UNIVERSITY OF SOUTHAMPTON

Probabilistic Leak Detection and Quantification Using Multi-output Gaussian Processes

by

Obaid Malik

A thesis submitted in partial fulfillment for the degree of Doctor of Philosophy

in the

Faculty of Physical Sciences and Engineering
School of Electronics and Computer Science
Agents, Interaction and Complexity Group

August 2016

UNIVERSITY OF SOUTHAMPTON

ABSTRACT

FACULTY OF PHYSICAL SCIENCES AND ENGINEERING
SCHOOL OF ELECTRONICS AND COMPUTER SCIENCE

Doctor of Philosophy

by Obaid Malik

A water distribution system WDS is often divided into smaller isolated and independent zones called district metering areas (DMA). A DMA can have anywhere from a few hundred to a few thousand properties. Normally only three locations within a district metering area are actively monitored for pressure or flow readings. These are the supply point pressure and flow and the critical point pressure which is the point of the lowest pressure in the DMA. As leakage rates are typically directly proportional to average pressures in the DMA, keeping the network pressure as low as possible while maintaining desired serviceability is an effective and widely used method for leak reduction. With advancement in technology this network pressure reduction is now done in real-time, where the network pressure is increased or decreased based on the demand. However, such real-time optimisation changes the DMA dynamics making it different from traditional unoptimised DMAs. We consider the problem of detecting and quantifying leaks in pressure optimised DMA, using only these three DMA-level hydraulic measurements. The DMA-level measurements represent the current aggregate water demand/consumption within the DMA. Detecting leaks at this point is challenging, particularly small leaks, as they do not produce a significant increase in the aggregated DMA-level measurements. Furthermore, the DMA-level data exhibits input signal dependence whereby both noise and leaks are dependent on the flow and pressure being measured, making leak detection task more difficult. To address this, we first propose a Gaussian process (GP) based approach that uses only the DMA-level flow to detect leaks (NSGP). We devise an additive diagonal noise covariance for the GP that is able to handle the input dependant noise observed in this setting. A parameterised mean step change function is used to detect and approximate leaks. As accurate leak data is often not available due to poor record keeping, we develop a detailed simulated model of a pressure optimised DMA and use it for analysing proposed leak detection methods. We show that active pressure optimisation changes the dynamics of a DMA. In light of the change in DMA dynamics, we proposed a domain specific, data driven, multi output gaussian process model, to detect and quantify leaks in pressure optimised DMAs (SMOGP). The novelty of the model is, firstly its ability to use all available available information from a DMA to detect leaks, secondly the ability to model the pressure dependant leak process mathematically within the GP framework. We compare the performance of the proposed methods with the current state of the art leak detection method. We show that our proposed method out perform other approaches considerably both in terms of the accuracy of leak detection and leak magnitude estimation.

Contents

List of Figures

List of Tables

Declaration of Authorship

I, Obaid Malik, declare that the thesis entitled *Probabilistic Leak Detection and Quantification Using Multi-output Gaussian Processes* and the work presented in the thesis are both my own, and have been generated by me as the result of my own original research. 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;
- none of this work has been published before submission
- parts of this work have been published as: (Malik et al., 2015)

Signed:.....

Date:.....

Dedicated to the loving memory of my mother.

You smile down upon me from the heavens

for a dream that could not come true in your life,
has finally been realised.

Acknowledgements

First and foremost, I would like to thank my supervisor, Alex Rogers, for his excellent advice, constant support and for being an endless source of motivation throughout the course of my PhD.

I would like to thank my wife Mehreen Malik for always standing by me even when times were hard. Completing a PhD feels like quite a journey and I am grateful she was a major part of it.

To my friends, Sasan Maleki, Lampros Stavrogiannis, Zeeshan Ziaf, Ali Khan Khattak, Kamran Yousaf and Nauman Mehmood, I am eternally grateful for their unconditional, unwavering friendship and for always being there when I needed them.

A special thanks to my brother and sisters, Avais Malik, Arooj Saqib, Kanwal Nadeem, Aqsa Aamir and Sana Khan, for being a constant source of motivation.

I am not sure there are enough words of thanks to express my gratitude for my father and mother in law, Amber Malik and Zia Ullah Malik. They have supported, motivated, helped and believed in me more than I did myself. I can only pray that Allah gives me the opportunity to reciprocate the love, support and care I received from them. I would also like to thank my wife's aunt and uncle, Chanda Chaudrey and Zameer Chaudrey for all their help throughout my time here.

Finally, I would like to acknowledge the financial support provided by I20-Water Ltd and EPSRC through the ALADDIN and ORCHID projects, without which my PhD would not have been possible.

Chapter 1

Introduction

1.1 Motivation

Water is one of the most important commodities in the world. Clean water has been a primary source of concern for civilisations since ancient times. However, after centuries, even with the major advances in science and technology, many parts of the world still suffer from clean water shortage. Like all other natural resources there is a limited amount of fresh water on earth, which is only 3 % of the total water body. Furthermore, 77 % of this fresh water is locked up in icecaps and glaciers (U.S. Geological Survey, 2010). With the ever growing population and demand for clean water, the current water supply is running short. The World Health Organisation (2014) estimates that 783 million people (i.e., one in ten) in the world are deprived from safe water. According to United Nations by 2050 the world population is expected to grow to 9 billion while the water volume available would approximately be the same. Furthermore, factors such as rapid urbanisation, unsustainable consumption patterns, pollution of clean water resources and the environmental impact of global warming, add to the already strained water resources.

In light of this clear scarcity, to maintain the same living standards for generations to come, the major challenge posed to the world today is to find innovative ways to develop, manage and utilise water resources. An obvious solution to meeting this increasing demand, is developing new clean water resources. However, given the limited amount of clean water available and the increased cost of constructing new water storage and treatment plants, this is often not feasible. Desalination is another way to tap in to the massive amount of sea water available, however, desalination is very expensive, thus most countries cannot rely on this technology as an alternate source of fresh water supply. Given this, the focus of the governments and water industry has been shifting to better utilisation and conservation of the already available resources, particularly better management of the current water distribution systems (WDS).

Water distribution systems (WDS) are responsible for delivering water from aquifers and other sources to consumers. These systems form a complex network of interdependent components e.g., reservoirs, treatment plants, pumping stations, valves etc. Most of this infrastructure has evolved over decades. The size, cost and complexity of these evolved water distribution means that both, making major changes and applying modern techniques is difficult.

In most water distribution systems a large percentage of water is lost while in transit from treatment plants and reservoirs to consumers. Unaccounted water is usually attributed to several causes including leakage, metering errors, and theft. Of these, leakage is the major cause[1] (Hunaidi, 2000). In addition to environmental and economic losses caused by leaks, leaky pipes pose a public health risk, as leaks are potential entry points for contaminants if a pressure drop occurs in the system.

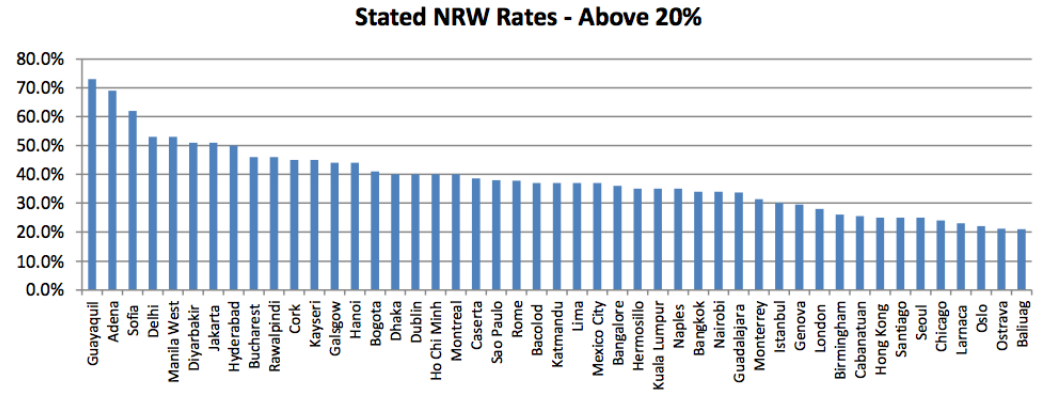

FIGURE 1.1: Cities around the world with NRW greater than 20%. Taken from (SWAN, 2011)

For water companies leaks result in revenue loss and are classified as non revenue water (NRW). Reduction of NRW is the primary step towards addressing the imminent water crisis. The Smart Water Networks Forum (SWAN, 2011) estimates that NRW in current water systems around the world is as high as 70 % in some cities. Figure 1.1 details cities with NRW greater than 20 %, showing NRW rates in London to be approximately 30% of the total water supplied. Since leaks are the primary source of NRW, leakage detection, location, and management is of paramount importance.

Although reducing the NRW caused by leaks is financially important for water companies, it is not the only factor that motivates companies to pursue timely detection and location of leaks. Water companies are not only responsible for optimally maintaining the WDS (i.e., meeting the government required standards while minimising the operational costs) but also efficiently managing leak situations. It is important to note that leaks are stochastic in nature and are impossible to predict beforehand (Hunaidi, 2000).

[1] Although leaks are often attributed as small pipe fractures resulting in smaller water losses where as larger fractures are termed as bursts, in this thesis the term leak or leakage will be used to refer to both and any distinction based on the leak size shall be made explicitly

Thus, effective management of leak events is critical to water companies, as both the public and regulatory authorities judge them on how well or poorly they perform during such events. In recent years, tolerance for supply interruption or damages caused by leaks has reduced significantly as more focus is applied to this area.

Given the aforementioned, timely leak detection can result in both financial and social cost savings. Timely leak detection can result in actions being taken before a leak grows large enough to cause significant damage or long term supply interruptions. This avoids fines imposed by UK Office of Water Services (OFWAT) for violating the Service Incentive Mechanism (SIM) and Serviceability measures and also payments compensating owners for infrastructure damage caused by leaks. Similarly, timely leak detection results in lower water supply, road traffic and fire fighting interruptions, reducing the social cost incurred by the companies, while also improving their image and customer trust. All these factors are the main drivers for the research carried out in this thesis.

1.2 Leak Detection in Water Distribution Systems

For ease of monitoring and serviceability, a WDS is often divided in to smaller zones called district metering areas, DMA, which can consist of a one hundred to a few thousand properties. *A DMA is an area of the distribution system that can be isolated by valves and for which the quantities of water entering and leaving can be metered* (Trifunovic, 2006). All leak detection/management activities in WDS systems are normally carried out at the DMA level. Given this, leakage detection and management can be classified in to three main categories based on the problem being addressed: leakage detection, quantification, and localisation. Leakage detection refers to the detection of leaks in the system without giving any information about the location of the leak in the network. Leakage quantification aims to estimate the amount of water lost in the system. Leakage localisation aims to identify and sometimes prioritise the area/location of the leaks within a DMA. In this research we will only address the problems of leak detection and quantification at the DMA-level. Since leakage detection is the logical first step before leaks can be quantified and located, early and accurate detection of leaks is of critical importance.

Currently, a wide range of techniques exist, which employ various principles, to detect leaks (Colombo et al., 2009; Puust et al., 2010; De Silva et al., 2011). Specialised hardware equipment such as acoustic sensors, ground penetrating radar (GPR), noise correlators etc., have already been used in the water industry. With continuous advancements in the hardware technology these methods offer increasingly accurate leak detection (Puust et al., 2010). However, hardware based methods are invasive and very costly, both in terms of the hardware costs and the manual labour required to effectively install and maintain them. These techniques are often slow to perform as they require

hardware installation and constant manual inspection and monitoring in different areas of the WDS. Furthermore, acoustic devices, noise correlators and GPR are prone to environmental noise (e.g., traffic, water use, wind and ground movement), pipe material variations (e.g., plastic pipes attenuate acoustic signals) and operator error (Hunaidi and Wang, 2006).

Given the above, in terms of practical or financial feasibility for long term deployment or real time leak detection, the techniques currently applied in the water industry are at best limited. Even now, in most water companies, leaks are only highlighted after a customer complaint or a visible water loss report. Given this, water companies are increasingly investing resources and manpower to pursue modern leak detection methods. Specifically, water companies are particularly interested in non-intrusive approaches that are low cost, automated, require minimal WDS information and can detect leaks in real time.

Owing to this, a considerable amount of research has been done towards hydraulic measurement based leak detection methods, which rely on the pressure and flow measurements in the WDS. Generally, hydraulic methods can be classified as model based or measurement based. Model based techniques try to form a numerical hydraulic model (NHM) of the whole WDS and then use this to detect leaks. This is done either by using the difference in the estimated and observed values of hydraulic measurements from sensors or by looking for transient signals caused by leaks in sections of the WDS. Some notable techniques in this category are steady state analysis (Pudar and Liggett, 1992; Mukherjee and Narasimhan, 1996), transient analysis (Liggett and Chen, 1994), negative pressure wave analysis (Misiunas et al., 2003, 2005; Srirangarajan et al., 2013). It must be noted that most model based approaches presented in the literature have been limited to very simple pipe systems or single pipe scenarios. Furthermore, most of these approaches assume ideal noise-free pipe or distribution networks and also require an accurate estimate of the type and configurations of pipe, joints, valves, pipe material, roughness coefficients etc., which are mostly not known or very difficult to estimate in reality. Another drawback of model based methods is the requirement to monitor pressure and flow at every node in the network, thus requiring sensors to be installed, maintained and monitored continuously which is often financially impractical for water companies to do. Even in cases when complete network information is at hand, incorporating all these factors in a mathematical model for the whole network is very difficult and computationally expensive. A select few attempts, (Covas et al., 2005), have been made to validate these approaches in field tests which are close to real life systems with little success.

In contrast to hydraulic model based methods, several measurement based methods exist that endeavour to offer cost effective, fast leak detection solutions. Out of these, of particular interest to this research are methods that use statistical and artificial intelligence techniques for automated, real time leak detection. Practical applications of intelligent

data analysis methods have only recently found their way to WDS problems. A variety of techniques including statistical analysis (Buchberger and Nadimpalli, 2004), artificial neural networks (ANN) (Khan et al., 2008; Mounce et al., 2002, 2007), generalised likelihood ratio method (Mukherjee and Narasimhan, 1996), principal component analysis (PCA) (Gertler et al., 2010), support vector machines (SVM) (Mashford et al., 2009) and Kalman filtering (Ye and Fenner, 2010) have been proposed to address the leakage detection, quantification and location problems. These methods offer several advantages over other leak detection methods. In particular, they do not require accurate knowledge of the WDS or specialised hardware and they rely solely on the empirical measurements. Furthermore, these methods require pressure and flow measurements to be sampled at a much lower frequency (e.g., 15 minutes) than model based methods (e.g., transient analysis).

In light of this, statistical/AI based approaches provide an attractive way forward, as they not only provide a way to automate the data analysis process, but can also efficiently handle sparse and noisy sensor data. This sparsity is important, since in a typical DMA, only three locations are commonly monitored for hydraulic readings. These are, the supply point pressure (P2), and flow (Q) (the point where the DMA connects to the rest of the network) and a point of the lowest pressure (P3) (critical point) as shown in Figure 1.2. Throughout this thesis we shall use the term DMA-level measurements/readings to refer to these three hydraulic time-series or data streams. Given this, statistical/AI based methods, that employ only DMA-level measurements, from the already installed sensors, are most appealing to water companies, as they require no additional infrastructure changes or investment.

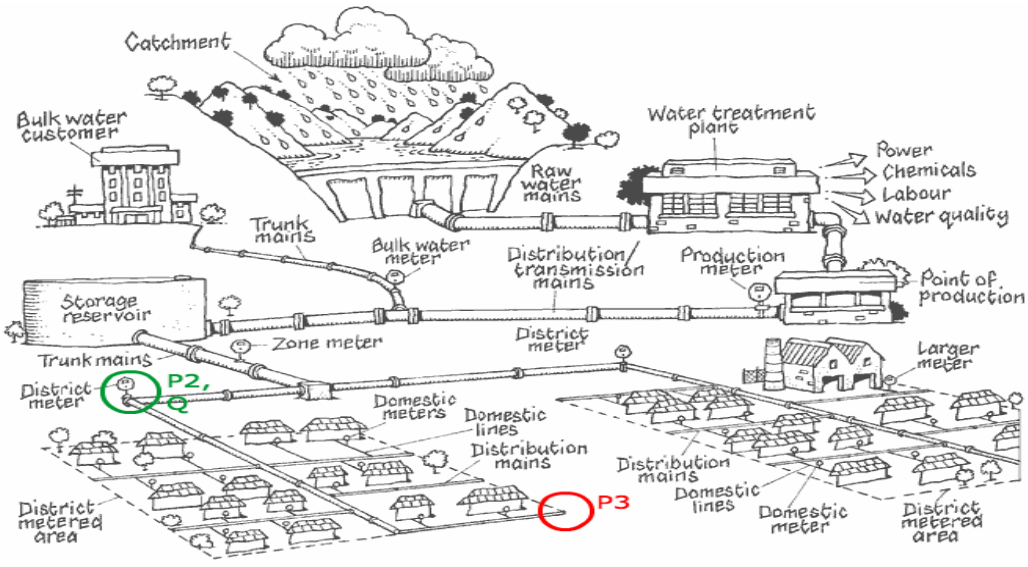

FIGURE 1.2: A typical water distribution system. Taken from PacificWater SOPAC (2012). Supply point pressure and flow, P2 and Q, marked with a green circle. Critical point pressure, P3, marked with a red circle.

Conversely some hydraulic measurement based methods require pressure or flow sensors either at every node, or at multiple locations within a WDS (Buchberger and Nadimpalli, 2004; Khan et al., 2008; Gertler et al., 2010). This may be feasible for conducting short term engineering studies or field tests. However, long term deployment of sensors, particularly ones that transmit data in near real time (e.g. 15 mins), is often not financially feasible for water companies. This is because in practice, the financial benefit gained from installing a large number of hydraulic sensors in the network is often far less than the costs associated with installing and maintaining them (e.g. infrastructure, sensor installation, maintenance, data transmission, labour training, battery life issues and replacement costs etc.) (McHenry, 2013). While customer meters can provide data for every node in the WDS, most customer meters installed in UK are analog flow meters which do not have the capabilities to provide data in real time. In fact the job of reading these meters, once a month, is often left to the customers, who then inform the water company by updating the meter readings on the company website. To follow this up, normally, every three months or so, a company employee is sent to take readings from all customer meters and update the company database. Smart autonomous, real time metering for WDS is an active area of research (Marvin et al., 1999; Gurung et al., 2015; Kashi, 2016). However, various cost benefit studies on smart meter technology have highlighted the large investment costs of installing smart meters as a critical barrier to their industry wide adaption (SWAN, 2010; McHenry, 2013). It must be noted that making smart meters technology affordable is an active area of research (Molina-Markham et al., 2012; Batista et al., 2013) and with companies like WavIot and Libelium already investing heavily in this area, cheap smart meters can be a reality in the near future (further discussion in Section 6.2). However, until end to end smart metering technology becomes affordable for the water industry, either by technological advancements or government incentives/subsidiaries, accurate detection and quantification of leaks when only DMA-level hydraulic data is available, is an area that entails further improvement.

1.3 Problem Description

In the past decade the rapid advancement in technology and outsourcing has found its way to the water industry. There are now a number of specialised firms offering both hardware and software services that are particularly tailored for the water industry. This research conducted in this thesis, is in collaboration with one such firm, I2O-Water Ltd, which offers pressure optimisation solutions to water companies. These solutions include both hardware (e.g., remotely controllable pressure reducing valves (PRV), data loggers etc.) and software (e.g., automatic network optimisation, data management, remote/online monitoring and configuration etc.). Indeed, the advantages of network pressure optimisation (e.g., leak reduction, energy saving, pipe asset lifetime extension)

have been known and proven in numerous studies (Germanopoulos and Jowitt, 1989; Lambert, 2001).

Since leakage rates are typically directly proportional to average pressures in the DMA, keeping the network pressure as low as possible while maintaining desired serviceability is an effective, encouraged and widely used way of proactive leak reduction. As the critical point pressure (P3) is representative of the lowest serviceable pressure in a DMA, pressure optimisation, in an I2O-Water managed DMA, is achieved by keeping the DMA supply pressure (P2) as low as possible while maintaining the P3 as close to the guaranteed standards scheme (GSS) set by OFWAT as possible (i.e. seven metres static head[2]). To monitor and control the pressure in real time specialised hardware (i.e., remotely controllable PRV, data loggers) are installed at the supply and critical point respectively. Despite this, leaks still occur and cannot be completely eliminated from WDS. Given this, the problem addressed in this work is how leaks can de detected and quantified in real time using only DMA-level hydraulic measurements in pressure optimised DMA's.

The DMA-level measurements represent the current aggregate water demand/consumption within the DMA. Detecting leaks at this point is challenging, particularly small leaks, as they do not produce a significant increase in the aggregated DMA-level measurements (e.g., a pumping station or water reservoir, even when there is a leak may still be able to provide the required pressure to sustain the water demand in a DMA). Also, flow and pressure within the DMA are not only correlated but also affect the noise levels associated with the readings being taken. For example, an increase in the demand (as consequentially, an increase in flow) decreases pressure, and an increase in the supplied pressure increases the flow within the pipes. This increases the noise levels associated with the readings being taken and also the leak magnitude (if a leak is present in the system at that time). Thus, both leaks and the noise in the observed readings in the system are dependent on the inputs. This input dependent and time varying noise is often referred to as heteroscedasticity. Besides leaks, genuine customer consumption can also cause an unexpected increase in the flow e.g. fire fighting, or high demand during a festival. Therefore, the duration of the unexpected increase must be taken into account before attributing it to a leak, as attributing legitimate short term increases in demand to leaks can result in unnecessary and costly inspections by engineers. This is particularly important for companies such as I2O-Water where any false positive leak reported to their client water companies can result in the clients losing trust and doubting the reliability of the services offered to them. Apart from detecting the presence of a leak, for financial accountability, quantification of the water loss during a leak, is of equal importance for the water companies and should be an integral part of any leak detection

[2]Meter head: A concept that relates the energy in an incompressible fluid to the height of an equivalent static column of that fluid. The pressure at any point in WDS is measured in meter water column (m) and is often referred as meters of pressure head or simply pressure head.

solution. As the water demand in a DMA changes with time, the pressure characteristics also change. Since leaks are pressure dependent, the amount of water loss due to leaks at any given time may also change (depending on the location of the leak in the DMA), this makes accurate estimation of water loss during a leak challenging. Given the aforementioned, having highlighted the research problem and challenges, we shall now list the conditions a leak detection solution must fulfil to be practically viable for industrial deployment.

1.3.1 Industrial Requirements for a Leak Detection Solution

In light of the nature of the leaks and monitoring requirements in a DMA, after discussion with subject matter experts from I2O-Water, the requirements for a practical and industrially viable leak detection solution can be listed as:

1. Timely detection of leaks when only limited water network data is available i.e using only aggregated DMA-level measurements.
2. Distinction of legitimate short term increase in water consumption from leaks.
3. Estimation of the leak magnitude to quantify water loss.

1.4 Research Scope and Objectives

In light of the problem, challenges and requirements described in Section 1.3, the scope of the research work done in this thesis is to propose and test a leak detection method that can detect and quantify leaks in real-time, using only DMA-level hydraulic measurements. In doing so, the proposed method would facilitate timely action and repair in case of leak events. Thereby providing water companies the opportunity to attain higher customer service and operational standards while reducing costs in terms of the energy and water lost due to leaks.

The research scope detailed above is achieved through the following objectives:

Obj 1: To investigate the feasibility of using DMA-level hydraulic sensor measurements as the sole source of information pertaining to the normal and abnormal (when there is a leak) operational behaviour of a DMA.

Obj 2: To explore the possibility of applying statistical/AI methods to model and learn the patterns of a DMA when there is no leak in the system i.e. the normal operation model (NOM).

Obj 3: To investigate and develop a methodology to detect and learn the deviations from the NOM, resulting from leaks and use them to detect and quantify leaks.

Obj 4: To research the feasibility of using Bayesian reasoning to incorporate multiple sources/sensors measurements as evidence of a leak i.e. all DMA-level measurement data streams.

Obj 5: To validate the proposed novel approach by testing its performance against the current state of the art leak detection methods based on DMA-level hydraulic measurements.

1.5 Research Contributions

Against the research scope and objectives detailed in the previous sections this thesis makes the following contributions:

1. **Leak Detection and Step Change Magnitude Approximation using Noise Scaled Semi Parametric Gaussian Process (NSGP) (Chapter 3)**: Initially we propose a method that uses the flow data stream, from all three DMA-level data streams, to detect and estimate leaks. In doing so we make the following contributions:

   (a) We propose a Gaussian process (GP) based approach, Noise Scaled GP (NSGP) that can efficiently model the normal operating DMA-level flow patterns even in presence of heteroscedasticity (satisfying Obj 1 and 2). We do this by first computing the weekly mean from historic flow data and then using it as GP prior mean. We model the heteroscedasticity by defining and using a diagonal noise covariance function that is based on the observed weekly historic flow variance.

   (b) Given the above, we parameterise the proposed GP in a way that allows us to use WDS domain knowledge to dynamically bound the leak detection parameters, at each step, resulting in a fast leak detection and quantification method that can manage sustained leaks with high detection accuracy (satisfying Obj 3). Specifically, we do this by defining an additive parametric step change mean function. The parameterization captures the properties of the leak we are interested in learning e.g. leak start time and leak magnitude, which allows us to not only detect a leak, but also, approximate the leak magnitude. The resulting model has the advantage of modelling a leak as being independent of the underlying NOM. This leads to a more accurate depiction of the physical leak process, since leaks result in sudden sustained additions uncorrelated to the original underlying consumption patterns. For each new observed flow reading, based on our knowledge of the WDS, we place dynamic bounds and then discretise the leak magnitude parameter (details in Chapter 3). This allows us to automatically prune our leak parameters search space

resulting in a robust and fast detection approach that can be used in an online setting.

(c) By analysing the performance of the proposed NSGP, against both, the state of the art in GP based fault detection algorithms and current applied leak detection approaches in the water industry, on real data with simulated time varying leaks, we show that the proposed NSGP outperforms the state-of-the-art considerably (satisfying Obj 5).

2. **Probabilistic Leak Detection And Characterisation Using Efficient Semi-Parametric Multi Output Gaussian Process. (SMOGP) (Chapter 5):** We extend the NSGP by proposing a second approach, Semi-Parametric Multi Output Gaussian Process (SMOGP), which incorporates all DMA-level readings allowing the leak detection decision to be based on all available data from a DMA. In detail we make the following contributions:

(a) We extend the NSGP by proposing a multi output Gaussian process model that can learn and model the normal operating patterns of the DMA based on all available DMA-level hydraulic time-series (satisfying Obj 1 and 2). We do this by modelling the cross-correlation between the hydraulic time-series in a DMA, as a Hadamard product of a covariance function over time and a covariance function over the time-series labels (satisfying Obj 3 and 4).

(b) We extend the additive parametric fault detection mean functions in two ways. Firstly, we extend it to incorporate multiple time series by redefining it as a Hadamard product of two labeled mean functions, which capture the leak time and magnitude for each of the time series. Secondly, instead of the simple step change leak magnitude approximation, we model pressure leak relationship within the mean function. This allows us to accurately quantify pressure dependant time varying leaks while keeping the leak detection mean function parameters limited to a computationally acceptable level (satisfying Obj 3 and 4).

(c) We propose a computationally efficient dual optimisation and learning mechanism for the hyper parameters of the SMOGP, which we divide in to two sets, NOM and leak hyper parameters. The NOM hyper parameters capture the characteristics of the DMA under normal no leak conditions and are optimised using conjugate gradient decent (the standard likelihood based optimisation method). While the leak hyper parameters are learnt by maximising the sum of the posterior predictive probabilities, using a bounded search process. As in case of the NSGP, at each step of the leak hyper parameters optimisation process we dynamically bound the possible hyper parameter values using WDS domain knowledge resulting in a domain specific, computationally efficient multi output Gaussian process model.

(d) We do a comparative analysis of the proposed novel SMOGP against the state of the art algorithms we compared the NSGP with. Additionally, we also analyse the performance of SMOGP against the proposed NSGP. To achieve this we model a synthetic pressure optimised DMA using the latest EPANET[3] extension WaterNetGen (Muranho et al., 2014), which adds the capabilities of running pressure dependant demand and leak simulation to EPANET (details in Chapter 4) (satisfying Obj 6).

Summary Of Contributions:

Given the above, in this thesis we propose two novel GP based leak detection algorithms NSGP and SMOGP. By proposing the NSGP we extend the state of the art in applied ML based anomaly detection methods for single time-series. In particular, firstly the NSGP proposes a way to model heteroscedastic periodic time series in a GP framework which is computationally more efficient and accurate than existing GP models (Kersting et al., 2007; Lazaro-Gredilla and Titsias, 2011). Secondly, it uses an efficient dual optimisation mechanism to detect sequential anomalies, which is far more accurate that existing GP based anomaly detection methods (Garnett et al., 2010; Osborne et al., 2012). For both, heteroscedasticity modelling and sequential anomaly detection the comparative analysis results (Chapter 3) show that the proposed method is computationally 5 times faster than current state of the art GP based fault detection methods with a detection accuracy of 99%. Although the NSGP has been applied to leak detection in WDS it can be easily applied to other domains where heteroscedasticity data is common e.g. energy prediction (BREHM et al., 2012) and wind speed measurements (Tol, 1997).

Our second contribution to applied ML is the SMOGP algorithm. In this method we propose a way to model linearly correlated time-series in a GP framework which results in a computational efficient covariance structure. In particular the SMOGP can model the cross-correlations between multiple linearly correlated time-series using only one additional hyperparameter in contrast to existing heavily parameterised multi-output GP models (Osborne et al., 2008; Alvarez and Lawrence, 2009) where to model the cross-correlations the required hyperparmeter, at best, are in order of the number of inputs i.e., O(N). Additionally the SMOGP, proposes a novel method to incorporate evidence from multiple time-series to detect sequential anomalies which to our current knowledge hasn't been done in any multi-output GP based approach.

In addition to our contributions to GP based applied machine learning anomaly detection algorithms, the proposed methods extend the state of the art in leak detection in two ways. Firstly, in this thesis we analyse the effects of pressure optimisation (Chapter 4) on leak detection and show that the behaviour of the WDS differs drastically in

[3]EPANET is software that models water distribution piping systems. EPANET is public domain software that may be freely copied and distributed. EPANET performs extended period simulation of the water movement and quality behaviour within pressurised pipe networks.

comparison to a normal unoptimised DMA. To our knowledge to this date, the effects of pressure optimisation on leak detection, specially leak detection using both DMA-level pressure and flow measurements, haven't been studied before. In light of this, the proposed method extends the state of the art in leak detection by proposing an efficient leak detection method that can detect leaks using all available DMA-level measurements. Additionally, the proposed method also gives a best effort estimate of both the leak coefficient and exponent. The leak coefficient and exponent provide an indication to both, a ratio of how the water loss increases with a unit increase in the supplied DMA-level pressure and the type of the pipe in which the leak occurred (for details see Chapter 5). The term best effort here means that given the amount of data available, after a leak has been detected, the algorithm gives a best effort estimate (not guaranteed optimal) for the values of the leak coefficient and exponent. As more data becomes available the estimates become more accurate (detailed discussion on benefits and limitations of the estimated in Chapter 5). To our knowledge, this is the first attempt by any leak detection approach, to firstly address the issue of leak detection in pressure optimised DMA's, secondly use multiple DMA-level data streams as evidence to improve the detection accuracy and thirdly to make an attempt to approximate the leak coefficient and exponent, all in one computationally efficient algorithm.

Having outlined the research contributions, we list the manuscripts that have been published or are in preparation for submission in support of these contributions:

1. Malik, O., Ghosh, S. and Rogers, A., 2015. A Noise Scaled Semi Parametric Gaussian Process Model for Real Time Water Network Leak Detection in the Presence of Heteroscedasticity. In Computational Sustainability Workshop at the Twenty-Ninth AAAI Conference on Artificial Intelligence.

2. Malik, O. and Rogers, A., Time-Delayed Probabilistic, Near Real Time Leak Detection In Water Distribution Systems Using Multi-Output Gaussian Process, In preparation for ACM Transactions on Intelligent Systems and Technology (ACM TIST)

Having summarised the research contributions, we now describe the structure of this thesis.

1.6 Thesis Structure

The remainder of the thesis is organised as follows:

1. In Chapter 2, we initially provide a brief overview of the workings of a typical water distribution systems and the challenges faced in managing them. This is

followed by a review of the current available leakage detection techniques. Here we highlight the advantages and disadvantages of each technique in light of the research requirements detailed earlier.

2. In Chapter 3, we give a description of the proposed NSGP model for leak detection using DMA-level flow readings. Following this, we explain the experimental setup and performance metrics used in the comparative analysis. Next, we briefly describe the state of the art leak detection algorithms used in the comparative analysis which is followed by a discussion on the results of the analysis.

3. In Chapter 4, we start by detailing the challenges faced in acquiring usable truth data for leaks. This is followed by a description of the simulated DMA in WaterNetGen. We then outline the strategy used to implement, I2O-Water like, pressure optimisation in the simulated DMA, which is followed by a discussion on the method used to simulate leaks. In the end we give a brief discussion on the effect of active pressure optimisation on DMA dynamics and leak detection methods.

4. In Chapter 5, we introduce the novel multi output Gaussian process model for leak detection and quantification. First, we show how we can efficiently use multiple hydraulic data streams to develop a NOM for a DMA. We then detail how leak can be detected and quantified by parameterising the proposed GP. This is followed by a description of the dual optimisation process used to learn the NOM and leak parameters and the approach used to assign a probabilistic time varying confidence value to the detected leaks. We then detail the experimental setup and data used in the comparative analysis. We end with a discussion on the results of the analysis.

5. Finally, Chapter 6 gives a summary of the research presented in this thesis and the conclusions that can be drawn from each chapter. We also discuss limitations of our methods and outline the direction for future work.

Chapter 2

Background

This chapter presents the background on water distribution systems (WDS) and a review of existing leak detection approaches in literature. In Section 2.1 we begin by giving a description of the structure and working of a typical WDS and its various components. This is followed by a general overview of leakage management in Section 2.2, briefly describing the traditional reactive and the modern proactive approach to leakage management. We classify the leak detection approaches found in literature in to three categories, hardware based, hydraulic model and hydraulic measurement based methods. We discuss each of these categories in Sections 2.3, 2.4 and 2.5 respectively, highlighting the strengths and limitations of each approach. Following this in Section 2.6 we give a detailed description of three state of the art leak detection methods from literature used as benchmark in this thesis. This is followed by a discussion of our data acquisition methodology in Section 2.7, where we give a very brief overview of the WDS simulation software, EPANET and WaterNetGen, used to generate DMA-level hydraulic measurements and leak data. Finally, in Section 2.8 we conclude by giving a summary of the chapter highlighting the gaps in the current literature.

2.1 The Structure of Water Distribution Systems

Water distribution systems collect water from source sites and deliver it to the consumers. Often the water from the sources is first stored in storage facilities like dams and reservoirs and transported from there to cities. Urban water supplies to consumers require high quality water, however natural surface or underground water sources often don't meet the requirements for domestic and industrial use. In such condition water is treated in a treatment plat before being delivered to end consumers.

Pipes are the backbone of water transportation systems. Based on the purpose pipes can be classified as follows Trifunovic (2006):

1. Trunk Main: Trunk mains are large pipes that supply water from the treatment plants to the distribution area. Their diameters can vary from 100 millimetres to a few meters depending upon the capacity being supplied.

2. Secondary Mains: Secondary mains link the main components in the WDS e.g. water sources, reservoirs, pumping station etc. They are also used to provide customers with high water demand e.g. factories.

3. Distribution Mains: These pipes are laid alongside roads and streets to supply water from secondary mains to the consumers.

4. Service Pipes: Service pipes supply water from the distribution mains directly to the customers.

Based on the layout of pipes, transmission and distribution systems can be looped or branched. Looped systems offer multiple paths for the water to reach the end consumer where as branched or tree systems have only one possible path to the end user. Coupled with sufficient valving mechanisms looped systems are preferable as the provide extra level of reliability, lower flow velocities and higher capacity. A typical WDS normally consists of a combination of both looped and branched pipe layouts, often trading off between reliability and infrastructure cost. Figure 2.1 gives an example of looped and branched network where the end consumption points are represented by nodes.

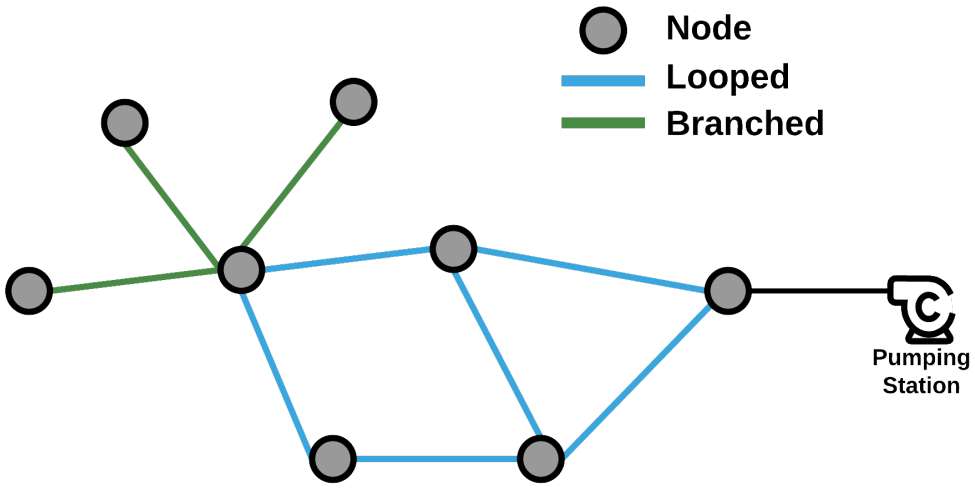

FIGURE 2.1: A simple hypothetical water network showing looped (blue) and branched (green) pipe configurations.

Water distribution schemes are based on the topographical conditions of the area. They can be classified based on the way water is supplied (Trifunovic, 2006). When the water sources are at a higher elevation than the distribution area there may be no need for pumping water. Gravity schemes rely on gravity to supply water directly to

the consumers. Where gravitational schemes can not be used direct pumping schemes are used. In such schemes a pumping station is used to pumps water directly to the distribution area. To cater for varying demand patterns the pumping stations are often connected to water storages such as towers, ponds etc. Such schemes are called combined schemes. For ease of monitoring and serviceability, a WDS is often divided in to smaller zones called district metering areas, DMA, which can consist of a one hundred to a few thousand properties. *A DMA is an area of the distribution system that can be isolated by valves and for which the quantities of water entering and leaving can be metered.* All leak detection/management activities in WDS systems are normally carried out at the DMA level. Figure 2.2 illustrates a typical water distribution system. Having given a basic

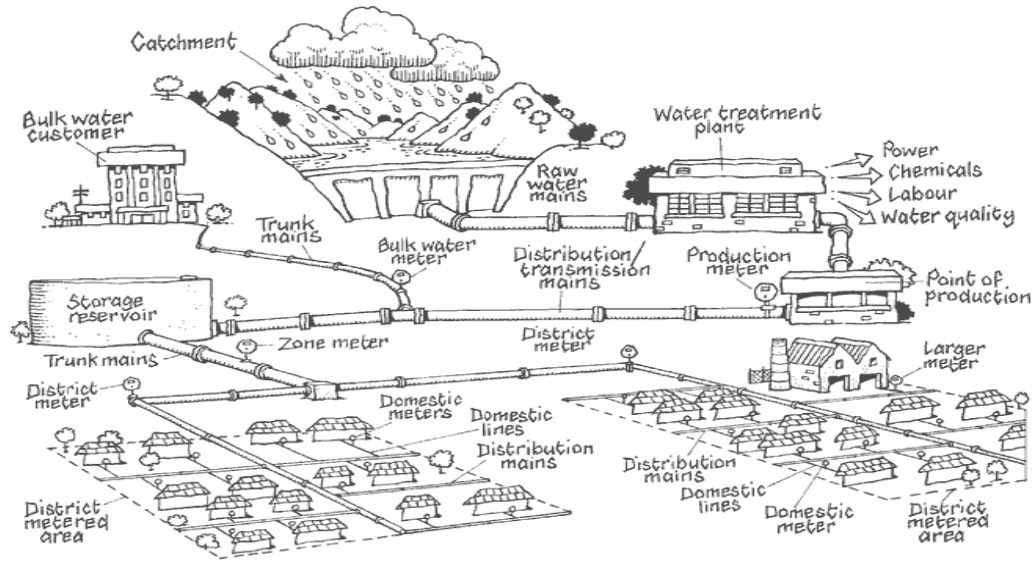

FIGURE 2.2: A typical water distribution system. Taken from PacificWater SOPAC (2012)

description of the structure and working of a typical WDS and its various components, in the next section we give a general overview of leakage management.

2.2 Leakage Management

Leaks have been a major cause of concern for the water industry since the earliest WDSs were made. As such, leak detection and management is an interdisciplinary challenge spanning multiple fields such as environmental control, network design/optimisation, water quality, sensor networks etc. Owing to this this is a vast amount of literature available where different tools and backgrounds have been explored to propose leak management strategies/systems. An extensive review of the these approaches can be found in (Al-Dhowalia and Shammas, 1991; De Silva et al., 2009; Colombo et al., 2009; Puust et al., 2010; Bieupoude et al., 2012). All these leak management approaches can be generally classified as either proactive or reactive.

Proactive leak management entails implementation of mechanisms for active control and preventions of leaks. Such mechanisms include active pressure management and real time leak detection and control. Active pressure management entails keeping the network pressure as low as permissible while maintaing the supply standards. This results in significant reduction of the number of new leaks, as leaks are directly related to the supplied pressure (Germanopoulos and Jowitt, 1989; Lambert, 2001). In case of a leak, resulting from a crack or aperture in a pipe, the higher the pressure the greater would be the water loss. Furthermore, the aperture in a pipe can enlarge as the pressure in the pipe increases resulting in even larger water losses. In this context, it must be noted that pressure optimisation is not only beneficial to companies like I2O-Water but can result in leak reduction and energy savings for all water industry. However, pressure optimisation also changes the underlying behaviour of the WDS (detailed discussion in Chapter 4), in particular in a pressure optimised WDS in case of a leak the overall pressure in the network would go down, in such case the pressure optimisation mechanism would increase the supply pressure of the network to cater for this reduction. In case of a leak the pressure increase in the network will result in larger water losses. Thus, currently industry wide used leak detection methods, such as the nightline analysis, which rely only on one minimum variance measurements during the night would no longer be effective. For example, by detecting leaks using only one night time flow measurement the nightline analysis method would not be able to detect a leak that occurred the following morning until night the next day. In case of a pressure optimised DMA the pressure optimisation might increase the pressure to cater for a pressure drop resulting in higher water loss. Thus, pressure optimised DMAs require specialised methods that can detect leaks in real-time so to avoid increasing water losses.

In addition to pressure optimisation, active replacement or rehabilitation of deteriorating assets, although costly but is another effective mechanism for reducing the occurrence of new leaks. Given this, pressure management and asset renewal are mechanisms that every modern company should practice for effective leak management.

Traditionally water companies only react and repair leaks when they become evident (e.g. water escaping from ground) or are found after inspection resulting from a customer complain (e.g. low pressure or no supply at customer end). Thus, all leak management approaches where the companies wait for a leak to occur and repair or take actions on it thereafter, fall under reactive leakage management. Reactive approaches can broadly be classified in to three main categories, hardware based methods which employ specialised hardware devices to detect and locate leaks, hydraulic model based methods which use the hydraulic behaviour of a WDS (pressure and flow measurements) to construct a numerical hydraulic model (NHM) of the WDS and use the NHM to detect and locate leaks and hydraulic measurement based methods which rely solely on the hydraulic behaviour of a WDS (pressure and flow measurements) to detect leaks. A detailed

description of the hardware, hydraulic model and hydraulic measurement based methods is given in the Sections 2.3, 2.4 and 2.5 respectively.

2.3 Hardware Based Leak Detection Methods

Hardware based methods employ use of specialised devices to detect and locate leaks in a WDS. These devices can range from acoustic listing devices, thermal imaging cameras to electromagnetic pulse generators. Based on the scientific principle employed for leak detection, these techniques can be classified in to the following categories:

- Sounding methods
- Gas based methods
- Temperature change based methods
- Inline pipe inspection
- Radar based methods

We give a description and review of each of the above mentioned category in Sections 2.3.1 to 2.3.5 highlighting the advantages and disadvantages of each approach.

2.3.1 Sounding Methods

As water escapes a leaking pipe, it emits an acoustic signal, this signal traverses through the pipe walls and its close surroundings. Sounding methods use specialised acoustic equipment to detect the leak generated acoustic signal to confirm the presence of a leak. Sounding methods are perhaps the most commonly employed methods for leak detection in the water industry. Based on the particular type of acoustic device employed these methods can be categorised in to three main types namely listening rods, noise loggers and leak noise correlators. We give a brief review of each type.

2.3.1.1 Listening Rods

These devices use modern equipment such as microphones coupled with signal amplifiers and hardware or software noise filtering mechanisms to improve the leak induce acoustic signal detection Pilcher et al. (2007). The use normally entails leak engineers walking along the sections of the WDS and placing the listening rod on them i.e, performing sounding surveys. The advantage of devices are ease of use. However, the detection ability of listening rods is prone to human error, due to fatigue and noise interference from surroundings e.g. traffic, wind etc.

2.3.1.2 Noise Loggers

Noise loggers are data loggers with specialised acoustic sensors that are deployed (placed on pipes or left in close proximity of them) in the WDS to record acoustic signals typically during the night time (3:00 to 4:00 AM). Leaks are then detected by comparing the recorded noise signals at different sections. As noise loggers typically work during night time, they are less affected by the high surrounding noise during the day. Furthermore, these devices can be deployed for an extended period. However, Sánchez et al. (2005) in results from a six month study, point out that the high cost associated with installing these devices as a major disadvantage in their long term use.

2.3.1.3 Leak Noise Correlators

Leak noise correlators, as the name suggests, detect leaks based on the cross correlation between the leak induced noise at two ends locations of a pipe. Noise correlators were first introduced by Water Research UK and were found be to too expensive and time consuming to be used commercially (Cole et al., 1979). These methods have the advantage of providing a greater resilience from environmental noise and human error. However their high cost remains the major reason why most companies do not employ them.

Additionally, Hunaidi and Wang (2006) point out factors that influence leak detection accuracy of sounding methods. In particular, these methods do not work well with plastic pipes, due to high noise attenuation, or pipes with large diameters that results in reduced sound propagation.

2.3.2 Tracer Gas Based Methods

These methods are based on the principle that water insoluble, non toxic gases, such as helium, when inserted in to an isolated WDS section will escape through leaky pipes. Scanning the WDS section for the presence of the gas, using gas sensor, can then confirm the presence of a leak. These method have the advantage of being able to detect leaks in plastic pipes where acoustic equipment based methods do not work well. Furthermore, these methods can also detect multiple leaks along the same pipeline. Despite this, tracer gas methods have several drawbacks (Hunaidi et al., 2000). As the detected area for the escaping gas is small i.e., the immediate vicinity of the pipe, leaks are often missed if the scanning is done at the correct distance from the pipe. Also, if the pipe is buried to deep underground or the leak is not at the top of the pipe the gas may not escape. Apart from the high personal and equipment cost, another drawback of these methods is the requirement to isolate sections of the WDS being surveyed which can hinder day to day WDS operations.

2.3.3 Temperature Change Based Methods

Leaks cause water to escape from a pipe into the ground surrounding the pipe. The resulting excess water in the ground, surrounding a leaky pipe, causes a temperature difference in the soil. This temperature difference can be detected using thermal imagery. Temperature change based methods employ specialised cameras that measure the thermal infrared radiations at the ground surface to construct thermal images. Computer analysis of the colour changes or tone difference in these thermal images is used to detect and locate leaks. These methods have been shown to give accurate leak detection results (Wirahadikusumah et al., 1998).

Despite the accuracy of these methods Fahmy and Moselhi (2009) show that these methods are greatly affected by environmental conditions such as cloud coverage, humidity, ground and thermal noise.

2.3.4 Inline Pipe Inspection Methods

With the advancements in technology specialised devices have been developed that can be inserted directly in to a WDS pipe. The flow in the pipe propagates these devices, which equipped with sensors record information on the go. With particular sensors, these devices can be used to detect and locate leaks. Based on the objective the sensors in these devices can range from microphones (for leak detection using noise), to wave generators or ultrasound sensors (for finding cracks in a pipe wall) (Fletcher and Chandrasekaran, 2008).

Despite their technological advantages, these devices are invasive and can affect the quality of water being supplied. Furthermore, these methods have limited applicability in practical WDS due to the presence of valves, joints and pipes of different sizes, all of which can hinder or totally obstruct the movement of these devices.

2.3.5 Radar Based Methods

Similar to the temperature profile based methods, radar based methods construct an image profile of the subsurface by sending high frequency pulsed electromagnetic waves in to the ground. The velocity of these waves is altered when passing through different materials. This changes in velocity causes a time lag in the reflection of the waves, which is captured at the surface. The time lag caused by, water saturated soil or circulating water around leaky pipes results in noticeable areas in the subsurface map, which are then sued to detect and confirm the presence of leaks.

Radar based techniques have been shown to effectively and accurately detect and locate leaks (Stampolidis et al., 2003; Farley, 2008). The main disadvantage of this method is

the loss of electromagnetic signal strength in conductive materials such as conductive clay or saturated soil. Furthermore, metallic objects underground can result in these methods giving incorrect results.

2.4 Hydraulic Model Based Methods

Hydraulic model based techniques try to form a numerical hydraulic model (NHM) of the whole WDS and then use this to detect leaks. Two prominent approaches in model based techniques are steady state and transient analysis which are discussed in the Sections 2.4.1 and 2.4.2.

2.4.1 Steady State Analysis-Based Methods

In stead state analysis a NHM of the WDS is made, where some aspects of the WDS such as pressure, demand and flows are modelled from actual WDS sensor measurements, while parameters such as pipe roughness, leaks, friction factors etc. are unknown. The NHM is then used to simulate leaks that match the observed measurements. Once a potential match, that minimises the difference between the observed and simulated values, is found the inverse problem of finding the unknown parameters is solved. (Pudar and Liggett, 1992) proposed the inverse steady state analysis methods by presenting a simple pipeline network hydraulic model where the inverse problem was solved by minimising the squared difference between the measured and simulated values. Following this several advancements were made to this method. Wu and Sage (2006, 2007); Wu et al. (2009) proposed methods to incorporate leak detection in the NHM parameter tuning process, where emitters where used to simulate leaks and the optimisation was done using genetic algorithms (GA). For their analysis, the effectiveness of the NHM solution was compared with field deployed pressure sensors, with a resolution of one sensor per two hundred properties. Although their analysis produced satisfactory results, it highlighted the need of setting precise and accurate demand and pressure values for multiple nodes, to obtain satisfactory results.

Given this, the main disadvantage of these methods is the sole reliance on accurate NHM of the WDS, which is often difficult to achieve. Indeed, NHM are rarely used by water companies, even in cases where they exist they are not maintained. Another drawback of these methods is the long computation time they require when solving the inverse problem. Owing to these drawbacks, these methods have not found any practical implementation in the industry.

2.4.2 Transient Analysis Methods

Transient flows, in WDS, are short duration unsteady flows in a pipe, caused by a sudden changes in pressure. These changes can be cased by a sudden leak or other operations such as pump startup or valve closure. Transient flows generate waves that traverse away from the source. A transient wave loses energy and produces a reflection wave when it encounters a leak in a pipe. The loss of energy produces a damping effect in the transient wave. This wave damping and the reflective wave can be measured and used to detect leaks. As transient wave speeds can be as high as 1200 meters per second, high frequency pressure measurements are required to detect the transient effects.

Given the above, transient analysis methods work by introducing transient waves in WDS sections and then using high frequency pressure sampling to detect and locate leaks. Two main categories of transient analysis techniques are inverse transient analysis (ITA) (Liggett and Chen, 1994) and frequency domain transient analysis (Lee et al., 2006). ITA methods analyse the transient signal in the time domain and use the measured transient data to calibrate a NHM which is then used to detect leaks. In contrast, frequency domain transient approaches analyse the transient signal in the frequency domain for calibration of the NHM.

Similar to steady state based methods, the main draw back of transient analysis methods is their requirement of precise information about friction factors, pipe roughness, WDS geometry etc., for forming an accurate NHM. Such factors are mostly now known and very difficult to estimate. Owing to this a select few attempts have been made to validate these methods in field test with little success (Covas et al., 2005).

2.5 Hydraulic Measurement Based Methods

Hydraulic measurement based methods rely solely on the pressure and flow measurements from the WDS to detect leaks. These methods offer several advantages over other leak detection methods. In particular, they do not require accurate knowledge of the WDS and they rely solely on the empirical measurements. Furthermore, these methods require pressure and flow measurements to be sampled at a much lower frequency (e.g., 15 minutes) than model based methods (e.g., transient analysis). The main categories of measurement based techniques discussed in this section are step tests (Section 2.5.1), water audits (Section 2.5.2) and statistical/AI based leak detection methods (Section 2.5.3).

2.5.1 Step Tests

Step testing procedures work by isolating sections of a DMA in a stepwise manner by closing valves temporarily. The flow rates when each section is isolated are noted. A larger than anticipated decrease in the measured flow, gives an indication of the presence of a leak in the section which has just been isolated. To minimise supply interruptions step tests are usually carried out during the minimum night line flow period. Step tests are generally performed in one of two ways (Pilcher et al., 2007). Firstly, by closing valves to isolate section and then reopening them once the flow measurements have been taken. In the second way, valves farthest away from the main supply point or flow meter, are closed first, moving sequentially to the supply point, until the measured flow drops to zero. The main drawback of these methods is the supply interruptions and the risk of infiltration as a result of network de-pressurisation due to valve closure.

2.5.2 Water Audits

Water audits are two part procedures that aim to find and account for all the water flowing in and out if a WDS. International Water Association (IWA) has published guidelines for conducting water audits (IWA, 2000). For ease of accountability, these guidelines give a breakdown the total amount of water supplied to a WDS which is shown in the Figure 2.3.

In the first part of the water audit, using WDS records (customer bills, reports etc), data is complied that breaks down the total supplied water in to categories listed in Figure 2.3. This part of the audit is often referred as top down loss assessment. The second part of the water audit aims to minimise the discrepancies in the loss figures estimated in the top down loss assessment. This is done by computing the real losses using the minimum nightline flow (MNF)[1]. As the water consumption is minimal during the night, the measurements at this time represent flows with the highest level of real losses. Subtracting the estimated consumption from the measured MNF gives an account of the real losses at night time, which are then extrapolated to obtain daily loss figures. Water audits can be system wide (for the whole WDS) or district wide (at DMA level)

Given this, water audits provide companies with a picture of the WDS efficiency and indication of areas that require further improvement. However, it must be noted that water audits give a snapshot of the WDS and are normally performed, at best, on a monthly or quarterly scale. Thus, although water audits are very useful they are cannot be used to detect leaks in real time.

[1] In literature two descriptions of the minimum nightline flow for a day are used. In the first description the MNF is defined as the minimum out of all the average hourly readings taken from 00:00 to 23:45 inclusive. The second description gives the MNF as the average flow between 02:00 to 04:00.

<table>
<tr><td rowspan="10">System Input Volume</td><td rowspan="2">Authorized Consumption</td><td rowspan="2">Billed Authorized Consumption</td><td>Billed Metered Consumption (including water exported)</td><td rowspan="2">Revenue Water</td></tr>
<tr><td>Billed Unmetered Consumption</td></tr>
<tr><td rowspan="2">Unbilled Authorized Consumption</td><td>Unbilled Metered Consumption</td><td rowspan="8">Non-Revenue Water</td></tr>
<tr><td>Unbilled Unmetered Consumption</td></tr>
<tr><td rowspan="6">Water Losses</td><td rowspan="3">Apparent Losses</td><td>Unauthorized Consumption</td></tr>
<tr><td>Customer Metering Inaccuracies</td></tr>
<tr><td>Systematic Data Handling Errors</td></tr>
<tr><td rowspan="3">Real Losses</td><td>Leakage on Transmission and Distribution Mains</td></tr>
<tr><td>Leakage and Overflows at Utility's Storage Tanks</td></tr>
<tr><td>Leakage on Service Connections up to point of Customer Metering</td></tr>
</table>

FIGURE 2.3: International Water Association (IWA) breakdown the total amount of water supplied to a WDS (IWA, 2000).

2.5.3 Statistical / Artificial Intelligence Based Methods

Statistical / Artificial Intelligence Based methods rely solely on the measurements from the WDS to detect leaks. These method employ various techniques from statistics and artificial intelligence algorithms like, artificial neural networks (ANN), support vector machines (SVM) etc. to the problem of leak detection. In this section we highlight some of the prominent AI based methods employed to the leak detection problem.

Mounce et al. (2002) proposed a ANN based method to learn and predict future values from a DMA sensor. The prediction and the actual observed values were then analysed by a classification module to indicate the presence of a leak. Later Mounce et al. (2007) introduced a fuzzy inference system in place of the classification module and further improved the detection by providing estimated average leak flow. However, their method used a 12 or 24 hour window for leak detection and reporting, furthermore this methods was not evaluated in active pressure optimised DMAs. To improve the 12 to 24 hour window Mashford et al. (2009) proposed a SVM methodology which was shown to generate faster detection results. In their analysis they generated leak data by using WDS simulation software EPANET. Pressure measurements from every node in the simulated network, both under leak and no leak conditions, were used to train the a binary SVM classifier to distinguish between leaky and non-leaky pipes to detect and locate leaks. Their results reports 76.8% detection accuracy with a false alarm rate of 23.2% in the

best case. However, the proposed SVM classifier required pressure measurements at every node in the network, furthermore the authors highlighted that practical applicability of the method is dependant on instrumentation and relative sensitivity of the pressure measurements. They concluded that pressure changes of 0.1 Pa or 0.00001m need to be registered in order to detect leaks of 1.5 litres per min.

An adaptive Kalman filter based methodology was presented by Ye and Fenner (2010) for analysing both the DMA flow and pressure data. In this approach weekly data for both pressure and flow from a DMA sampled at 15 minute intervals was modelled using one Kalman filter per time slot in a week i.e., 672 Kalman filters. Leak detection was based on the normalised residuals obtained from the multiple filter setup. The proposed method was tested on data from engineered leak events.

Palau et al. (2011) proposed a principle component analysis (PCA) based method where the hourly flow data from a DMA was analysed to detect leaks. However, this method considered flow reading at the night time only. In case a leak occurs during the day this method will not be able to detect it until the night time readings are observed. Thus, at best this methods can detect leaks the next day.

In light of the above, statistical / AI based methods offer several advantages over other leak detection methods. In particular, they do not require accurate knowledge of the WDS or specialised hardware and they rely solely on the empirical measurements. Furthermore, these methods require pressure and flow measurements to be sampled at a much lower frequency (e.g., 15 minutes) than model based methods (e.g., transient analysis). In light of this, statistical/AI based approaches provide an attractive way forward, as they not only provide a way to automate the data analysis process, but can also efficiently handle sparse and noisy sensor data. This sparsity is particularly important, given our research requirements of using only DMA-level hydraulic measurements.

Given this, in the next section we will detail the background fundaments and workings of selected state of the art leak detection methods, which we will use as a benchmark to test our proposed algorithms. We shall also highlight the reason for choosing these methods.

2.6 Selected Approaches for Comparative Analysis

To do an effective comparative analysis it is necessary to consider all of the following, the state of the art in GP based fault/leak detection, current applied leak detection approaches in the industry and the most recent advancements in leak detection applied and tested in a real life scenario. Our first choice is the nightline analysis (Muncke, 2011) algorithm, which is the defecto leak detection approach used in the water industry. The second chosen approach is a recent advancement by Ye and Fenner (2010) proposing a

multiple kalman filter based model for leak/burst detection in water distribution systems which was applied and tested in a real WDS setting with engineered leaks. In GP based fault/detection methods the fault bucket algorithm proposed by Osborne et al. (2012), which is an extension of the previous works in GP based change-point/fault detection (Garnett et al., 2010; Osborne et al., 2010), represents the current state of the art and is thus our method of choice. The following sections present a brief description of the selected approaches.

2.6.1 Mean Shift Nightline Algorithm

One of the most common approach used to detect anomalous flow events in the DMA is by analysing the nightline flow readings. The flow reading at certain times in the night represent the minimum legitimate consumption with minimum variations, therefore the nightline across days remains fairly consistent. Unexpected increase in the nightline flow can be attributed to leaks or bursts. However in making this attribution the duration of the unexpected increase must be taken into account, as short duration increase in flow can occur due to legitimate use e.g. fire fighting, or high demand during a festival. Muncke (2011) formulated this problem, from a mathematical standpoint, as that of detecting a mean shift in nightline discrete time series data. If x_t and m_t represents the fluctuations (noise) and the mean point of the nightline at time t then the nightline flow can be modelled by a discrete random variable F_t as:

$$F_t = x_t + m_t \tag{2.1}$$

The author makes the assumption that x_t is normally distributed. Using this assumption, an iterative algorithm is proposed which performs two one-tailed tests on every new value x_t, at time step t, to see if it belongs to the same distribution as x_1 to x_{t-1}. The mean point, m_t, is recalculated, calculated at each step based on the results of the one-tailed tests. If x_t belongs to the same distribution the calculated mean is kept unchanged otherwise a warning is raised and the mean is recalculated using readings starting from time step t. This introduces a calibrated mean that can adjust to leaks and bursts.

To test if x_t belongs to the same distribution predictive hypothesis tests are done. If there are n past errors x_1 to x_n then the null hypothesis is that x_{n+1} comes from the same distribution i.e. $H_0 : x_{n+1} \sim N(\mu, \sigma^2)$ and $H_1 : x_{n+1} \nsim N(\mu, \sigma^2)$. An alternative and more convenient formulation, based on the assumption that errors x are normally distributed, is: if the past n errors i.e. x_1 to x_n are normally distributed with a mean

$\bar{x}_n$ and standard deviation s_n, then the above hypothesis can be written as:

$$H_0 : \frac{x_{n+1} - \bar{x}_n}{s_n\sqrt{1+\frac{1}{n}}} \sim T^{n-1}$$

and

$$H_1 : \frac{x_{n+1} - \bar{x}_n}{s_n\sqrt{1+\frac{1}{n}}} \nsim T^{n-1}$$

Where T^{n-1} denotes a Student's t-distribution with $n-1$ degrees of freedom. Based on the above the null hypothesis is rejected at a chosen confidence α if:

$$x_{n+1} > \bar{x}_n + s_n\sqrt{1+\frac{1}{n}}.T_\alpha^{n-1}$$

or

$$x_{n+1} < \bar{x}_n - s_n\sqrt{1+\frac{1}{n}}.T_\alpha^{n-1}$$

The complete algorithm is detailed below:

Set $s = 1$ the starting time;
for $i = 1$ *to* n **do**
  if *($i <= 2$)* **then**
    $m_1 = F_1$ and $m_2 = \frac{F_1+F_2}{2}$
  else
    if *($x_i = F_i - m_i$ belongs to the same distribution as x_1 to x_{i-1})* **then**
      $m_{i+1} = mean(F_1 \text{ to } F_i)$;
    else
      Set $s = i$;
      Raise warning (yellow, orange or red);
      $m_{i+1} = mean(F_s \text{ to } F_i)$;
    end
  end
end

Algorithm 1: Nightline algorithm

The confidence interval α sets the trade off between the false positive rate and the false negative rate. Increasing α decreases the false positive rate but increases the false negative rate. To cater for legitimate short time increases in demand a threshold number of days of high flow can be set that distinguish between a burst and legitimate use. This threshold was set to three days with a three step warning system. The first true positive is flagged yellow, on the consecutive day the next true positive if flagged orange, if the anomaly continues for a third day it is flagged red and is categorised a burst or leak in the system. Furthermore the algorithm can adjust to seasonal variation by introducing a moving window of 90 days on the past data being used.

2.6.2 Burst Detection Using KF Residual (Neptune Project)

In research conducted under the EPSRC funded Neptune project Ye and Fenner (2010) propose a simple application of Kalman filtering for burst detection in water networks. The filter is used to estimate the normal hydraulic parameters (flow or pressure) in the distribution system. The residual from the estimates represents excessive variations from normal patterns and are considered bursts. If we consider X to be the hydraulic parameter being estimated then the state of X at time k is only dependant on its previous state at time $k-1$ and can be modelled by the equation:

$$X(k) = X(k-1) + Q \tag{2.2}$$

where Q is the the Gaussian process noise. The observed value of the parameter can be modelled as:

$$Z(k) = X(k) + R \tag{2.3}$$

where R is the the Gaussian observation noise, which models inaccuracies in measurements. The filter works in five steps which are as follows:

1. Predict the flow or pressure

$$X(k) = X(k-1) \tag{2.4}$$

2. Predict the flow or pressure covariance

$$p(k|k-1) = p(k-1|k-1) + q(k) \tag{2.5}$$

3. Calculate the Kalman gain

$$g(k) = \frac{p(k|k-1)}{p(k|k-1) + r(k)} \tag{2.6}$$

4. Update estimate

$$X(k|k) = X(k|k-1) + g(k)[Z(k) - X(k|k-1)] \tag{2.7}$$

5. Update covariance

$$p(k|k) = [1 - g(k)]p(k|k-1) \tag{2.8}$$

The parameters $q(k)$ and $r(k)$ are adaptively computed, in a process originally developed by Mehra (1970), using the filters innovation sequence (the difference between and

estimate and the measurement) as:

$$s(k) = Z(k) - X(k|k-1) \tag{2.9}$$

$$q(k) = g^2(k)c(k) \tag{2.10}$$

$$r(k) = c(k) + p(k|k-1) \tag{2.11}$$

Here $c(k)$ is the innovation covariance which is calculated using a moving window of size M as (Price, 2005):

$$c(k) = \frac{1}{M} \sum_{i=k-M+1}^{k} s^2(i) \tag{2.12}$$

The normalised residual of the filter given by equation (2.13) is used to indicate the occurrence of a burst. Positive values of the residual represent burst while negative values indicate that the actual demand is less than the predicted estimate.

$$R_f(k) = \frac{Z(k) - X(k)}{Z(k)} \tag{2.13}$$

The flows in a DMA follow a diurnal pattern with two peaks, in the morning and evening, where as the Kalman filter and the model presented in equation (2.2) assumes a linear stochastic system. Adjacent data points in the flow time series are therefore not linear. To overcome this limitation the author uses the fact that flow readings are more likely to be similar to the readings the previous week at the same time i.e., the consumption on Monday 8:15 A.M. would be more similar to Monday 8:15 previous week than at 8:00 AM on the same day. Incorporating this fact, there can be a separate Kalman filter for each time slot across a 1 week window. If readings are taken every 15 minutes then 672 Kalman filters would be required to cover the whole week. The residual of each filter at each time slot indicates whether there is a burst in the DMA at that particular time or not.

2.6.3 GP based Fault Bucket (FB) Algorithm

In this section we start by giving a brief introduction to GP's which is followed by a description of the GP based fault detection algorithm Fault Bucket (FB).

2.6.3.1 Introduction To Gaussian Processes (GP):

Formally a GP can be defined as a stochastic process defining a distribution over functions $H \rightarrow \mathbb{R}$ such that the distribution over any finite subset $F \subset H$ is a multivariate gaussian distribution. A GP can be completely defined by a mean function, $m(.)$ and a positive semi-definite covariance function, $k(.,.)$. Given the observed values, $y = y_1 \dots y_n$, of a function f at a set on inputs $x_1 \dots x_n$, the observed sample can be

thought as being drawn from a multivariate Gaussian distribution.

$$\begin{bmatrix} f(x_1) \\ \vdots \\ f(x_n) \end{bmatrix} \sim \mathcal{N}\left(\begin{bmatrix} m(x_1) \\ \vdots \\ m(x_n) \end{bmatrix}, \begin{bmatrix} k(x_1,x_1) & \dots & k(x_1,x_n) \\ \vdots & \ddots & \vdots \\ k(x_n,x_1) & \dots & k(x_n,x_n) \end{bmatrix} \right)$$

Assuming a GP prior on f the prior distribution can be given as:

$$p(f|x,\theta) = \mathcal{N}\left(m(x;\theta), k(x,x;\theta)\right)$$

where $k(x,x;\theta)$ is the positive semi-definite covariance matrix and θ is the set of hyper-parameters that characterise the mean and covariance functions. The covariance function is often defined by one or more kernel functions. A good example which has become the de-facto default kernel for GPs and SVMs is the squared exponential kernel (SE) which is given as:

$$KSE(x,x') = \sigma^2 \exp -\frac{(x-x')^2}{2l^2}) \tag{2.14}$$

Here the hyper-parameters, θ, for kSE are the length scale, l, that controls the horizontal length scale over which the function varies, and output variance, σ^2, controls the vertical variation. Often the observed data, y, is a noise corrupted version of the true underlying function f which is not known. In such cases, if the noise is assumed to be zero mean Gaussian i.i.d, $\epsilon \sim \mathcal{N}(0,\sigma^2)$, then the prior distribution can be written as:

$$p(y|x,\theta) = \mathcal{N}\left(m(x;\theta), k(x,x;\theta) + \sigma^2 I\right) \tag{2.15}$$

It can be shown (Rasmussen and Williams, 2006) that the posterior predictive density given a set of test points x_* can be given as:

$$\begin{bmatrix} y \\ y_* \end{bmatrix} \sim \mathcal{N}\left(\begin{bmatrix} m(x) \\ m(x_*) \end{bmatrix}, \begin{bmatrix} K_y + \sigma^2 I & K(x_*,x) \\ K(x_*,x)^T & K_(x_*,x_*) \end{bmatrix} \right)$$

$$p(y_*|x_*,x,y,\theta) = \mathcal{N}\left(y_*; m(x_*|y,x,\theta), \Sigma(x_*|y,x,\theta)\right) \tag{2.16}$$

$$m(x_*|y,x,\theta) = m(x_*) + K(x_*,x)K_y^{-1}(y - m(x)) \tag{2.17}$$

$$\Sigma(x_*|y,x,\theta) = K(x_*,x_*) - K(x_*,x)^T K_y^{-1} K(x_*,x) \tag{2.18}$$

$$K_y \triangleq K(x,x) + \sigma^2 I \tag{2.19}$$

The hyper-parameters, θ, can be optimised by maximising the log marginal likelihood (Murphy, 2012) since:

$$\begin{aligned} p(y|x) &= \int p(y|f,x)p(f|x)df \\ p(f|x) &= \mathcal{N}(f|m,K) \\ p(y|f) &= \prod_i \mathcal{N}\left(y_i|f_i,\sigma^2\right) \end{aligned}$$

the log marginal likelihood is given as:

$$\begin{aligned} \log p(y|x) &= \log \mathcal{N}(y|m,K)) \\ &= \underbrace{-\frac{1}{2}(y-m)K^{-1}(y-m)}_{\text{Data fit term}} \quad \overbrace{-\frac{1}{2}\log|K|}^{\text{Model complexity term}} \quad \underbrace{-\frac{N}{2}\log(2\pi)}_{\text{Constant}} \end{aligned} \tag{2.20}$$

The first term is a data fit term, the second term is a model complexity term, and the third term is just a constant. Considering a SE kernel in $1D$, if the length scale, l, is varied while keeping σ^2 fixed. For short length scales, the data fit term will be small producing a good fit but the model complexity will be high as K will be almost diagonal as most points will not be considered near any others. Similarly, for long length scales, the fit will be poor but the model complexity will be low as K will have non diagonal elements. The partial derivative of the this marginal likelihood equation with respect to each hyper-parameter θ_i in θ can be derived as:

$$\begin{aligned} \frac{\partial}{\partial\theta_i}\log p(y|x) &= \frac{1}{2}(y-m)^T K^{-1}\frac{\partial K}{\partial\theta_i}K^{-1}(y-m) - \frac{1}{2}tr\left(K^{-1}\frac{\partial K}{\partial\theta_i}\right) \\ &= \frac{1}{2}tr\left((\alpha\alpha^T - K^{-1})\frac{\partial K}{\partial\theta_i}\right) \end{aligned} \tag{2.21}$$

where $\alpha = K^{-1}y$. Here $\partial K/\partial\theta_i$ depends on the kernel being used and the parameters with respect to which the derivative is being taken. With the expression for the marginal likelihood (2.20) and its derivative (2.21) any standard gradient-based optimiser can be used to estimate the kernel parameters. However gradient-based optimisers will not always yield the global optimal. In such cases alternative sampling approaches, such as Monte Carlo or quasi-Monti Carlo, can be used.

2.6.3.2 Fault Bucket (FB) Algorithm

The FB algorithm is based on the expectation that points that are more likely to be generated by noise with wide variance, than under the normal predictive model of a GP, are likely to be faulty. The data available at any point in time is partitioned into old and

new halves. The approach uses four keys assumptions which are listed below (Osborne et al., 2012):

1. **Fault Bucket:** *Faulty observations are assumed to be generated from a Gaussian noise distribution with a very wide variance.*
2. **Single Gaussian marginal:** *A mixture of Gaussians, weighted by the posterior probabilities of faultiness of old data, is approximated as a single moment-matched Gaussian.*
3. **Old/new noise independence:** *The noise contributions are assumed to be independent, the contributions for new data are independent of old observations.*
4. **Affine precision:** *The precision matrix over both old and new halves is assumed to be affine in the precision matrix over the old half.*

The model is formalised by choosing an observation noise distribution which is independent but not i.i.d. with separate variance for non-fault and fault cases as:

$$p(y|f,t,\rceil\text{fault},(\sigma^n)^2) = \mathcal{N}\left(y; f, (\sigma^n)^2\right)$$
$$p(y|f,t,\text{fault},(\sigma^f)^2) = \mathcal{N}\left(y; f, (\sigma^f)^2\right)$$

where fault $\epsilon\{0,1\}$ is an indication of fault presence or absence in observation $y(t)$ and $\sigma^f > \sigma^n$ is the standard deviation around the mean of the fault. Both σ^f and σ^n form part of the hyper-parameters of the model.

Since, a priori, it cannot be know if the observations is faulty or not, assumption 2 is used to approximately marginalise the faultiness of old observations as a single Gaussian. Formally, given available data, D, of all observations y, then for each y if the possible values of the noise variances, σ, are indexed by i, where each σ_i represents a different combination of faultiness over D ($\sigma_t^0 = \sigma^n$ for not fault and $\sigma_t^1 = \sigma^f$ for fault), the predictive distribution for the latent function, f, is found by marginalising over predictions indexed by i i.e collapsing the weighted sum of all predictions to a single Gaussian prediction.

$$\begin{aligned} p(f*|y) &= \sum_i p(\sigma^i|y)p(f*|y,\sigma^i) \\ &= \sum_i p(\sigma^i|y)\mathcal{N}\left(f*; m(f*|y,\sigma^i), K(f*|y,\sigma^i)\right) \end{aligned} \tag{2.22}$$

The observed data is partitioned in to two portions, old observations $D_a = (x_a, y_a)$ and new observations $D_b = (x_b, y_b)$ where the possible values of σ_a are indexed by i while those for σ_b are indexed by j. The covariance matrices for the old (V_a^i), new (V_b^j) and

complete data ($V^{i,j}_{a,b}$) are given as:

$$V^i_a = K_{a,a} + \text{diag}\sigma^i_a$$
$$V^i_b = K_{b,b} + \text{diag}\sigma^j_b$$
$$V^{i,j}_{a,b} = K_{\{a,b\},\{a,b\}} + \text{diag}\{\sigma^i_a, \sigma^j_b\}$$

Simple calculations and rearrangement shows that the approximation of equation 2.22 as a single Gaussian, requires the expected values of V^{-1} and $V^{-1}yy^TV^{-1}$ (the expectations with respect to $p(\sigma^i|y)$). Given D_a these matrices are denoted by M_a and referred to as the marginal set. The values of M_a are stored and simple updates are made to it to arrive at matrix $M_{a,b}$ which is the marginal set given $D_{a,b}$. The update is done based on assumption 3 i.e. it is assumed that the faults do not persist longer than D_b. Mathematically this can be stated as:

$$p(\sigma^{i,j}_{a,b}, y_{a,b}) \simeq p(\sigma^i_a)p(y_a|\sigma^i_a)p(\sigma^j_b)p(y_b)p(y_b|\sigma^{i,j}_{a,b}, y_a)$$

Assumption 3 and 4 ($V^{i,j}_{a,b}{}^{-1}$ is affine in $V^{i}_{a}{}^{-1}$) allow efficient updates to the marginal set M. When new data, D_c, is received the existing data is treated as the old set and the iteration is performed again. At each iteration the algorithm gives the predictions $p(f*|y_{a,b})$ and the posterior probability of the new observations faultiness $p(\sigma_b|y_{a,b})$. As analytical marginalization of the hyperparamters is not possible Bayesian Monte Carlo (Rasmussen and Williams, 2006; Garnett et al., 2010) is used to approximate them.

Having given a detailed description of the state of the art leak detection approaches used as a benchmark in this thesis, we shall now briefly describe the methodology used to acquire WDS and leak data for analysing the performance of the algorithms proposed in this thesis.

2.7 Data Acquisition Methodology

Currently most companies rely on either, visible burst reports or customer complaints to record the occurrence of a leak. Even in these cases, the exact accurate time a leak started in never known or found. Also, there may be multiple leak complaints, reported at different dates, at a given time in different parts of a DMA, in such scenarios it becomes even harder to determine when a particular leak started. Similarly, when a leak is fixed only the repair date is recorded, which in cases of multiple repairs on the same date makes the task of finding the exact leak start, end times even harder. Even in cases when leak reporting and repair dates are available, it is often impossible to determine the exact leak times and magnitudes from such data. This is mainly because this data is not only susceptible to human errors but also ongoing network optimisation, seasonal, weather changes and DMA level adjustments e.g., repairs, modifications, firefighting,

addition of new properties etc. Often such changes go un-recorded, thus making it impossible to separate out leak effects from the effects caused by such changes.

Given the above challenges and the difficulties in acquiring verifiable accurate real leak data, a large number of leak detection methods in literature (Mounce et al., 2002, 2007; Mashford et al., 2009; Palau et al., 2011) proceed by testing the algorithms on simulated data generated from WDS simulation tools such as EPANET. *EPANET is a public domain, water distribution system modelling software package developed by the United States Environmental Protection Agency's (EPA) Water Supply and Water Resources Division. It performs extended-period simulation of hydraulic and water-quality behaviour within pressurised pipe networks.*

However, one limitation of EPANET is that it can only simulate demand-driven WDS models i.e., models where the outflows in the network are assumed to be constant irrespective of the available pressure. Such models do not consider the impact of pressure changes on the flow supplied to the network. However, given the objectives of this research, we are particularly interested in pressure optimised DMAs i.e., DMA's where the supply point pressure is dynamically changes so that the critical point pressure is as close to a set serviceability threshold as possible. Therefore, to implement pressure optimisation it is crucial to have a hydraulic model that takes in to account the effects of pressure changes on the link and nodal flows. Given the limitation of EPANET, in this thesis we use a enhanced version of EPANET, namely WaterNetGen that, among other enhancements, adds the capability to run pressure-driven analysis to EPANET (for details see Muranho et al. 2012).

Given the above, it must be noted that although EPANET generated WDS models have been extensively used in literature to generate data to verify leak detection algorithms. Due to the limitations of EPANET, mentioned above, such datasets or models are not suitable for analysis of pressure optimised DMAs, and thus, cannot be directly used to in context of the research problem addressed in this thesis.

2.8 Summary

In this chapter we began by giving a brief description about the structure and working of a typical water distribution and listing its major components. We then gave an overview of the general are of leak management. We highlighted the proactive and reactive leak management approaches, emphasising the need and importance of proactive leak management. Following this, we detailed the general categories of leak management approaches. We then reviewed various hardware based leak detection methods highlighting their advantages and disadvantages. Following this we described the two main type of hydraulic model based techniques. We highlighted the requirement of precise network parameters such as, friction factors, pipe roughness etc, as their main

drawback. Following this we detailed various hydraulic measurement based methods, their advantages and disadvantages. We reviewed various statical and AI based methods where we pointed out the shortcoming of each of the approach.

Chapter 3

Leak Detection and Step Change Magnitude Approximation using Noise Scaled Semi Parametric Gaussian Process (NSGP)

In Chapter 2, we described the existing works related to leak detection and quantification. We also discussed how these works fall short in addressing the research requirements and objectives detailed in Sections 1.3 and 1.4. To address these shortcomings, this chapter proposes a Gaussian process (GP) model with an additive diagonal noise covariance that is able to handle the input dependent noise observed in this setting. A parameterised step change mean function is used to detect leaks and to estimate their size. Using prior water distribution systems (WDS) knowledge we dynamically bound and discretise the detection parameters of the step change mean function, reducing and pruning the parameter search space considerably. We evaluate the proposed noise scaled GP (NSGP) against both the latest research work on GP based fault detection methods and the current state of the art and applied leak detection approaches in water distribution systems. We show that our proposed method outperforms other approaches, on real water network data with synthetically generated time varying leaks.

The rest of this chapter is organised as follows. In Section 3.1, we start by describing how we model the DMA-level flow using a GP to construct a NOM. We also detail the diagonal noise covariance function used to model input dependent noise. This is followed by a description of the step change mean function used to detect and approximate leaks (Section 3.2). Next in Section 3.3, we describe the dual optimisation mechanism used to learn the NOM and leak hyperparameters. Following this, in Section 3.4 we give a brief description of the DMA flow data setup and the leak simulation model used for the comparative analysis. This is followed by a brief description of the performance metrics

used in the evaluation. We present the results of the comparative analysis in Section 3.5. We end by giving a summary of the chapter and the conclusions drawn from it, in Section 3.6.

3.1 DMA Flow Modelling Using A GP

Water consumption in a DMA is periodic, at seasonal, monthly and weekly scales and is mostly measured at fixed discrete intervals, usually 15 mins. On a weekly scale, the consumption at a particular time in a day, is very similar to the historic readings at the same time in the previous weeks. To elaborate this, Figure 3.1 shows the DMA-level flow readings, from a UK DMA, over a period of four weeks. Here the mean for each discrete time slot is calculated by taking an average of the historic readings across four weeks at the same time slot.

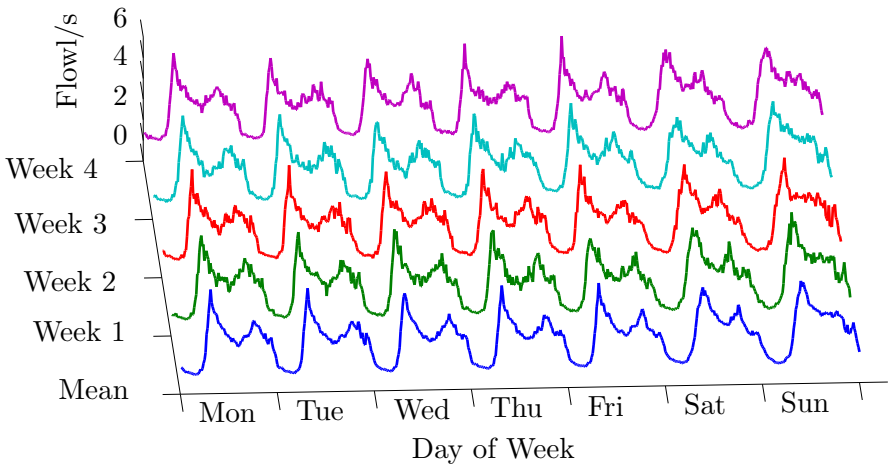

FIGURE 3.1: Four week flow data showing similarity in weekly consumption patterns and the calculated weekly mean. Flow measured in litres per second (l/s).

Given this, the consumption at a particular time in a day can be modelled as a function of the historic readings at the same time in the previous weeks. Similarly the variations in the readings at a particular time in the week can also be modelled as a function of the historic variations. To incorporate this in the GP framework, we firstly define a weekly moving window on the data being modelled. Secondly, we define two vectors, the weekly mean (m_w) and weekly variance (σ_w^2). Assuming sensor data is discrete (with readings taken at fixed time intervals), the weekly means and variances vectors consist of a mean and variance value for each reading time slot, t, in a week i.e. $m_w(t)$ and $\sigma_w^2(t)$.

Figure 3.2 shows the computation of the mean values, $m_w(t)$, for one day of a week (Monday). Here each mean and variance value is computed from the historic flow data,

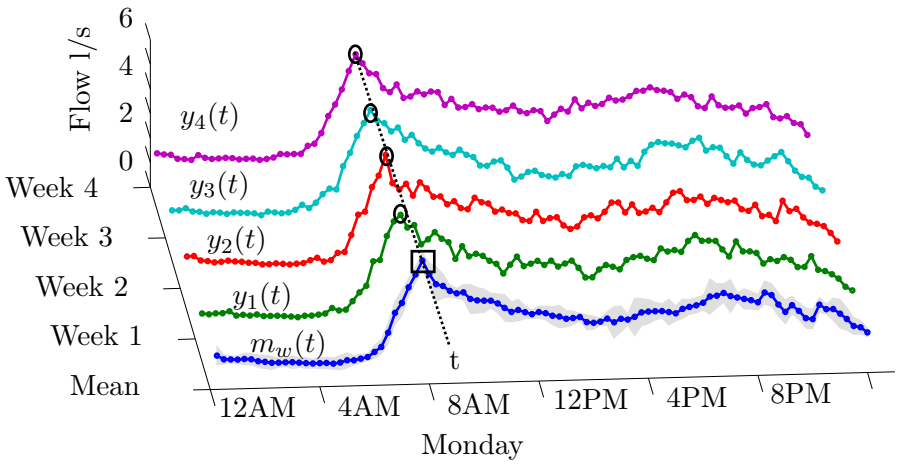

FIGURE 3.2: Showing consumption patterns on Monday across four weeks and the computation of weekly mean and variance vectors.

represented as $y(t)$, using a moving average window of four weeks as:

$$m_w(t) = \frac{1}{4}\sum_{i=1}^{4} y_i(t)$$

$$\sigma_w^2(t) = \frac{1}{4}\sum_{i=1}^{4}(y_i(t) - m_w(t))^2$$

The historic weekly mean can be easily modelled as a fixed deterministic mean function in a GP (Rasmussen and Williams, 2006). On the other hand, heteroscedastic noise in GP models results in a non-closed form likelihood expression. This is due to the fact that the noise function, σ^2, is no longer a constant value but represents input-specific noise rates that are typically unknown and must be calculated by some mechanism. Various approaches have been proposed to incorporate heteroscedastic noise in GPs. This is mostly done by modelling the latent function[1] in one GP, while the input dependant noise is modelled as a separate additive GP. However, unlike a single GP, this methodology results in analytically non-tractable inference. To address this, multiple approximate inference approaches, for heteroscedastic GPs, have been proposed. In particular these are, Markov Chain Monte Carlo (MCMC) (Goldberg et al., 1997), Expectation Maximisation, EM-like, procedure (Kersting et al., 2007) and variational Bayes approximation (Lazaro-Gredilla and Titsias, 2011). However, these approximations are computationally expensive. Furthermore, it has been shown that variational based approaches in GP result in under estimation of the posterior variance (Kuss and Rasmussen, 2005), which in the case of leak detection would result in false positives (as underestimation of the variance, particularly at times with high variability in readings, can result in outliers

[1]In statistics, latent variables as opposed to observable variables, are variables that are not directly observed but are rather inferred (through a mathematical model) from other variables that are observed (directly measured). In GPs the observations are assumed to be a noise corrupted version of a hidden function i.e. the latent function.

that would otherwise be considered as normal behaviour). Given this, in case of our problem, two observations greatly simplify the inference which are as follows:

1. Firstly, in our model the noise variance at each time slot are known, as we compute the noise variance vector from the historical data (with the assumption that the noise variance at a particular time slot is a function of the historical variances at the same time slot).

2. Secondly, in our settings, it is reasonable to assume that the noise rates at each time slot are independent of the noise rates at other time slots. For example, there is minimum variance in the readings at 3:00 AM in the morning as there is minimal water usage throughout the week at this time, where as on weekdays, at 8:00 AM the water usage varies a lot as most people are preparing to get ready for work. Given this, even for 8:00 AM (i.e. 15 mins after the 7:45 AM reading) it is reasonable to assume (for computational efficiency) that the variations at the previous time slot (7:45 AM) are unrelated to the next time slot (8:00 AM) or any other time slot in the readings.

Assuming the aforementioned, a tractable exception can be derived with the assumption of independence between noise variances (Venanzi et al., 2013). Specially, since the weekly observed variance in our model is calculated over a moving window of one month, the heteroscedasticity is inherently captured outside the GP framework. Given this, if we represent the observed DMA-level flow readings at time steps, t, as y and the computed historic weekly mean as, $m_w(t)$, then the computed variance as, $\sigma_w^2(t)$ is incorporated in our GP model as a diagonal noise matrix $\Sigma = \sigma_w^2(t)I$. Using this, we can define NOM for a DMA's flow patterns using a GP as:

$$p(y|t, \theta_{cov}) = \mathcal{N}\left(f; m_w(t), k(t, t; \theta_{cov}) + \Sigma)\right) \tag{3.1}$$

Here, $k(t, t; \theta_{cov})$ is the function that models the covariance. In our GP model we use the standard squared exponential function to model the covariance between time slots, e.g. t_i, t_j as:

$$k_{SE}(t_i, t_j) = \sigma^2 \exp(\frac{-||t_i - t_j||}{2l^2}) \tag{3.2}$$

Thus in our case, $k(t, t; \theta_{cov}) = k_{SE}(t, t; l, \sigma^2)$, where the the length scale, l, and the output variance, σ^2, form the hyper parameters, θ_{cov}, of the covariance function. The length scale determines the smoothness of the function i.e., small length scale values represent a function that can change quickly while large values represent functions that change slowly. Similarly, the the output variance parameter determines the variation of function values from its mean. Given this, for a set of new flow readings, y_*, at time steps, t_*, the posterior predictive distribution, given in equation (2.16), can be re-written

as:

$$p(y_*|t_*, t, y, \theta_{cov}) = \mathcal{N}\left(y_*; m_w(t_*|y, t), \Sigma(t_*|y, t, \theta_{cov})\right) \tag{3.3}$$

where :

$$m_w(t_*|y, t) = m_w(t) + K(t_*, t)K_y^{-1}(y - m_w(t)) \tag{3.4}$$
$$\Sigma(t_*|y, t, \theta_{cov}) = K(t_*, t_*) - K(t_*, t)^T K_y^{-1} K(t_*, t) \tag{3.5}$$
$$K_y \triangleq K(t, t) + \Sigma \tag{3.6}$$
$$\Sigma = \sigma_w^2(t)I \tag{3.7}$$

Using this, the posterior distribution can be derived as the combination of the GP kernel and the diagonal noise matrix. Also, given the above we can obtain a closed form equation of the marginal likelihood, just as in case of a normal GP (Rasmussen and Williams, 2006). Using the marginal likelihood equation (2.20) of a standard GP and substituting the values of m, K and σ^2 to m_w, $K(t, t)$ and σ_w^2 we obtain:

$$\log p(y|t) = -\frac{1}{2}(y - m_w)^T[K(t, t) + \sigma_w^2 I]^{-1}(y - m_w) - \frac{1}{2}|K(t, t) + \sigma_w^2 I| - \frac{n}{2}log(2\pi) \tag{3.8}$$

The diagonal noise matrix inherently adds a bias to the diagonal of the covariance structure that reflects the actual variations in data and thus providing better handling of the input dependent noise in the underlying system. Having defined the NOM of DMA-level flow readings using a GP, in the next section we define our methodology for leak detection and quantification.

3.2 Leak Detection and Quantification by Parameterising the GP

We now return to the problem of leak detection and quantification. Leaks are non-deterministic and cause an increase in the measured flow thus changing the characteristic of the underlying system for the duration of the leak. This deviation from normal behaviour can be used to detect leaks in the system. Recent works on GP based fault/change-point detection (Garnett et al., 2010) propose various kernels to detect different types of faults. Out of these, of particular relevance to leak detection problem is the bias fault kernel (Garnett et al., 2010). We use the same approach, however, instead of modelling a bias kernel in the GP covariance structure, we propose a step change mean function, m_l, parameterized by $\theta_l = \{\theta_{time}; \theta_{mag}\}$. Here, the set of hyper parameters represented by θ_l, capture the properties of the leak we are interested in. Particularly, θ_{mag} captures the magnitude of the leak and θ_{time} represents the starting

time slot of the leak, in the one week data modelling window. The step change mean is then represented as:

$$m_l(t;\theta_l) = \begin{cases} \text{zero} & \text{if } t < \theta_{time} \\ \theta_{mag} & \text{if } t \geq \theta_{time} \end{cases}$$

We incorporate this in our GP model by defining the GP prior mean as a composite mean function, m_c, which is a sum of the historic weekly mean, m_w, and the step change mean function, m_l, resulting in a semi-parametric GP (Murphy, 2012). Given this, the posterior predictive distribution in equation (3.3) can be rewritten as:

$$p(y_*|t_*,t,y,\theta_{cov}) = \mathcal{N}\left(y_*; m_c(t_*|y,t), \Sigma(t_*|y,t,\theta_{cov})\right) \tag{3.9}$$

where :

$$\begin{aligned} m_c(t_*|y,t) &= m_c(t) + K(t_*,t)K_y^{-1}(y - m_c(t)) \\ m_c(t) &= m_w(t) + m_l(t;\theta_l) \end{aligned} \tag{3.10}$$

Here, by modelling the leak as a parametric mean function we have explicitly made the leak independent of the correlations in NOM. This leads to a better modelling of the actual physical leak process, as leaks are un-correlated additions to the underlying latent process. Furthermore, when learning the leak parameters, this results in the marginal likelihood giving an estimate of the leak based on deviation from the NOM. Thus, any variations in the observation noise due to the leak are also captured in the leak parameters, giving a more accurate estimate of the leak effect without altering the observation noise model.

Having defined the leak detection and quantification methodology, in the next section we define the dual hyper parameter optimisation scheme used to learn the NOM and leak hyper parameters.

3.3 Hyper parameter Learning

In this section we describe how the two sets of hyper parameters i.e. the covariance hyper parameters, θ_{cov}, and the mean hyper parameters, θ_l, are learned and how the two learning schemes are interwoven. In our model, θ_{cov}, represents the correlations in the DMA-level flow/demand when there is no leak in the system. As we use a squared exponential kernel, $k_{SE}(t_i,t_j) = \sigma^2 \exp(\frac{-||t_i - t_j||}{2l^2})$, to model these correlations (for a discussion of kernel types see Rasmussen and Williams, 2006) the NOM hyper parameters in, θ_{cov}, are the length scale, l, and the output variance, σ^2_{out}. We find the optimal values for these by setting $m_l(t;\theta_l)$ to zero and then using gradient descent search

to train the NSGP on four weeks of DMA-level flow data (Rasmussen and Williams, 2006).

Once the GP is trained, for new observed flow values, detection is done over a moving window of one week by learning the step change mean function hyper parameters, $\theta_l = \{\theta_{time}; \theta_{mag}\}$, using a bounded search process. At each time step, t, both θ_{mag} and θ_{time} are bounded based on the observed flow. The leak starting time parameter, θ_{time}, is bounded by our data modelling window of one week. For a new observed flow value, $y(t)$, at time, t, the leak magnitude will always be between $m_w(t)$ and $y(t)$. The leak at a particular time can not be more than the observed flow reading, $y(t)$, and less than the historic weekly mean, $m_w(t)$. An observed flow value less than the historic mean suggests a reduction in water flow whereas a leak always produces an increase. Thus, at a particular time slot, t, if $y(t) - m_w(t) < 0$, then search for θ_{time} at that particular time slot, t, can be pruned. These bounds reduce the leak parameter search space considerably. In case $y(t) - m_w(t) > 0$ or a set min threshold value, we discretise the possible θ_{mag} values by sampling five consecutive equidistant values between the upper and lower bounds, $y(t)$ and $m_w(t)$. Exhaustive search is then used to find both a single θ_{mag} (out of the discretised possibilities) and a θ_{time} value (out of all possible time slots in a week) that yields the lowest negative log likelihood.

Having defined the leak detection process we now refer back to the previously mentioned requirement of distinguishing between legitimate short term increase in flow and leaks. Depending on the specific reliability requirements of a utility company, a duration, T, of continuous increased flow can be defined for an anomaly to be considered a leak. In such cases θ_{time} will always be between the current time slot t and $t - T$. This not only allows continuous monitoring and record keeping of the leak but also leak correction, in the observed flow data, based on any previously confirmed leaks.

Having described the proposed NSGP algorithm, in the next section we shall detail the dataset and leak model used in the performance evaluation of all benchmark methods previosly described in Section 2.6 and performance metrics we consider for the analysis.

3.4 Data Setup and Leak Simulation Model

To evaluate our model we use five weeks of real flow data from a UK DMA in Arnesby Village, near Leicestershire, with readings taken at 15 minutes interval. Four weeks data is used to calculate the weekly mean and variance vectors. We add synthetic time varying leaks at different time slots in the fifth week. We introduce time varying leaks using the following formula.

$$L_s^n = B + \alpha y_s^n \tag{3.11}$$

Where L represents the leak, s is the starting time slot of the leak in the one week window, n is the total number of time slots in a week (672 time slots for readings taken at 15 mins interval), B is a constant base magnitude of the leak, α is a scaling factor (explained later) and y is the observed flow. This allows us to simulate leaks that are dependant on the observed flow. Ideally, leaks should be modelled in a pressure dependant way. Most leak simulations in literature (which normally employ specialised WDS modelling software such as EPANET) leaks are normally modelled as emitters, which model leaks as the flow through a nozzle or orifice using the equation, $q = Cp^{\gamma}$. Here q = emitter flow rate, p = pressure at the emission point, C = discharge coefficient, and γ = pressure exponent. The discharge coefficient in the equation represents the rate at which water is lost through an aperture per unit change in the pressure (units l/s per meter head of pressure). The pressure exponent models the effect of different aperture shapes on leak e.g., a large horizontal crack in pipe or a small hole (a value of 0.5 is mostly used for circular apertures). However, modelling a leak in this manner requires the pressure, flow, pipe, head loss e.t.c. characteristics at, at least, the location of the leak in a DMA. Although this might be easy to do in WDS modelling software, it is not possible in our case, when using only real DMA-level flow data. Particularly, the emitter methodology models a leak based on the pressure and other characteristics of the pipe the leak occurs in (Mashford et al., 2009; Muranho et al., 2012). Where as, given DMA-level flow measurements, we wish to introduce an increase in the flow that mimics the effect of a leak at an arbitrary, unknown location in a DMA. Despite this, we can still model the pressure dependance of a leak, at least to a certain extent, given only the DMA-level flow. Since in a pressure optimised DMA as the flow (water demand) increases the inlet pressure will increase too so as to keep the critical point pressure serviceable. Thus, in such case the pressure patterns would increase and decrease in conformance with the flow patterns, in turn any leak, dependant on pressure, would also follow these patterns. Keeping this in mind, in equation (3.11) we define a base leak magnitude, B, which defines the minimum magnitude of the simulated leak, where as α defines a scaling factor, which when multiplied with the observed flow y, attempts to simulate the effect of the changes in flow and pressure by mimicking the flow patterns. The larger a leak, the greater effect it will have on the pressure and flow within a DMA, thus keeping this in mind, in our simulations, we set the α parameter to $B/8$. After running simulations This allows the resultant simulated leak to mimic flow patterns based on how big or small the base leak is e.g., a small leak with a base magnitude of 0.1 l/s will have a scaling factor of 0.0125, showing minimal effect on the final simulated leaky flow, where as a large leak with a base magnitude of 1.5 l/s will have a larger scaling factor of 0.1875 and thus a greater effect on the final flow.

Figure 3.3 shows the one week test data used in the performance evaluation, both under normal operating conditions and with a simulated leak using base magnitude of 0.5 l/s.

In literature, the performance of leak detection solutions is normally evaluated using the

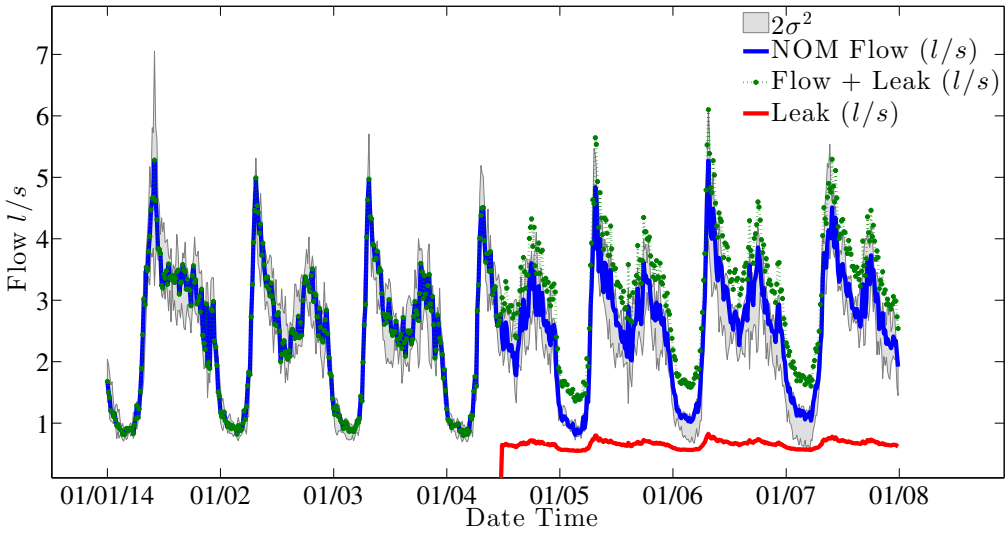

FIGURE 3.3: Showing: One week of NOM flow data (blue) with the uncertainty (shown in grey) computed using four weeks of historic flow readings. The same one week flow with an added simulated leak of base magnitude 0.5 l/s (green). The simulated leak is shown in red

statistical accuracy metric, which is computed using both the correctly and incorrectly identified leaks out of the total observed or simulated leaks (Mounce et al., 2002; Gertler et al., 2010; Ye and Fenner, 2010). In cases where simulated data from a WDS modelling software e.g. EPANET is used, leaks of varying magnitude are simulated at different locations within the network and then used in performance evaluation (Mashford et al., 2009). We employ the same methodology to evaluate the performance of the proposed NSGP in comparison with the benchmark approaches detailed in section 2.6. In particular, we test the detection accuracy of all approaches on 100 different leaks. We select 10 different time slots in the week between the fourth and seventh day. At each selected time slot we simulate 10 varying magnitude leaks (base magnitude of 0.15 l/s to 1.5 l/s). Doing this allows us to see how well each approach does for very small to very large leaks. Also, having leaks at different time slots e.g., fourth day of week at 12:00 PM and seventh day of the week at 21:00 PM, allows us to measure the accuracy of each algorithm when there is varying levels of leaky readings (evidence) available.

Additionally, in our experiments we consider the following performance metrics for evaluation:

1. **True Positive Rate (TPR):** The TPR defines how many correct positive results occur among all positive samples available during the test.

$$TPR = \frac{TP}{TP + FN} \tag{3.12}$$

where:

(a) TP The number of correctly identified time slots with a leak.

(b) FN The number of incorrectly identified time slots with a leak.

2. **False Positive Rate (FPR):** FPR, defines how many incorrect positive results occur among all negative samples available during the test.

$$FPR = \frac{FP}{FP + TN} \tag{3.13}$$

where:

(a) FP The number of time slots without a leak incorrectly identified as leaky.

(b) TN The number of time slots without a leak correctly identified as normal i.e. without a leak.

3. **Accuracy (ACC):** Accuracy is the proportion of true results (both true positives and true negatives) in the population. It is the degree of closeness of the number of detected leaks to the actual number of leaks in the data i.e.

$$ACC = \frac{TP + TN}{TP + FP + TN + FN} \tag{3.14}$$

4. **F1 Score (F1):** The F1 score, measures accuracy using the statistics precision P and recall R. Precision is the ratio of true positives (TP) to all predicted positives (TP + FP). Recall is the ratio of true positives to all actual positives (TP + FN). The F1 score is given by:

$$F1 = 2\frac{P \times R}{P + R} \tag{3.15}$$

where:

$$P = \frac{TP}{TP + FP}$$
$$R = \frac{TP}{TP + FN}$$

5. **Root Mean Square Error (RMSE):** To test the accuracy of the detected magnitude of a leak, we compute the root mean square error using the detected leak, $\hat{L_d}$, and the actual simulated leak, L_s as:

$$RMSE = \sqrt{\frac{1}{N}\Sigma_i^N(\hat{L_d} - L_s)^2}$$

In the comparative analysis, for all GP based methods we use the standard squared exponential kernel with the same learned hyperparmeters (length scale, $l = 125$, and output variance, $\sigma^2 = 0.2$). The only difference being FB algorithm use a constant gaussian noise model where as in NSGP we use the diagonal noise covariance kernel in conjunction with our mean functions.

3.5 Results of Comparative Analysis

Given this, our main objectives are to analyse that performance of the proposed NSGP, in comparison with other approaches. In doing so we wish to see how well each approach performs on leaks of different magnitudes. Furthermore, how well does each method perform as more and more leaky data becomes available i.e., as the available evidence of a leak increases. Thus to achieve this, in addition the aggregate overall performance results of all approaches on all 100 simulated leaks, we consider two special cases in our experiments. The first selected time slot e.g., fourth day of week 12:00 PM, as a case when maximum evidence of a leak is available i.e. three and a half days of leaky readings. Similarly, the seventh day of the week at 21:00 PM, when minimum evidence of a leak is available i.e., only two hours of leaky readings. We shall present and discuss the results for each of these cases separately. All algorithms were executed on a Mac book Pro 2.5 Ghz computer with 16 GB random access memory and the execution time, in seconds, for each algorithm is shown in the results table for each case.

Table 3.1 details the averaged detection results over the 10 leaks (base magnitude of 0.15 l/s to 1.5 l/s) simulated at the fourth day at 12:00 PM. It can be seen from the results that the NSGP performs better than other approaches with the highest TPR, 0.993, and lowest FPR, 0.000. In leak magnitude estimation the proposed model has the lowest root mean square error (RMSE) of 0.120 l/s.

	Reporting Time	**TPR**	**FPR**	**ACC**	**F1**	**RMSE**	**Execution Time (secs)**
		Averages over 10 different leaks					
NL	2 days	0.967	**0.000**	**0.999**	0.980	-	**0.245**
NKF	15 mins	0.773	0.015	0.879	0.834	0.652	0.381
FB	15 mins	0.291	0.003	0.643	0.406	0.611	94.878
NSGP	15 mins	**0.993**	**0.000**	0.996	**0.996**	**0.120**	31.534

TABLE 3.1: Showing averaged detection results over the 10 leaks with base magnitudes of 0.15 to 1.5 l/s simulated on 12:00 PM on the fourth day.

Although it may seem that the NL algorithm performs as good as the proposed NSGP it must be noted that the NL algorithm addresses a much simpler problem, as it uses one minimum hourly average flow reading per day i.e. the night line flow (which is the reason why accurate leak magnitude estimation particularly for time varying leaks is not possible for the NL algorithm in the current experiments and RMSE estimation for it is not shown in table). Thus, even though the NL algorithm shows high performance on the night line flow, it always takes at least two days for the NL algorithm to confirm and report a leak. Owing to this the NL algorithm is not included in Figure 3.4, which shows the leak detection results of the selected algorithms over the smallest, 0.15 l/s and the largest, 1.5 l/s, simulated leak at 12:00 PM on the fourth day. Given this, Figure 3.4 highlights both NKF and FB leak predictions are irregular, with false negatives within

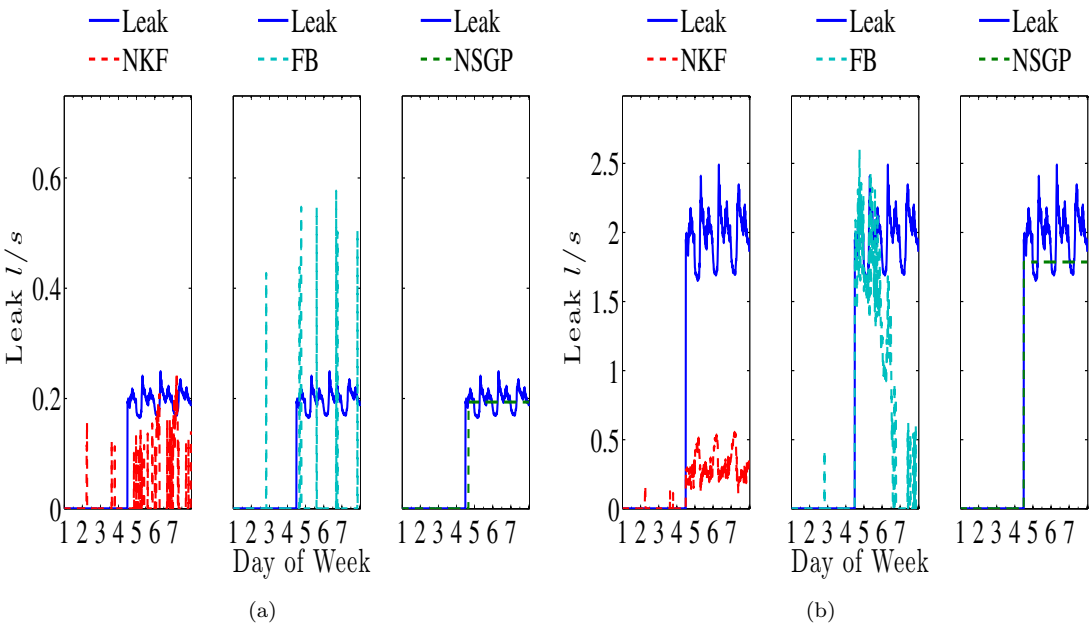

FIGURE 3.4: (a) The detected leak by all approaches for a leak starting at 12:00 PM on the 4th day with a base magnitude of 0.15 l/s. (b) The detected leak by all approaches for a leak starting at 12:00 PM on the 4th day with a base magnitude of 1.5 l/s.

the leak duration and false positives before the leak starting time. This is mainly due to two reasons, firstly both the FB and NKF approaches assume a constant gaussian noise and secondly these methods assume the leak at a particular time to be independent of the leak at other times, thus ignoring the sustained nature of leaks and using point estimates to represent leaks. Elaborating on this, Gaussian noise assumptions when the underlying noise is heteroscedastic results in over estimation of the noise at low variance times and under estimation at times of high variance. Thus, gaussian noise assumption with point estimation of leaks may give good detection results in cases where the leak magnitude is greater than the maximum observed heteroscedastic noise in the data. However, for smaller leaks, that lie within the bounds of the max and min observed variance, these approaches result in false negatives. Similarly an opposite effect of assuming a constant noise variance and leak independence between observations, is false positives observed at times with high observation noise. These effects can be seen in both cases with the smallest, Figure 3.4(a), and the largest leak, Figure 3.4(b).

Having detailed the detection results in the maximum leak evidence case we know look at the case when only two and a half hours i.e, ten leaky readings, have past since the leak occurred. Table 3.2 shows the aggregate performance results of all algorithms on ten leaks with base magnitudes ranging from 0.15 l/s to 1.5 l/s, all starting on the seventh day at 21:00 PM.

The results show that even when minimum evidence of a leak is available the NSGP performs better than the state of the art with the highest TPR, 1.00, lowest FPR,

	Reporting Time	**TPR**	**FPR**	**ACC**	**F1**	**RMSE**	**Execution Time (secs)**
		Averages over 10 different leaks					
NL	2 days	0.000	0.010	0.990	0.000	-	**0.231**
NKF	15 mins	0.808	0.050	0.947	0.362	0.127	0.371
FB	15 mins	0.838	0.009	0.988	0.692	0.052	96.807
NSGP	15 mins	**1.000**	**0.000**	**1.000**	**0.993**	**0.013**	5.667

TABLE 3.2: Showing averaged detection results over the 10 leaks with base magnitudes of 0.15 to 1.5 l/s simulated on 21:00 PM on the seventh day.

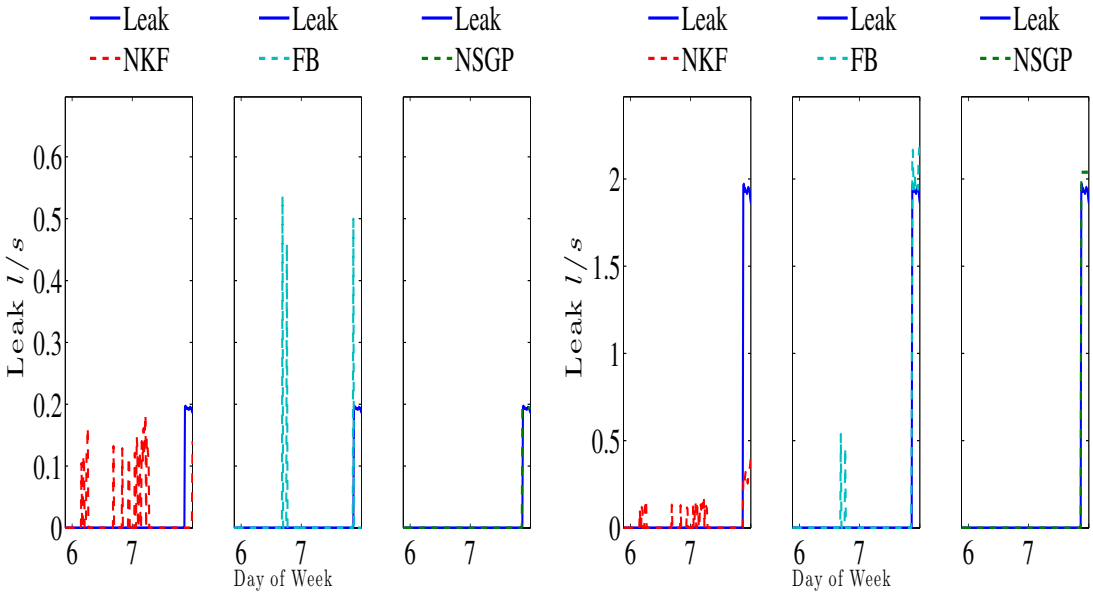

FIGURE 3.5: (a) The detected leak by all approaches for a leak starting at 21:00 PM on the 7th day with a base magnitude of 0.15 l/s. (b) The detected leak by all approaches for a leak starting at 21:00 PM on the 7th day with a base magnitude of 1.5 l/s.

0.00 and the lowest root mean square error (RMSE) of 0.120 l/s in leak magnitude estimation. The NKF and FB algorithm results again show false positives and false negatives, the result of Gaussian noise assumptions and point estimation of leaks. As in the maximum evidence case, these effects can be seen in Figure 3.5, which shows the leak detection results of the selected algorithms over the smallest, 0.15 l/s and the largest, 1.5 l/s, simulated leak at 21:00 PM on the seventh day. In contrast to FB and NKF, the NSGP models the sustained nature of leaks by re-evaluating the leak parameters over the whole week window with each new observation. This may seem computationally expensive however, with our dynamic bounds on the leak parameters we considerably reduce this cost. In the 100 leak detection experiments conducted, out of the 672 leak starting time slots in a one week window, on average using the dynamic bounds the proposed NSGP was able to prune 492 time slots, a 73.21% reduction in the total searchable space. Comparing the computation time of both the GP based

methods i.e., FB and NSGP, it can be seen that proposed NSGP is computationally less expensive. In the maximum evidence case, Table 3.1, the NSGP takes 31.534 seconds to compute where as the FB algorithm take 94.878 seconds. In contrast for the minimum evidence case, Table 3.2, the NSGP takes only 5.667 seconds to compute where as the FB algorithm computation time remains approximately the same, 96.807 seconds. This is because in minimum evidence case only 10 out of the 672 readings are leaky and most of the normal readings fall within the dynamic bounds i.e., the observed value minus the historic mean is less than zero, resulting in that particular time slot to be dropped from the possible leak start time search space. In the two leak detection experiments detailed in Figure 3.5(a) and 3.5(b) out of the possible 672 starting time slots in a week, the NSGP was able to prune 640 and 628 time slots respectively.

	Reporting Time	**TPR**	**FPR**	**ACC**	**F1**	**RMSE**	**Execution Time (secs)**
		Averages over 100 different leaks					
NL	2 days	0.793	**0.002**	**0.998**	0.796	-	**2.134**
NKF	15 mins	0.777	0.026	0.922	0.748	0.445	3.363
FB	15 mins	0.487	0.007	0.836	0.544	0.339	957.942
NSGP	15 mins	**0.999**	0.006	0.994	**0.986**	**0.066**	190.077

TABLE 3.3: Detection results of all approaches averaged over 100 leaks ranging from 0.15 to 1.5 l/s

Given this, in Table 3.3 we present the aggregate performance results of all algorithms on all 100 simulated leaks, which show that overall the NSGP has the highest TPR, 0.999, one of the lowest FPR, 0.006 and an almost perfect F1 Score of 0.986. In leak magnitude estimation the NSGP overall has the lowest RMSE of 0.066, where as in comparison to the GP based FB method it is 5 time computationally more efficient. Having detailed the results of the experiments the next section summarises the findings of this chapter.

3.6 Summary

In this chapter, we began by proposing a GP based method to model the DMA-level flow to construct a NOM, detailing the diagonal noise covariance function used to model input dependant noise. We then proposed a step change mean function and showed how it can be used to detect and approximate leaks. Following this we described how the parameters of the NOM and step change mean function can be optimised in a computationally efficient manner. We also described in detail the experimental setup, the leak models used and the performance metrics used in our evaluations. Discussing the results, we showed that our proposed approach outperform other leak detection methods. Furthermore, we showed that our proposed dynamic bounds for the step change mean hyper parameter learning, makes the proposed NSGP at least 5 time faster than the current state of the art GP based methods.

Given the above, the NSGP presents promising results, however, it leaves grounds for further improvement. Particularly, testing the proposed method on verifiable accurate real leak data and investigation of whether incorporating other DMA hydraulic parameters can enhance leak detection in any way. However, such analysis requires accurate truth data for leaks, specifically the exact starting date time and magnitude of a leak, for a pressure optimised DMA. However, such leak data is often not available, most water companies only record the date a leak was reported and fixed providing no information about the actual leak start time or the magnitude of the leak. Thus to do this, particularly in light of the pressure optimisation schemes used at I2O Water, their effects on leaks and the unavailability of verifiable accurate real leak data, we advance by using the industrial standard simulation software EPANET to model a pressure optimised DMA, the details of which will follow in the next chapter.

Chapter 4

WDS Modelling With EPANET Extension - WaterNetGen

In this chapter we start by giving a detailed discussion of the challenges faced in acquiring accurate leak data (Section 4.1). In the absence of real verifiable leak data we proceed by simulating a pressure optimised DMA using an enhanced version (WaterNetGen) of the widely used WD modelling software EPANET. Given this, in Section 4.2, we provide a brief description of EPANET software, its limitations and a brief description of WaterNetGen EPANET extension (Muranho et al., 2014) that solves these limitations. Following this, we describe how a synthetic DMA can be modelled and simulated in WaterNetGen. After this, in Section 4.3 we detail how control rules in WaterNetGen can be used to implement pressure optimisation in a simulated DMA. Next in Section 4.4, we detail the leak model used to generate pressure dependent leaks, we also discuss the effects of pressure optimisation on the DMA dynamics, particularly the implications of pressure optimisation on any leak detection solution. Finally, in Section 4.5 we give a summary of the chapter and our findings from the simulation.

4.1 Challenges in Leak Data Acquisition

Leaks are non deterministic and a source of concern for the water companies. Despite this, there is a clear lack of applied leak detection and record keeping mechanisms in the industry. As mentioned before in Section 2.7, currently most companies rely on either, visible burst reports or customer complaints to record the occurrence of a leak. Even in these cases, the exact accurate time a leak started in never known or found. Also, there may be multiple leak complaints, reported at different dates, at a given time in different parts of a DMA, in such scenarios it becomes even harder to determine when a particular leak started. Similarly, when a leak is fixed only the repair date is recorded,

which in cases of multiple repairs on the same date makes the task of finding the exact leak start, end times even harder.

Given the above, the task of finding true leak times becomes even harder for companies like I20 Water. This is because, firstly as I20 Water provides network optimisation solutions to water companies, they do not have direct access to the recorded leak reporting and repair data. Secondly, this data is often private and sensitive, as water companies can incur fines from Ofwat based on this, therefore, such data has to be requested from the client companies. These requests sometimes require legal formalities such as signing a non disclosure agreement (NDA), which makes this process long and time consuming.

In context of this research, even in cases when leak reporting and repair dates are available, it is often impossible to determine the exact leak times and magnitudes from such data. This is mainly because this data is not only susceptible to human errors but also ongoing network optimisation, seasonal, weather changes and DMA level adjustments e.g., repairs, modifications, firefighting, addition of new properties etc. Often such changes go unrecorded, thus making it impossible to separate out leak effects from the effects caused by such changes.

Given the above challenges and the unavailability of verifiable accurate real leak data, we proceed by using and enhanced version of the WDS simulation software EPANET, namely WaterNetGen, to model a pressure optimised DMA. In the next sections we briefly describe EPANET its limitation and how they are addressed by WaterNetGen. Following this, we describe how a DMA can be modelled in EPANET to obtain simulated data for evaluating the performance of the leak detection approaches.

4.2 DMA Modelling With WaterNetGen

EPANET is a free open source WDS modelling software, where the WDS models are built based on a link-node formulation. Links join together various components of the WDS and are normally of three types. Namely, pipes that convey water within the network, pumps that increase the hydraulic head by imparting energy and valves that limit the flow or pressure at a particular point within the network. Water sources such as tanks, rivers, lakes etc. are modelled using nodes. Similarly, junctions are used to represent water entry and exit points in the network. Therefore, they are normally used to model customer demand or consumption, in such cases a junction can represent a single or multiple consumers.

EPANET allows defining various properties for each type of component in the WDS e.g., pipe length, diameter, material, pump speed, power, demand pattern at each junction, ground elevation, pressure head etc. All these parameters, once set, are then used by the EPANET hydraulic solver to compute the pressure head at junctions, link flows,

tank levels etc. at each time step in the simulation, where the time step and the total simulation length are configurable by the user. The hydraulic solver computes the solution by iteratively solving the non linear conservation of flow equations and the head loss relationship across each link in the network.

However, one limitation of EPANET is that it can only simulate demand-driven WDS models i.e., models where the outflows in the network are assumed to be constant irrespective of the available pressure. Although, demand-driven models are accurate and have been used in planning, operation and analysis in WDS (Ingeduld et al., 2006; Yu et al., 2010; Alexander and Boccelli, 2010), they only work when the WDS is assumed to be in normal conditions i.e., with sufficient pressure available in the network. Such models do not consider the impact of pressure changes on the flow supplied to the network. However, given the objectives of this research, we are particularly interested in pressure optimised DMAs[1]. Therefore, to implement pressure optimisation it is crucial to have a hydraulic model that takes in to account the effects of pressure changes on the link and nodal flows. Given the limitation of EPANET, we proceed by using a enhanched version of EPANET, namely WaterNetGen that, among other enhancements, adds the capability to run pressure-driven analysis to EPANET (for details see Muranho et al. 2012). Having given a brief description of EPANET, its limitations in terms of this research and the reason for using WaterNetGen, we know briefly describe the details of modelling a DMA in WaterNetGen.

WaterNetGen, allows design of a WDS as a whole, the design and isolation of metering areas are left to the user, i.e., these softwares have no inbuilt mechanisms or conception of DMA. However, multiple DMAs can be modelled by isolating different parts of the WDS using valves. In our case, we are only interested in modelling a single pressure optimised DMA, which is identical to modelling a small isolated WDS. To do this, we firstly choose the Highfield area in the University of Southampton from google maps and lay pipes in line with roads, connecting them with junctions. We add a reservoir and a pumping station at the DMA supply point, which supplies the required pressure to the network (Figure 4.3). The next step in the model design involves setting parameters for the various DMA components, we highlight the most important of these.

To model the demands in the DMA, as close to a real life DMA as possible, we assume the DMA caters for five thousand inhabitants, where the consumption per person is 150 litres per day i.e., the UK wide average water consumption per person as per day (Ofwat, 2015). As there is always some background water loss in real WDS from pipes, joints etc., we compensate for this by adding a 10% background water loss to the overall demand. Figure 4.1(a) shows the DMA-level demand setup for five thousand inhabitants in WaterNetGen. This overall DMA-level demand is equally distributed to all junctions

[1]Pressure optimised DMA: A DMA in which the network pressure is kept as low as possible while maintaining the desired serviceability standards. This is achieved by varying the supply point pressure, P2, such that the critical point pressure, P3, stays within a desired rage or close to a predefined value

in the DMA, where at the junction level, WaterNetGen allows to setup time based consumption patterns of how the DMA-level distributed demand changes over time at a particular juncation. This is done by defining demand patterns and assigning them to each junction. Demand patterns are a list of time based multiplier which when multiplied with the constant distributed demand at each junction yield time varying consumption patterns. To model the demands as close as possible to real life consumption patterns and to model both daily and weekly variations in data, we define a five week time pattern. The five week pattern is computed by normalising five weeks of normal operation flow data from the Arnesby Village, Leicestershire, UK DMA, previously shown in Figure 3.3. Using actual normalised data as a demand pattern allows us to mimic both the daily diurnal demand characteristics and weekly consumption variations observed in real life DMAs. Figure 4.1(b) shows the 5 week demand pattern setup in WaterNetGen. Considering a 15 minute measurement recording interval, a 5 week demand pattern would have 5 (weeks) x 7 (days in a week) x 24 (hours in a day) x 4 (readings in an hour) = 3360 multipliers, where each multiplier represents the consumption pattern for a particular 15 minute interval in the 5 weeks.

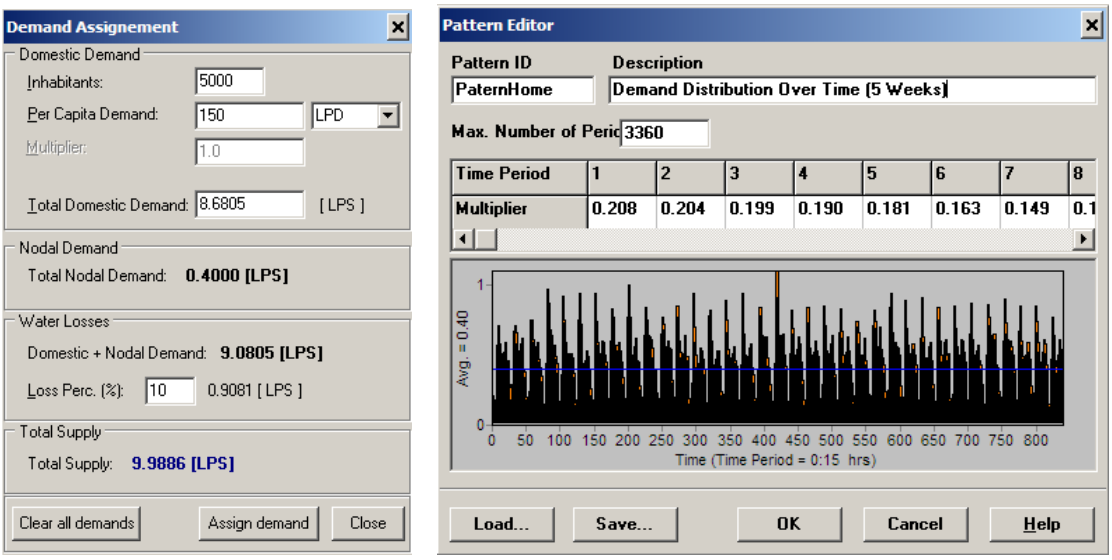

(a) (b)

FIGURE 4.1: (a) The DMA-level demand setup for five thousand inhabitants in WaterNetGen. (b) A five week demand pattern setup in WaterNetGen

Having defined the demand modelling parameters, we now define properties of components that govern the suppliable pressure and head loss characteristics of the DMA. The first property that must be set for every node and junction is the ground elevation. The elevation, measured in meters, is used to compute the pressure at junctions and the suppliable head at nodes. For example, in a gravitational water supply scheme a water tower with a suppliable pressure head of 30m at an elevation of 50m, will be able to supply a total pressure head of 70m to all connect nodes which are at an elevation of 10m i.e., $50 + 30 - 10$ (not considering frictional losses in pipes). Similarly, if a junction

is at a higher elevation, say 60m, than the water source e.g. 30m, than a pumping source would be required to supply the additional 30m head to get the water from the source to the junction. In the current DMA model we assume the DMA to be in a flat area with a constant elevation of 50m for all components in the DMA. However, to define a critical pressure point in the DMA we set one of the junction, J57, to a higher elevation of 60m.

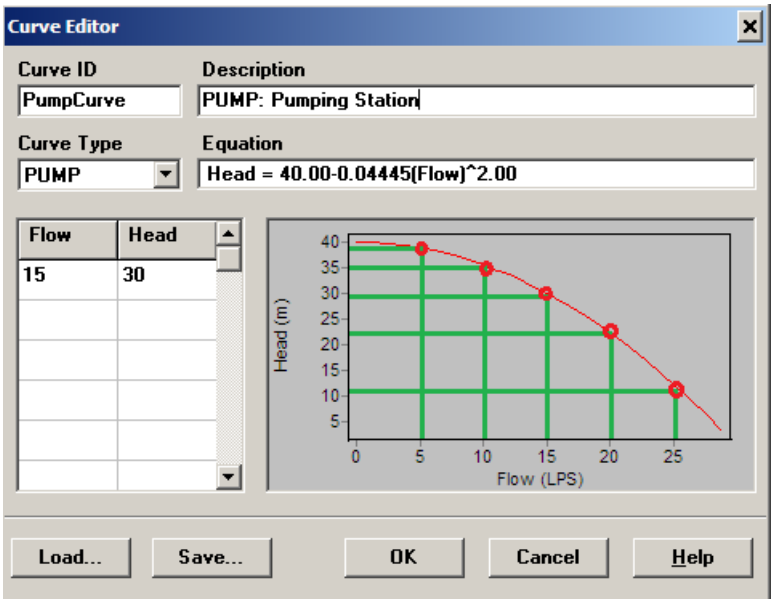

FIGURE 4.2: Pump curve for the pumping station. The pump is able to supply any head flow combination that lies under the area of the curve. The green lines map the max head the pump can supply at the given flow rates.

The final step in the DMA model design is to define the pump characteristics i.e., how much pressure head can the pump deliver given a particular litres per second demand. In WaterNetGen this is done by defining a pump curve. A pump curve represents the relationship between the head and flow rate that a pump can deliver. The choice of a pump for a WDS is determined by the demand and required head in the network. However, often, to compensate for seasonal variations, DMA expansion, fire fidgeting etc, pumps capable of supplying a higher head flow ration than the demand requirements are chosen. Keeping this in mind, in our current DMA model we select a pump curve that can supply slightly higher head flow ratio than required by our demand of approximately 10 litres per second. We assume the pump to be a variable speed pump, i.e. the speed of the pump can be changed to increase or decrease the output pressure head. Given this, the modeled pump can provide any head flow combination that lies under the pump curve area. Figure 4.2 shows the pump cure modelled in WaterNetGen, here the green lines in the pump curve map the max head the pump can supply at the given flow rates. The final modelled DMA with all the defined components, overlaid on Google maps is shown in Figure 4.3. The supply point pressure and flow, P2 and Flow, at junction J1 are marked with a green circle, where as the critical point pressure, P3, at junction J57 is marked with a red circle.

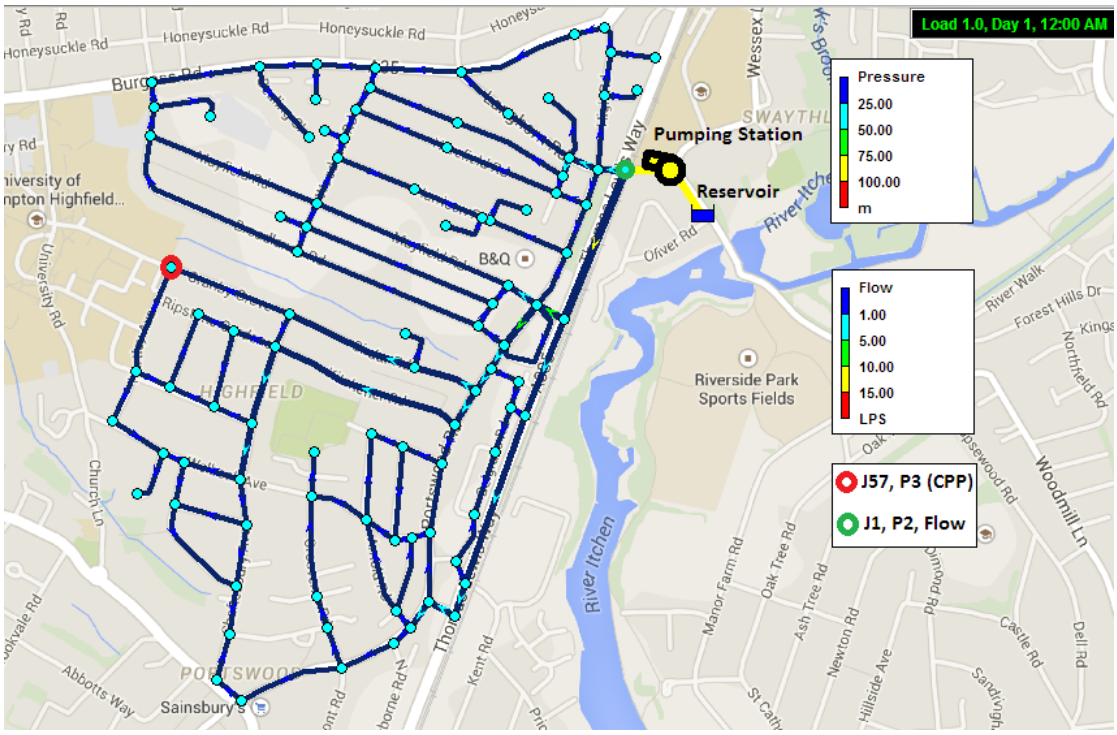

FIGURE 4.3: The modelled DMA overlaid on google maps. The inlet pressure and flow points, P2 and Flow, at junction J1 are marked with a green circle, where as the critical pressure point, P3, at junction, J57, is marked with a red circle.

Having defined the DMA model, in the next section we detail the mechanism used in implementing I20 style pressure optimisation in the modelled DMA.

4.3 Implementing Pressure Optimisation

Pressure optimisation in a DMA involves keeping the network pressure as low as possible while maintaining the desired serviceability standards. At I20 water, this is achieved by varying the supply point pressure, P2, such that the critical point pressure, P3, stays within a desired rage or close to a predefined value. In WaterNetGen this can be achieved by using rule based controls, which allow changes to be made to network components based on a combination of conditions that might exist in the network during the course of the simulation. Given this, we define a min range of 9m to 10m head for the critical point pressure, P3. The defined range is higher than the 7m head Ofwat requirement for household supply standards. This is because, often the critical point does not represent a single consumer, but rather a point in the network that supplies multiple consumers. Also, the critical point pressure is normally kept slightly higher than the required value to compensate for sudden increases in demands such as during firefighting. Thus, we maintain P3 within the 9 to 10m head rage by varying the pump speed in small increments. In detail, if the P3 value falls below 9m head we keep increasing the pump speed in small increments of 0.2% till P3 is back in range. Similarly, if P3 pressure goes

above the 10m head threshold we decrease the pump speed incrementally by 0.2% till P3 is back within range. In WaterNetGen the original pump curve supplied to the program has a relative speed setting of 1 i.e. operating at its maximum speed. For a pressure optimised DMA, assuming the DMA is optimised such that the pressure requirements are met when the pump is operating at half its capacity, then the relative setting would be 0.5. Given this we use the following algorithm to generate 500 rules to implement pressure optimisation in WaterNetGen.

Input1: NodeP3; //Critical point pressure Node
Input2: P; //Pump
Define P3min= 9; // minimum allowed Critical point pressure value
Define P3max= 10; // maximum allowed Critical point pressure value
Define Smin= 0.5; // minimum speed pump can operate on
Define Smax= 1; // maximum speed pump can operate on
Define ΔS= 0.002; // step change to introduce in pump speed
Define Rcounter= 1; // Rule number counter
Define i =Smax;
while $i <$*Smax* **do**
    Print Rule Rcounter
    Print IF NODE NodeP3 PRESSURE > P3max
    Print AND PUMP P SETTING = i
    Print THEN PUMP P SETTING = i $-\Delta$S
    Rcounter=Rcounter+1
    Print Rule RCounter
    Print IF NODE NodeP3 PRESSURE < P3min
    Print AND PUMP P SETTING = i $-\Delta$S
    Print THEN PUMP P SETTING = i
    $i = i - \Delta$S
end

Algorithm 2: WaterNetGen Pressure Optimisation Rules algorithm

After describing the DMA optimisation mechanism, we now briefly look at one week of the simulated DMA-level hydraulic readings of the modelled DMA with and without the pressure optimisation in place. Figure 4.4 and 4.5 shows the supply point pressure and flow, P2 and Flow and the critical point pressure, P3, for the unoptimised and optimised DMA respectively.

In the unoptimised DMA the P2 and P3 values are around 40m and 30m head respectively. Where as in the optimised DMA the P3 value is controlled and kept within the range of 9 to 10m head, resulting in the reduction of supply point pressure from 40m head at peak time to around 20m head. Thus, the optimisation results in approximately 50% pressure reduction in the simulated DMA. The stepwise reduction of the pressures in the DMA can be seen in the initial readings on 01/01/14 in Figure 4.5 where the rule based controls are incrementally adjusting the pump speed to regulate the critical point pressure in the DMA.

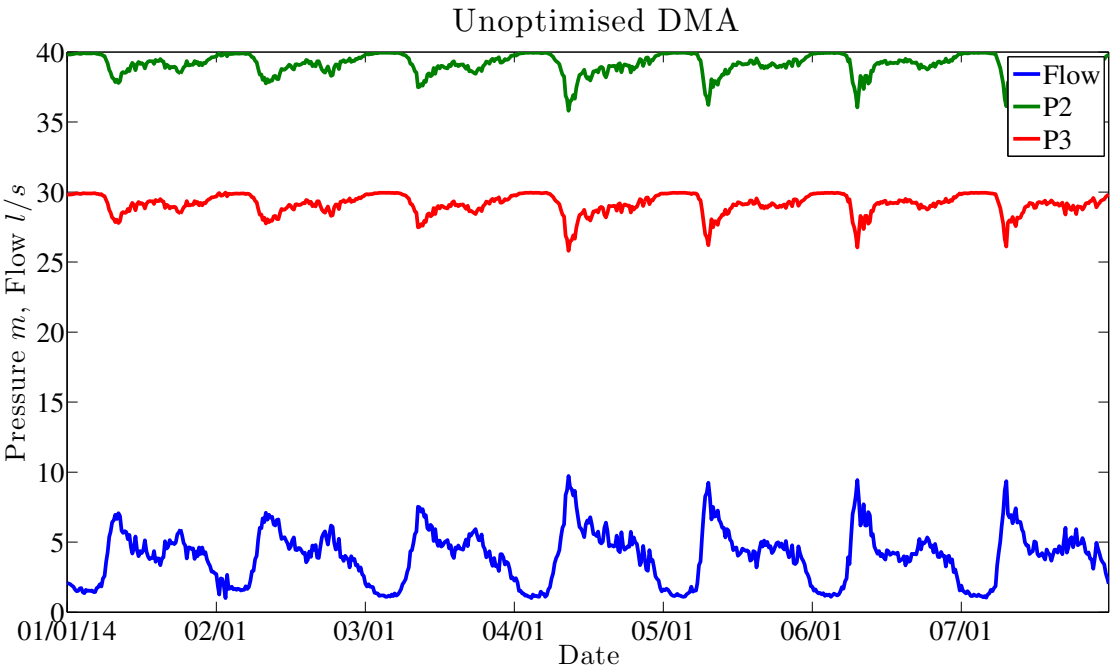

FIGURE 4.4: The DMA-level hydraulic readings without active pressure optimisation.

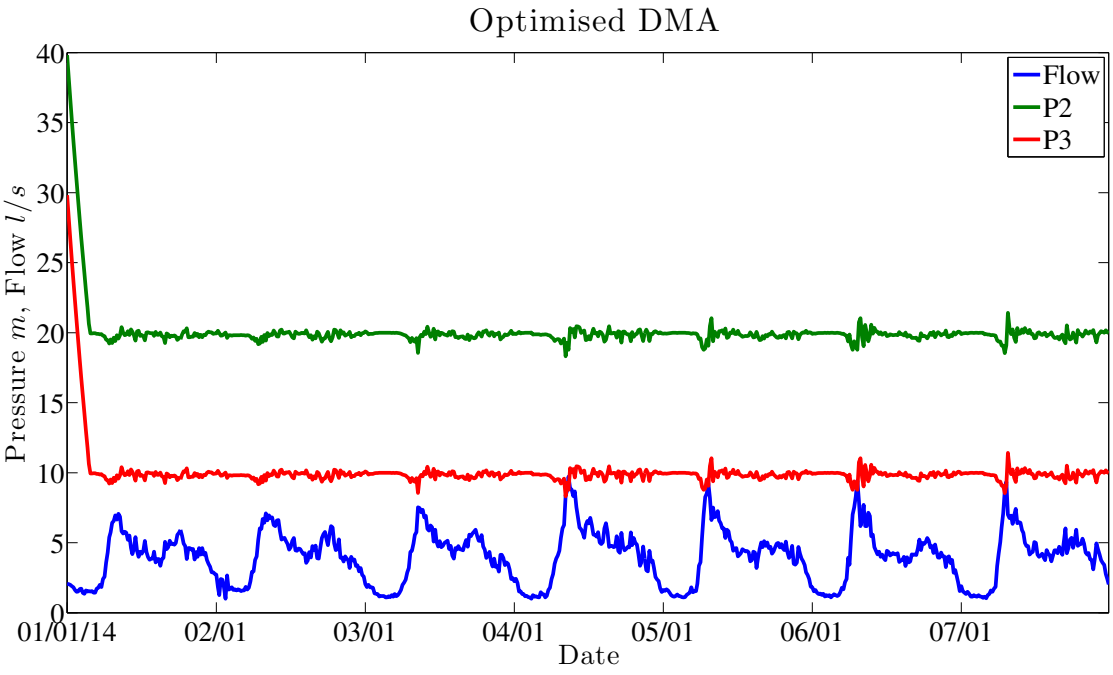

FIGURE 4.5: The DMA-level hydraulic readings with active pressure optimisation.

Having detailed the pressure optimisation mechanism and the resulting simulated DMA-level hydraulic readings, in the next section we will describe the leak simulation model used to generate leak in the DMA. We shall also look at how the pressure optimisation changes the dynamics of the DMA in case of leak events and briefly discuss its implication of leak detection solutions.

4.4 Pressure Dependent Leak Simulation

Unlike EPANET, WaterNetGen adds the capability of modelling leaks in individual pipes using the classical orifice flow formula, $q = Cp^{\gamma}$, as described earlier in Section 3.4. WaterNetGen assumes leakage increases continuously with pressure and the pressure-leak relationship for a given pipe, k, in the network is given as:

$$q_k^{leak}(P_k) = \begin{cases} C_k(P_k)^{\alpha_k} & \text{if} \quad P_k > 0 \\ 0, & \text{otherwise} \end{cases} \tag{4.1}$$

Here, given a pressure of P_k, the total leakage in pipe k is given by q^{leak}, where as C_k and α_k are the leak coefficient and exponent respectively. WaterNetGen computes the pressure P_k in pipe k as the mean of the pressure values of its end nodes.

For remainder of the thesis, we use the above formulation to simulate leaks in the DMA. In all leak simulations we assume a leak exponent of 0.5 (used for a circular aperture), where as we model varying leak magnitudes by using different values for the leak coefficient. We now look briefly look at the DMA dynamics under leak conditions for both the optimised and unoptimised DMA. We simulate a leak on 04/01/2014 at 12:00PM with a leak coefficient of 0.5 litres per second per meter head change in pressure.

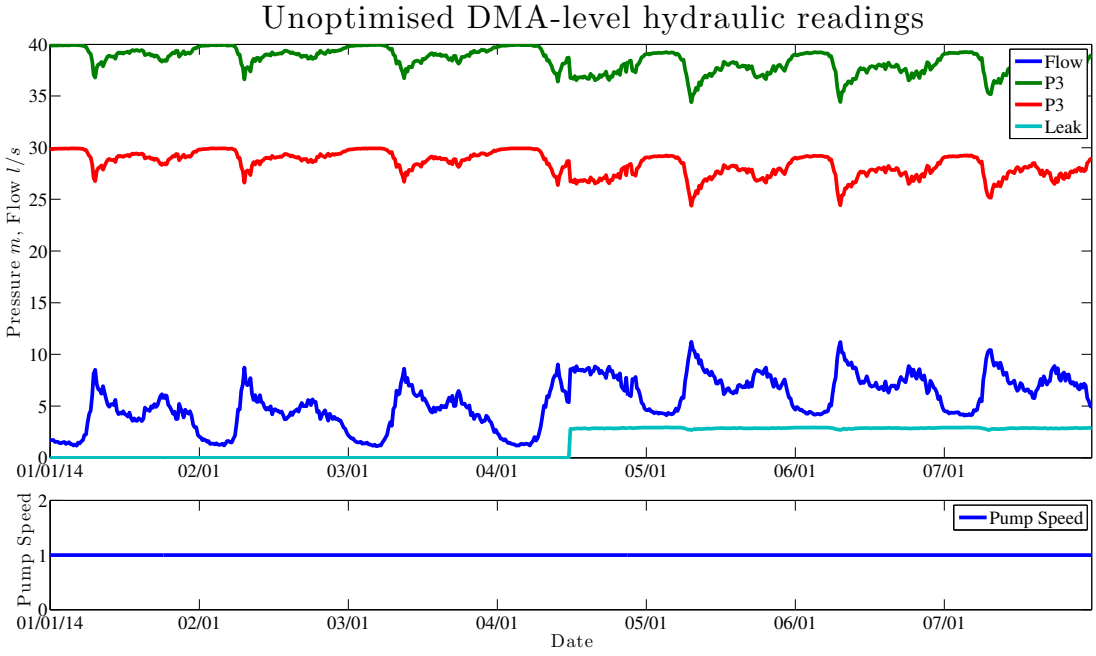

FIGURE 4.6: The DMA-level hydraulic readings for the modelled DMA without active pressure optimisation for a simulated leak with a leak coefficient of 0.5.

Figure 4.6 shows one week of the resulting DMA-level hydraulics readings and the pump speed for the unoptimised DMA under the simulated leak conditions. In the unoptimised DMA the leak affects all the DMA-level hydraulic time series. The leak increases the

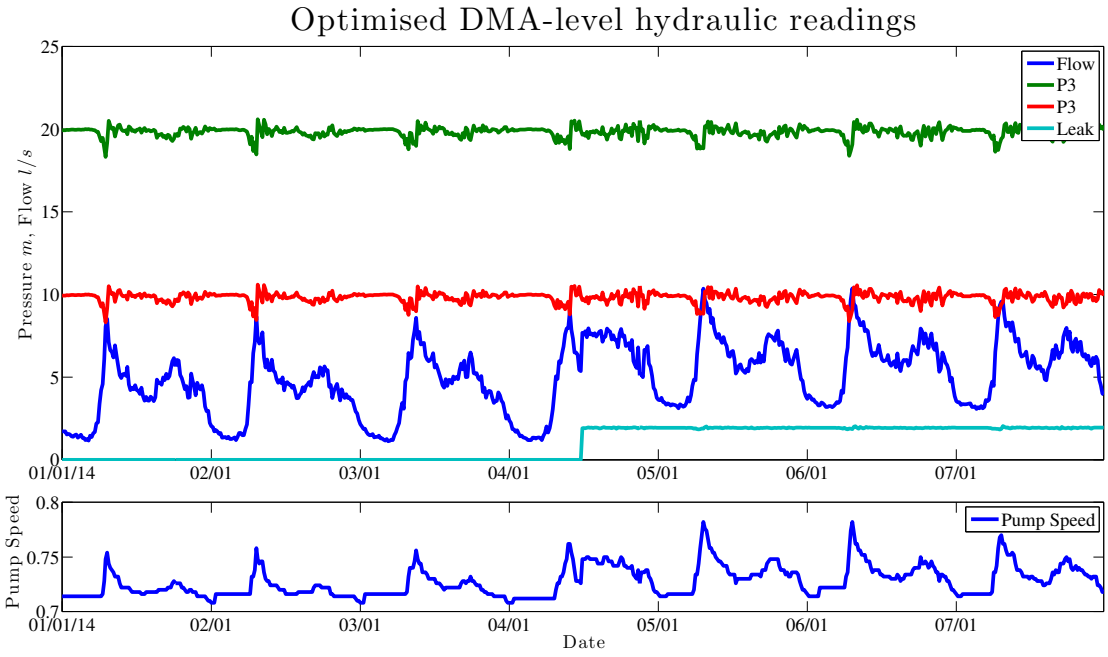

FIGURE 4.7: The DMA-level hydraulic readings for the modelled DMA with active pressure optimisation for a simulated leak with a leak coefficient of 0.5.

DMA-level flow which in turn, given the fixed pump speed, results in the reduction of the supply point pressure, P2, and the critical point pressure P3.

In contrast to the unoptimised DMA, the dynamics of the optimised DMA, shown in Figure 4.7, under leak conditions are the opposite. As the pump can supply higher level of pressure, as the pressure in the DMA drops due to the leak, the DMA level optimisation keeps the supply point and critical point pressures, P2 and P3, at the same optimised levels as when there was no leak in the system Figure 4.5. Thus, under leak conditions, pressure optimisation has a profound effect on the information that can be gathered from different DMA-level hydraulic readings. In particular, for pressure optimised DMAs, leak only affects the DMA-level flow, as the supply and critical point pressures remain mostly unchanged. Thus, they yield no extra information that could aid in leak detection. However, pressure optimisation adds another hydraulic DMA-level property to the available DMA-level data, namely the pump speed. In case of a pressure optimised DMA the pump speed keeps changing to keep the overall network pressures at desirable rates, since the overall consumption patterns in a DMA follow a daily and weekly trend and directly affect the pressures in the DMA. This results in the pump speed having a pattern similar to the DMA-level flow. Furthermore, in case of a leak the the pressure optimisations results in the pump to operate at a higher speed deviating from its normal pattern. Both these factors allow an additional set of information that can be used in conjunction with the DMA-level flow measurements to aid leak detection. Having detailed the effects of pressure optimisation on the DMA dynamics, particularly in light of leak detection, in the next section we give a brief summary of this chapter.

4.5 Summary

In this chapter we started by giving a description of the various challenges faced in acquiring real and accurate leak data. We highlighted that poor record keeping and on going unrecorded DMA level changes make deduction of leaks from real DMA-level hydraulic measurements very difficult. Given the unavailability of real leak data we then proceed by simulating a DMA in a WDS simulation software EPANET. We give a brief description of EPNAET software highlighting its limitations, particularly in context of running pressure dependant simulations, which are inherently required to implement a pressure optimised DMA. We highlighted that WaterNetGen WDS modelling software, which is based on EPANET, addresses this limitation and adds the capability of running pressure dependant analysis to EPANET. Following this we gave a detailed explanation of the steps involved in modelling a DMA in WaterNetGen. In particular we showed how demand pattern computed from real DMA data can be used to model close to real life DMA-level demands that exhibit both, the daily diurnal flow and weekly variations observed in real life DMAs. In the following section we described in detail the process of using rule based controls to model active pressure optimisation in the DMA. In detail, we described how control rules can be used for a variable speed pump to increase or decrease its speed based on the critical point pressure, to implement I20 Water like pressure optimisation. The resulting DMA-level hydraulic readings for both optimised and unoptimised DMA showed that the proposed optimisation resulted in 50% reduction of the overall network pressure. Following this we briefly descried the pressure dependent leak model used to simulate leaks in the pressure optimised DMA. After analysis of both, the optimised and unoptimised DMAs, under leak conditions, we discussed how pressure optimisation changes the DMA dynamics. In particular, we highlighted that in pressure optimised DMAs a leak only affects the DMA-level flow, where as the pressure optimisation keeps the supply and critical point pressures at the same levels as under no leak conditions. In contrast in an unoptimised DMA as the flow increases due to a leak, it results in the overall network pressure to drop. Thus, affecting all the three DMA-level hydraulic measurements. Given this, we concluded that in pressure optimised DMAs under leak conditions, the supply point and critical point pressures provide no useful information pertaining to the leak. However, as the optimisation needs to keep track of not only the pressures in the network but also the pump speed, it provides an additional data stream i.e. the pump speed, that is directly affected by leaks. Under leak conditions, in pressure optimised DMAs, the pump speed increases to compensate for the pressure drop in the DMA, thus deviating from its normal pattern. This deviation can be detected and can provide an extra source of information pertaining to a leak for any leak detection solution.

Given the above, in the next chapter we propose a novel multi-output GP based model, that incorporates the information from not only the all DMA-level pressure and flow measurements but also the pump speed, to detect and quantify leaks in a DMA.

Chapter 5

Probabilistic Leak Detection And Characterisation Using Efficient Semi-Parametric Multi Output Gaussian Process (SMOGP)

In this chapter we begin by proposing a multi output GP based method (SMOGP) to model the DMA-level flow and pump speed. In Section 5.1 we describe how the two time series and their correlations can be modelled using labeled inputs. We then used this formulation to define the covariance structure for NOM of a pressure optimised DMA. Following in Section 5.2 we describe the parametric leak mean function, where we model pressure dependent leaks using the flow through an orifice model. Next, in Section 5.3 we detail the dual optimisation mechanism used to learn the NOM and leak hyper parameters. Following this in Section 5.4 we briefly describe the simulate leak data used to evaluate the performance of the proposed approach. The results of the performance analysis are then given in Section 5.5. Following this, in Section 5.6 we give a summary of the chapter.

5.1 Multi output GP Model For A DMA

Most GP implementations model only a single output variable. Numerous ways have been proposed to model multiple correlated outputs in a GP. Boyle and Frean (2005) and Alvarez and Lawrence (2009) propose modelling multiple outputs using convolution kernels. In particular, these methods model each output as a convolution between a smoothing kernel and a latent function, where all outputs share the same latent function which is assumed to be white noise GP. As convolution on a function is a linear operation,

the convolution outputs are expressed as a jointly distributed GP i.e. the multi-output GP.

Modelling multiple outputs in GPs poses challenges as this requires computation of the cross covariance between model outputs. Although, the resulting models can improve the prediction accuracy of each output given the other, (as the correlations between them as modelled), such models increase the computational complexity considerably. In particular, as these approaches aim at providing a generalised multi-ouput GP method that can handle both linear or non-linearly correlated outputs. They result in heavily parameterised complex covariance structures which have high computation and storage overheads (Boyle and Frean, 2005; Alvarez and Lawrence, 2009). However, if the outputs being modelled have a high linear correlation then the complexity in representing the cross covariance can be greatly simplified by representing it as a linear function of the covariance of either of the outputs (Osborne et al., 2008; Rogers et al., 2011).

In the context of this research, our objective is to use the available DMA-level measurements to detect a leak and provide as much information as possible about it. In light of this, from the modelled DMA simulation in Chapter 4 we obtained four time series, i.e., the supply point pressure and flow, P2 and Flow, the critical point pressure, P3 and the pump speed which we represent as S. The simulation showed that the pump speed and flow follow a similar pattern and are highly correlated. A Pearson's correlation coefficient (a standard measure of the linear correlation between two variables) value of 0.8945, computed using five weeks of NOM pump speed and flow, shows that both the time series have high positive linear correlation. In light of this, we model both the pump speed and flow in a multi-output GP where the cross covariance between the pump speed and the flow are represented as a simple linear multiple of the flow covariance. In contrast to the NSGP in which the leak start time decision is solely based on the deviations in the flow time series we aim to make a more informed decision based on the information gathered from both the time series. Having described the objective in using a multi output GP, we know describe the multi output GP model for the DMA under normal operating conditions.

We model the DMA-level pump speed, S, and flow, F, as a concatenated output vector $y_{i=1}^{n\times m} = [S_1 \dots S_n, F_1 \dots F_m]$. The inputs, x, to the GP are also represented as a concatenated vector of the measurement times for each individual time-series with an additional dimension specifying a label, l, to distinguish each time-series e.g. $x = [(t_1, l = S), \dots (t_n, l = S), (t_1, l = F) \dots (t_m, l = F)]$. In our case our label set, L, consists of two values, S and F, to identify the pump speed and flow time-series i.e., $L = [S, F]$. The correlation between the time-series are then modelled as a Hadamard product of the covariance over time and a covariance over the time-series labels as:

$$K([t, l], [t^{'}, l^{'}]) = K_L(l, l^{'}) \odot K_T(t, t^{'}) \tag{5.1}$$

Where the covariance function over the time-series labels is given as:

$$K_L(l, l^{'}) = \left\{ \begin{array}{ll} 1 & \text{if } l = l^{'} \\ \rho & \text{otherwise} \end{array} \right\} \tag{5.2}$$

Here $K_L(l, l^{'})$ is 1 if both the if both inputs are from the same time series where the hyper parameter ρ determines the cross-correlation between the time series. Thus, for an input vector $x = [(t_1, l = S), \ldots (t_n, l = S), (t_1, l = F) \ldots (t_m, l = F)]$, the time-series label covariance, $K_L(l, l^{'})$, would produce a matrix where the top left and bottom right blocks, of sizes $n \times n$ and $m \times m$ respectively, would have all elements with a value of one and the top right and bottom left blocks, of sizes $n \times m$ and $m \times n$ respectively, would have all elements having a value of ρ.

In case of pressure optimised DMA, as seen in Figure 4.7, the data has three distinct patterns we wish to model. Firstly, we aim to model the weekly periodic nature of the DMA-level hydraulic data, we achieve this by using a periodic covariance kernel k_P given as:

$$k_P(t, t^{'}) = \sigma_{f_P}^2 \exp(-\frac{\sin^2 \pi(t - t^{'})}{2l_P^2}) \tag{5.3}$$

Here, $\sin^2 \pi(t - t^{'})$, determines the distance between repetitions of the function, the length scale, l_P^2, controls the horizontal length scale over which the function varies where as the output variance, $\sigma_{f_P}^2$, controls the vertical variation. In addition to the weekly periodicity, we also wish to model the changes within a day, we model these using the matérn, k_M, covariance kernel which is given as:

$$k_M(t, t^{'}) = \sigma_{f_M}^2 \left(1 + \frac{\sqrt{3}(t - t^{'})}{l_M}\right) \exp(-\frac{\sqrt{3}(t - t^{'})}{l_M}) \tag{5.4}$$

Similar to the periodic kernel, here the length scale, l_M, controls the horizontal length scale over which the function varies where as the output variance, $\sigma_{f_M}^2$, controls the vertical variation. The third pattern we wish to model is the noise in the data. We model this using the a gaussian noise kernel, k_N given as:

$$k_N(t, t^{'}) = \sigma_n^2 \delta_{t,t'} \tag{5.5}$$

Here δ is the Kronecker delta and σ_n^2 represents the noise variance. In addition to this, we also want the correlation between data points which are close to each other in time (i.e., 8:00 AM and 8:15 AM) to be stronger as compared to data points which are further apart in time. We achieve this by modelling our temporal covariance as, $K_T(t, t^{'}) = k_M(t, t^{'}) + k_P(t, t^{'}) + k_N$, as a sum of the matérn, k_M, periodic, k_P, and a gaussian noise kernel, k_N kernels. In general, by adding together kernels, the resulting temporal covariance will have a high value if the base kernels have a high value (for

a discussion of kernel types see Rasmussen and Williams, 2006). Given this, by using time-series labels we define a constant mean function, m_{nom}, which represents the mean of each time-series in the NOM model as:

$$m_{nom}(l; l') = \begin{cases} c_S & \text{if } l = S \\ c_F & \text{if } l = F \end{cases} \tag{5.6}$$

Here, the NOM mean function has two hyper parameters, $[c_S, c_F]$, which represent the constant mean values of the pump speed and flow time-series. These constant mean parameters are computed from the NOM data. For notational brevity we represent the NOM hyper parameters for the mean and covariance jointly as $\theta_{nom} = [c_P, c_F, \rho, \sigma^2_{f_M}, l_M, \sigma^2_{f_P}, l^2_P, \sigma^2_n]$. Having defined the basic multi output GP structure, in the next section we will describe the proposed parametric leak detection and quantification mechanism.

5.2 Pressure Dependent Leak Detection and Quantification Using an Additive Parametric Mean Function

Leaks are pressure dependent and increase continuously with an increase in the supplied pressure. This pressure dependant relationship is usually modelled using the equation of flow through an orifice i.e., $q_{leak} = C(P)^{\alpha}$. Here the leak exponent, α, and the leak coefficient, C, determine the characteristics of a leak at the given pressure P. The leak exponent represents the flow through different types of fractures or apertures in pipes and is dependant on both the pipe material and the aperture shape. Field tests and experimental studies conducted by Greyvenstein and Van Zyl (2007) show that the typical leak exponents values range from 0.4, for uPVC pipe with circumferential crack, to 2.30, for corroded cluster holes in steel pipes. The leak coefficient represents the water loss per unit change in pressure given a particular type or aperture e.g., for a cement pipe with longitudinal crack with a leak exponents of 1, with the supplied pressure of 10m head, a leak coefficient of 0.5 would mean mean that a 1m increase in pressure would cause the leak to increase by 0.5 l/s, from 5 l/s ($0.5 \times 10^1 = 5$) to 5.5 l/s ($0.5 \times 11^1 = 5$). Thus, if these variables are known or estimated from a leak, they can provide valuable information, particularly for pressure optimised DMAs where the optimisation can increase the pressure in the network resulting in larger leaks.

Given this we aim to infer these variables in addition to detecting and quantifying leaks. We do so by defining a leak mean function, m_L, that models the pressure dependant leak relationship using the orifice flow equation. In particular, similar to the covariance function, we define the leak mean function as a Hadamard product of two labeled mean function, m_{f_t} and m_{f_m}, defining the leak start time and magnitude for each time-series (pump speed and flow), respectively i.e. $m_F = m_{f_t} \odot m_{f_m}$. As our simulations show, a

leak will cause a positive change in both the pump speed and the flow at the same time when the leak occurs. If we represent the starting time of the deviation as τ then the leak time mean function, m_{f_t}, is given as:

$$m_{f_t}(t,l) = \begin{cases} 1 & \text{if } l = S \quad \& \quad t \geq \tau \\ 1 & \text{if } l = F \quad \& \quad t \geq \tau \\ 0 & \text{otherwise} \end{cases} \tag{5.7}$$

Here, the leak time mean function, m_{f_t}, is parameterised by τ which represents the time each of the modelled time series started to deviate from the NOM. The leak magnitude mean function, m_{f_m}, is used to quantify the deviation in both time-series. We model the deviation in the pump speed as a simple step change function, similar to the NSGP step change mean function, where we define a hyper parameter θ_S to represent the step size. We model the deviation in the flow time-series i.e, the leak, by modelling the orifice flow equation, where the leak coefficient and exponent, C and α, form hyper parameters of the mean function. As at the DMA-level we do not know the exact pipe and pressure at the point of the leak, we use the average of the supply and critical point pressure i.e. average zone pressure (AZP), as a representative of the pressure in the DMA. For notational brevity we represent the AZP as P. Given this, the leak magnitude mean function, m_{f_m}, parameterised by θ_S, C, α is given as:

$$m_{f_m}(P;t,l) = \begin{cases} \theta_S & \text{if } l = S \\ C(P)^{\alpha} & \text{if } l = F \end{cases} \tag{5.8}$$

where P is computed from the observed pressures in the DMA and supplied as an exogenous input to the leak magnitude mean function. Ideally one would like to learn the exact time varying deviation pattern for the pump speed, however, our main goal is detecting, quantifying and obtaining information about leaks where the pump speed deviation serves as an additional layer of confirmation for the leak starting time. Given our leak model, based on the flow through an orifice equation, the pump speed does not influence the leak estimation, thus learning the exact time varying pump speed deviation pattern serves only to increase the computation complexity without any other resulting gain. Given this, the hyper parameters for the overall fault mean function, m_F, are jointly represented as $\theta_L = [\tau, \theta_S, C, \alpha]$, where the final mean function for the SMOGP is a sum of the NOM mean function, m_{nom}, and the leak mean function, m_F, and given as $m = m_{nom} + m_F$.

Having defined the leak detection and quantification methodology, in the next section we define the dual hyper parameter optimisation scheme used to learn the NOM and leak hyper parameters.

5.3 Hyperparameter Learning

In this section we describe how the two sets of hyper parameters i.e. the normal operation, $\theta_{nom} = [c_P, c_F, \rho, \sigma^2_{f_M}, l_M, \sigma^2_{f_P}, l^2_P, \sigma^2_n]$, and leak mean hyper parameters, $\theta_L = [\tau, \theta_S, C, \alpha]$, are learned and how the two learning schemes are interwoven. In our model θ_{nom} represents the DMA characteristics when there is no leak in the system. To learn the optimal values for the NOM hyper parameters we first fix all leak mean hyper parameters, to zero i.e., $\theta_L = [\tau = 0, \theta_S = 0, C = 0, \alpha = 0]$. We use four weeks of normal DMA data, from the simulated DMA described in the previous chapter, as the training set. We define a one week sliding data window for the SMOGP. For each sliding window we find the optimal set of hyper parameters by minimising the negative log likelihood using conjugate gradient decent (Rasmussen and Williams, 2006). After the minimisation we record the likelihood value and computed hyper parameter for each sliding window. From all the resulting negative log likelihood values and hyper parameter sets we select the set of parameters with the lowest negative log likelihood value.

Once the multi output GP has been trained, subsequent leak detection is done solely based on the SMOGP predictions over a one week moving window (prediction window). Thus, requiring the computation of the covariance matrix and its inverse only once every week (re-training window), after which the NOM hyper parameters are re-optimised on the most recent one week data using the previously optimised hyper parameters as priors.

Most GP based change point or fault detection methods (Garnett et al., 2010; Osborne et al., 2010, 2012), where the fault is modelled in the covariance, require recomputing the covariance matrix and its inverse (or updating the cholesky factorisation) for learning the fault hyper parameters after each new observation. Thus, incurring $\mathcal{O}(N^3)$ or $\mathcal{O}(hN^2)$ computation cost for each new observation (where h is the number of fault hyper parameters). We propose an alternative computationally efficient method, facilitated by our choice of modelling leaks as a parametric mean function, where we learn the leak parameters by maximising the sum of the posterior predictive probabilities. Elaborating on this, for a set of new observations, $y^* = [S^*, F^*]$, at input times $x^* = [t^*, l]$, the SMOGP gives us a predictive mean, $\hat{\mu^*} = [\hat{\mu^*_S}, \hat{\mu^*_F}]$, and variance, $\hat{\sigma^2_*}$, then we can evaluate the quality of predictions in several ways. The standardised mean squared error (SMSE), $(y^* - \hat{\mu^*})^2/\hat{\sigma^2_*}$, is one of the simplest ways. Using this and producing a predictive distribution at each input we can evaluate the log probability of the predictions given the training data, $D = [y, x]$, and our model, $M = [\theta_{nom}, \theta_L]$ as (Rasmussen and Williams, 2006):

$$\log p\left(y^* | x^*, D, M\right) = -\frac{1}{2}\log\left(2\pi\sigma^2_*\right) - \frac{\left(y_* - \hat{\mu^*}\right)^2}{2\hat{\sigma^2_*}} \tag{5.9}$$

Given the above, leak detection in the proposed SMOGP is done using a moving input data window of one week. Assuming sensor readings 15 minutes apart, the one week input data window has 672 DMA-level hydraulic readings i.e., for a new observation at time t, the inputs to the SMOGP are $x^* = [(t-672, l = S), \ldots, (t, l = S), (t-672, l = F), \ldots, (t, l = F)]$ and $y^* = [S_{t-672}, \ldots, S_t, F_{t-672}, \ldots, F_t]$ which we represent compactly as $x^* = [t^*, l]$ and $y^* = [S^*, F^*]$. Given this, for each new observation, we first make a prediction using the SMOGP and assuming a NOM by setting $\theta_L = 0$ which gives the predictive mean $\hat{\mu^*} = [\hat{\mu^*_{S_{t-672}}}, \ldots, \hat{\mu^*_{S_t}}, \hat{\mu^*_{F_{t-672}}}, \ldots, \hat{\mu^*_{F_t}}]$ and variance $\hat{\sigma^2_*} = [\hat{\sigma^2_{*_{t-672}}}, \ldots, \hat{\sigma^2_{*_t}}, \hat{\sigma^2_{*_{t-672}}}, \ldots, \hat{\sigma^2_{*_t}}]$. For notational brevity we represent these as $\hat{\mu^*} = [\hat{\mu^*_S}, \hat{\mu^*_F}]$, and $\hat{\sigma^2_*}$ respectively.

Following this, to learn the leak hyper parameters, we maximise the sum of the predictive probability in equation (5.9) using a bounded search process where we dynamically bound each leak hyper parameter based on the observed values, y^*, and NOM predictive mean, $\hat{\mu^*}$, and variance, $\hat{\sigma^2_*}$. The leak starting time parameter, τ, is bounded by our data modelling window of one week. For observations $y^* = [S^*, F^*]$ of the pump speed and flow, at input times $x^* = [t^*, l]$, the deviation caused by a leak will always be between $\hat{\mu^*}$ and y^*. The deviation cause by a leak in both the pump speed and flow is always positive and can not be greater than their observed values. Thus, at a particular, t, in the data window, if the residuals of either of the modelled time series are less than zero i.e., $S^*_t - \hat{\mu^*_{S_t}}, < 0$ or $F^*_t - \hat{\mu^*_{F_t}}, < 0$, we assume this to be an indication that the leak did not start at that particular time i.e, $\tau \neq t^*$, and search for other leak hyper parameters, θ_S, C and α, at that particular time can be pruned. This search procedure has two advantages, firstly it reduces the leak parameter search space considerably, secondly the leak start time parameter, τ, in this case is based on evidence from both the pump speed and the flow. Given this, if there is a positive deviation in both the modelled time series we search for the optimal values of the the remaining leak parameters, θ_S, C and α, based on dynamic bounds placed on each of them.

Elaborating on this, we know that the leak coefficient, C, is always positive (as a negative leak coefficient would result in a negative value for the leak flow). Similarly, we also know that α is always positive and typically in the range of $0.4 - 2.3$ (Greyvenstein and Van Zyl, 2007). In light of this, we set the lower bound for both C and α as 0.01. We believe this to be a suitable choice, as when both parameters are at their minimum value the resulting leak would be approximately 0.01 l/s even for a very high pressure value of 1000 m head ($0.01 \times 1000^{0.01} = 0.01$). A leak value lower than this, would produce negligible effects on the over DMA-level hydraulic data, nor would it be significant enough to warrant any action from water companies.

Given this, for a particular time, t, and the flow residuals, $R^*_t = F^*_t - \hat{\mu^*_{F_t}}$, we compute the max values C and α can take, using the orifice flow equation, $q_{leak} = C \times P^{\alpha}$, and

bound them as follows:

$$0.01 < C < C_{max} = \frac{R_t^*}{(P_t^*)^{\alpha=0.01}} \tag{5.10}$$

$$0.01 < \alpha < \alpha_{max} = \frac{\log R_t^* - \log(C = 0.01)}{\log P_t^*} \tag{5.11}$$

where P^* is the average zone pressure at input time x^*. We then discretise the possible values of C and α, by sampling ten consecutive equidistant values between the upper and lower bounds.

Let $\theta_{\mathbf{nom}} = [c_P, c_F, \rho, \sigma^2_{f_M}, l_M, \sigma^2_{f_P}, l^2_P, \sigma^2_n]$
Set $\theta_{\mathbf{L}} = [\tau, \theta_S, C, \alpha] = \mathbf{0}$
for *each new observation* $y = [S, F]$ *at time*$= t$ **do**
  #Make Predictions
  Compute $\hat{\mu^*} = [\hat{\mu^*_S}, \hat{\mu^*_F}]$, and $\hat{\sigma^2_*}$ #GP NOM Predictions
  Set $R^*[R^*_S, R^*_F] = y[S, F] - \hat{\mu^*}[\hat{\mu^*_S}, \hat{\mu^*_F}]$ #Compute Residuals
  Set $Mps = -\infty$ #Max Probability sum
  Set $C_{min} = 0.01$ #Min bound for C
  Set $\alpha_{min} = 0.01$ #Min bound for α
  Set $\theta_{L_{Opt}} = [\tau, \theta_S, C, \alpha] = \mathbf{0}$ #Vector to record optimal leak hyper-parameters
  #For each reading in the 1 week data window (n=672 readings)
  for $t = 1$ *to* n **do**
    if $(R^*_F(t) > 0)$ **then**
      Set $\tau = t$
      #Compute upper bounds for leak hyper parameters:
      Set $C_{max} = \frac{R^*_F(t)}{(P^*(t))^{\alpha_{min}}}$
      Set $\alpha_{max} = \frac{\log R^*_F(t) - \log(C_{min})}{\log P^*(t)}$

      #Sample 10 equidistant values between upper and lower bounds for C, α:
      Sample **Cvec** $\sim$ Uniform$([C_{min}, C_{max}])$ #Vector of 10 C values
      Sample α**Vec** $\sim$ Uniform$([\alpha_{min}, \alpha_{max}])$ #Vector of 10 α values
      for $j = 1$ *to* $length$(**Cvec**) **do**
        for $k = 1$ *to* $length(\alpha$ ***Vec***$)$ **do**
          Set $\theta_L = [\tau, \theta_S, C, \alpha] = [t, R^*_S(t), CValues(j), alpha\text{Values}(k)]$
          #Compute sum of the log posterior predictive probability (Lps) of each leak:
          Compute $Lps = \sum_{t=1}^{n} -\frac{1}{2}\log\left(2\pi\hat{\sigma^2_*}(t)\right) - \frac{\left(y(t) - \hat{\mu^*}(t)\right)^2}{2\hat{\sigma^2_*}(t)}$
          if $(Lps > Mps)$ **then**
            Set $\theta_{L_{Opt}} = \theta_L$
            Set $Mps = Lps$
          end
        end
      end
    end
  end
end

Algorithm 3: SMOGP leak detection

Similarly, in case of a leak, the deviation in the pump speed cannot be greater than the current observed value of the pump speed or less than the NOM prediction. An observed speed lower than the predictions would indicate that the pump is operating at a reduced capacity which is in contradiction to the pump behaviour under leak conditions, suggesting that there is no leak in the system. In light of this for time, t, we set $\theta_S = S_t^* - \hat{\mu}_{S_t}^*$ i.e, the residual. These steps are detailed in Algorithm 3.

Given the above, starting from time $t = 1$, in the input data window, the residuals of both the modelled time series are computed based on the SMOGP NOM predictions and the observed values. If both residuals are positive, the bounds for the leak parameters are computed using equations 5.10 and 5.11 and the leak starting time parameter, τ, is set to equal to t. Based on the bounds, the possible values of C and α are discretised. Using the discretised values, the sum of log predictive probabilities for each combination of C and α is computed using equation 5.9. The combination that result in the highest sum of log predictive probabilities is selected as the optimal parameter set, and the highest probability sum is recorded as the maximum probability sum. The above process is repeated for subsequent times, $t^* = 2$ to n. The optimal parameter set is updated, with a new one from a subsequent time t^*, only if the new one yields a probability sum greater than the maximum probability sum. In which case the leak starting time parameter, τ, is set to t^* and the maximum probability sum is set to the probability sum yielded by the new set. When the whole data window has been exhausted, the posterior predictive probability using the final optimal leak parameters are computed using 5.9. The final probability of a leak, for each time t in the data window, is then given as the normalised difference between the SMOGP NOM predictive probabilities and the posterior predictive probability using the final optimal leak parameters.

Having defined the hyper parameter learning schemes, in the next section we will briefly describe our simulated data set used for performance evaluation of the proposed approach.

5.4 Simulated Leak Dataset

In line with the methodology found in literature and previously discussed in Section 3.4, we evaluate the performance of the proposed method on leaks of varying sizes simulated at different location within the DMA. In particular, using the pressure optimised DMA modelled in Chapter 4 we simulate five weeks of data. In the fifth week, at different times, we simulate leaks at two different locations in the DMA. For the first location we simulate a leak in pipe, P52, close to the supply point, while for the second location we simulate a leak close to the critical point, in pipe P68. Figure 5.3 shows the selected locations in the DMA. At each location we simulate three different leaks, where we vary

the leak magnitude using leak coefficients values of 0.2, 0.3 and 0.4. For all the leak simulations we use a leak exponent of 0.5.

To evaluate the performance of the proposed SMOGP, we compare it against both the NSGP and the selected approaches listed in Section 2.6. For the evaluation we use the same performance metrics we used earlier in the comparative analysis for the NSGP (listed in Section 3.4).

For the SMOGP and FB algorithms we use a sum of a matérn, k_M, periodic, k_P, and a gaussian noise kernel, k_N, i.e, $K_T(t,t') = k_M(t,t') + k_P(t,t') + k_N$. We use the same optimised hyper parameters values for both algorithms. For the NSGP we use the diagonal noise covariance kernel in conjunction with the weekly and step change mean functions (described in Section 3.1).

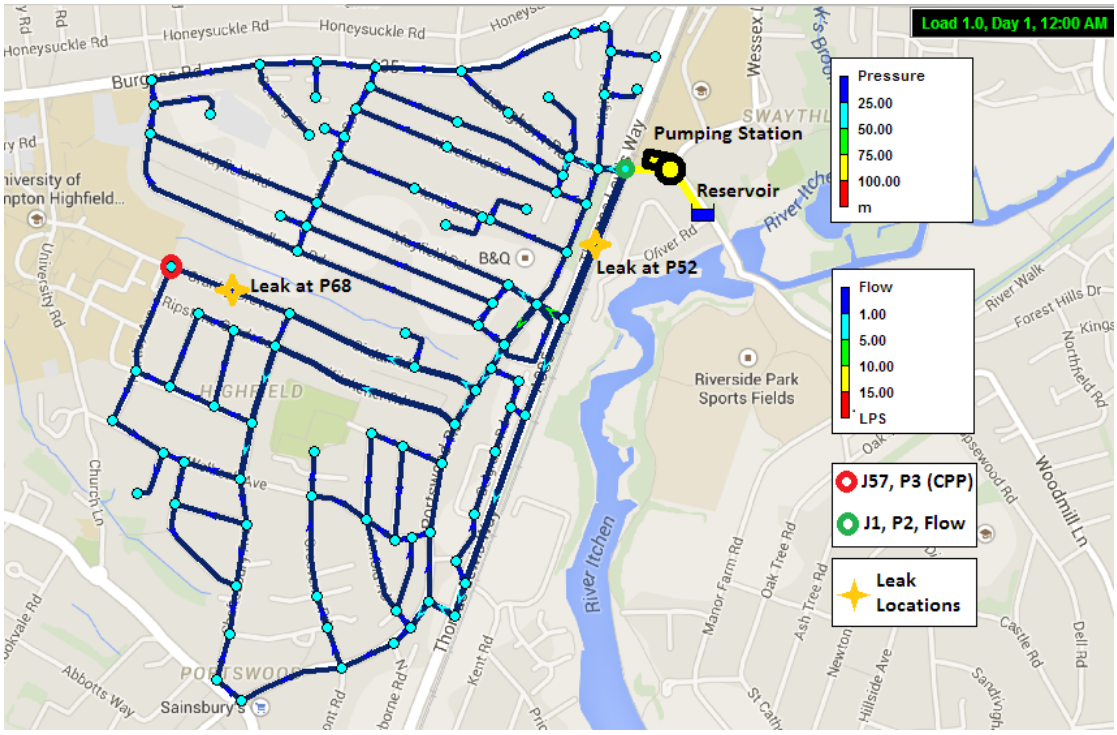

FIGURE 5.1: The modelled DMA overlaid on google maps. The supply pressure and flow points, P2 and Flow, at junction J1 are marked with a green circle, where as the critical pressure point, P3, at junction, J57, is marked with a red circle. The yellow stars indicate the two pipes where leaks of different magnitudes are simulated. At each selected pipe three leaks with leak coefficients of $0.2, 0.3$ and 0.4 are simulated.

5.5 Results And Discussion

Our main objectives are to analyse that performance of the proposed SMOGP, in comparison with other approaches. In doing so we wish to see how well each approach performs on leaks of different magnitudes. Furthermore, how well does each method perform as more and more leaky data becomes available i.e., as the available evidence of

a leak increases. Thus to achieve this, we discuss two cases in detail, where as the detection results of all methods on the six simulated leaks are listed in Appendix A. The first case we select is a leak simulated on the fifth day at 8:30 AM, in pipe P68 close to the critical point, with a leak coefficient of 0.2 and a leak exponent of 0.5. The second case we will discuss is a leak simulated in pipe P52 close to the supply point, on the sixth day at 20:00 PM. Given this, we also wish to see the effectiveness of the proposed pressure dependant leak detection mean function. In particular, for the SMOGP, we expect that modelling orifice flow based leak equations in the mean function would result in better prediction of the variations in the leak due to demand based pressure changes. We also expect that using the pump speed deviations in conjunction with the flow deviations to detect leaks would result in a higher detection accuracy and lower false positive rate in comparison to the NSGP and other approaches. Finally, we also wish to analyse the accuracy of the SMOGP estimated leak coefficients and leak exponents for each leak. For this analysis, all algorithms were executed on a Mac book Pro 2.5 Ghz computer with 16 GB random access memory and the execution time, in seconds, for each algorithm is shown in the results table for each case.

Figure 5.2 shows the detected leaks by all approaches for the first case. The detected leak by the proposed SMOGP follows the follows the same patterns as the actual simulated one resulting in a very low RMSE. Against the simulated leak coefficient and exponent of 0.2 and 0.5 (both unitless), the SMOGP estimates values were found to be 0.2369 and 0.1725 respectively. It must be noted that the simulated coefficient value are based on the pressure supplied at pipe P68, where as the SMOGP estimates are based on the average zone pressure. Thus, the best way verify these values is to simulate the average zone pressure at pipe P68, then set the leak coefficient and exponents to the values estimated by the SMOGP, and then compare the resultant leak with the one predicted by the SMOGP. However, such an analysis is not possible with WaterNetGen.

	Reporting Time (Min)	**TPR**	**FPR**	**ACC**	**F1**	**RMSE**	**Execution Time (secs)**
NL	2 days	**1.000**	**0.000**	**1.000**	**1.000**	**0.009**	**0.010**
NKF	15 mins	0.956	0.423	0.667	0.576	0.284	**0.024**
FB	15 mins	0.157	0.078	0.741	0.223	0.208	5.530
NSGP	15 mins	**1.000**	**0.004**	**0.997**	**0.994**	0.064	0.775
SMOGP	15 mins	**1.000**	**0.004**	**0.997**	**0.994**	**0.028**	24.838

TABLE 5.1: Detection results of all approaches for a leak simulated in pipe P68 on the 5th day at 8:30 AM with a leak coefficient of 0.20.

Given this, looking at Table 5.1 the overall performance of the NSGP and SMOGP, apart from the leak magnitude estimation is the same. The SMOGP is computationally the most expensive. Both NKF and FB leak predictions are irregular, with false negatives within the leak duration and false positives before the leak starting time. As stated

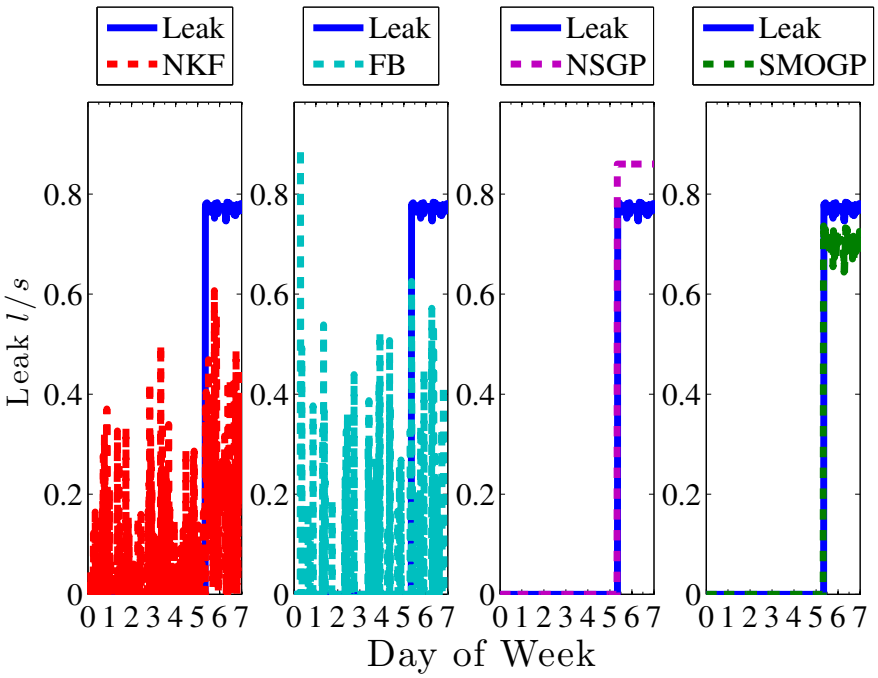

FIGURE 5.2: Detection results of all approaches for a leak simulated in pipe P68 on the 5th day at 8:30 AM with a leak coefficient of 0.20.

earlier, this is due to the fact that these methods do point estimation of leaks, thus not taking in to consideration their sustained nature.

The detection results for the second case are presented in Table 5.2 where the performance of the SMOGP is superior to all other methods. In this case only three hours of leaky readings are available for all algorithms to detect and quantify leaks. In contrast to the NKF and FB methods both NSGP and SMOGP methods show good results. The SMOGP leak magnitude estimation is better than all other methods, except the NL, with the lowest RMSE error of 0.047 and zero false positive rate. In this case against the simulated leak coefficient and exponent of 0.4 and 0.5 the SMOGP estimate are found to be 0.0100 and 1.0626. This offset in estimation is due to the fact that there are only four hours of leaky readings, in this case most of the one week data window will not be searched. Thus, their leak parameter optimisation will not have enough examples to find the optimal value for the leak coefficient and exponent.

Having detailed the results of the experiments the next section summarises the findings of this chapter.

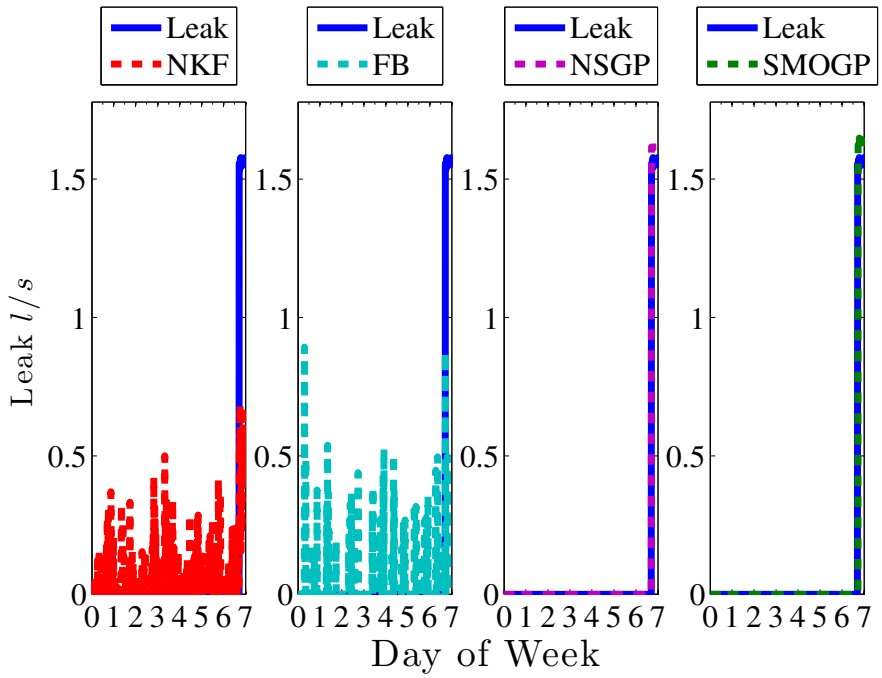

FIGURE 5.3: Detection results of all approaches for a leak simulated in pipe P52 on the 6th day at 20:00 PM with a leak coefficient of 0.40.

	Reporting Time (Min)	**TPR**	**FPR**	**ACC**	**F1**	**RMSE**	**Execution Time (secs)**
NL	2 days	**1.000**	**0.000**	**1.000**	**1.000**	**0.000**	**0.011**
NKF	15 mins	**1.000**	0.451	0.568	0.167	0.247	**0.025**
FB	15 mins	0.172	0.086	0.882	0.112	0.190	5.621
NSGP	15 mins	0.862	**0.000**	0.994	0.926	0.120	0.464
SMOGP	15 mins	0.931	**0.000**	**0.997**	**0.964**	**0.047**	18.979

TABLE 5.2: Detection results of all approaches for a leak simulated in pipe P52 on the 6th day at 20:00 PM with a leak coefficient of 0.40.

5.6 Summary

In this chapter we began by proposing a multi output GP based method to model the DMA-level flow and pump speed. We detailed how the two time series and their correlations can be modelled using labeled inputs. In particular we highlighted that in case of linearly correlated inputs the cross correlations in a multi output GP can be modelled as a Hadamard product of covariance over time and a covariance over time series labels. We then used this formulation to define the covariance structure for NOM of a pressure optimised DMA. Following this we described the parametric, pressure dependant leak mean function, in which we modelled leaks using the flow through an orifice model. Next, we described how the two sets of hyper parameters i.e., the NOM and leak hyper

parameter are learnt. We described in detail the dual optimisation mechanism used. Here we described how we define dynamic bound for the leak hyper parameters based on the observed hydraulic readings and the SMOGP predictions, at each step of the leak detection process. Using the pressure optimised DMA modelled using WaterNetGen, in the previous chapter, we generated six leaks in two distinct pipes in the modelled DMA. Using this data we evaluated the performance of the proposed SMOGP against both the state of the art leak detection methods and the previously proposed NSGP. The results indicated that pressure dependant leak modelling in the SMOGP produced better estimates of the leak magnitude. Although, overall the SMOGP showed the best detection accuracy among all real time algorithms (excluding the NL), the additional computational costs of the SMOGP warrants further improvement, particularly in the search mechanism used to learn the leak coefficient and exponent parameters.

Chapter 6

Conclusions, Limitations and Future Work

This thesis has described two efficient leak detection approachs for pressure optimised DMAs using only DMA-level hydraulic readings. We now summarise the contributions of this work, highlight its limitations and give directions for future work.

6.1 Conclusions

We first defined the problem of leak detection in pressure optimised DMAs in Chapter 1. We identified the key requirements that must be fulfilled in order to realise a realistic solution to this problem. The requirements stated that the solution must be able to detect and quantify leaks in real-time. Most importantly, the solution must not require any additional data other than the DMA-level hydraulic measurements available from the already installed sensors. This requirement is crucial for any leak detection approach to be accepted and adopted by the industry.

In Chapter 2, we then provided a background of existing work in the area of leak detection. We categorised leak detection approaches from literature in to three main categories, hardware based and hydraulic model based and hydraulic measurement based methods. With in each category we discussed the advantages and disadvantages of the various types of approaches used for leak detection. We showed that hardware based methods are invasive and very costly, both in terms of the hardware costs and the manual labour required to effectively install and maintain them. Within the category of hydraulic model based methods, we showed that methods that require monitoring of pressure and flow at every node in the network are financially impractical for water companies. In hydraulic measurement based methods we highlighted that statistical and AI based methods that use only DMA-level hydraulic measurements for leak detection

are best suited to our research requirements. However, each existing approach in this category had discrepancies. In particular, all of these approaches used only one of the DMA-level data streams to detect leaks, ignoring the leak evidence that can be gathered from the other data streams. Furthermore, all of these approaches assume the dynamics of unoptimised DMAs where a leak results in the DMA pressure to decrease. None of these approaches have been tested in active pressure optimised DMA setting where the DMA dynamics change under active pressure optimisation

Chapter 3, presents the first major contribution of this thesis, in which we gave a description of the proposed noise scaled semi-parametric gaussian process (NSGP) model for leak detection using only DMA-level flow readings. We described the how the historic DMA flow measurements can be used to from a NOM model in GP where the proposed diagonal noise covariance can efficiently and effectively model the input dependant noise. We then described the proposed step change mean function parametrisation that captures the leak magnitude and start time. Following this we described the efficient dual optimisation process used to learn the NOM and leak hyper parameters. Next, we gave a brief description the selected leak detection approaches from literature that we use to evaluate the performance of the proposed NSGP. Following this, we explain the experimental setup and performance metrics used in the comparative analysis. In the end we described the results of the comparative analysis which highlighted the superior performance of the proposed method. We concluded the chapter highlighting the need for a more accurate close to real life pressure optimised DMA dataset with verifiable leaks.

Chapter 4, represent the second contribution of this thesis, in which we analyse the dynamics of pressure active optimised DMAs under leak conditions. We start the chapter by detailing the challenges faced in acquiring usable truth data for leaks from actual DMAs. We highlight that poor record keeping, human errors and lack of leak detection, reporting and recording mechanisms as the main reason why such data is not available. We then provided a brief description of EPANET software, its limitation and a brief description of WaterNetGen EPANET extension that solves these limitations. In the next section we gave a detailed description on the mechanism used to model a close to real life DMA in WaterNetGen. This was followed by a description of the methodology used to implement active pressure optimisation used rule based control. We showed that the implemented rule based optimisation greatly reduced the DMA pressures. We then described the pressure dependant leak model used to generate synthetic leak events in the modelled DMA. Next with simulated examples we showed that the dynamics of a pressure optimised DMA under leak conditions are different from that of an unoptimised DMA. We also highlighted the implications of pressure optimisation on any leak detection solution.

Chapter 5, represents the major and final contribution of the thesis. We introduce the novel semi parametric multi output Gaussian process model (SMOGP) for leak detection and quantification in active pressure optimised DMA. First, we show how we

can efficiently model multiple hydraulic data streams in a GP by using labeled inputs and modelling the cross correlations as a linear function of the temporal covariance. We detail how this is used to construct a NOM model for the DMA. Following this, we detail the proposed implementation of the pressure dependant leak mean function and how we incorporate the orifice flow model of leaks in the GP framework as a parametric mean function. Next, we detail our novel dual optimisation mechanism for the NOM and leak hyper parameter learning. For leak detection we detail how we dynamically define bounds for the leak hyper parameters based on the observed hydraulic readings and the SMOGP predictions. We then detail the experimental setup and simulated DMA data used to compare the performance of the proposed SMOGP against both the current state of the art leak detection approaches and the proposed NSGP. We end with a discussion on the results of the analysis which highlight the superior performance of the proposed approach.

6.2 Limitations and Future Work

Although, the proposed methods have been successfully applied and evaluated in active pressure optimised DMA scenarios as per the research requirements. There are a number of areas and factors that leave considerable room for improvement. We categorise them in to the following areas:

NSGP Limitations

The most important limitation of the proposed NSGP is approximation of leak as a step change in the DMA-level flow. Although in most scenarios a step change leak estimation would provide good results both in terms of detection accuracy and approximate estimation of the leak magnitude. In case of leaks that increase rapidly, such as cluster erosions of metallic pipes, the shape of the leak cannot be approximated efficiently approximated with a simple step change. In such scenarios the NSGP leak detection and quantification would be compromised.

Limitations of the Simulated DMA and WaterNetGen

Although, different leaks have been simulated in WaterNetGen for a pressure optimised DMA. Certain scenarios could not be simulated owing to the limitations of the WaterNetGen. In particular, if the simulated leak is large enough to produce a pressure drop in the system that surpasses the pumps head providing capacity WaterNetGen aborts the simulation. Additionally, the only way to simulate leaks in WaterNetGen is by using the provided flow through an orifice model at pipe level. The model does not allow for

the leak coefficient or exponent to be altered during the simulation, which would allow to model leaks that increase over time.

SMOGP Limitations

Although the current bounded search for leak hyper parameters produces good results, the computation time for the proposed SMOGP can be greatly improved by investigating more efficient search mechanisms. Also, given the limitations of the simulated DMA and leaks in WaterNetGen, the proposed method needs to be evaluated in more realistic scenarios, particularly in scenarios where the leaks increase over time. However, this would require a WDS simulation software that allows changing the leak coefficient values at runtime i.e. during the simulation. Another way to get data for increasing leaks would be by conducting engineering tests, where increasing leaks can be simulated by opening valves or hydrants in increments over a predetermined duration of time.

Although the proposed algorithms donot address leak localisation, however the estimates of the leak exponent values from the SMOGP can be used to approximate the material of the pipe in which the leak occurred. In cases where there is sufficient information about the deployed pipes in the DMA e.g. for new DMAs, the estimated leak exponent value from the SMOGP can be used to focus on areas of the DMA where the desired pipe types are installed.

Future Work

Given the above the, one area of focus for future work will be investigation of more efficient search mechanisms to reduce computation cost of the proposed SMOGP. Another interesting area of future work is extending the proposed approach to learn correlations between multiple DMAs to verify leaks and further reduce false positives. As one of main concerns in leak detection techniques is decreasing the false positive leak predictions. As previously mentioned, a sudden increase in the water consumption may not always be a result of a burst or leak e.g. increase in demand for a few hours due to an anticipated football game. In such scenarios a leak detection algorithm will flag this demand increase as leak in a DMA. However, if a similar demand increase pattern is observed in multiple neighbouring DMA's at the same time then it can safely be attributed to unprecedented demand change, since the probability of multiple DMAs having a leak at the same time is very low.

Real-time smart water meters provide another way to distinguish between legitimate demand increase and leaks. Indeed, if the flow/pressure at each end node in a DMA is known then the change in flow/pressure across nodes can be used to not only detect and quantify but also locate leaks. In the scenario presented earlier, a short term

increase in demand in case of a festival or anticipated football game would be visible as a demand increase in all end nodes. In contrast a leak would only affect the end nodes connected to the pipe in which the leak occurs. Thus, for a leak the change in flow/pressure between two end nodes can be used to detect, quantify and locate leaks within DMA. As mentioned in Section 1.2, smart autonomous, real time metering for WDS is an active area of research (Marvin et al., 1999; Gurung et al., 2015; Kashi, 2016). However, various cost benefit studies on smart meter technology have highlighted the large investment costs of installing smart meters as a critical barrier to their industry wide adaption (SWAN, 2010; McHenry, 2013). It must be noted that making smart meters technology affordable is an active area of research (Molina-Markham et al., 2012; Batista et al., 2013) and with companies like WavIot and Libelium already investing heavily in this area. In particular, WavIot has recently developed a cheap real-time smart water meter ($35) capable of sending readings over distances as large as 16km (using LoraWan technology) while providing a deployment battery life of 10 years. There is a deployment study underway to test the proposed solution (Link to study). In light of these advancements, the future of leak detection and management in WDS will be shaped heavily by smart water metering technology.

Appendix A

Detection Result Tables

This appendix provides the leak detection of all approaches for the six simulated leaks from Chapter 5.

	Reporting Time (Min)	**TPR**	**FPR**	**ACC**	**F1**	**RMSE**	**Execution Time (secs)**
NL	2 days	0.500	**0.000**	0.964	0.667	**0.035**	**0.004**
NKF	15 mins	**1.000**	0.416	0.695	0.636	0.591	**0.020**
FB	15 mins	0.168	0.077	0.722	0.243	0.414	5.432
NSGP	15 mins	**1.000**	0.004	0.997	0.994	0.086	0.976
SMOGP	15 mins	**1.000**	**0.000**	**1.000**	**1.000**	**0.043**	27.559

Table A.1: Detection results of all approaches for a leak simulated in pipe P68 with a leak coefficient of 0.40

	Reporting Time (Min)	**TPR**	**FPR**	**ACC**	**F1**	**RMSE**	**Execution Time (secs)**
NL	2 days	**1.000**	**0.000**	**1.000**	**1.000**	**0.008**	**0.002**
NKF	15 mins	0.993	0.434	0.652	0.534	0.391	**0.021**
FB	15 mins	0.170	0.078	0.771	0.230	0.277	5.568
NSGP	15 mins	**1.000**	0.004	0.997	0.993	0.068	0.838
SMOGP	15 mins	**1.000**	**0.000**	**1.000**	**1.000**	**0.029**	25.189

Table A.2: Detection results of all approaches for a leak simulated in pipe P68 with a leak coefficient of 0.30.

	Reporting Time (Min)	**TPR**	**FPR**	**ACC**	**F1**	**RMSE**	**Execution Time (secs)**
NL	2 days	**1.000**	**0.000**	**1.000**	**1.000**	**0.009**	**0.010**
NKF	15 mins	0.956	0.423	0.667	0.576	0.284	**0.024**
FB	15 mins	0.157	0.078	0.741	0.223	0.208	5.530
NSGP	15 mins	**1.000**	**0.004**	**0.997**	**0.994**	0.064	0.775
SMOGP	15 mins	**1.000**	**0.004**	**0.997**	**0.994**	**0.028**	24.838

TABLE A.3: Detection results of all approaches for a leak simulated in pipe P68 with a leak coefficient of 0.20.

	Reporting Time (Min)	**TPR**	**FPR**	**ACC**	**F1**	**RMSE**	**Execution Time (secs)**
NL	2 days	**1.000**	**0.000**	**1.000**	**1.000**	**0.000**	**0.011**
NKF	15 mins	**1.000**	0.451	0.568	0.167	0.247	**0.025**
FB	15 mins	0.172	0.086	0.882	0.112	0.190	5.621
NSGP	15 mins	0.862	**0.000**	0.994	0.926	0.120	0.464
SMOGP	15 mins	0.931	**0.000**	**0.997**	**0.964**	**0.047**	18.979

TABLE A.4: Detection results of all approaches for a leak simulated in pipe P52 with a leak coefficient of 0.40.

	Reporting Time (Min)	**TPR**	**FPR**	**ACC**	**F1**	**RMSE**	**Execution Time (secs)**
NL	2 days	**1.000**	**0.000**	**1.000**	**1.000**	**0.000**	**0.002**
NKF	15 mins	**1.000**	0.457	0.552	0.080	0.160	**0.019**
FB	15 mins	0.154	0.088	0.897	0.055	0.137	5.281
NSGP	15 mins	0.846	**0.000**	0.997	0.917	0.075	0.399
SMOGP	15 mins	**1.000**	**0.000**	**1.000**	**1.000**	**0.009**	18.083

TABLE A.5: Detection results of all approaches for a leak simulated in pipe P52 with a leak coefficient of 0.30.

	Reporting Time (Min)	**TPR**	**FPR**	**ACC**	**F1**	**RMSE**	**Execution Time (secs)**
NL	2 days	0.500	**0.000**	0.964	0.667	**0.035**	**0.004**
NKF	15 mins	**1.000**	0.416	0.695	0.636	0.591	**0.020**
FB	15 mins	0.168	0.077	0.722	0.243	0.414	5.432
NSGP	15 mins	**1.000**	0.004	0.997	0.994	0.086	0.976
SMOGP	15 mins	**1.000**	**0.000**	**1.000**	**1.000**	**0.043**	27.559

TABLE A.6: Detection results of all approaches for a leak simulated in pipe P52 with a leak coefficient of 0.40.

Bibliography

Al-Dhowalia, K. and Shammas, N. (1991). Leak detection and quantification of losses in a water network. *International Journal of Water Resources Development*, 7(1):30–38.

Alexander, M. T. and Boccelli, D. L. (2010). Field verification of an integrated hydraulic and multi-species water quality model. In *Water Distribution Systems Analysis 2010*, pages 687–696. ASCE.

Alvarez, M. and Lawrence, N. D. (2009). Sparse convolved gaussian processes for multi-output regression. In *Advances in neural information processing systems*, pages 57–64.

Batista, N., Melício, R., Matias, J., and Catalão, J. (2013). Photovoltaic and wind energy systems monitoring and building/home energy management using zigbee devices within a smart grid. *Energy*, 49:306–315.

Bieupoude, P., Azoumah, Y., and Neveu, P. (2012). Optimization of drinking water distribution networks: Computer-based methods and constructal design. *Computers, Environment and Urban Systems*, 36(5):434 – 444.

Boyle, P. and Frean, M. (2005). Dependent gaussian processes. *Advances in neural information processing systems*, 17:217–224.

BREHM, M.-A., Scholtes, V. A., Dallmeijer, A. J., Twisk, J. W., and Harlaar, J. (2012). The importance of addressing heteroscedasticity in the reliability analysis of ratio-scaled variables: an example based on walking energy-cost measurements. *Developmental Medicine & Child Neurology*, 54(3):267–273.

Buchberger, S. and Nadimpalli, G. (2004). Leak estimation in water distribution systems by statistical analysis of flow readings. *Journal of Water resources Planning and Management*, 130(4):321–329.

Cole, E. S. et al. (1979). Methods of leak detection: An overview (pdf). *Journal-American Water Works Association*, 71(2):73–75.

Colombo, A. F., Lee, P., and Karney, B. W. (2009). A selective literature review of transient-based leak detection methods. *Journal of Hydro-environment Research*, 2(4):212–227.

Covas, D., Ramos, H., Young, A., Graham, N., and Maksimovic, C. (2005). Uncertainties of leak detection by means of hydraulic transients: From the lab to the field. In *Proc., 8th Int. Conf. on Computing and Control for the Water Industry (CCWI 2005)*, pages 143–148. Univ. of Exeter Exeter, UK.

De Silva, D., Mashford, J., and Burn, S. (2009). *Computer Aided Leak Location and Sizing in Pipe Networks*. Urban Water Security Research Alliance.

De Silva, D., Mashford, J., Burn, S., Alliance., U. W. S. R., for a Healthy Country Flagship (Program), W., University., G., of Queensland., U., and CSIRO. (2011). *Computer aided leak location and sizing in pipe networks / Dhammika De Silva, John Mashford and Stewart Burn*. Urban Water Security Research Alliance, City East, Qld. :, 2nd ed. edition.

Fahmy, M. and Moselhi, O. (2009). Detecting and locating leaks in underground water mains using thermography. In *Proceedings of the 26th International Symposium on Automation and Robotics in Construction (ISARC?09)*, pages 61–67.

Farley, M. (2008). Finding the ?difficult? leaks. *International Water Association Specialist Group-Efficient Operation and Management*.

Fletcher, R. and Chandrasekaran, M. (2008). Smartball?: A new approach in pipeline leak detection. In *2008 7th International Pipeline Conference*, pages 117–133. American Society of Mechanical Engineers.

Garnett, R., Osborne, M. A., Reece, S., Rogers, A., and Roberts, S. J. (2010). Sequential bayesian prediction in the presence of changepoints and faults. *The Computer Journal*, 53(9):1430–1446.

Germanopoulos, G. and Jowitt, P. (1989). Leakage reduction by excess pressure minimisation in a water supply network. In *ICE Proceedings*, volume 87, pages 195–214. Thomas Telford.

Gertler, J., Romera, J., Puig, V., and Quevedo, J. (2010). Leak detection and isolation in water distribution networks using principal component analysis and structured residuals. In *Control and Fault-Tolerant Systems (SysTol), 2010 Conference on*, pages 191 –196.

Goldberg, P. W., Williams, C. K., and Bishop, C. M. (1997). Regression with input-dependent noise: A gaussian process treatment. *Advances in neural information processing systems*, 10:493–499.

Greyvenstein, B. and Van Zyl, J. (2007). An experimental investigation into the pressure-leakage relationship of some failed water pipes. *Journal of Water Supply: Research and Technology-AQUA*, 56(2):117–124.

Gurung, T. R., Stewart, R. A., Beal, C. D., and Sharma, A. K. (2015). Smart meter enabled water end-use demand data: platform for the enhanced infrastructure planning of contemporary urban water supply networks. *Journal of Cleaner Production*, 87:642–654.

Hunaidi, O. (2000). *Detecting leaks in water distribution pipes*, volume 40. Institute for Research in Construction, National Research Council of Canada.

Hunaidi, O., Chu, W., Wang, A., and Guan, W. (2000). Leak detection for plastic water distribution pipes. *Journal AWWA*, 92(2):82–94.

Hunaidi, O. and Wang, A. (2006). A new system for locating leaks in urban water distribution pipes. *Management of Environmental Quality: An International Journal*, 17(4):450–466.

Ingeduld, P., Svitak, Z., Pradhan, A., and Tarai, A. (2006). Modelling intermittent water supply systems with epanet. In *8th annual water distribution systems analysis symposium, Cincinnati*, pages 27–30.

IWA (2000). Manual for best practice: performance indicators for water supply services. *International Water Association, London, UK*.

Kashi, H. (2016). Smart energy and water meter (sewm): An innovative approach towards groundwater monitoring and management. In *2016 NGWA Groundwater Summit*. Ngwa.

Kersting, K., Plagemann, C., Pfaff, P., and Burgard, W. (2007). Most likely heteroscedastic gaussian process regression. In *Proceedings of the 24th international conference on Machine learning*, pages 393–400.

Khan, A., Widdop, P. D., Day, A. J., Wood, A. S., Steve, Mounce, R., and Machell, J. (2008). Artificial neural network model for a low cost failure sensor: Performance assessment in pipeline distribution. *International Journal of Mathematical, Computational, Physical, Electrical and Computer Engineering*, 2(9):690 – 696.

Kuss, M. and Rasmussen, C. E. (2005). Assessing approximations for gaussian process classification. In *Advances in Neural Information Processing Systems*, pages 699–706.

Lambert, A. (2001). What do we know about pressure-leakage relationships in distribution systems. In *IWA Conf. n Systems approach to leakage control and water distribution system management*. Citeseer.

Lazaro-Gredilla, M. and Titsias, M. (2011). Variational heteroscedastic gaussian process regression. In Getoor, L. and Scheffer, T., editors, *Proceedings of the 28th International Conference on Machine Learning*, pages 841–848, New York, NY, USA.

Lee, P. J., Lambert, M. F., Simpson, A. R., Vítkovskỳ, J. P., and Liggett, J. (2006). Experimental verification of the frequency response method for pipeline leak detection. *Journal of Hydraulic research*, 44(5):693–707.

Liggett, J. and Chen, L. (1994). Inverse transient analysis in pipe networks. *Journal of Hydraulic Engineering*, 120(8):934–955.

Malik, O., Ghosh, S., and Rogers, A. (2015). A noise scaled semi parametric gaussian process model for real time water network leak detection in the presence of heteroscedasticity. In *Workshops at the Twenty-Ninth AAAI Conference on Artificial Intelligence*.

Marvin, S., Chappells, H., and Guy, S. (1999). Pathways of smart metering development: shaping environmental innovation. *Computers, Environment and Urban Systems*, 23(2):109–126.

Mashford, J., De Silva, D., Marney, D., and Burn, S. (2009). An approach to leak detection in pipe networks using analysis of monitored pressure values by support vector machine. In *Network and System Security, 2009. NSS'09. Third International Conference on*, pages 534–539. IEEE.

McHenry, M. P. (2013). Technical and governance considerations for advanced metering infrastructure/smart meters: Technology, security, uncertainty, costs, benefits, and risks. *Energy Policy*, 59:834–842.

Mehra, R. (1970). On the identification of variances and adaptive kalman filtering. *Automatic Control, IEEE Transactions on*, 15(2):175–184.

Misiunas, D., Lambert, M., Simpson, A., and Olsson, G. (2005). Burst detection and location in water distribution networks. *Water Science and Technology: Water Supply*, 5(3-4):71–78.

Misiunas, D., Vitkovsky, J., Olsson, G., Simpson, A., and Lambert, M. (2003). Pipeline burst detection and location using a continuous monitoring technique. In *Advances in Water Supply Management: Int. Conf. on Computing and Control for the Water Industry (CCWI)*, pages 89–96.

Molina-Markham, A., Danezis, G., Fu, K., Shenoy, P., and Irwin, D. (2012). Designing privacy-preserving smart meters with low-cost microcontrollers. In *International Conference on Financial Cryptography and Data Security*, pages 239–253. Springer.

Mounce, S., Boxall, J., and Machell, J. (2007). An artificial neural network/fuzzy logic system for dma flow meter data analysis providing burst identification and size estimation. *Water Management Challenges in Global Change*, pages 313–320.

Mounce, S., Day, A., Wood, A., Khan, A., Widdop, P., and Machell, J. (2002). A neural network approach to burst detection. *Water Ccience and Technology: A Journal of the International Association on Water Pollution Research*, 45(4-5):237.

Mukherjee, J. and Narasimhan, S. (1996). Leak detection in networks of pipelines by the generalized likelihood ratio method. *Industrial & Engineering Chemistry Research*, 35(6):1886–1893.

Muncke (2011). Burst detection by nightline analysis. *I2O Water Internal company report*.

Muranho, J., Ferreira, A., Sousa, J., Gomes, A., and Marques, A. S. (2012). Waternetgen: an epanet extension for automatic water distribution network models generation and pipe sizing. *Water Science and Technology: Water Supply*, 12(1):117–123.

Muranho, J., Ferreira, A., Sousa, J., Gomes, A., and Marques, A. S. (2014). Pressure-dependent demand and leakage modelling with an epanet extension–waternetgen. *Procedia Engineering*, 89:632–639.

Murphy, K. (2012). *Machine learning a probabilistic perspective*. MIT Press, Cambridge, Mass.

Ofwat (2015). On Serving Water ofwat.

Osborne, M., Garnett, R., and Roberts, S. (2010). Active data selection for sensor networks with faults and changepoints. In *24th IEEE International Conference on Advanced Information Networking and Applications (AINA)*, pages 533–540.

Osborne, M., Garnett, R., Swersky, K., and de Freitas, N. (2012). Prediction and fault detection of environmental signals with uncharacterised faults. In *26th AAAI Conference on Artificial Intelligence*, pages 533–540, Toronto, Ontario, Canada.

Osborne, M. A., Roberts, S. J., Rogers, A., Ramchurn, S. D., and Jennings, N. R. (2008). Towards real-time information processing of sensor network data using computationally efficient multi-output gaussian processes. In *Proceedings of the 7th international conference on Information processing in sensor networks*, pages 109–120. IEEE Computer Society.

PacificWater SOPAC (2012). The water supply system @ONLINE.

Palau, C., Arregui, F., and Carlos, M. (2011). Burst detection in water networks using principal component analysis. *Journal of Water Resources Planning and Management*, 138(1):47–54.

Pilcher, R., Hamilton, S., Chapman, H., Field, D., Ristovski, B., and Stapely, S. (2007). Leak location and repair guidance notes. *Guideline of IWA Water Loss Task Force*.

Price, E. (2005). *The Use of Residuals for Adaptive Signal Processing*. University of Oxford.

Pudar, R. and Liggett, J. (1992). Leaks in pipe networks. *Journal of Hydraulic Engineering*, 118(7):1031–1046.

Puust, R., Kapelan, Z., Savic, D., and Koppel, T. (2010). A review of methods for leakage management in pipe networks. *Urban Water Journal*, 7(1):25–45.

Rasmussen and Williams, Carl Edward, C. K. I. (2006). *Gaussian Processes for Machine Learning (Adaptive Computation and Machine Learning)*. The MIT Press.

Rogers, A., Maleki, S., Ghosh, S., and Jennings, N. R. (2011). Adaptive home heating control through gaussian process prediction and mathematical programming. *In International Workshop on Agent Technology for Energy Systems (ATES)*.

Sánchez, E., Ibáñez, J., and Cubillo, F. (2005). Testing applicability and cost effectiveness of permanent acoustic leakage monitoring for loss management in madrid distribution network. *Leakage 2005*, pages 12–14.

Srirangarajan, S., Allen, M., Preis, A., Iqbal, M., Lim, H. B., and Whittle, A. J. (2013). Wavelet-based burst event detection and localization in water distribution systems. *Journal of Signal Processing Systems*, 72(1):1–16.

Stampolidis, A., Soupios, P., Vallianatos, F., and Tsokas, G. (2003). Detection of leaks in buried plastic water distribution pipes in urban places-a case study. In *Advanced Ground Penetrating Radar, 2003. Proceedings of the 2nd International Workshop on*, pages 120–124. IEEE.

SWAN (2010). Smart water metering cost benefit study. Technical report, Smart Water Networks Forum (SWAN).

SWAN (2011). Stated NRW (Non-Revenue Water) Rates in Urban Networks. Technical report, Smart Water Networks Forum (SWAN).

The World Health Organisation (2014). *UN-water global analysis and assessment of sanitation and drinking-water (GLAAS) 2014*. Investing in water and sanitation: increasing access, reducing inequalities. World Health Organisation.

Tol, R. S. (1997). Autoregressive conditional heteroscedasticity in daily wind speed measurements. *Theoretical and applied climatology*, 56(1-2):113–122.

Trifunovic, N. (2006). *Introduction to Urban Water Distribution: Unesco-IHE Lecture Note Series*. UNESCO-IHE Delft Lecture Note Series. Taylor And Francis.

U.S. Geological Survey (2010). How much water is there on, in, and above the earth?

Venanzi, M., Rogers, A., and Jennings, N. R. (2013). Crowdsourcing spatial phenomena using trust-based heteroskedastic gaussian processes. In *First AAAI Conference on Human Computation and Crowdsourcing*.

Wirahadikusumah, R., Abraham, D. M., Iseley, T., and Prasanth, R. K. (1998). Assessment technologies for sewer system rehabilitation. *Automation in Construction*, 7(4):259–270.

Wu, Z. and Sage, P. (2007). Pressure dependent demand optimization for leakage detection in water distribution systems. In *9th International Conference on Computing and Control in the Water Industry, Leicester, UK.* Taylor & Francis London.

Wu, Z. Y. and Sage, P. (2006). Water loss detection via genetic algorithm optimization-based model calibration. *ASCE 8th Annual International*, 5.

Wu, Z. Y., Sage, P., and Turtle, D. (2009). Pressure-dependent leak detection model and its application to a district water system. *Journal of Water Resources Planning and Management*, 136(1):116–128.

Ye, G. and Fenner, R. (2010). Kalman filtering of hydraulic measurements for burst detection in water distribution systems. *Journal of Pipeline Systems Engineering and Practice*, 2(1):14–22.

Yu, T., Liya, M., Xiaohui, L., and Yunzhong, J. (2010). Construction of water supply pipe network based on gis and epanet model in fangcun district of guangzhou. In *Geoscience and Remote Sensing (IITA-GRS), 2010 Second IITA International Conference on*, volume 2, pages 268–271. IEEE.