce pressure-deficient conditions and were repeated multiple times. All simulations were 24 hours long and used the same default HFR parameter values used in the sensitivity analysis. The attacks were implemented as follows:

- Scenario 1: Tank T2 was prevented from refilling by targeting the communication between the T2 level sensor and PLC3, replacing the sensor's level readings with a constant high value. (The readings received by the PLC, as well as the actual level, are shown in Figure 5.3.) This prevented valve V1 from opening, cutting off Districts 2 and 3 from the source. When the tank level in T2 dropped below 0.1m, the attack was stopped. This simulation had a hydraulic time step of 5 minutes and was carried out using the original DDA model and the new PDA model. The simulation was also repeated without the attack, for comparison.

- Scenario 2: Similar to attack scenario 1, but tank T4 was instead prevented from refilling by targeting the communication between the T4 level sensor and PLC6, replacing the sensor's level readings with a constant high value. This switched off pumps PU6 and PU7, cutting off District 2 from the source.

- Scenario 3: Same as scenario 1, but the attack lasted for 12 hours, from $t = 8$ hrs to $t = 20$ hrs, i.e. the attack continued even when tank T2 was completely empty. This scenario was repeated 100 times with randomised initial tank levels and demand patterns, and had a hydraulic time step of 15 minutes to reduce computation time.

- Scenario 4: Same as scenario 2, but the attack lasted for 12 hours, from $t = 8$ hrs to $t = 20$ hrs, i.e. the attack continued even when tank T4 was completely empty. This scenario was also repeated 100 times as above.

- Scenario 5: The actuators controlled by PLC1 were directly attacked, switching off pumps PU1, PU2, and PU3. This completely cut off the network from the water source. This attack could equally have been achieved by attacking the communication between the PLC and the actuators, or by a complete takeover of the SCADA system. It was similar to the attack used in the sensitivity analysis, but in this case the pumps were switched off at a random start time, for a random duration. Both values were selected from a normal distribution with the parameters shown in Table 5.2, below. This scenario was repeated 1,000 times with randomsed initial tank levels and demand patterns, and had a hydraulic time step of 15 minutes.

	Mean, μ	Standard Deviation, σ
Start time (hrs)	8	4
Duration (hrs)	10	4

TABLE 5.2: Distribution of attack scenario 5 parameters

5.4.2 Results

No convergence issues were experienced during any of the five attack scenarios.

Scenarios 1 & 2

Attack scenario 1 closed valve V1, preventing Tank 2 from filling. Attack scenario 2 shut off pumps PU6 and PU7, preventing Tank 4 from filling. Figures 5.3 and 5.4 each show the result of 4 simulations: using the modified *epanetCPA* toolkit (results labelled "PDA") both with and without an attack, and using the original *epanetCPA* toolkit (results labelled "DDA") both with and without an attack. Each figure displays the levels in the tanks targeted by the corresponding attack scenario. Both scenarios exhibit near perfect agreement between the pressure-driven and demand-driven models. This should be expected, as these parts of the network experienced only pressure-sufficient conditions. Also shown are the tank levels that were reported back to PLC3 and PLC6, including the erroneous values reported during the attack that prevented the tanks from refilling.

The levels in the tanks that were not targeted were also compared, and the only significant deviations between PDA and DDA simulations were observed in Tank 7. (See Figures B.1 & B.2.) Further investigation revealed two nodes near Tank 7 in District 4 where the demanded water was not being supplied, even when attacks were not implemented. These two nodes had elevations only slightly below Tank 7: 97.07 m and 99.05 m, compared to 102.00 m. The head difference between Tank 7 and the nodes was thus less than the desired pressure, P_{des}. This may indicate that pressures in the network are inadequate for these two nodes, and so at least in a small section of the network, a DDA model will produce inaccurate results. This is an example of the problem discussed in Section 3.3.1.

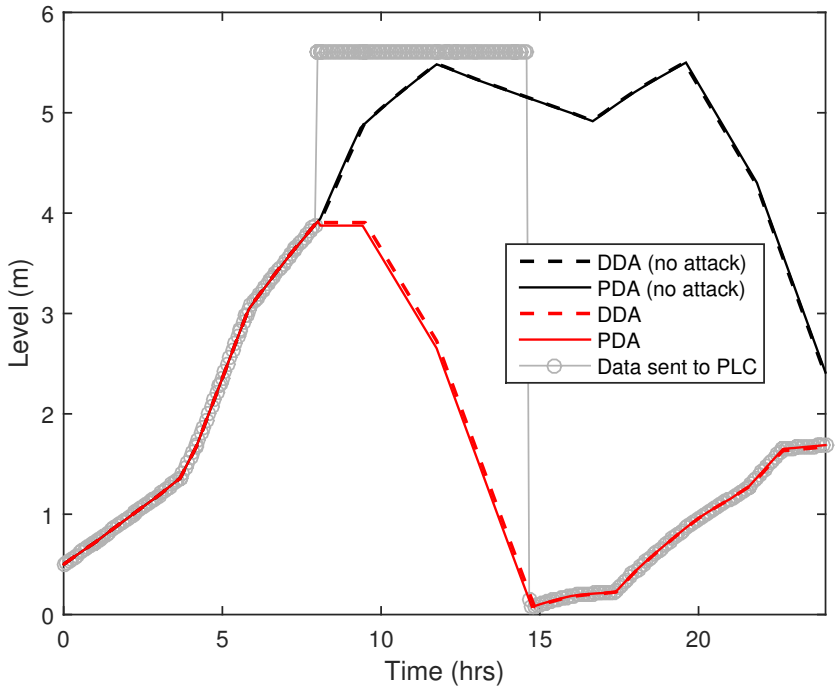

FIGURE 5.3: Attack scenario 1 – Tank 2 level over time

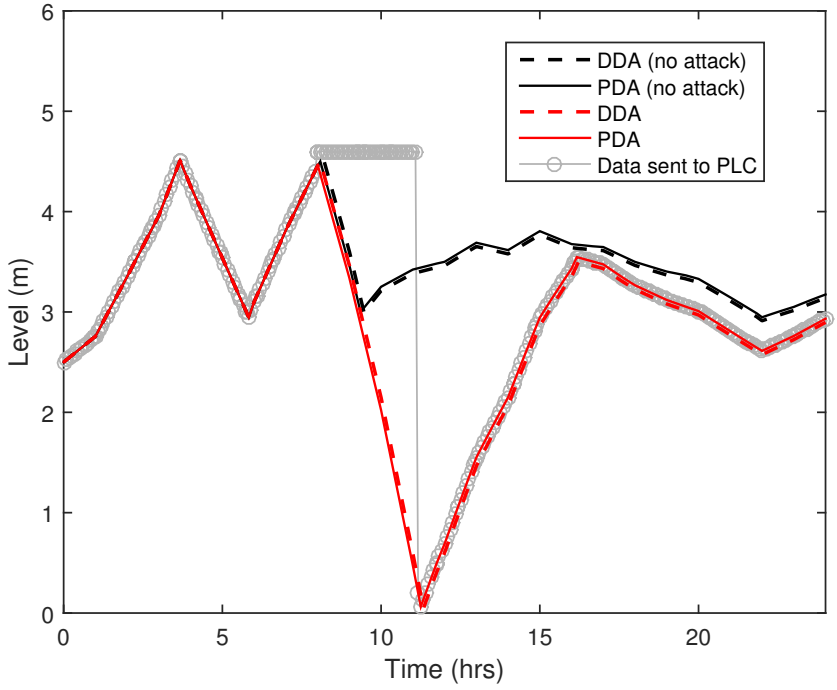

FIGURE 5.4: Attack scenario 2 – Tank 4 level over time

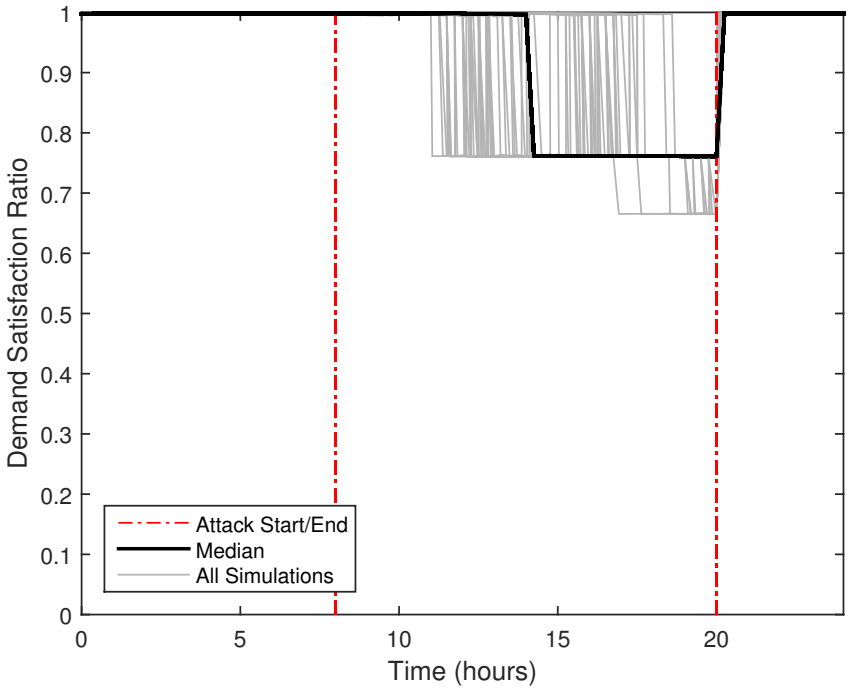

FIGURE 5.5: Attack scenario 3 – Network-wide DSR over time

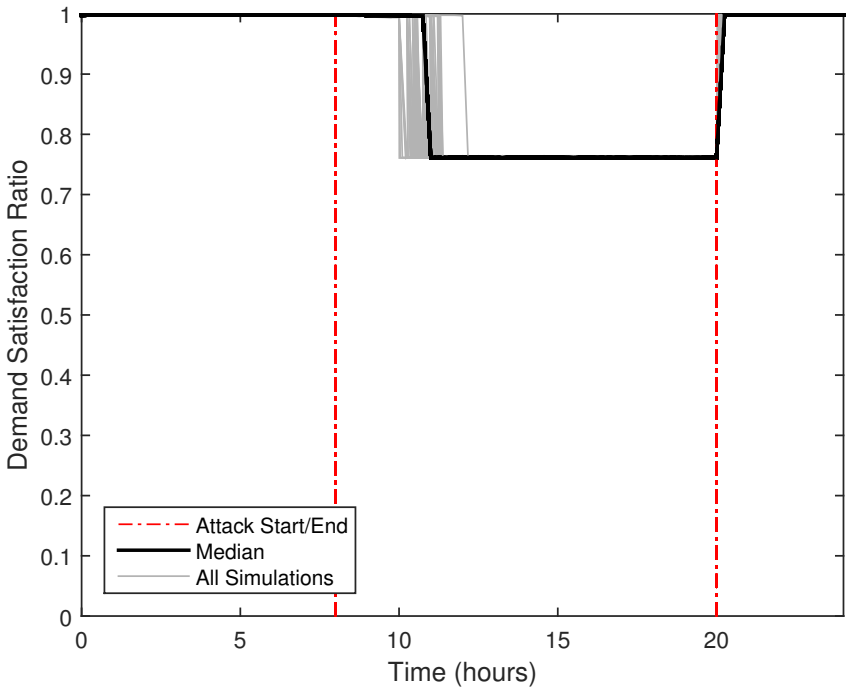

FIGURE 5.6: Attack scenario 4 – Network-wide DSR over time

Scenarios 3 & 4

Attack scenario 3 closed valve V1, cutting off Districts 2 and 3. Figure 5.5 shows the demand satisfaction ratio, DSR, for the entire network over time under attack scenario 3. Each of the grey lines represent the results from a single simulation, while the black line indicates the median value at each point in time. Note that though the attack always started at $t = 8$ hrs, a drop in the DSR was universally not experienced until at least 3 hours later. All of the simulations show a cut off in water supplied to District 2 (the first drop), and some show a cut off in supply to District 3 (the second, smaller drop). In every simulation, the water supply went from 100% to 0 almost immediately for the affected district once the corresponding tank ran dry. Similarly, the DSR immediately returned to 1.0 when the attack was stopped and the valve was reopened. (See also Figure B.3, which shows DSR over time at the district level.)

Attack scenario 4 shut off pumps PU6 and PU7, cutting off District 2. Figure 5.6 shows the equivalent plot for attack scenario 4. In this case, there is no second drop in DSR because only District 2 is affected. Compared to scenario 3, the drop in DSR tends to occur sooner, and there is less variation in the timing. This is because Tank 4 is much smaller than Tank 2, and so there is less water to continue supplying demand once District 2 is cut off from the source. (See also Figure B.5.)

The combined resilience-failure index, I_{rf}, (Equations 2.13 and 2.14) was also calculated, but it did not effectively illustrate the impact of the attack. When the nodal outflows in District 2 went to zero (lowering the numerator), the flows from Tank 4 also went to zero (lowering the denominator). At the network level, the product of nodal outflows and heads ($q_{user}H$) still exceeded the product of demand and desired head (dH_{des}), meaning that the resilience index dominated the failure index. These factors meant that the effect of the attack was only visible at the district level. Additionally, the variations in demand and flow led to variations in the value of I_{rf}, leading to a noisier dataset (see Figures B.4 & B.6).

Scenario 5

Attack scenario 5 shut off pumps PU1, PU2, and PU3, cutting off the entire network. The attack start times and durations were randomised. Figure 5.7 shows the network-wide Demand Satisfaction Ratio over time for all 990 simulations, with the DSR at the end point of the attack represented by a dot. (In 10 simulations the attack start time was randomly selected to be greater than 24 hours.) Predictably, longer attacks tended to result in a lower DSR. It is notable that no drop in supply was experienced before 1 hour after the attack started and significant drops in supply were not experienced until at least 2 hours into the attack. Conversely, there were some attacks up to nearly 8 hours in length that had no effect on DSR. The stepwise clustering observed in the final DSR values was due to the fact that the DSR in each district tended to go from 1.0 to 0.0 very rapidly when the corresponding tank ran dry. This is the same effect observed in attack scenarios 3 and 4, but applied to more districts.

Figure 5.8 shows the minimum DSR experienced during the attack for all 5 districts,

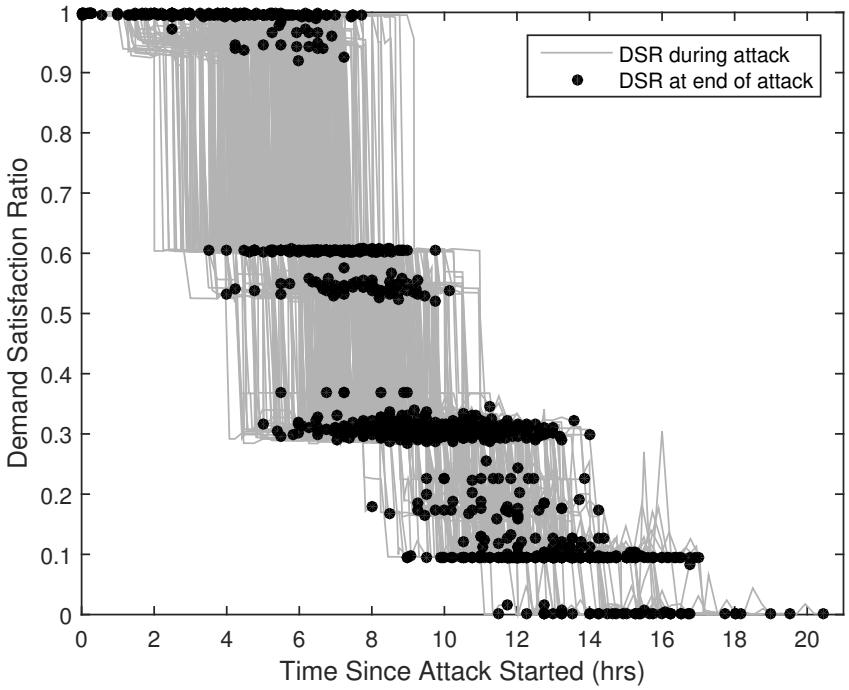

FIGURE 5.7: Attack scenario 5 – network-wide DSR over time during the attack for all simulations.

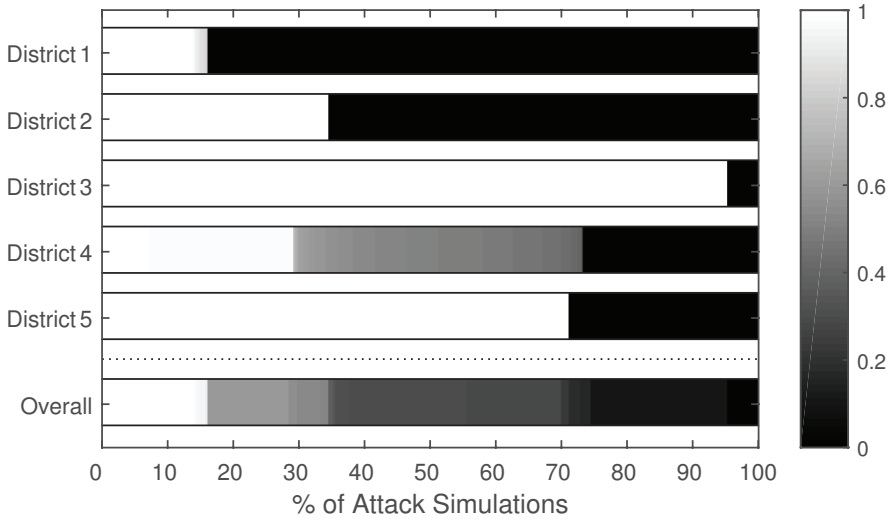

FIGURE 5.8: Attack scenario 5 – minimum DSR experienced during the attacks for each district and across the entire network.

as well as for the overall network. White indicates that the area experienced full supply (a DSR of 1.0) for the entire simulation, while black indicates that supply completely cut out at some point (a DSR of 0.0). For districts 1, 2, 3, and 5, there were few or no simulations where an intermediate level of supply was experienced; that is, the water supply tended to cut out very suddenly in these districts. From this plot it is clear that the most vulnerable district is District 1, where the water was completely cut off (herein referred to as complete failure) during 83.6% of attacks. Districts 1 and 2 were likely the most vulnerable because they contain tanks that also supply water to other districts. District 3 was the least vulnerable; it retained full supply during all but 4.5% of attacks. District 2 (65.2% complete failure rate) was more vulnerable than District 5 (28.7%), while District 4 was arguably somewhere in the middle; it completely failed 26.3% of the time and retained partial supply during a significant number of simulations. Partial supply was experienced in District 4 when tank T7 ran dry; tank T6 could only supply a fraction of the demand.

5.4.3 Discussion

Attack scenarios 1 and 2 showed excellent agreement between the modified, pressure-driven toolkit and the existing *epanetCPA* toolkit when pressure-sufficient conditions were simulated. This agreement was observed for simulations with and without attacks. It can thus be concluded that the PDA toolkit does not adversely affect the robustness of the hydraulic model. That is, the modified toolkit should be a reliable tool in simulating a range of scenarios both pressure-sufficient and pressure-deficient. Furthermore, it can be used to check whether there are any nodes that may be experiencing pressure-deficient conditions, as was observed in District 4. While the approach introduced in this work is compatible with reduced or skeletonised network models, using this kind of model could preclude nodal-level insights such as this.

The results of attack scenarios 3 and 4 showed that in a system like C-Town, there may be little warning before water is cut off entirely to a district. Because the tanks in C-Town had sufficient head to service typical demand right up until they were empty, consumers did not experience a marked decrease in pressure before the water supply cut out entirely. It would not be possible to know this using a traditional DDA model. Attack scenario 5 confirmed that this occurred in Districts 1, 2, 3, and 5. This tendency for sudden drops in supply could be improved by implementing emergency controls, such as flow control valves, that are activated when a tank reaches a critically low level. In such a case, the supply may not meet demand, but there would still be some supply available for critical customers, such as hospitals. Perhaps even more importantly, if an attack detection method is to just look for unusual flows in the network, then an attack may not even be detected until the water supply is completely depleted. Attack detection algorithms should be designed to detect such events before they become so critical. Attack detection is an area of active research that could be aided through the use of the modelling approach proposed in this work (Taormina et al., 2016).

Attack scenario 5 revealed two key observations. One is that some districts were clearly more vulnerable to attacks than others. Performing this kind of analysis on real-world networks will help identify the most critical areas for network improvements, such as

installing additional storage capacity or additional sensors that use a different communication protocol. In the latter case, if one sensor is compromised, the other could reveal the attack. Limited funds can thus have a bigger impact in reducing overall network vulnerability. Another key finding was identifying that there appeared to be a attack duration below which attacks had little to no impact (1-2 hours in this case). This can be seen as the operator's "window of opportunity" for identifying attacks and rectifying them before the network is affected. The shorter the window, the more crucial it is for water distribution system operators to have rapid response capabilities.

It is clear from these experiments that the effects of attacks will most likely not be shared equally by all customers across the network. Attacks scenarios 3, 4, and 5 demonstrated complete shutoffs of water supply to some districts, while the other districts were unaffected. This problem of equity was identified by Fujiwara and Li, 1998, and could be improved by incorporating redundant connections that link one district to another. These connections would normally be closed, to ensure that the districts could still operate independently, but they could be opened in the event that a pumping station breaks down. In such a scenario, the pumping station for one district could potentially supply water to two districts. This would probably result in some unmet demand across both districts, but it could mean a more equitable spread of an attack's impacts. Determining whether or not a system could cope with such modifications would require further modelling.

The combined resilience-failure index does not appear to be a useful metric when it comes to assessing the performance of water distribution systems subject to cyber-physical attacks. This metric was designed for comparing different network topologies, as opposed to measuring network performance over time. Instead, a metric such as demand satisfaction ratio should be used. Other metrics, such as the percent of customers experiencing less than a certain fraction of demand (which would be highly correlated with DSR), could also be used. Deciding which metric to use will depend on the purpose of the analysis and its intended audience.

Chapter 6

Conclusions

The overall goal of this work was to develop and test a modified version of the *epanetCPA* toolkit in order to simulate pressure-deficient conditions experienced due to cyber-physical attacks on water distribution systems. The experiments conducted in Chapters 4 and 5 showed that the toolkit presents an improvement over the previous DDA toolkit, producing more realistic simulation results. They also demonstrated that the toolkit can aid in understanding the behaviour of real-world systems, and so be a useful design and planning tool to improve network resilience and reliability. Specific findings and recommendations for future work follow.

6.1 Findings

Calibrating the toolkit against real-world flow data from the WADI network revealed useful findings for simulating cyber-physical attacks on water distribution systems.

- Any pressure-driven modelling approach, when using reasonable parameters, will produce more realistic results than a traditional demand-driven approach under pressure-deficient conditions. Most PDA models should also perform equally well as DDA models for pressure-sufficient conditions. This is a finding that was backed up by the simulations on C-Town.
- Approaches that incorporate a head-flow relationship that allows for intermediate flow rates yield better results than approaches that just use artificial components. This is dependent on obtaining reasonable values for the HFR parameters.
- For the HFR parameters, the value of γ should be based on how the consumers are expected to respond to low-pressure conditions. The values of P_{min} and P_{des} should be based on network topology and operating conditions.

Running attack simulations on the town-scale C-Town network produced results that have useful implications for water distribution system design and operation.

- The head-flow relationship (HFR) parameters used in this pressure-driven model did not have a large effect on network performance. If the goal of the simulations is to assess the effects of attacks at a large scale, say at the network or district level, then it should be adequate to use educated estimates of the HFR parameters. The initial levels of tanks in the network and the consumer demand patterns had a larger influence on network performance. It is more important to obtain accurate data on these factors than it is to calibrate the HFR parameters in order to get good results.

- This pressure-driven hydraulic modeling approach produced nearly identical results to the demand-driven *epanetCPA* toolkit under pressure-sufficient conditions, and so should be an appropriate approach to use for all conditions. It also revealed a section of the C-Town network experiencing pressure-deficient conditions during normal operations, highlighting a possible inaccuracy introduced by using a demand-driven model.
- Using this modeling approach revealed that the hydraulic response of the C-Town network to attacks was for flows to cut out very suddenly, without noticeable preceding drops in flow for most consumers. This has implications for attack detection (manual or automated) and the design of emergency controls that could help provide water to critical customers.
- In the C-Town network, certain districts were shown to be more vulnerable to attacks than others. This presents a problem to network operators who value equitable water supply. Recognising and quantifying these differences in vulnerability is an important first step in deciding how to best allocate resources to improve overall network reliability. This toolkit can thus be an important tool for network operators to aid in design and planning.
- When attacks were short enough, water was able to be supplied across the network without any service interruption. This modelling approach allowed for the identification of a "window of opportunity" for responding to an attack before it adversely affects water supply. Identifying this window could help network operators in developing cyber-physical attack response and contingency plans.

6.2 Recommendations for Future Work

One aspect of the modified toolkit not considered in this work is the effect that it has on simulation runtime. The modifications that allowed for pressure-driven modelling did significantly increase the time for a given simulation to run. However, *epanetCPA* is still undergoing improvements and refinements, and it is expected that as the PDA modifications are integrated into the latest version, then significant runtime reductions can be achieved.

For the course of this study, pipe leakage was not considered. Leakage is highly dependent on the characteristics of the network, and varies widely from location to location. Conventional approaches for modelling leakage in EPANET, such as adding emitters throughout the network, should have no trouble being integrated into the modified toolkit. Considering leakage in future studies would help ensure that the simulation results accurately predict real-world performance, but of course this depends on having a good understanding of leakage in the network being modelled.

Relatedly, inducing pressure-deficient conditions in a network may result in physical damage to network components. For example, pumps forced to continue running when the flow rate is below their specification (or indeed when there is no water supply) are liable to break and require repair before coming back online (McKee et al., 2011). Incorporating such factors into the toolkit for extended period simulations would require a

good understanding of different component failure modes, but could add greater realism to the results.

Now that the *epanetCPA* toolkit is capable of modelling pressure-deficient conditions, simulating a wide range of attack scenarios is possible. A natural extension of this work, then, would be to simulate the behaviour of attackers in a more mature way, incorporating the attackers' goals into the model and changing the attackers' behaviour based on the outcome of previous simulations. In this way, an "intelligent" attacker could be simulated that optimises its actions to maximise the impact on the network. Such an approach could go deeper than this present work to reveal yet further vulnerabilities in C-Town or any other network.

Another extension of this work would be to consider the broader, societal ramifications of cyber-physical attacks that target water distribution systems. Reductions in the quality and/or quantity of drinking water supplied to a population can have serious economic, environmental, and human health impacts. The tools developed in this work could form the basis of a model that quantifies these impacts (possibly in conjunction with EPANET's water quality modelling capabilities). Such a model would be of interest and use to a wide range of sectors from government to insurance.

Beyond these academic considerations, the modified *epanetCPA* toolkit would be most useful when used in the real world. Water distribution system operators can use this toolkit to test the resilience of their systems to a range of simulated cyber-physical attacks, potentially identifying vulnerabilities that would have otherwise gone unnoticed before it was too late. If this thesis can help utilities to prepare for and prevent the worst effects of such attacks, then it will have done its job.

Appendix A

Calibration Results

The following pages include figures showing the actual (post-processed) flows and the modelled flows from the WADI consumer tanks, using the optimal HFR parameters. Also shown are the total network flows (the sum of consumer tank outflows) and the level of water over time in one of the elevated reservoirs. The data for each of the four runs of WADI are shown on separate pages, and one of the three HFRs is shown per page, making for a total of 12 pages.

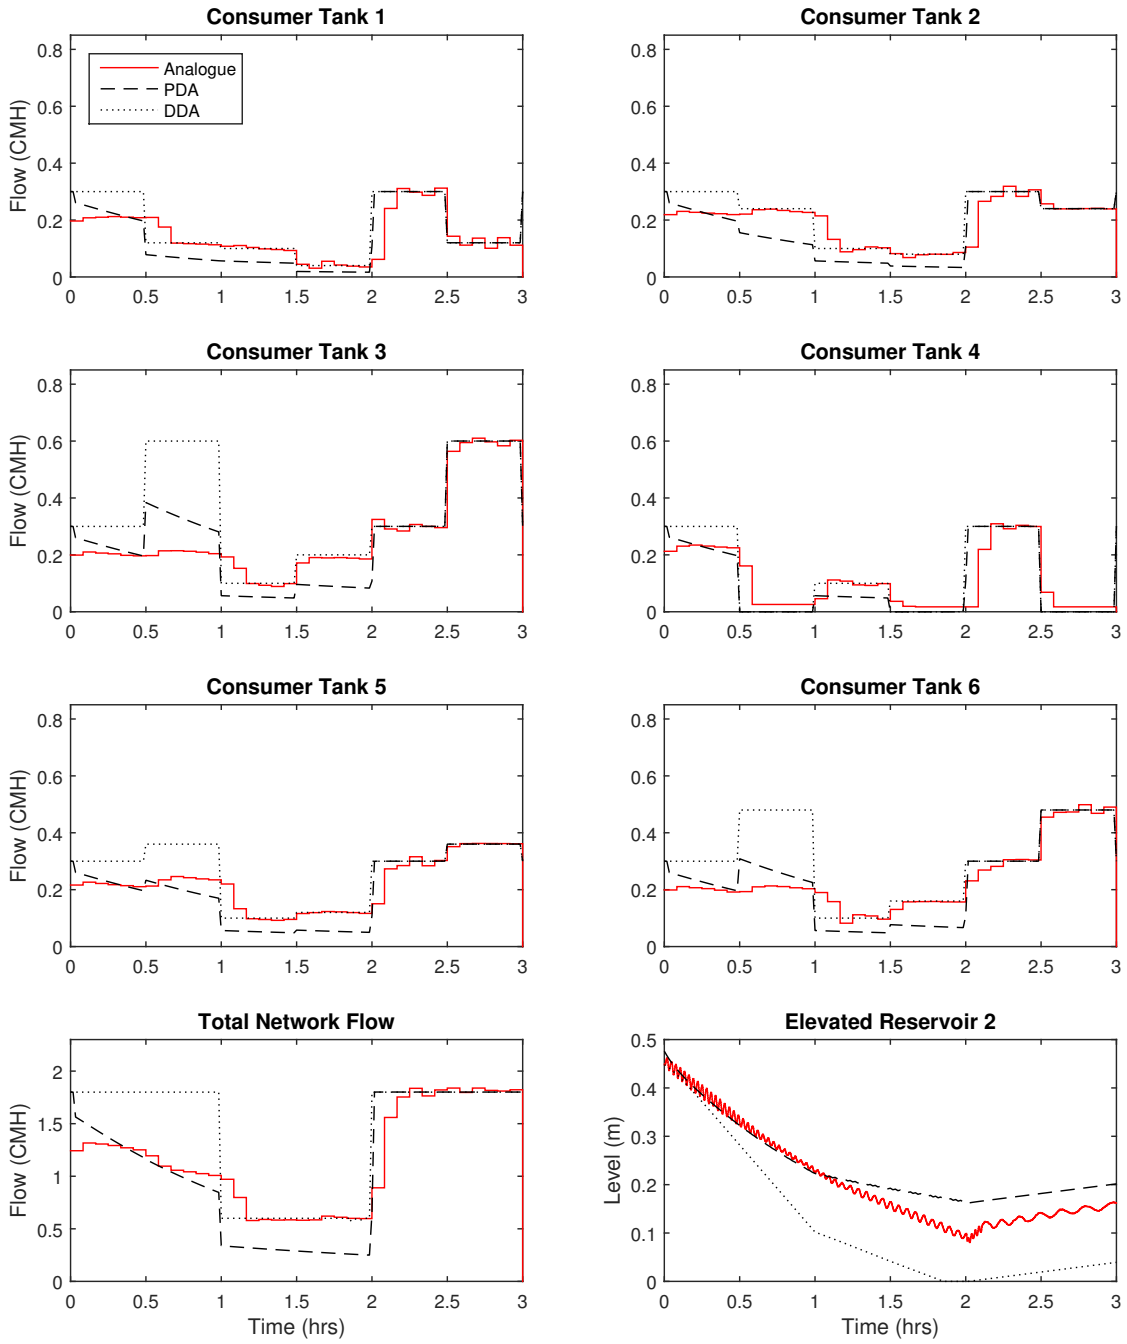

FIGURE A.1: Run 1 (Wagner HFR)

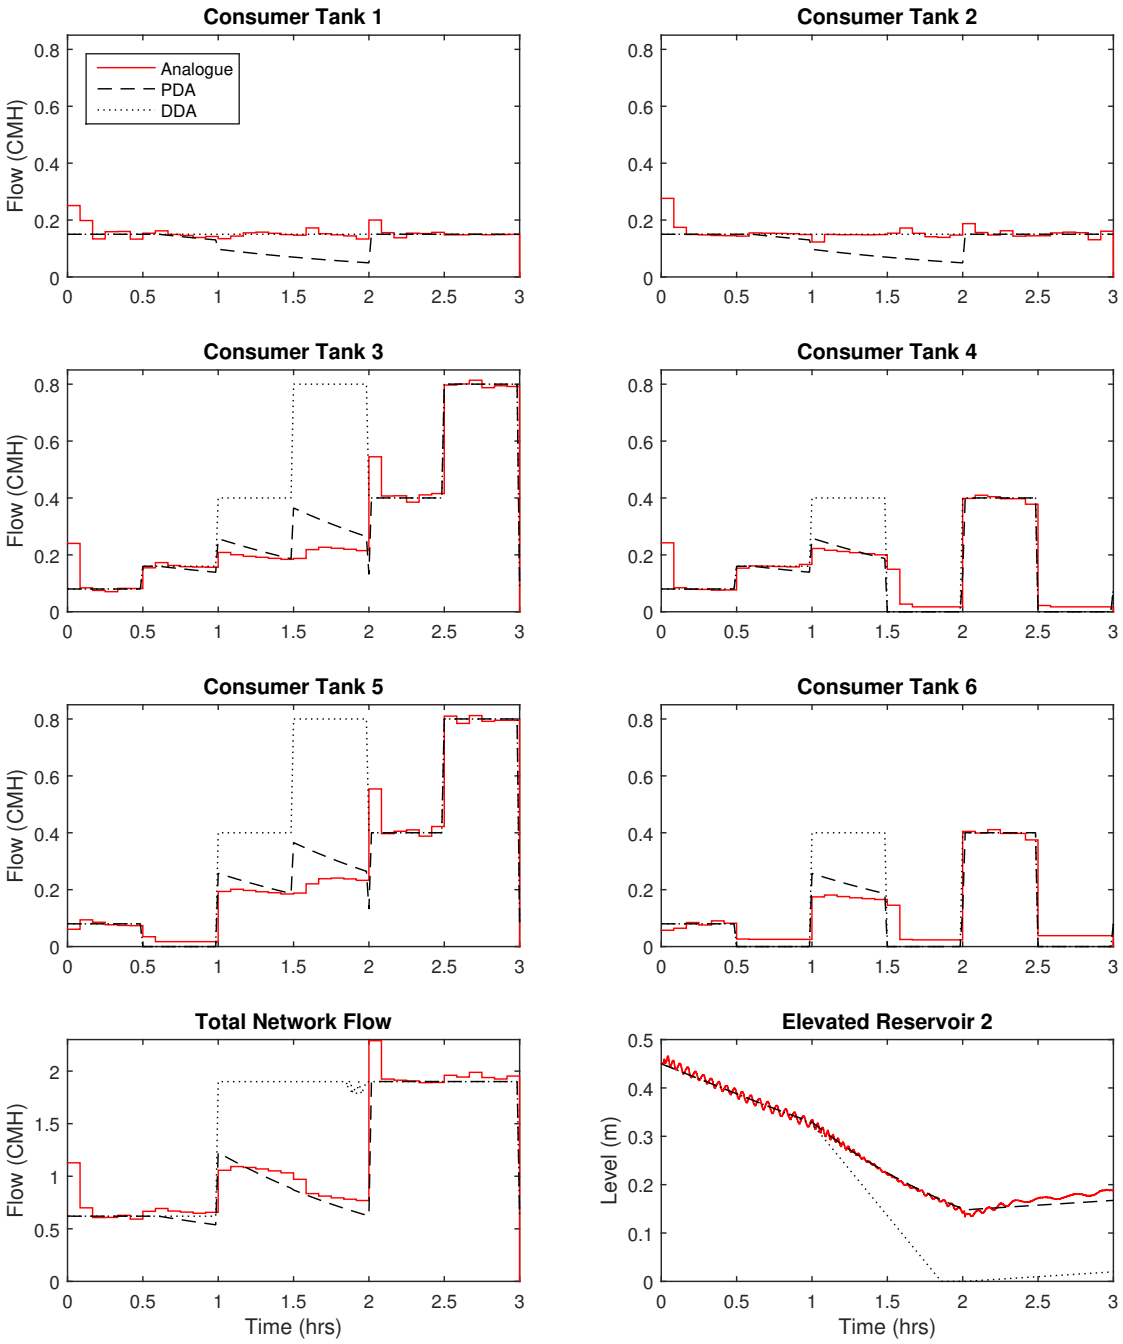

FIGURE A.2: Run 2 (Wagner HFR)

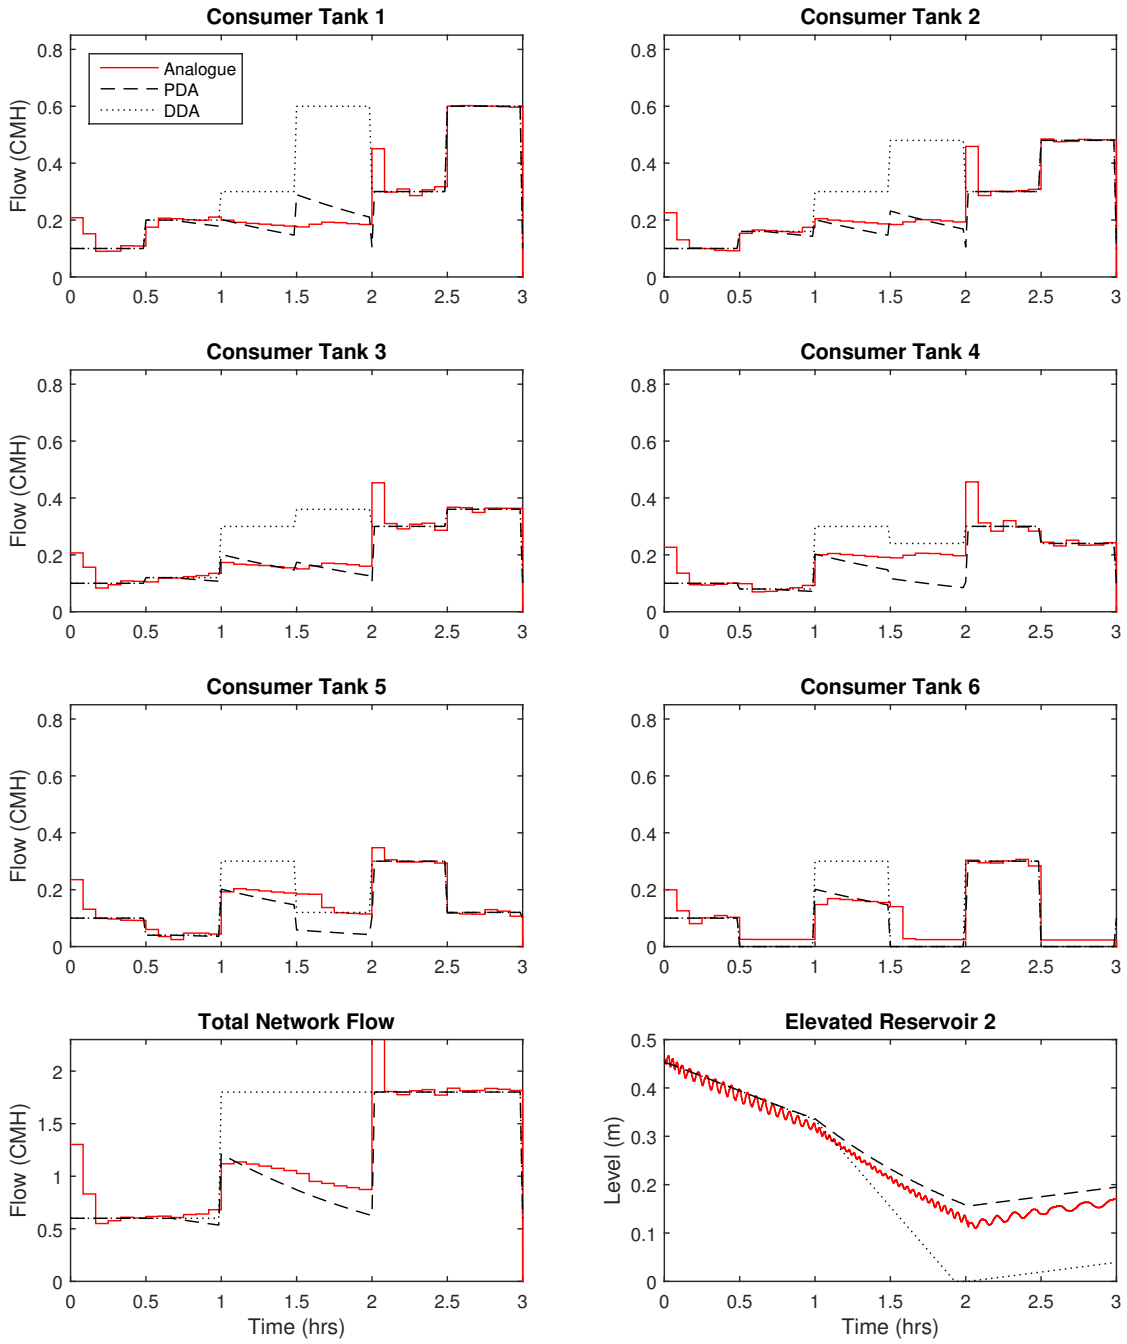

FIGURE A.3: Run 3 (Wagner HFR)

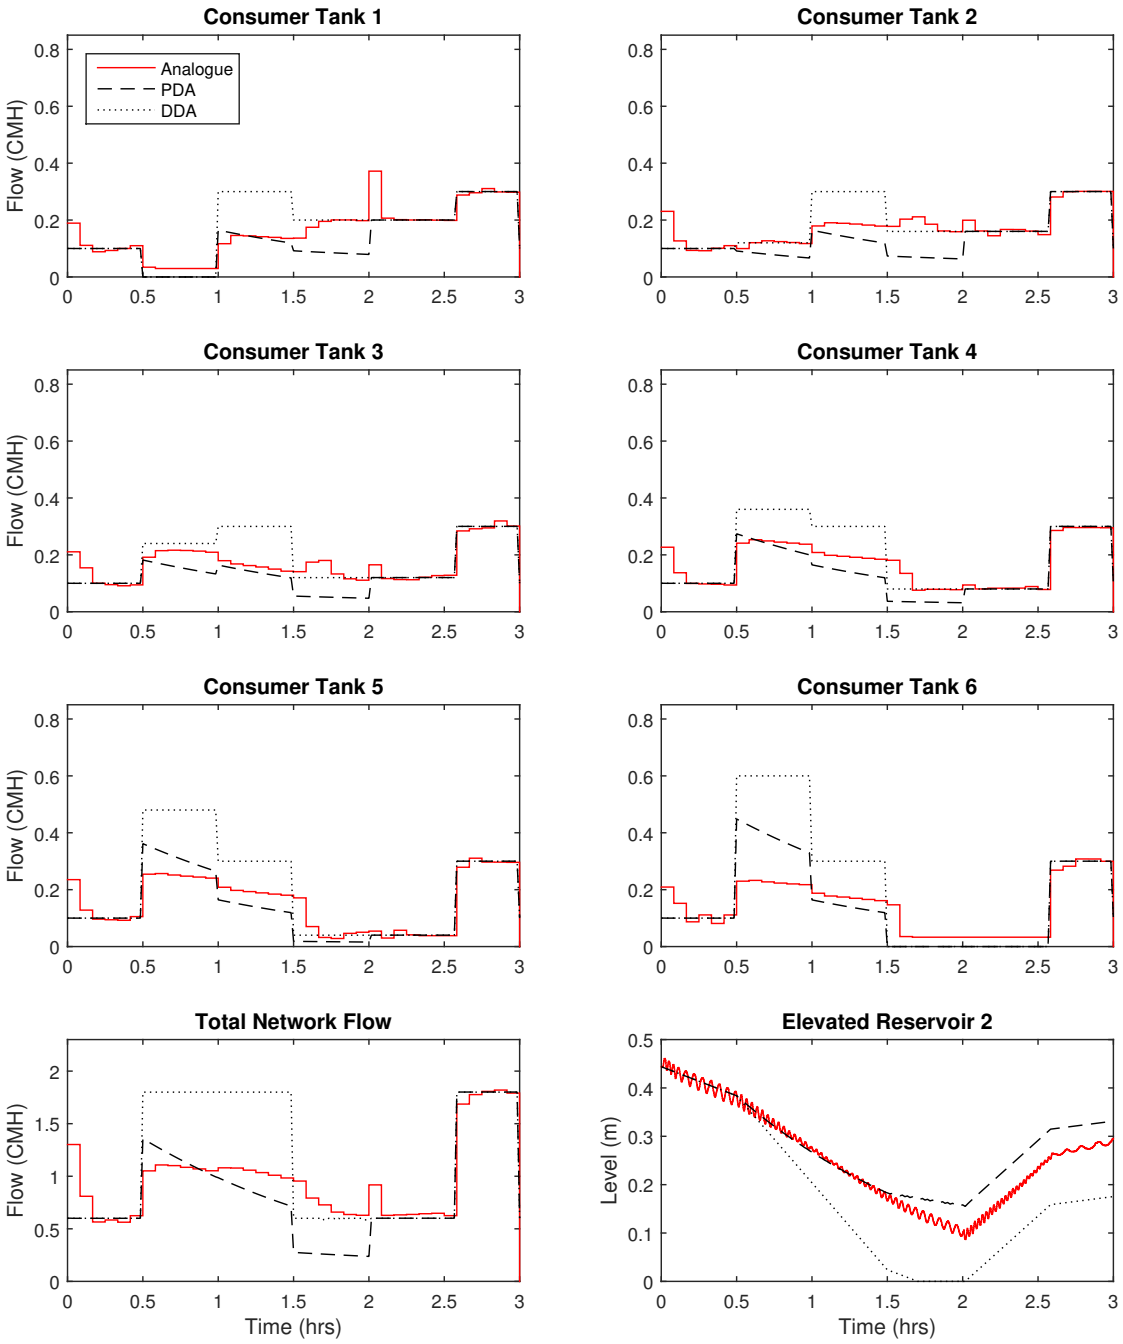

FIGURE A.4: Validation run (Wagner HFR)

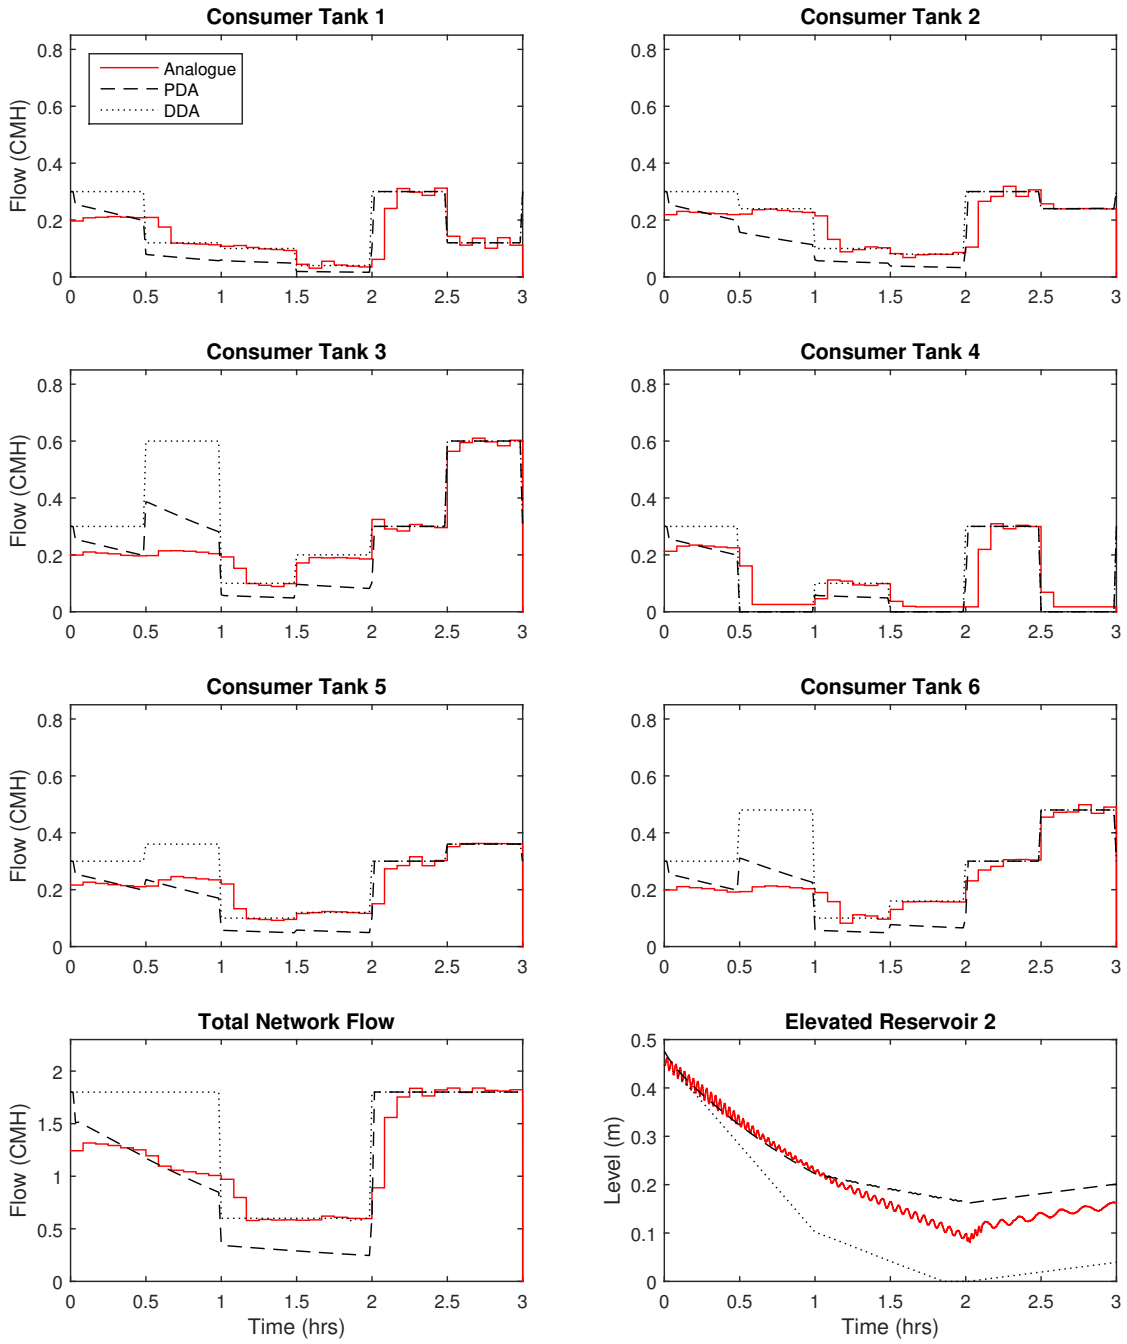

FIGURE A.5: Run 1 (Fujiwara HFR)

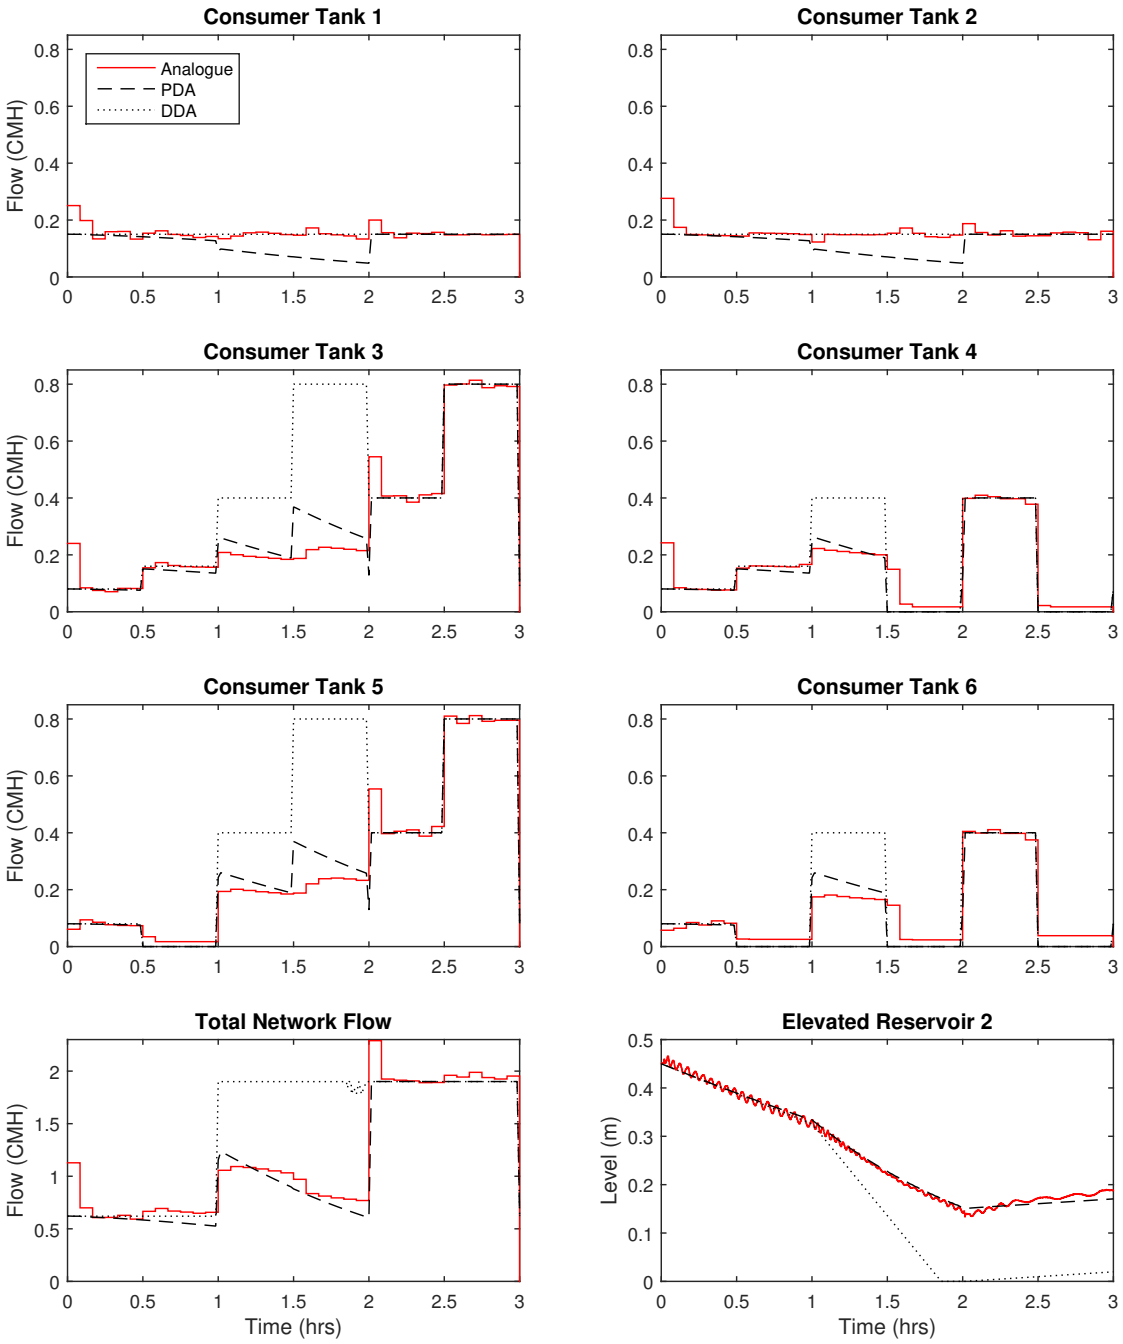

FIGURE A.6: Run 2 (Fujiwara HFR)

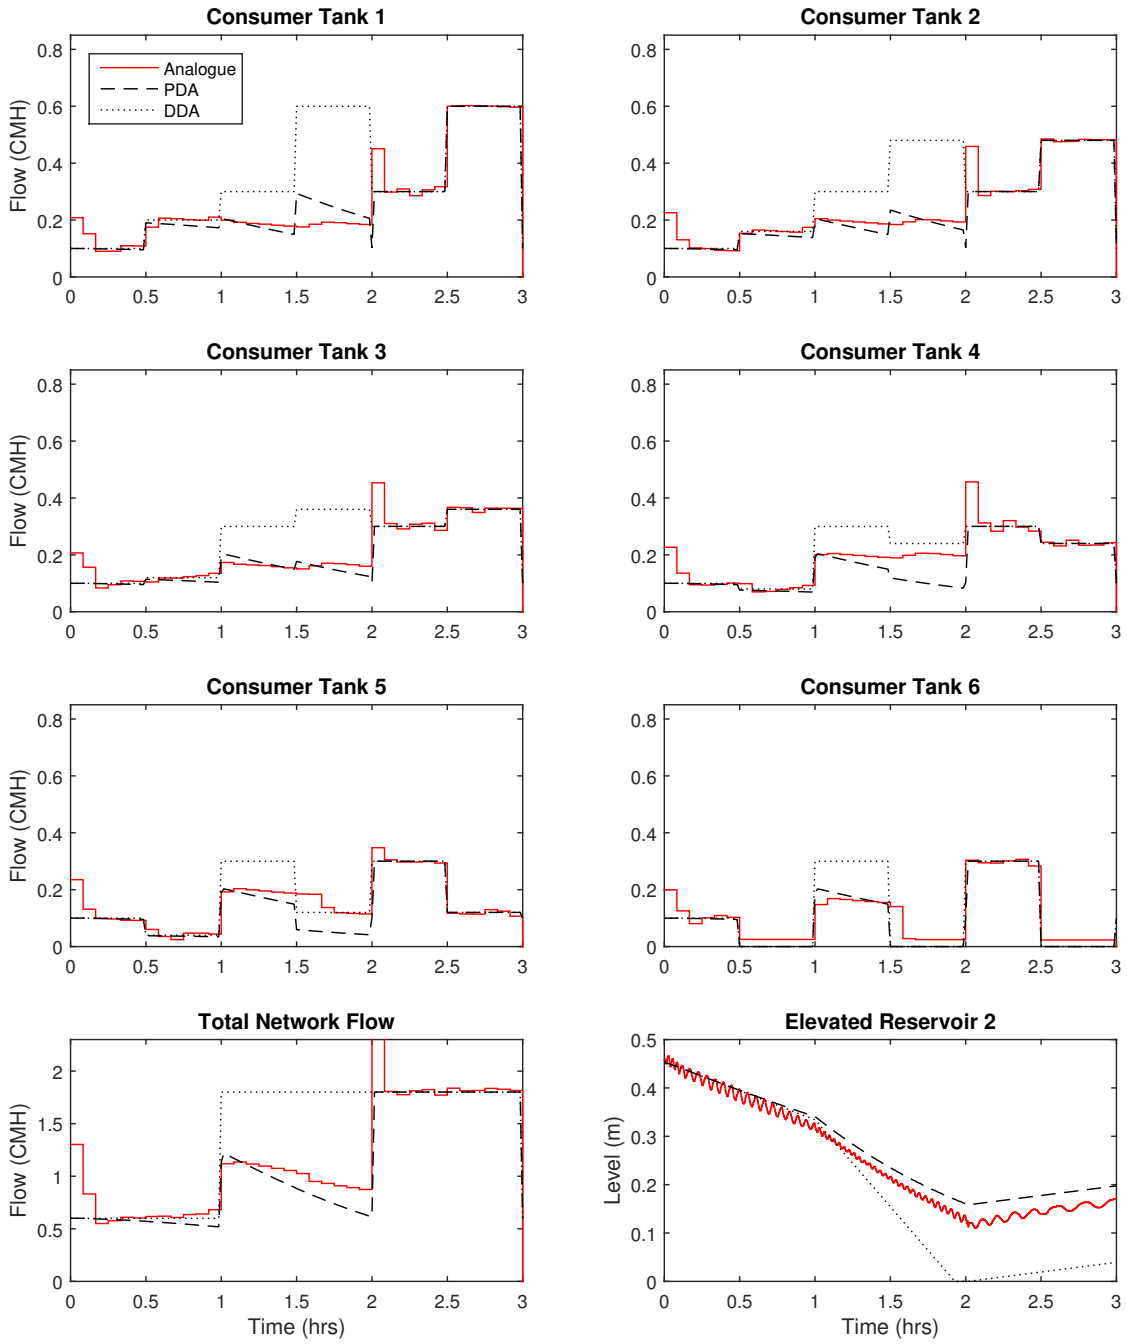

FIGURE A.7: Run 3 (Fujiwara HFR)

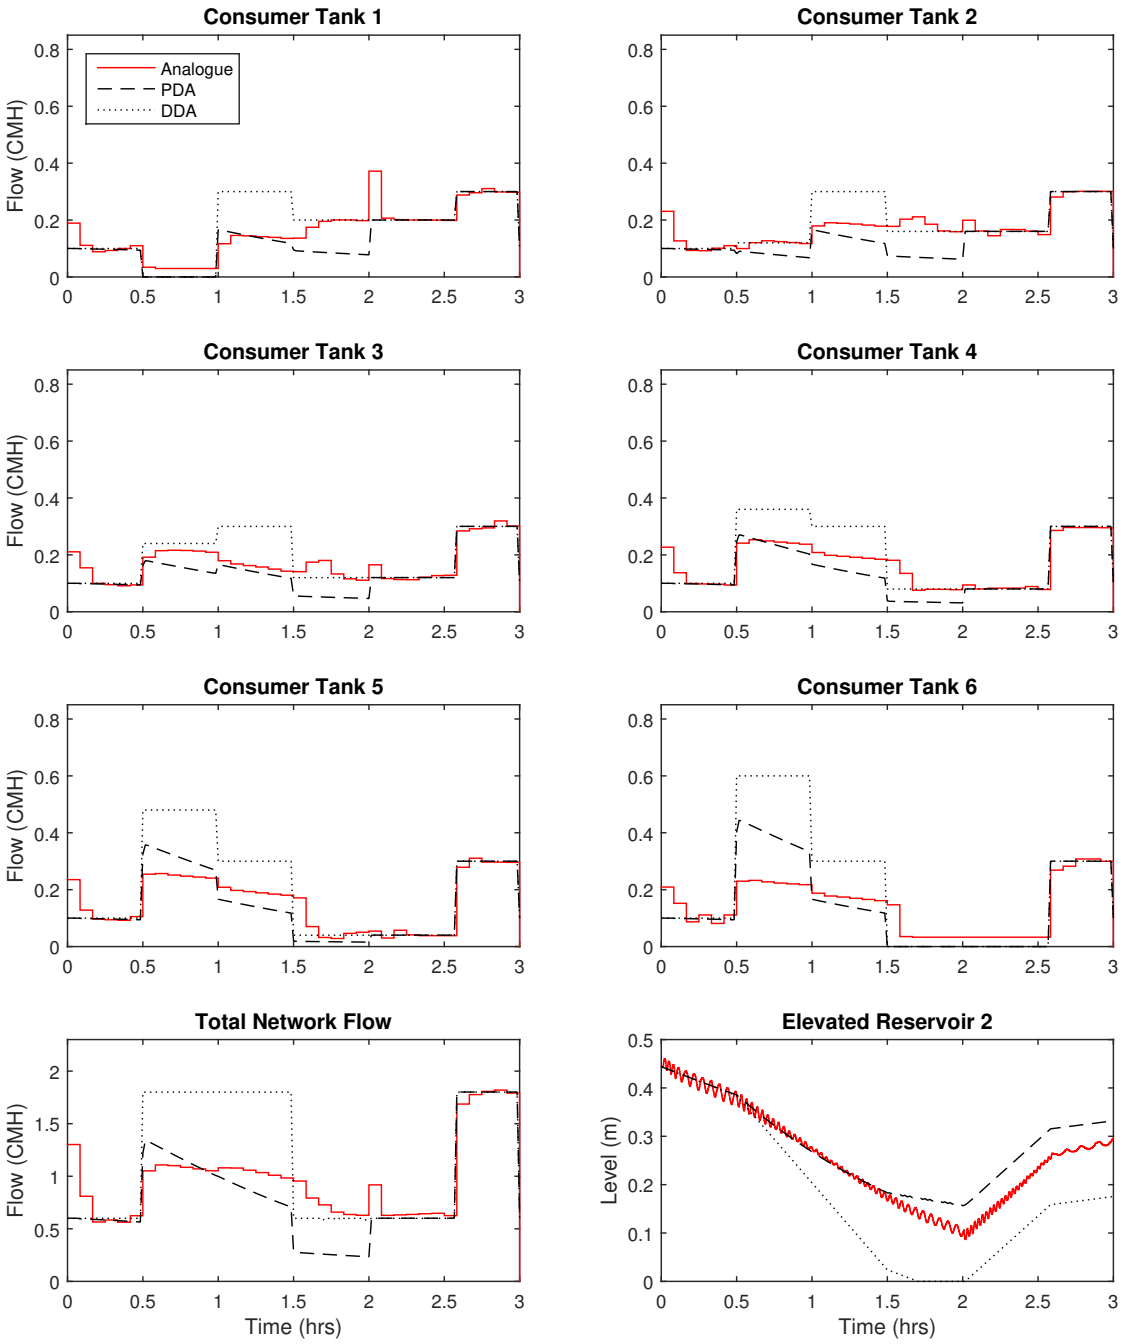

FIGURE A.8: Validation run (Fujiwara HFR)

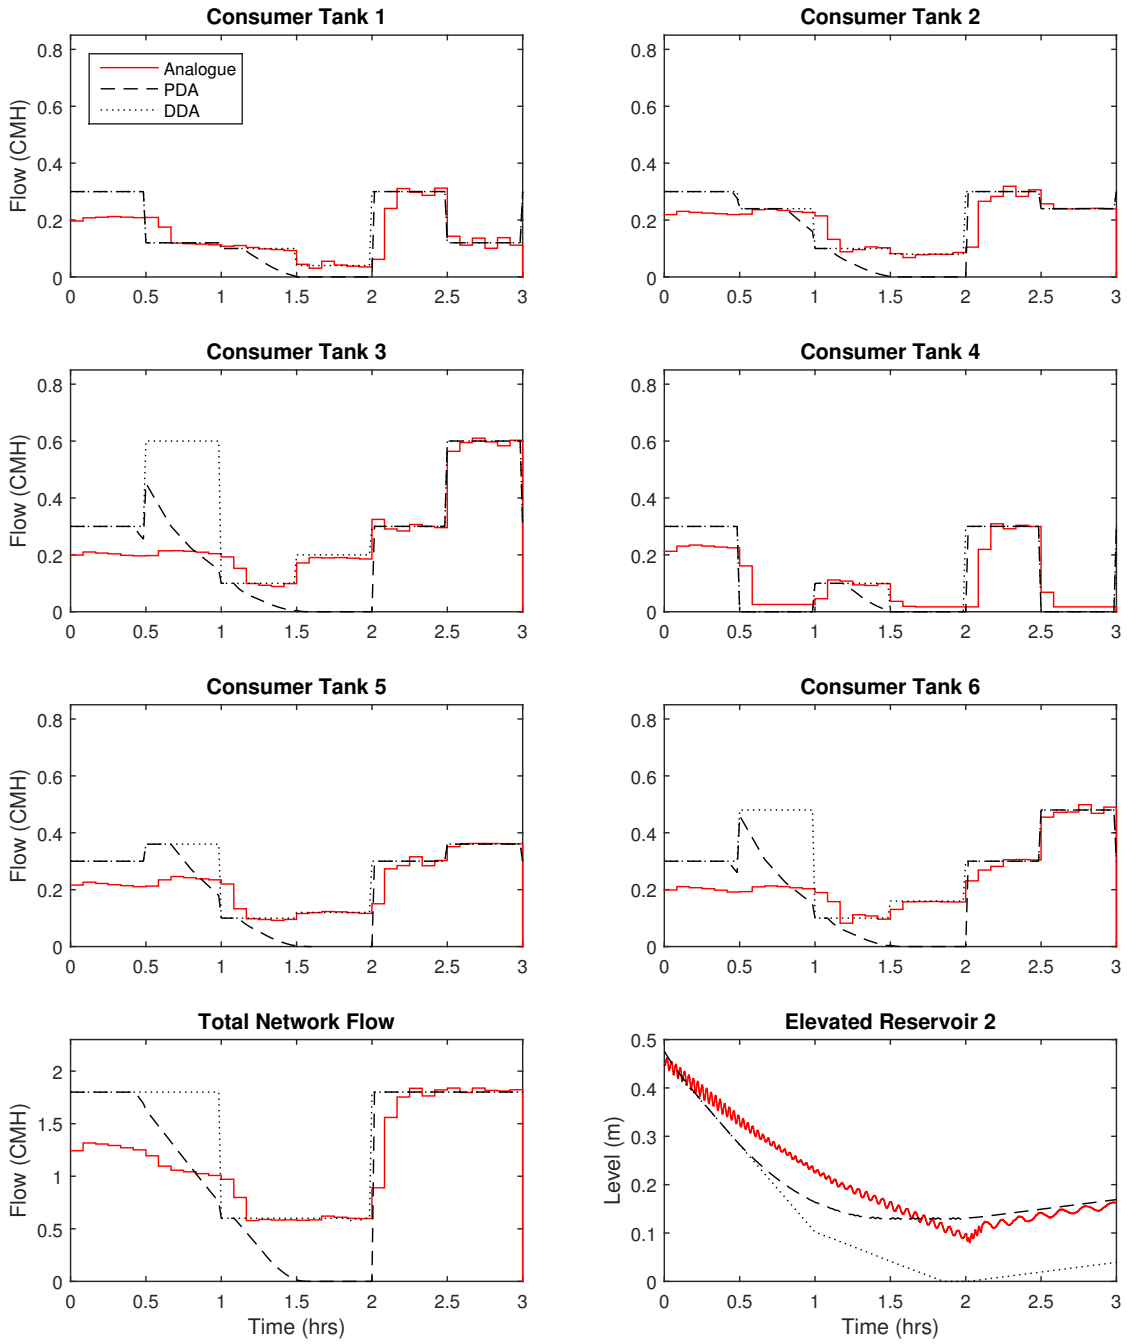

FIGURE A.9: Run 1 (Bhave HFR)

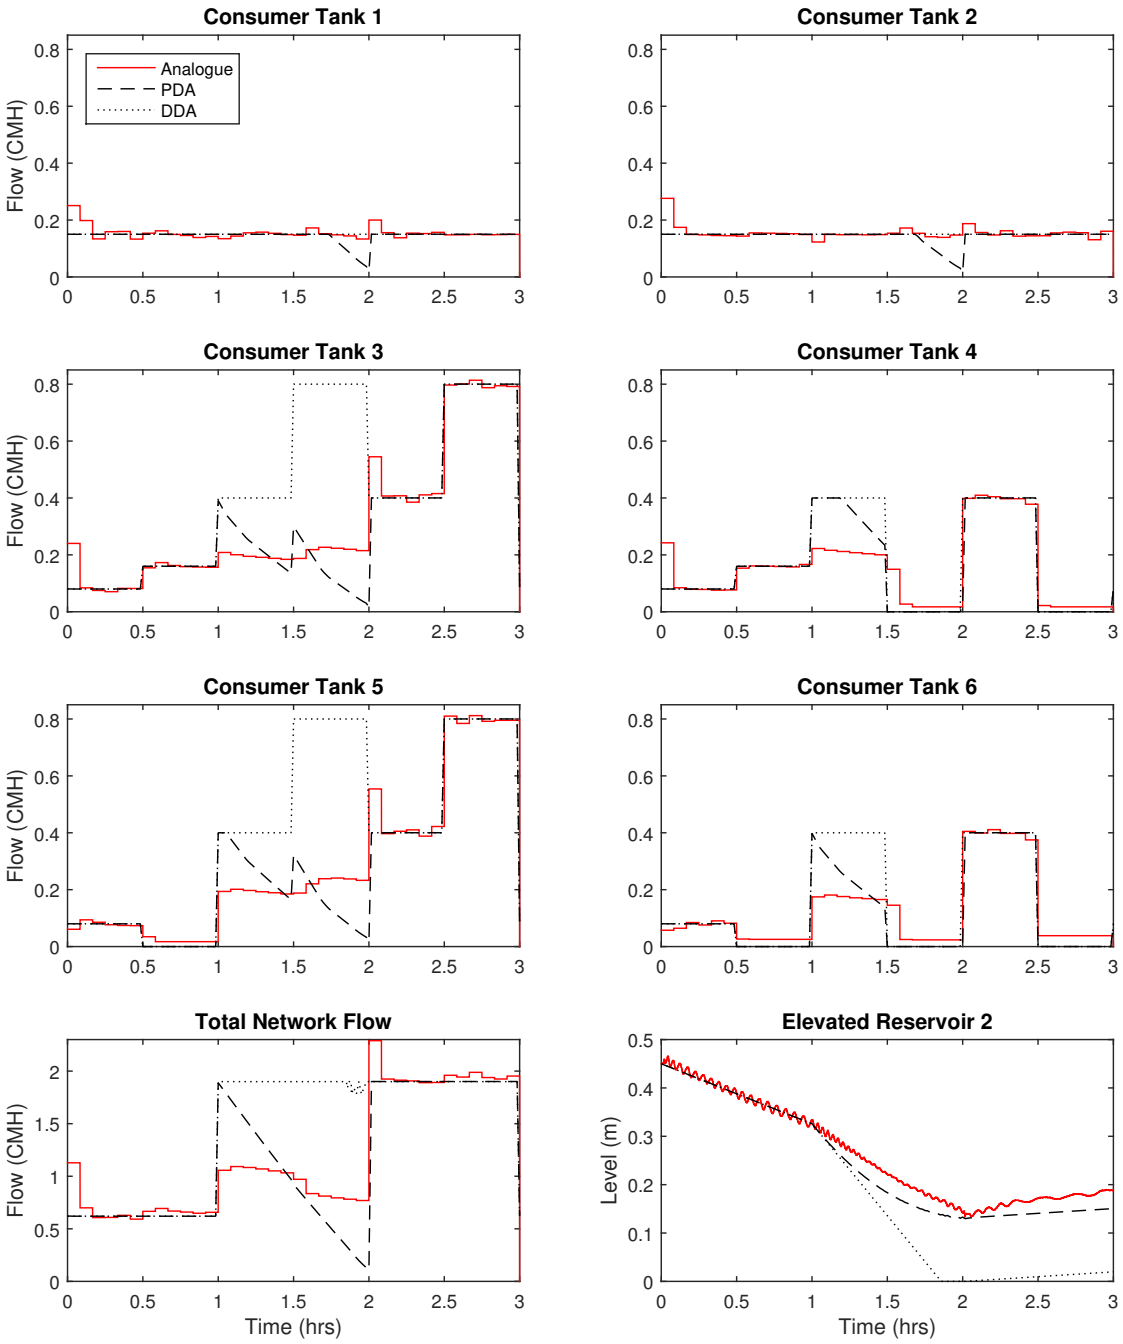

FIGURE A.10: Run 2 (Bhave HFR)

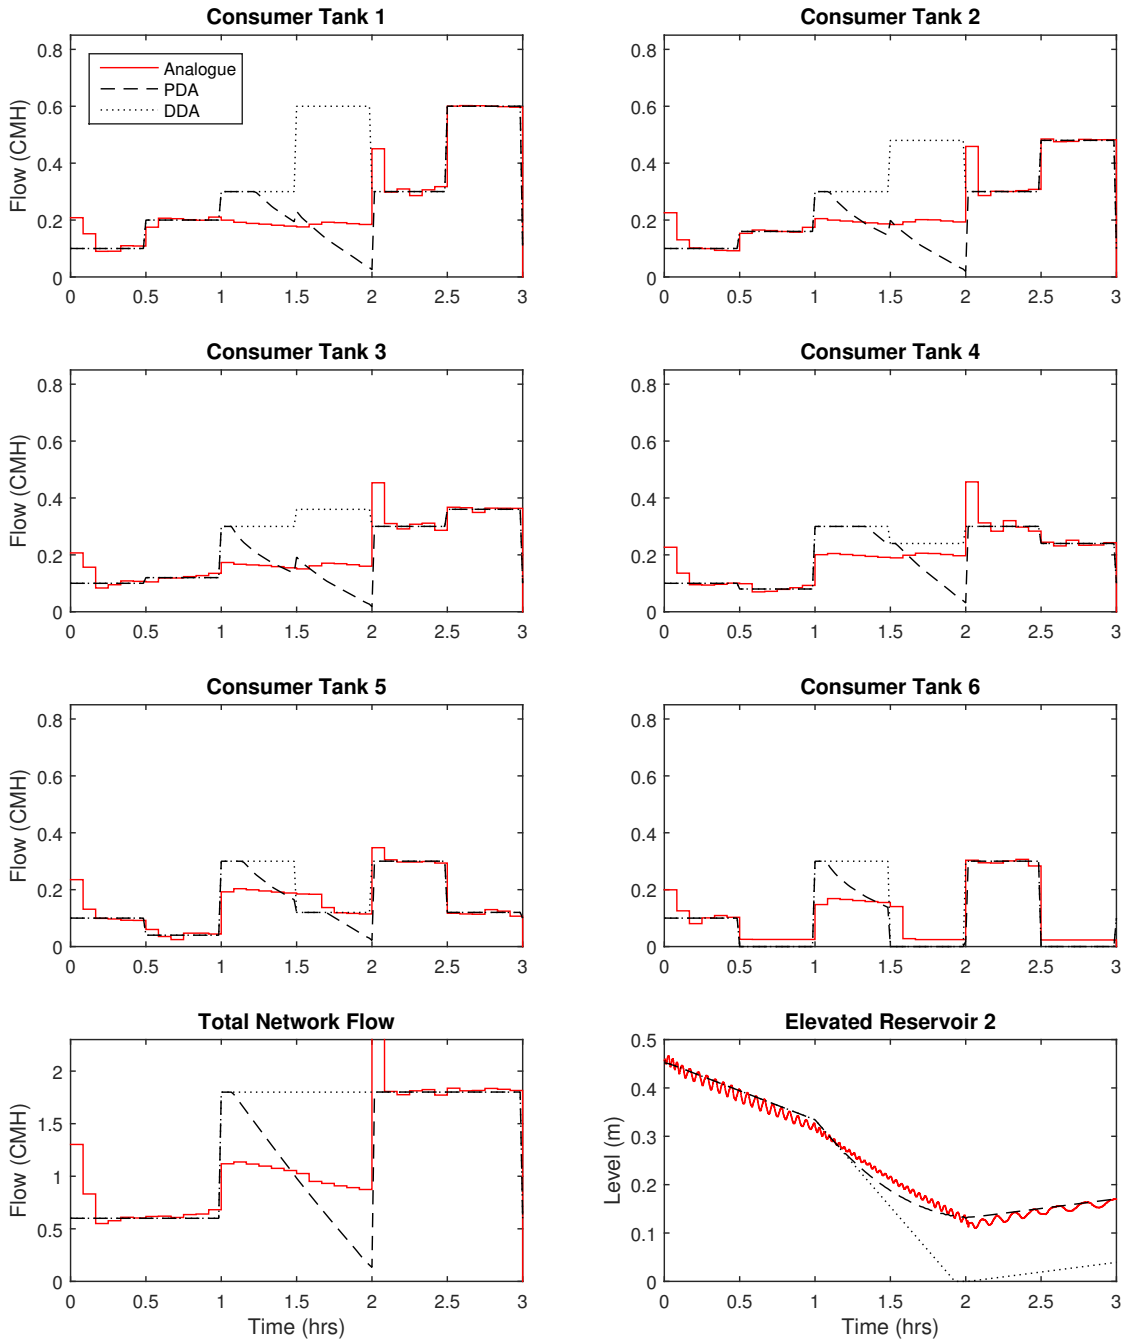

FIGURE A.11: Run 3 (Bhave HFR)

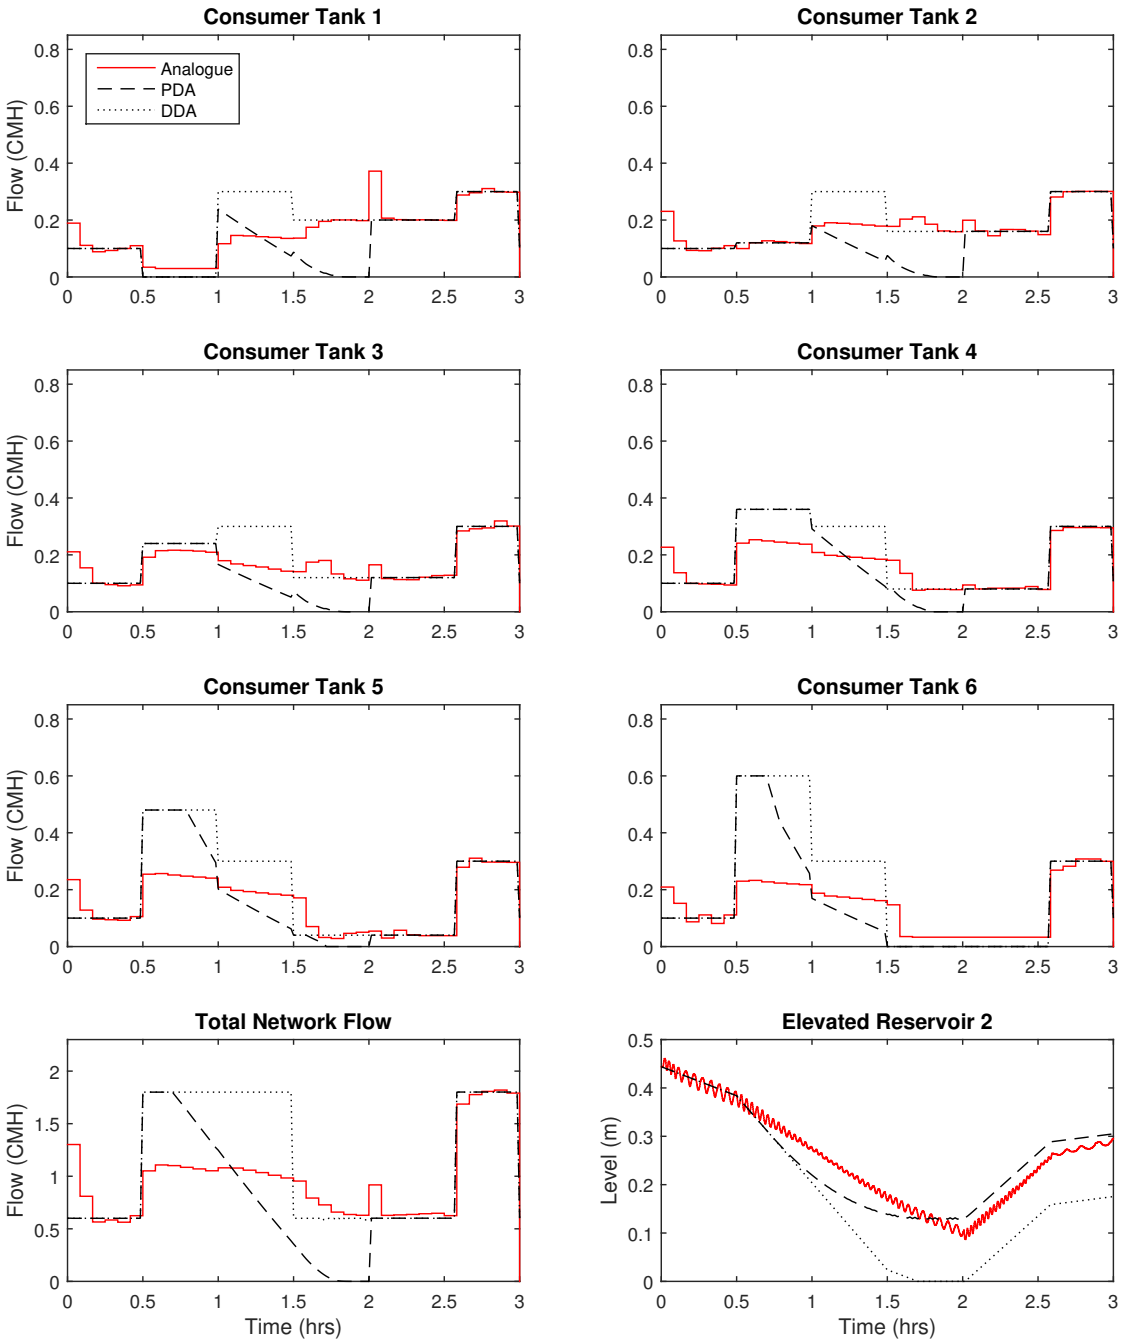

FIGURE A.12: Validation run (Bhave HFR)

Appendix B

C-Town Simulation Results

The following pages include figures showing the results of the different attack scenarios on the C-Town network described in Chapter 5. For attack scenarios 1 and 2, the tank levels in all seven tanks are shown. For attack scenarios 3 and 4, the demand satisfaction ratios and combined resilience-failure indices are shown for all five districts and the overall network.

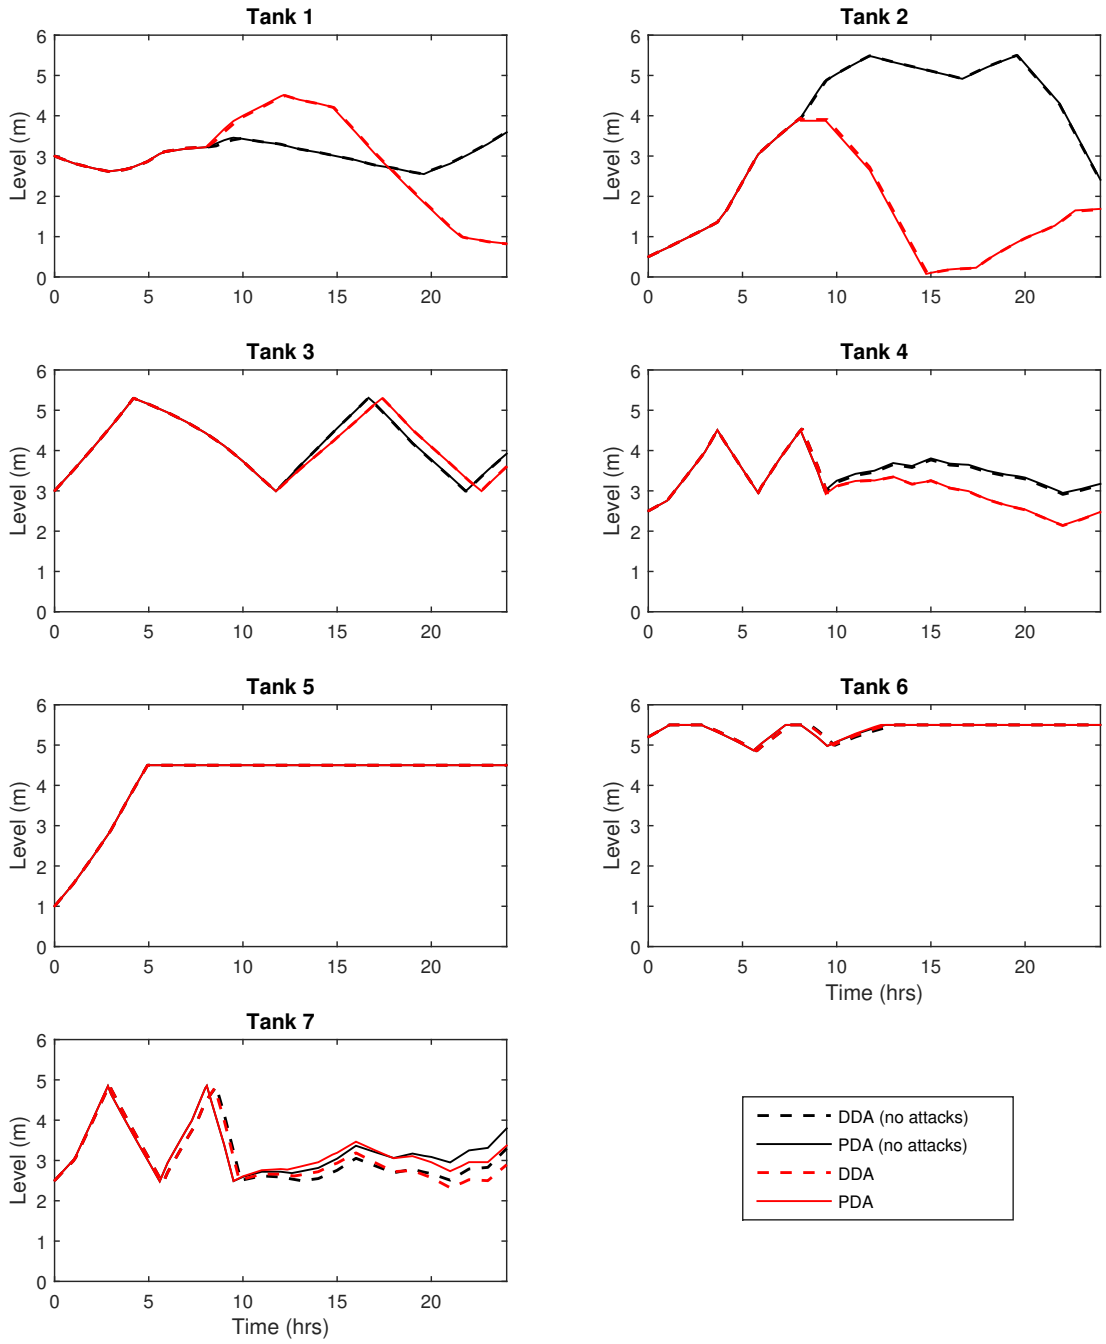

FIGURE B.1: Attack scenario 1 – Tank levels

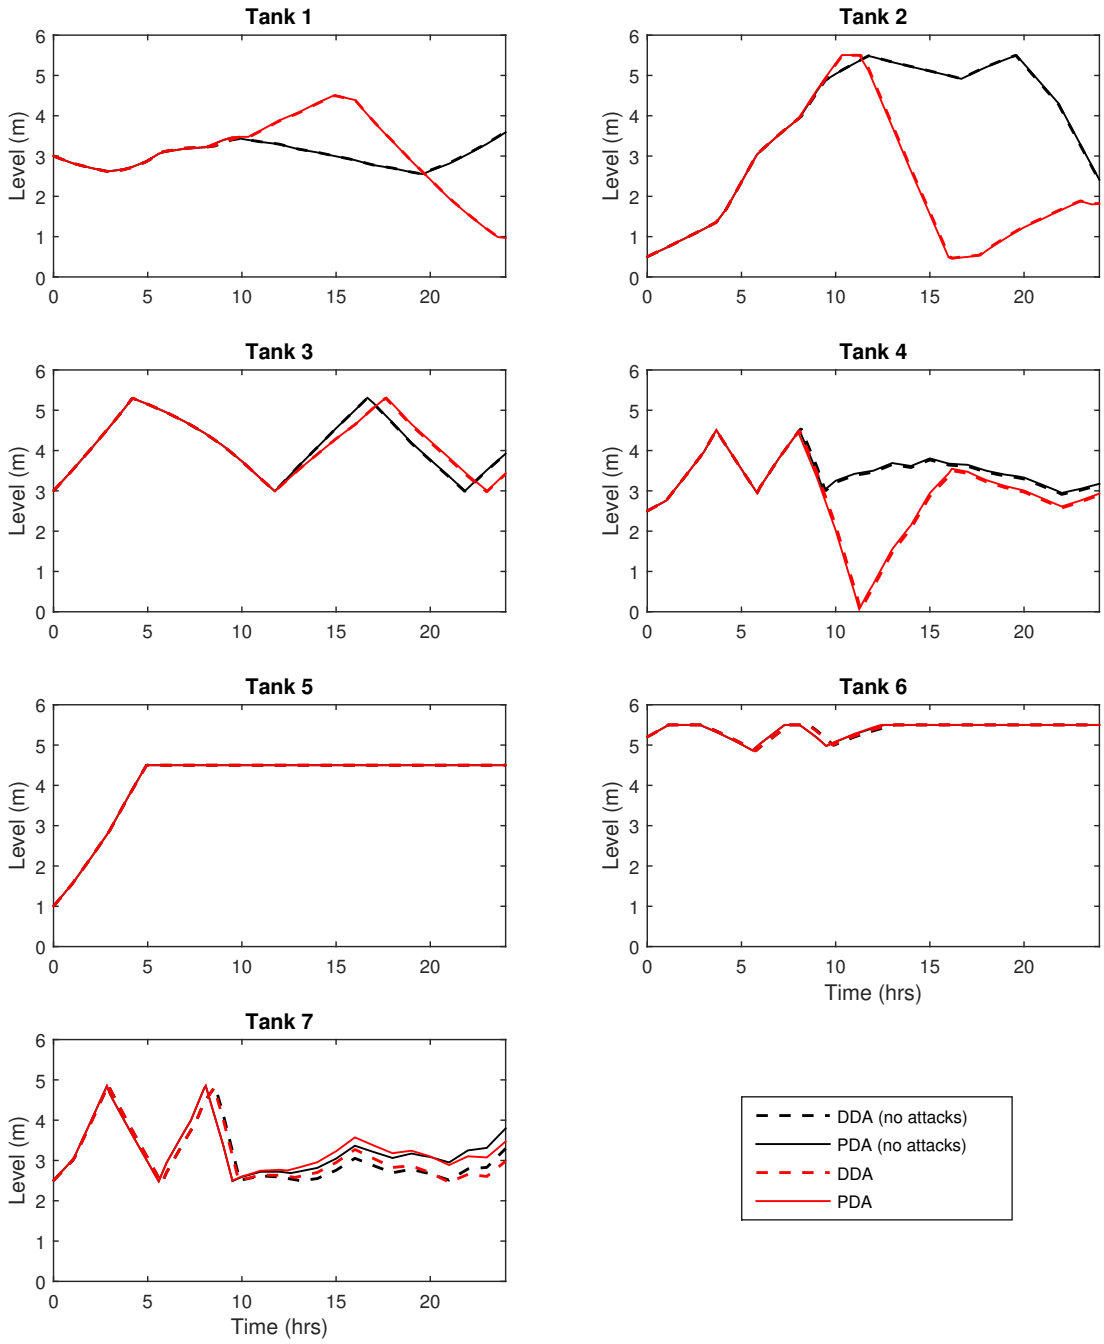

FIGURE B.2: Attack scenario 2 – Tank levels

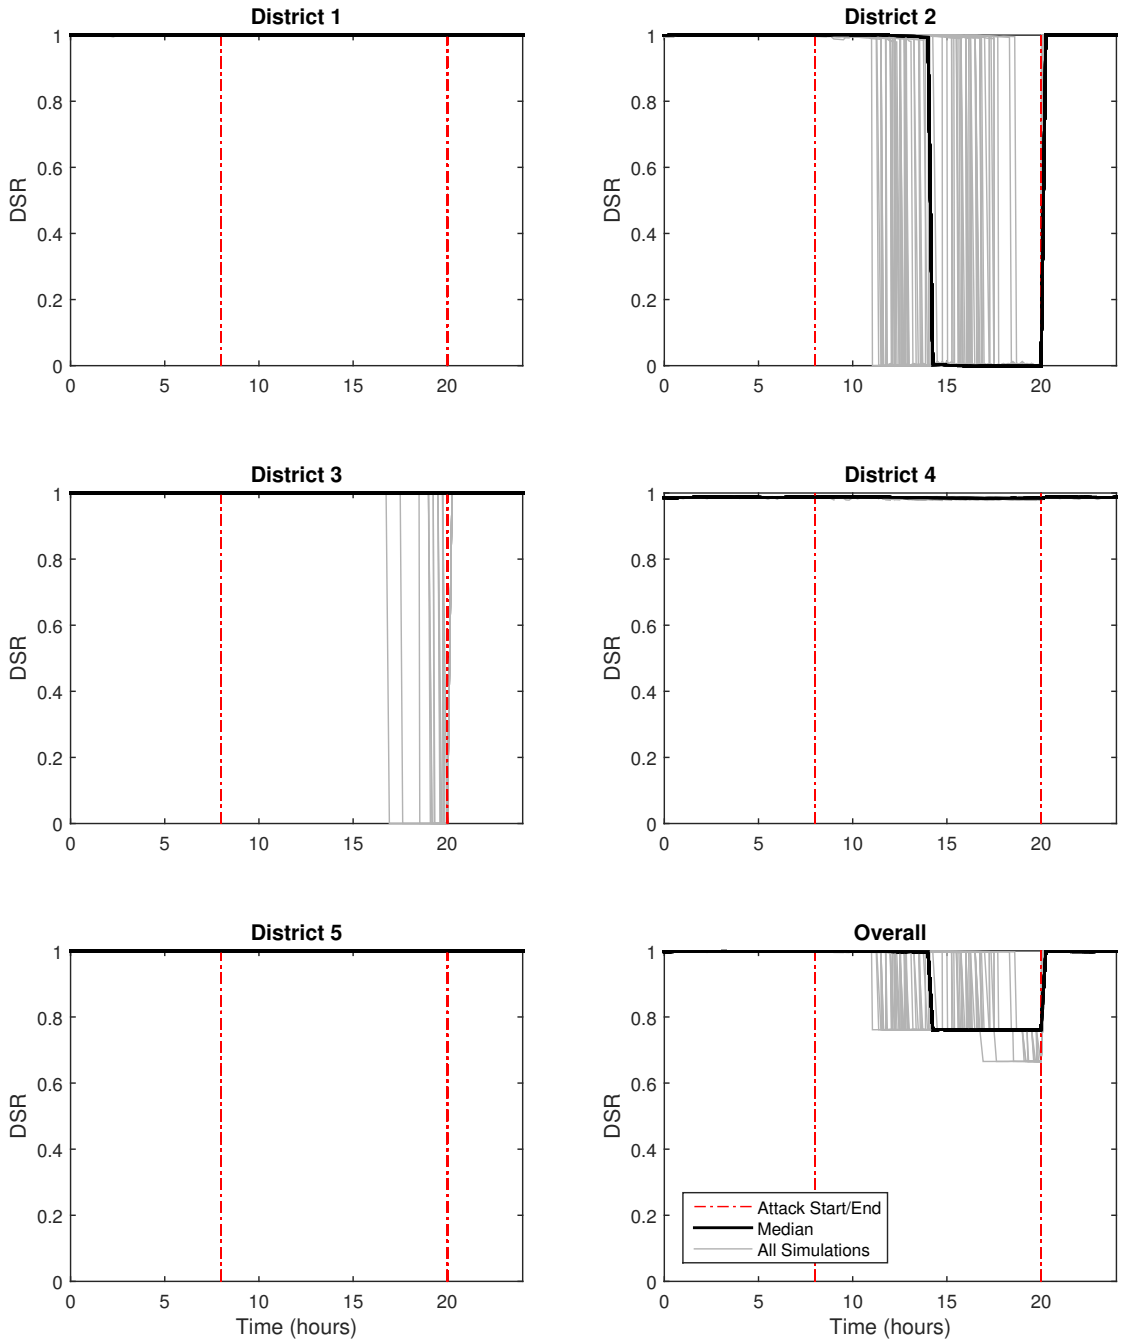

FIGURE B.3: Attack scenario 3 – Demand satisfaction ratio

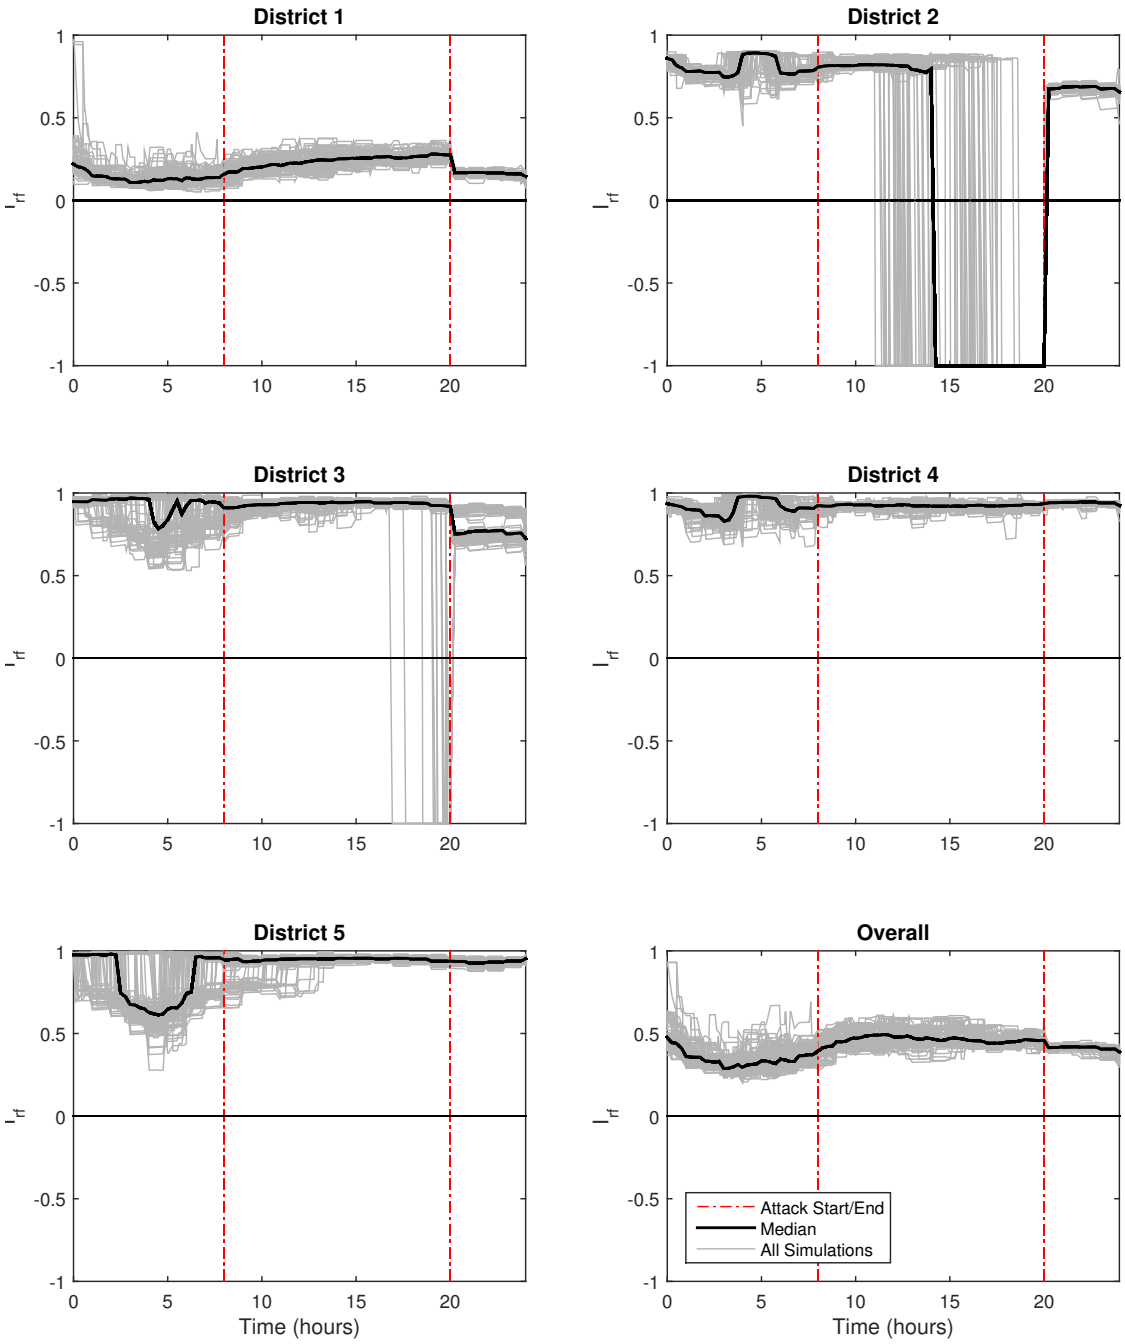

FIGURE B.4: Attack scenario 3 – Combined resilience-failure index

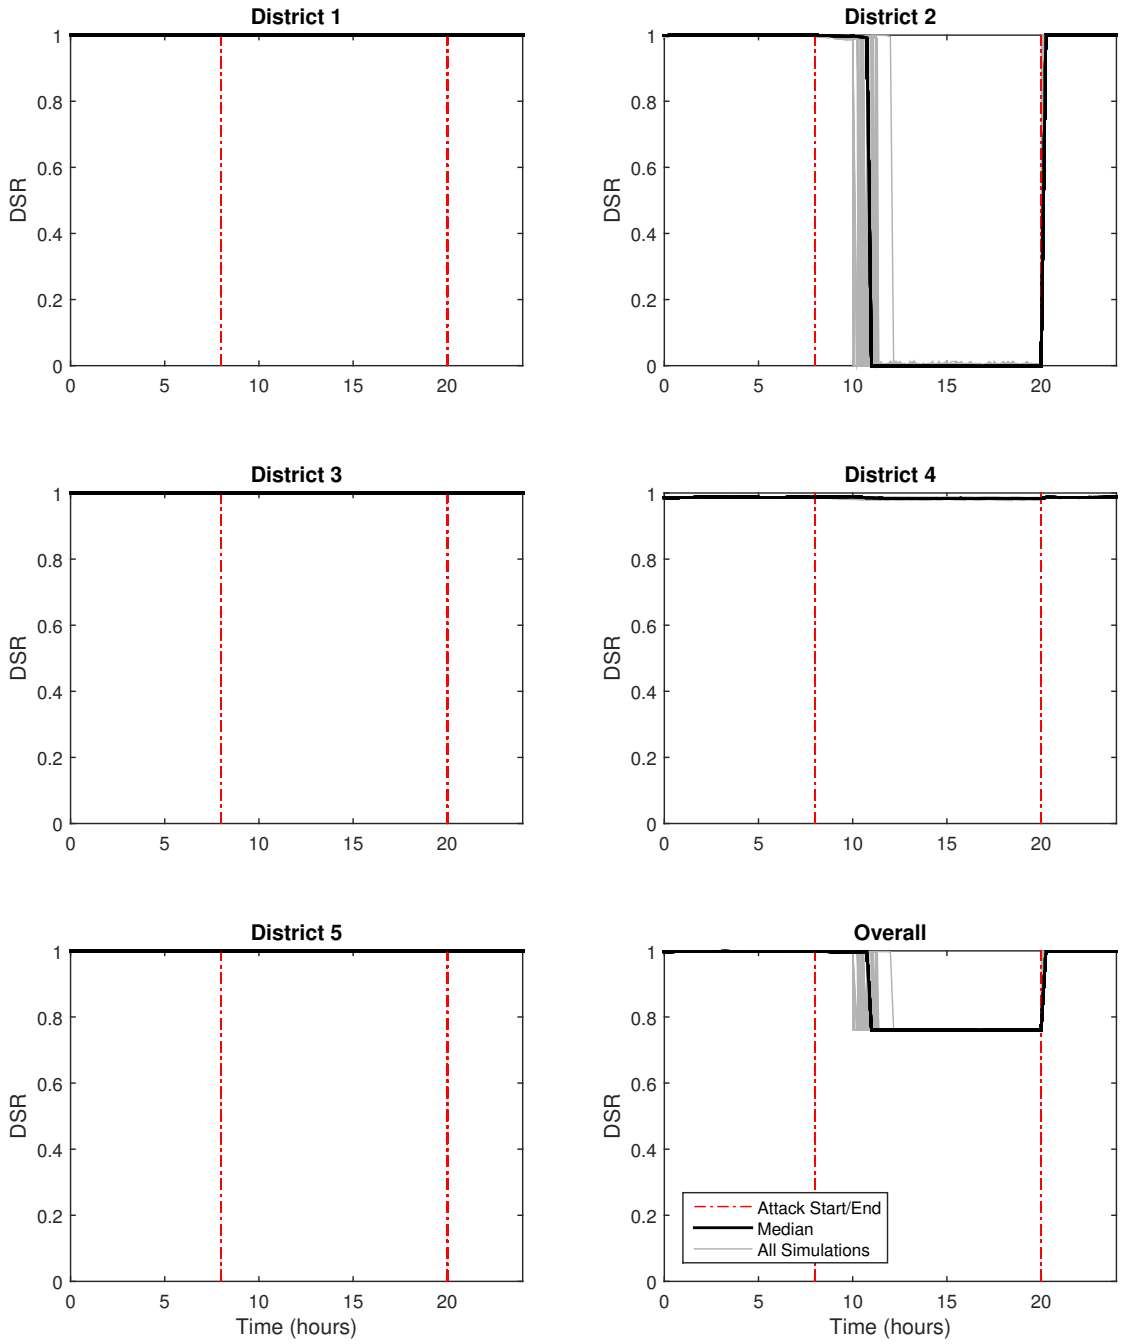

FIGURE B.5: Attack scenario 4 – Demand satisfaction ratio

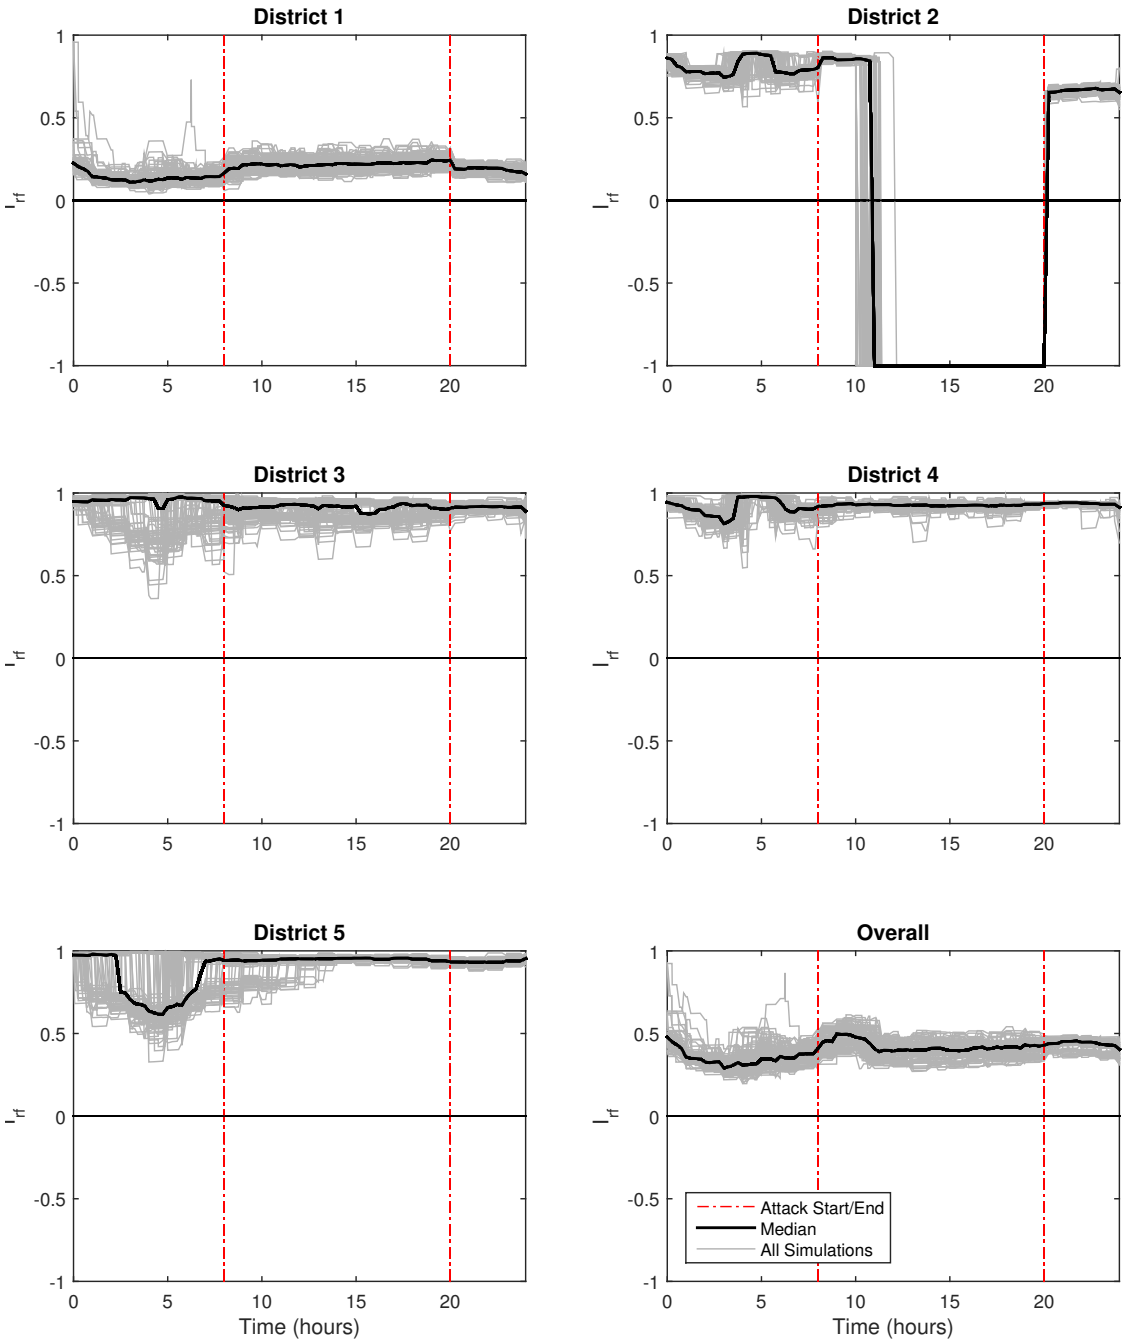

FIGURE B.6: Attack scenario 4 – Combined resilience-failure index

Bibliography

Abdy Sayyed, Mohd Abbas H., Rajesh Gupta, and Tiku T. Tanyimboh (2015). "Noniterative Application of EPANET for Pressure Dependent Modelling Of Water Distribution Systems". In: *Water Resources Management* 29.9, pp. 3227–3242. ISSN: 1573-1650. DOI: 10.1007/s11269-015-0992-0.

Ackley, J.R.L. et al. (2001). "Head-driven analysis of water distribution systems". In: *Water Software Systems: Theory and Applications*. Vol. 1. Research Studies Press Ltd. Chap. 3, pp. 183–192.

Ahmed, Chuadhry Mujeeb, Carlos Murguia, and Justin Ruths (2017). "Model-based Attack Detection Scheme for Smart Water Distribution Networks". In: *Asia Conference on Computer and Communications Security*. ACM, pp. 101–113. ISBN: 978-1-4503-4944-4. DOI: 10.1145/3052973.3053011.

Ang, Wah Khim and Paul W. Jowitt (2006). "Solution for Water Distribution Systems under Pressure-Deficient Conditions". In: *Journal of Water Resources Planning and Management* 132.3, pp. 175–182. DOI: 10.1061/(ASCE)0733-9496(2006)132:3(175).

Bertola, P and M Nicolini (2007). "Evaluating reliability and efficiency of water distribution networks". In: *Efficient management of water networks: design and rehabilitation techniques, Ferrara, May 2006*. FrancoAngeli, pp. 7–23.

Bhave, P. R. (1981). "Node flow analysis of water distribution systems". English. In: *Journal of Transportation Engineering* 107.4, pp. 457–467. ISSN: 0569-7891.

Buehren, Markus (2014). *Differential Evolution*. URL: https://www.mathworks.com/matlabcentral/fileexchange/18593.

Chan, Pui Wah (1983). "Country Report: Hong Kong". In: *Public Water Supply Leakage and Wastage Control*. World Health Organisation, pp. 70–75.

Cheung, P. B., J. E. Van Zyl, and L. F. R. Reis (2005). "Extension of EPANET for Pressure Driven Demand Modeling in Water Distribution System". In: *Computing and Control in the Water Industry (CCWI)*. Vol. 1. Centre for Water Systems.

Cominola, A. et al. (2015). "Benefits and challenges of using smart meters for advancing residential water demand modeling and management: A review". In: *Environmental Modelling & Software* 72, pp. 198–214. ISSN: 1364-8152. DOI: https://doi.org/10.1016/j.envsoft.2015.07.012.

Creaco, Enrico, Marco Franchini, and Ezio Todini (2016). "Generalized Resilience and Failure Indices for Use with Pressure-Driven Modeling and Leakage". In: *Journal of Water Resources Planning and Management* 142.8. DOI: 10.1061/(ASCE)WR.1943-5452.0000656.

Do, Van Long (2015). "Sequential Detection and Isolation of Cyber-physical Attacks on SCADA Systems". PhD thesis. Troyes, France: Université de Technologie de Troyes.

Elhay, Sylvan et al. (2016). "A Robust, Rapidly Convergent Method That Solves the Water Distribution Equations for Pressure-Dependent Models". In: *Journal of Water*

Resources Planning and Management 142.2. DOI: 10.1061/(ASCE)WR.1943-5452.0000578.

EPANET [Computer software]. U.S. Environmental Protection Agency, Cincinnati, OH.

Fontana, Nicola, Maurizio Giugni, and Gustavo Marini (2016). "Experimental assessment of pressure-leakage relationship in a water distribution network". In: *Water Science and Technology: Water Supply* 17.3, pp. 726–732. ISSN: 1606-9749. DOI: 10.2166/ws.2016.171.

Formiga, Klebber Teodomiro Martins and Fazal Hussain Chaudhry (2008). "Modelos de analise hidraulica de redes de distribuicao de agua considerando demanda dirigida pela pressao e vazamentos". In: *Engenharia Sanitaria e Ambiental* 13, pp. 153–162. ISSN: 1413-4152.

Fu, Guangtao et al. (2013). "Optimal Design of Water Distribution Systems Using Many-Objective Visual Analytics". In: *Journal of Water Resources Planning and Management* 139.6, pp. 624–633. DOI: 10.1061/(ASCE)WR.1943-5452.0000311.

Fujiwara, Okitsugu and Tharmarajah Ganesharajah (1993). "Reliability assessment of water supply systems with storage and distribution networks". In: *Water Resources Research* 29.8, pp. 2917–2924. ISSN: 1944-7973. DOI: 10.1029/93WR00857.

Fujiwara, Okitsugu and Jun Li (1998). "Reliability analysis of water distribution networks in consideration of equity, redistribution, and pressure-dependent demand". In: *Water Resources Research* 34.7, pp. 1843–1850. ISSN: 1944-7973. DOI: 10.1029/98WR00908.

Germanopoulos, George (1985). "A technical note on the inclusion of pressure dependent demand and leakage terms in water supply network models". In: *Civil Engineering Systems* 2.3, pp. 171–179. DOI: 10.1080/02630258508970401.

Gheisi, A., M. Forsyth, and Gh. Naser (2016). "Water Distribution Systems Reliability: A Review of Research Literature". In: *Journal of Water Resources Planning and Management* 142.11. DOI: 10.1061/(ASCE)WR.1943-5452.0000690.

Giustolisi, Orazio, Dragan Savic, and Zoran Kapelan (2008). "Pressure-Driven Demand and Leakage Simulation for Water Distribution Networks". In: *Journal of Hydraulic Engineering* 134.5, pp. 626–635. DOI: 10.1061/(ASCE)0733-9429(2008)134:5(626).

Gorev, Nikolai B. and Inna F. Kodzhespirova (2013). "Noniterative Implementation of Pressure-Dependent Demands Using the Hydraulic Analysis Engine of EPANET 2". In: *Water Resources Management* 27.10, pp. 3623–3630. ISSN: 1573-1650. DOI: 10.1007/s11269-013-0369-1.

Grafton, R. Quentin et al. (2011). "Determinants of residential water consumption: Evidence and analysis from a 10-country household survey". In: *Water Resources Research* 47.8. ISSN: 1944-7973. DOI: 10.1029/2010WR009685.

Gupta, Rajesh and Pramod R. Bhave (1996). "Comparison of Methods for Predicting Deficient-Network Performance". In: *Journal of Water Resources Planning and Management* 122.3, pp. 214–217. DOI: 10.1061/(ASCE)0733-9496(1996)122:3(214).

Haimes, Yacov Y. et al. (1998). "Reducing Vulnerability of Water Supply Systems to Attack". In: *Journal of Infrastructure Systems* 4.4, pp. 164–177. DOI: 10.1061/(ASCE)1076-0342(1998)4:4(164).

Huang, J., E. McBean, and W. James (2005). "A Review of Reliability Analysis for Water Quality in Water Distribution Systems". In: *Journal of Water Management Modeling* R223.07. DOI: 0.14796/JWMM.R223-07.

ICS-CERT (2014). *NCCIC/ICS-CERT Year in Review: FY 2013*. Tech. rep. 13-50369. Washington, D.C.: U.S. Department of Homeland Security – Industrial Control Systems-Cyber Emergency Response Team.

— (2015). *NCCIC/ICS-CERT Year in Review: FY 2014*. Tech. rep. 14-50426. Washington, D.C.: U.S. Department of Homeland Security – Industrial Control Systems-Cyber Emergency Response Team.

— (2016). *NCCIC/ICS-CERT Year in Review: FY 2015*. Tech. rep. 15-50569. Washington, D.C.: U.S. Department of Homeland Security – Industrial Control Systems-Cyber Emergency Response Team.

Jinesh Babu, K. S. and S. Mohan (2012). "Extended Period Simulation for Pressure-Deficient Water Distribution Network". In: *Journal of Computing in Civil Engineering* 26.4, pp. 498–505. DOI: 10.1061/(ASCE)CP.1943-5487.0000160.

Kalungi, P. and T. Tanyimboh (2003). "Redundancy model for water distribution systems". In: *Reliability Engineering and System Safety* 82.3, pp. 275–286. ISSN: 0951-8320. DOI: 10.1016/S0951-8320(03)00168-6.

Mahmoud, Herman A., Dragan Savić, and Zoran Kapelan (2017). "New Pressure-Driven Approach for Modeling Water Distribution Networks". In: *Journal of Water Resources Planning and Management* 143.8. DOI: 10.1061/(ASCE)WR.1943-5452.0000781.

MATLAB [Computer software]. Mathworks, Natick, MA.

McKee, K. et al. (2011). "A review of major centrifugal pump failure modes with application to the water supply and sewerage industries". In: *ICOMS Asset Management Conference Proceedings*. Asset Management Council. DOI: 20.500.11937/28560.

Mckenzie, R S and W Wegelin (2009). *Implementation of Pressure Management in Municipal Water Supply Systems*. Tech. rep. 0309. WRP Pty Ltd.

Morley, M.S. and C. Tricarico (2008). *Pressure Driven Demand Extension for EPANET (EPANETpdd)*. Tech. rep. 2008-02. UK: Centre for Water Systems, University of Exeter.

Muranho, J. et al. (2014). "Pressure-dependent Demand and Leakage Modelling with an EPANET Extension – WaterNetGen". In: *Procedia Engineering* 89, pp. 632–639. ISSN: 1877-7058. DOI: 10.1016/j.proeng.2014.11.488.

Ostfeld, Avi et al. (2011). "Battle of the Water Calibration Networks". In: *Journal of Water Resources Planning and Management* 138.5, pp. 523–532. DOI: 10.1061/(ASCE)WR.1943-5452.0000191.

Pacchin, E., S. Alvisi, and M. Franchini (2017). "Analysis of Non-Iterative Methods and Proposal of a New One for Pressure-Driven Snapshot Simulations with EPANET". In: *Water Resources Management* 31.1, pp. 75–91. ISSN: 1573-1650. DOI: 10.1007/s11269-016-1511-7.

Perelman, Lina and Saurabh Amin (2014). "A Network Interdiction Model for Analyzing the Vulnerability of Water Distribution Systems". In: *High Confidence Networked Systems*. ACM, pp. 135–144. ISBN: 978-1-4503-2652-0. DOI: 10.1145/2566468.2566480.

Rasekh, Amin and Kelly Brumbelow (2014). "Drinking water distribution systems contamination management to reduce public health impacts and system service interruptions". In: *Environmental Modelling & Software* 51, pp. 12–25. ISSN: 1364-8152. DOI: 10.1016/j.envsoft.2013.09.019.

Rasekh, Amin et al. (2016). "Smart Water Networks and Cyber Security". In: *Journal of Water Resources Planning and Management* 142.7. DOI: 10.1061/(ASCE)WR.1943-5452.0000646.

Rossman, L. A. (2000). *EPANET 2 Users Manual*. EPA/600/R-00/057. U.S. Environmental Protection Agency. Washington, D.C.

Sharoonizadeh, Shokofeh, Jafar Mamizadeh, and Javad Sarvarian (2016). "Comparison of solution methods for analyzing water distribution networks under pressure-deficient conditions". In: *Journal of Water Supply: Research and Technology - Aqua* 65.4, pp. 330–341. ISSN: 0003-7214. DOI: 10.2166/aqua.2016.084.

Siew, Calvin and Tiku T. Tanyimboh (2010). "Pressure-Dependent EPANET Extension: Pressure-Dependent Demands". In: *Water Distribution Systems Analysis 2010*, pp. 75–84. DOI: 10.1061/41203(425)9.

Slay, Jill and Michael Miller (2008). "Lessons Learned from the Maroochy Water Breach". In: *Critical Infrastructure Protection*. Springer US, pp. 73–82. ISBN: 978-0-387-75462-8. DOI: 10.1007/978-0-387-75462-8_6.

Tanyimboh, T., B. Tahar, and A. Templeman (2003). "Pressure-driven modelling of water distribution systems". In: *Water Science and Technology: Water Supply* 3.1-2, pp. 255–261. ISSN: 1606-9749.

Taormina, Riccardo et al. (2016). *BATtle of the Attack Detection ALgorithms (BATADAL): Detailed Problem Description and Rules*. Accessed 30-05-2017. URL: http://www.batadal.net/images/rules.pdf.

Taormina, Riccardo et al. (2017). "Characterizing Cyber-Physical Attacks on Water Distribution Systems". In: *Journal of Water Resources Planning and Management* 143.5. DOI: 10.1061/(ASCE)WR.1943-5452.0000749.

Thornton, J and A Lambert (2005). "Progress in practical prediction of pressure: leakage, pressure: burst frequency and pressure: consumption relationships". In: *Leakage 2005 - Conference Proceedings*. IWA. Halifax, Canada.

Todini, E (2003). "A more realistic approach to the "extended period simulation" of water distribution networks". In: *Advances in Water Supply Management*. Taylor & Francis. Chap. 19, pp. 173–183. ISBN: 978-90-5809-608-1. DOI: 10.1201/NOE9058096081.ch19.

Todini, E. and S. Pilati (1988). "A Gradient Algorithm for the Analysis of Pipe Networks". In: *Computer Applications in Water Supply: Vol. 1 - Systems Analysis and Simulation*. Research Studies Press Ltd. Chap. 1, pp. 1–20. ISBN: 0-471-91783-4.

Todini, Ezio (2000). "Looped water distribution networks design using a resilience index based heuristic approach". In: *Urban Water* 2.2. Developments in water distribution systems, pp. 115–122. ISSN: 1462-0758. DOI: 10.1016/S1462-0758(00)00049-2.

Tucciarelli, T., A. Criminisi, and D. Termini (1999). "Leak Analysis in Pipeline Systems by Means of Optimal Valve Regulation". In: *Journal of Hydraulic Engineering* 125.3, pp. 277–285. DOI: 10.1061/(ASCE)0733-9429(1999)125:3(277).

U.S. Department of Homeland Security (2017). *Personal correspondence with Antonio Soliz, External Affairs, Office of Cybersecurity and Communication on 31/03/2017*.

van Zyl, J. E. and C. R. I. Clayton (2007). "The effect of pressure on leakage in water distribution systems". In: *Proceedings of the Institution of Civil Engineers - Water Management* 160.2, pp. 109–114. DOI: 10.1680/wama.2007.160.2.109.

Wagner, M., Uri Shamir, and David H. Marks (1988). "Water Distribution Reliability: Simulation Methods". In: *Journal of the Water Resources Planning and Management Division, ASCE* 114.3, pp. 276–294.

Walski, Thomas M. et al. (2003). *Advanced Water Distribution Modeling and Management*. 1st. Haestead Press. ISBN: 9780971414129.

Wu, Z. Y. and T. M. Walski (2006). "Pressure dependent hydraulic modelling for water distribution systems under abnormal conditions". In: *Proceedings of the 5th IWA World Water Congress*.

Wu, Zheng Y. et al. (2009). "Extended Global-Gradient Algorithm for Pressure-Dependent Water Distribution Analysis". In: *Journal of Water Resources Planning and Management* 135.1, pp. 13–22. DOI: 10.1061/(ASCE)0733-9496(2009)135:1(13).