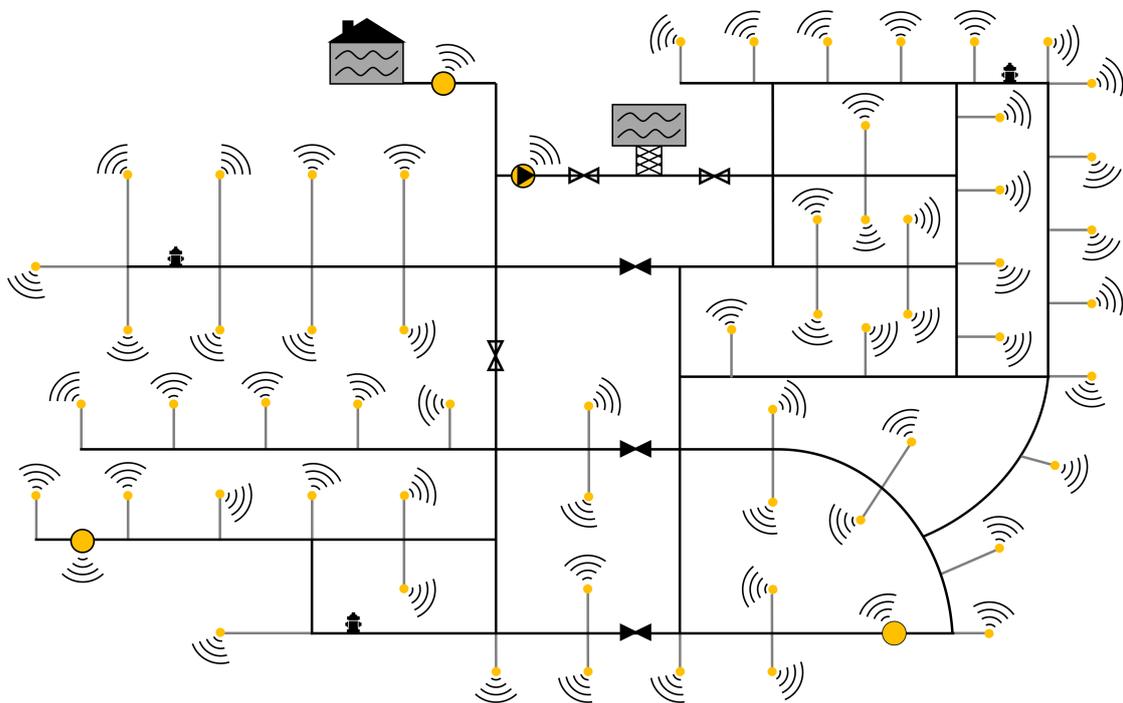


Data-driven water distribution system analysis - exploring challenges and potentials from smart meters and beyond



Jonas Kjeld Kirstein

PhD Thesis
February 2020

Data-driven water distribution system analysis – exploring challenges and potentials from smart meters and beyond

Jonas Kjeld Kirstein

PhD Thesis
February 2020

DTU Environment
Department of Environmental Engineering
Technical University of Denmark

Data-driven water distribution system analysis – exploring challenges and potentials from smart meters and beyond

Jonas Kjeld Kirstein

PhD Thesis, February 2020

The synopsis part of this thesis is available as a pdf file for downloading from the DTU research database ORBIT: <http://www.orbit.dtu.dk>.

Address: DTU Environment
Department of Environmental Engineering
Technical University of Denmark
Miljoevej, building 113
2800 Kgs. Lyngby
Denmark

Phone reception: +45 4525 1600

Fax: +45 4593 2850

Homepage: <http://www.env.dtu.dk>

E-mail: reception@env.dtu.dk

Preface

This PhD thesis was conducted at the Department of Environmental Engineering at the Technical University of Denmark (DTU). The work was guided by main supervisor Assoc. Prof. Martin Rygaard, co-supervisor Assoc. Prof. Morten Borup and co-supervisor Klavs Høgh from NIRAS A/S. The PhD study was supported by the LEAKman project and partners.

This thesis is based on six papers and these will be referred to in the text by their paper number in Roman numerals (I–VI).

- I Kirstein, J.K.,** Høgh, K., Rygaard, M. & Borup, M. (2019). Effect of data sampling resolution of smart meter readings in water distribution network simulations. *Manuscript in preparation.*
- II Kirstein, J.K.,** Høgh, K., Rygaard, M. & Borup, M. (2019). A semi-automated approach to validation and error diagnostics of water network data. *Urban Water Journal*, **16**(1), 1–10, doi:10.1080/1573062X.2019.1611884
- III Kirstein, J.K.,** Liu, S., Høgh, K., Borup, M. & Rygaard, M. (2019). Valve status identification by temperature modelling in water distribution networks. *Manuscript in preparation.*
- IV Kirstein, J.K.,** Høgh, K., Rygaard, M. & Borup, M. (2019). Using smart meter temperature and consumption data for water distribution system analysis. *Manuscript in preparation.*
- V Hubeck-Graudal, H., Kirstein, J.K.,** Ommen, T., Rygaard, M. & Elmegaard, B. (2019). Drinking water supply as low-temperature source in the district heating system: a case study for the city of Copenhagen. *Submitted.*
- VI Lund, N.S.V., Kirstein, J.K.,** Mikkelsen, P.S., Madsen, H., Mark, O. & Borup, M. (2019). Using smart meter water consumption data and in-sewer flow observations for model based sewer system analysis. *Submitted.*

In addition, the following publications and conference contributions, not included in this thesis, were also concluded during this PhD study:

- a) **Kirstein, J.K.**, Høgh, K., Borup, M. & Rygaard, M. (2016). Fra big data til smart data: driftsoptimering med højopløste sektionsdata i vandforsyningen. Casestudie: Halsnæs Forsyning. Dansk Vand Konference 2016, November 8-9, Århus, Denmark. *Abstract & oral presentation*.
- b) **Kirstein, J.K.**, Borup, M., Rygaard, M. & Høgh, K. (2016). Målerdata kan gemme på gratis informationer for forsyningerne: Et case studie fra Halsnæs Forsyning baseret på højopløste data. *DanskVand*, **84** (6), 50–51. *Article*.
- c) **Kirstein, J. K.**, Høgh, K., Borup, M. & Rygaard, M. (2018). Is your data correct? Validating and improving data collected in smart water networks. Nordic Drinking Water Conference 2018, June 11–13, Oslo, Norway. *Abstract & oral presentation*.
- d) Høgh, K. & **Kirstein, J.K.** (2018). ICT Frameworks – Moving Towards Smart Water Networks. IWA World Water Congress & Exhibition 2018, 16–21 September, Tokyo, Japan. *Abstract*.
- e) **Kirstein, J.K.**, Høgh, K., Borup, M. & Rygaard, M. (2018). Identifikation af ventilindstillinger fra temperaturmålinger. Dansk Vand Konference 2018, November 13–14, Århus, Denmark. *Abstract & oral presentation*.
- f) **Kirstein, J.K.**, Høgh, K., Borup, M. & Rygaard, M. (2019). Valve status identification by temperature modelling in water distribution networks. 13th Danish Water Forum Conference, January 31, Copenhagen, Denmark. *Abstract & oral presentation*.
- g) **Kirstein, J.K.**, Høgh, K., Rygaard, M. & Borup, M. (2019). Valve status detection using smart meter temperature and flow. 17th International Computing & Control for the Water Industry Conference, September 1–4, Exeter, UK. *Abstract & oral presentation*.
- h) Lund, N.S.V., Borup, M., **Kirstein, J.K.**, Mark, O., Madsen, H. & Mikkelsen, P.S. (2019). Lessons learned from comparing smart meter water consumption data with measured wastewater flow in the drainage system. 17th International Computing & Control for the Water Industry Conference, September 1–4, Exeter, UK. *Abstract*.

Furthermore, a meter data validation tool based on the tests presented in Paper II was made available online at <https://leakagemanagement.net/meter-validate>.

Acknowledgements

This PhD project was conducted at DTU Environment at the Technical University of Denmark in collaboration with the LEAKman project partners.

First, I should like to thank my two main supervisors, Martin Rygaard and Morten Borup. Their doors were always open to me during the project irrespective of timing, question or topic. They guided me when needed and I greatly appreciated their support when the project tended to take second place to my personal life. I will miss our supervisor meetings, at which they were at times entertainingly contradictory, ultimately increasing our mutual knowledge of water distribution system analyses.

Second, I should like to thank my supervisor, Klavs Høgh for sharing his professional knowledge and experience, often highlighting the gap between academia and the ‘real world’. Also, I should like to thank the VAF2 team at NIRAS A/S, in particular, Gaby, Gitte, Jesper, Adam, Andreas, Thomas, Lars, Jan, Anders, Janus, Nikolai, Rosa and Pia for help with all kinds of questions related to water supply and for always making me feel welcome at NIRAS.

Third, I should like to thank Professor Shuming Liu and his research team at Tsinghua University. They made me feel welcome from day one of my stay in China. I learned a lot from being a member of their group – from an academic as well as a personal point of view. I am sure this was not my last visit to China.

Next, I should like to thank Mette Fogtmann and Bo Lindhardt from Novafos and Per Grønvald, Henrik Horsholt Christensen and Thorkil Bartholdy Neergaard from Brønderslev Forsyning for sharing data and knowledge related to their water supply systems.

I should also like to thank the project partners of the LEAKman project, in particular Kamstrup A/S and HOFOR for their input and financial support to my PhD project.

I also want to praise and thank Frank for his help with my English during the project and, in particular, Nadia, for her never-ending support.

Last but not least, my final thanks go to my dad who, nine years ago, persuaded me to follow an education in environmental engineering instead of computer sciences. “You will end up in front of a computer anyway.” Well, he was right!

Summary

The availability of clean drinking water at any time is often taken for granted. However, the reliable supply of drinking water is, unfortunately, challenged by various threats, such as global warming and aging infrastructures, which put an immense pressure on the drinking water systems as we know them today. Consequently, utilities, technology providers and researchers seek to identify optimised and new approaches to maintaining and improving the quality of the delivered water. In an age of digitalisation, data-driven approaches are becoming increasingly important, as they have demonstrated various benefits for the operation and design of water distribution networks. However, the increased collection and application of the data also pose a major challenge to the water sector. This PhD developed methods to help utilities in validating and applying their data in novel ways by analysing ‘real world’ data obtained from five Danish utilities, applied in six case studies. The PhD study was sectioned into: 1) data collection; 2) data validation and reconstruction; and 3) data application.

Data-collection devices such as smart meters are increasingly deployed throughout water distribution systems. When utilities introduce smart meters, the selection of sampling resolutions has a trade-off between the applicability of the collected data and transmission costs. Analysis of a district metered area with smart meters installed revealed that common sampling resolutions of between 1 and 24 hours are sufficient for water loss assessments as long as utilities have representative demand patterns of their network available. However, sampling resolutions < 1 hour are potentially important to obtain reliable water quality simulations.

Automatic validation and reconstruction processing of the collected data are of paramount importance for utilities. The PhD project developed a systematic approach for categorizing anomalies. Four categories were introduced, with Type 0 describing system anomalies, Types 1 and 2 describing sensor data containing a low data quality, and Type 3 covering sensor data anomalies storing information about actual – though unusual – events appearing in the water distribution network. To identify anomalies of Types 1 and 2, seven validation tests were developed. Analysis of pressure and flow data sets from three Danish utilities revealed a large proportion of anomalies, with on average 10% missing data and up to 35% anomalies of Types 1 and 2 in a utility’s pressure data sets. These high numbers also emphasised the need for reconstruction processes to generate reliable data streams that are required in data applications. An example was presented whereby artificial neural networks were used to provide

missing data and to further validate dubious observations.

The collection of data from water distribution systems is not a new concept, but large amounts of data (such as temperature data) are often left unused due to a lack of evidence of successful applications. To show the benefits of temperature data, a temperature model and a hydraulic model were combined to identify the status and location of valves in the network. This novel approach and field tests in the network unexpectedly revealed various anomalies of Type 0 in the utility's asset database, ultimately casting doubt on the validity of the hydraulic model. As long as such anomalies prevail in the data sets, it is not possible to apply advanced data-driven applications successfully. Another issue in the case study included a low quantity of applicable temperature data. Smart meter temperature data, potentially available in each household, can be used to overcome this challenge. In another case study, the simulated temperature throughout a district metered area showed a satisfying resemblance to smart meter temperature data (average root mean square error of 0.9 °C). This highlights the potential of using smart meter temperature data for more advanced applications, such as leakage detection and valve status detection.

Silo thinking is traditionally a common feature of the water sector, and the value of water supply data is thus often overlooked in external applications. In a case study, the effect of deploying heat pumps on the water distribution network mains was assessed as a supplement to the district heating system of Copenhagen. A net heat extraction potential of 20.7 MW was estimated. Moreover, this caused the share of users complying with an upper temperature limit of 12 °C to increase from 41% to 81% during August. In another case study, smart meter water consumption data were linked to an urban drainage model to compare simulated wastewater flows with in-sewer observations. The in-sewer observations were found to be erroneous, and the smart meter data were deemed more valid in estimating dry weather flow than in-sewer observations.

Overall, the project showed that many anomalies prevailing in the utilities' asset databases and sensor data are first discovered through the application of the data. As long as utilities cannot maintain a high level of data reliability, it is doubtful whether more sensors will increase utilities' understanding of their systems. The true potential of the data will not be unlocked until a high level of data reliability is secured. In the coming years, utilities, technology providers and researchers should together identify methods for reducing the uncertainties prevailing in asset and sensor data, making it possible for the sector to reach higher levels of digital maturity.

Dansk sammenfatning

Tilgængelighed af rent drikkevand til enhver tid tages ofte som en selvfølge. Desværre er den høje forsyningssikkerhed truet af en række udfordringer, såsom global opvarmning og aldrende infrastruktur, hvilket lægger et enormt pres på drikkevandssystemerne, som vi kender dem i dag. Forsyningsselskaber, teknologileverandører og forskere arbejder derfor på at identificere optimerede og nye tilgange til at opretholde og forbedre kvaliteten af det leverede vand. I digitaliseringens tidsalder bliver datadrevne fremgangsmåder mere og mere vigtige, da disse har vist sig fordelagtige inden for drift og design af vandledningsnet. Den øgede indsamling og anvendelse af data udgør dog også en stor udfordring for vandsektoren. Ph.d.-afhandlingen udviklede derfor nye metoder til at hjælpe forsyninger med at øge valideringen og anvendelsen af deres data. Dette blev gjort ved at analysere rigtige data fra fem danske forsyningsselskaber fordelt over seks casestudier. Ph.d. afhandlingen blev opdelt i tre trin: 1) dataindsamling, 2) datavalidering og datarekonstruktion, og 3) brug af dataene.

Forsyninger installerer i stigende grad dataindsamlende enheder, såsom intelligente målere (*smart meters*). Når forsyningerne vælger at implementere smart meters, skal fordelene og ulemperne imellem en fin dataopløsning (større brugbarhed af de indsamlede data) og transmissionsomkostningerne opvejes mod hinanden. En analyse af et målerdistrikt med installerede smart meters viste, at typiske prøveopløsninger imellem 1 og 24 timer er tilstrækkelige til at vurdere vandtab, så længe forsyningsselskabet har repræsentative forbrugskurver fra netværket. En dataopløsning på < 1 time er dog potentielt nødvendigt for at opnå pålidelige vandkvalitetssimuleringer.

Automatiske validerings- og rekonstruktionsprocesser for de indsamlede data bør være af høj prioritet for forsyningerne. Ph.d.-projektet udviklede en systematisk tilgang til at kategorisere anomalier. Der blev introduceret fire kategorier, hvor Type 0 repræsenterede systemanomalier, Type 1 og 2 beskrev sensordata med en lav kvalitet, og Type 3 beskrev anomalier i sensordata baseret på rigtige, men usædvanlige, begivenheder i ledningsnettet. Der blev udviklet syv forskellige valideringstest til at bestemme Type 1 og 2 anomalier. En analyse af tryk- og flowdatasæt fra tre danske forsyningsselskaber afslørede et stort antal af anomalier, med et gennemsnit på 10% manglende data og op mod 35% anomalier i en af forsyningernes trykdata. Disse høje tal understreger behovet for rekonstruktionsprocesser til at generere de konsistente datastrømme, der kræves til en pålidelig anvendelse af data. Neurale netværk blev anvendt som et eksempel på en metode til at rekonstruere manglende data og yderligere

validere tvivlsomme observationer.

Indsamling af data fra vandforsyningssystemer er ikke et nyt koncept. Desværre lagres store mængder af data (såsom temperaturdata) ofte uden at blive brugt fordi der mangler eksempler på vellykkede anvendelser af dataene. For at demonstrere fordelene ved temperaturdata blev en temperaturmodel og en hydraulisk model kombineret for at identificere indstillingen og placeringen af ventiler i et netværk. Denne nyudviklede metode samt forsøg i felten førte utilsigtet til identifikation af forskellige Type 0 anomalier i forsyningens ledningsdatabase. Metoden har dermed sæt tvivl om pålideligheden af den hydrauliske model. Så længe anomalier gemmer sig i sensor- og ledningsdatabaser, er det ikke muligt succesfuldt at udføre mere avancerede datadrevne analyser. Et andet problem i casestudiet var det lave antal af relevante temperaturdata. Temperaturdata fra smart meters, som muligvis er tilgængelige fra hver husstand, kan bruges til at løse denne udfordring. I en anden forsynings målerdistrikt viste smart meter temperaturdata tilfredsstillende lighed med den simulerede temperatur (average root mean square error på 0,9 °C). Dette understreger potentialet ved at bruge smart meter temperaturdata til mere avancerede undersøgelser, såsom lækagesporing og identifikation af ventilindstillinger.

Den typiske silo-tankegang i vandsektoren resulterer ofte i, at man overser værdien af drikkevandsdata uden for vandforsyningssektoren. I et casestudie blev det undersøgt om det kan betale sig at implementere varmepumper på drikkevandsvandlejninger som et supplement til Københavns fjernvarmenet. Der blev fundet et varmpotentiale på 20,7 MW. Derudover steg andelen af drikkevandsforbrugere, hvor den øvre temperaturgrænse på 12 °C blev overholdt, fra 41% til 81% i august. I et andet casestudie blev smart meter vandforbrugsdata koblet til en afløbsmodel for at sammenligne den simulerede spildevandsstrøm med observationer fra afløbssystemet. Der blev identificeret fejl i målerne i afløbssystemet, og smart meter dataene blev derfor vurderet mere pålidelige til at estimere tørvejsflowet end observationerne fra disse målere.

Generelt viste ph.d.-projektet at mange fejl i forsyningernes ledningsdatabaser og sensordata kun opdages ved at anvende dataene. Så længe forsyningerne ikke kan opretholde en høj datapålidelighed, er det tvivlsomt om flere sensorer vil øge forsyningernes forståelse af deres system. Først når der kan sikres en højere pålidelighed af dataene, kan deres sande potentiale udnyttes. I de kommende år bør forsyningsselskaber, teknologileverandører og forskere identificere metoder til at reducere fejlene som florerer i forsyningernes lednings og sensordata, så forsyningerne kan nå et højere niveau af digital modenhed.

Table of contents

Preface	iii
Acknowledgements	v
Summary	vii
Dansk sammenfatning	ix
Table of contents	xi
1 Introduction	1
1.1 Challenges of water supply management	1
1.2 Digitalisation	2
1.3 The digital challenges	4
1.4 Objectives	6
1.5 Thesis structure	7
2 Digitalised water supply systems	9
2.1 Flow and pressure meters (at distribution level)	9
2.2 Water quality sensors	10
2.3 Noise loggers.....	11
2.4 The smart meter revolution.....	11
3 The urban water data warehouse	15
3.1 Sampling resolution and gap filling	15
3.2 Validating and reconstructing data	18
3.2.1 Need for automatic data cleansing	18
3.2.2 Anomaly types	19
3.2.3 Identification of anomalies	20
3.2.4 The value of anomalies	21
3.2.5 Higher-level validation of anomalies and reconstruction	22
4 Novel data applications and modelling for enhanced system analysis 25	
4.1 Improved system understanding through temperature modelling	26
4.1.1 Semi-synthetic case study	28
4.1.2 Real transportation network case study	28
4.2 Smart meter temperature data	31
4.2.1 Filtration and classification of smart meter temperature data	31
4.2.2 Temperature model analysis with smart meter data	33
4.3 Drinking water as a low-temperature source in district heating systems ...	35
4.3.1 Case study: Greater Copenhagen Utility, Denmark.....	37
4.4 Smart meter data as an add-on to urban drainage management	39
4.4.1 Case study: Utility Elsinore, Denmark	39
5 Conclusions	43

6	Perspective	47
7	References	49
8	Appendix	55
	A Scopus search query	55
	B Data reconstruction	55
9	Papers.....	59

1 Introduction

“Water and wastewater utilities must embrace digital solutions. There is really no alternative” (Sarni et al., 2019).

This quotation originates from a report about digital water by the International Water Association and sets the framework for this thesis. Sarni et al. (2019) state that “digital water is already here”, and that utilities are evolving from simple to complex and interconnected institutions. In general, digitalisation entails a long list of benefits for the water sector. But is digitalisation really a straightforward path to more success? The short answer is no, as each utility is unique in terms of levels of digital maturity and overall needs. However, to give a clearer and more elaborate answer to the above question it is important to understand the ongoing challenges of the water sector driving the transition towards digitalisation. Moreover, digitalisation requires that multiple obstacles be addressed properly before it is possible to refer to water systems as ‘truly smart’ (Moy de Vitry et al., 2019).

1.1 Challenges of water supply management

The United Nations Sustainability Goal 6 is about ensuring the availability and sustainable management of water and sanitation for all (United Nations, 2019). Around 10% of the global population lacks basic drinking water services, half of the global population lives in areas that experience water scarcity at least one month of the year, and rivers in Africa, Asia and South America are now more polluted than they were in the 1990s (United Nations, 2019). In addition, ever-increasing urbanisation, global population growth and climate change have put immense pressure on the drinking water systems as we know them today (for example, WWAP, 2019). These factors combined with socio-economic changes, will not only lead to changing consumption patterns but will also affect worldwide water consumption, which will continue to grow (about 1% each year since the 1980s), ultimately increasing global water stress (WWAP, 2019).

The summer of 2018 showed that temporal and regional water shortages occurred even in areas otherwise unfamiliar with water stress, such as Denmark and Germany. In the future, declining amounts and decreasing quality of water resources will lead to an increased competition between water users, demanding new ways of distributing water fairly between them (DANVA, 2018; German Environment Agency, 2019).

As a consequence of water stress, utilities need to pay increased attention to resource efficiency and the exploitation of alternative water resources (Rygaard et al., 2011). However, the introduction of alternative solutions, such as rainwater collection or wastewater reclamation, poses severe challenges that need to be addressed (Rygaard et al., 2011). Examples include thorough analysis of the solutions' energy requirements and effects on the quality of the delivered water (Rygaard et al., 2011). Thus, the overall complexity of drinking water supplies will continue to increase. This is also the case in Denmark, where, among other things, new and rising numbers of contaminants are being detected in the groundwater. The number of reported cases of pesticide levels exceeding guideline levels increased notably from 15 to 65 waterworks between 2013 and 2017 (Ministry of Environment and Food of Denmark, 2018a). This demands more advanced treatment processes than currently implemented.

Moreover, long-established utilities experience that a great proportion of their drinking water infrastructure is past its prime. Due to this unreliable infrastructure, not only the costs of operation and management (e.g. increased leakage) will grow, but also businesses and the standard of living will be affected negatively (ASCE, 2011). For example, investment in American water infrastructure does not keep up with need, reaching an estimated funding gap of up to \$144 billion by 2040 (ASCE, 2011). Thus, there is an increasing incentive to cope with the failing infrastructure in an efficient manner due to lack of funds (Eggimann et al., 2017). Here, 'efficiency' comprises better rehabilitation planning (asset management) and extending the expected system lifetime, as well as an optimised operation and design of urban systems. For example, Nguyen et al. (2018) state that coarse and out-of-date information is often used during the design and planning of urban water infrastructure, leading to inefficient management. This is also the case for Denmark, where Kirstein (2016) and Kirstein et al. (2016), by analysing district metered area (DMA) data showed that there were significant differences between the actual consumption and apparently out-of-date demand patterns provided by the Danish environmental protection agency (Watertech A/S, 2005). For example, application of out-of-date demand patterns may lead to unintentional network augmentation during design (Gurung et al., 2016) and doubtful hydraulic model simulations.

1.2 Digitalisation

Digitalisation is often pointed out as the process capable of mitigating and solving parts of the previously mentioned challenges. Gartner's IT glossary (Gartner, 2019) defines the term as: "the use of digital technologies to change

a business model and provide new revenue and value-producing opportunities; it is the process of moving to a digital business.”

In urban water management, digitalisation is envisioned to increase the effectiveness and flexibility of urban water systems and establish the opportunity for new services (Moy de Vitry et al., 2019). Such opportunities have boosted the interest of the research community on digitalisation in the water sector, notably since the start of the 2000s (Figure 1).

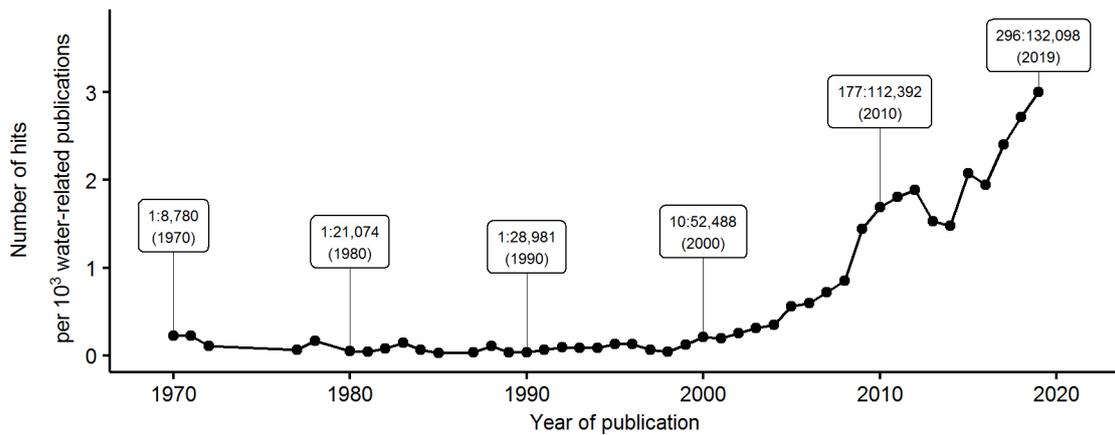


Figure 1. Normalised hits based on 1) publications concerning digitalisation in the water sector only and 2) publications related to water research in general. Number boxes display the actual hit ratio between these two categories. Scopus search (21/09/2019), Appendix A.

This boost in research is, among other things, driven by new technologies that are believed to benefit the water sector in multiple ways, providing more secure, resilient, reliable, efficient, cost-effective and innovative water solutions (Sarni et al., 2019). The literature has many promising examples justifying this claim (see, for example, Section 2). A comprehensive overview of the history and future of digitalisation in the urban water sector is given in Makropoulos and Savić (2019), indicating that the sector is far from being at the end of its digital transition journey. Sarni et al. (2019) outline a digital adaptation for utilities spanning steps of basic, opportunistic, systematic and transformational adaptation (Figure 2). Utilities around the world have already started the digital transformation, averaging an adoption level of ‘opportunistic’ according to a survey conducted in Sarni et al. (2019). Here, opportunistic accounts for, among other things, utilities with digital automation and control mechanisms and analytical tools during process optimisation.

The digital transformation is mainly enabled through recent advances in smart information and communication technology (ICT) and facilitated through applications and implementations of, for example, the Internet of Things (IoT)

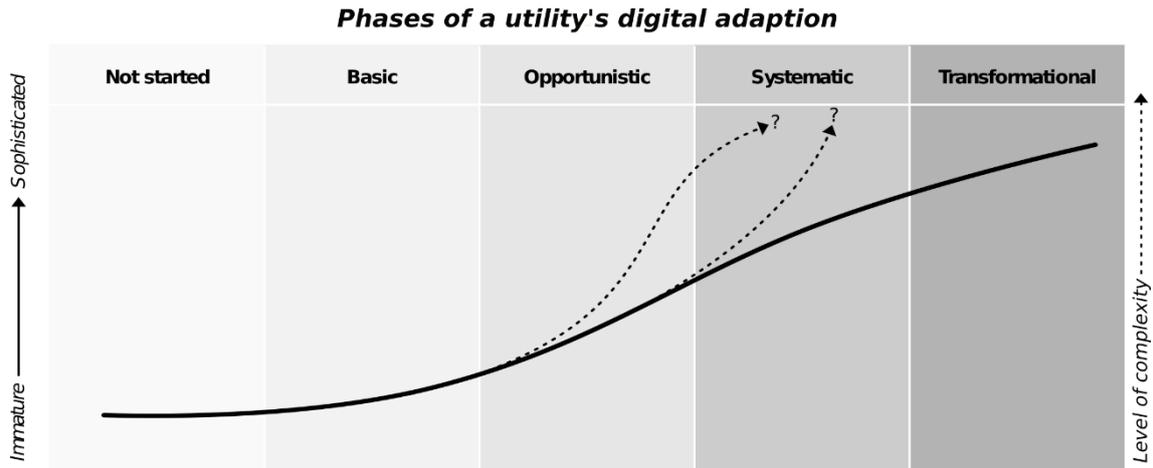


Figure 2. ‘Digital Water Adoption Curve’ adapted and modified from (Sarni et al., 2019).

and edge and cloud computing (Eggimann et al., 2017; Kulkarni and Farnham, 2016; Sun and Scanlon, 2019). In the water sector, these technologies have increased the speed of data collection and analytics notably over the last several decades in parallel with decreasing costs, ultimately opening up the world of ‘Big Data’ (e.g. Monks et al., 2019; Sun and Scanlon, 2019). An example of how fast the sector is evolving is highlighted in Cominola et al. (2015), where the installation of ultrasonic smart meters was described as too costly; now it is state-of-the-art, as shown, for example, by the deployment of 2,300 ultrasonic smart meters in Northern Copenhagen (LEAKman, 2018) or by the case study utilities in Papers IV and VI.

Digitalisation concerns the entire urban water sector, but this PhD study focuses only on data from water distribution systems. Whereas the collection and application of data from WDNs are not new to drinking water utilities, the increasing volumes of data introduce challenges and potentials within the field of water distribution system analysis. Here, the PhD study will to a large extent focus on the application of data from smart meters, being a prime example of digitalisation.

1.3 The digital challenges

Simply installing more devices and collecting more data do not result in a higher level of digital maturity for a utility. Due to the low usage of the collected data in many Danish utilities, I therefore speculate that reaching higher levels of digital maturity (e.g. systematic) is a more difficult (steeper) process than envisioned and may lead to unknown paths of rising complexity for utilities (dotted lines, Figure 2). Many important factors for becoming more *data-*

driven (i.e. relying on the collection and analysis of data) need to be considered, here divided into three steps:

- 1 Data collection.** It should be considered what and how much data should be collected. For example, Kulkarni and Farnham (2016) and Mekki et al. (2019) showed that the implementation of ICT comes at various costs and depends on multiple factors such as the selected data transmission technologies. Eggimann et al. (2017) pointed out that data-driven urban water management should aim to reach a region of optimal data availability, meaning that there is a point where more data is not necessarily better. Moreover, there is still a lack of evidence about what types and frequencies of information best suit the needs of utilities and consumers (Boyle et al., 2013).
- 2 Validation and reconstruction of data.** A reliability assessment of the collected data is necessary. I recognize ‘data reliability’ as a term covering the accuracy and completeness of collected data, which is secured only through forms of validation and reconstruction processes. With the increasing amount of data collected, it is of paramount importance that autonomous data collection and verification processes are in place to secure a high reliability of the data (Sun and Scanlon, 2019). In a survey of the urban water community about how ubiquitous sensing will shape the future of the sector, Blumensaat et al. (2019) showed that data validation and integrated management were ranked as the two most important topics among professionals. It is important to be able to trust the collected data with a very high reliability before it can be applied in decision-making processes (Blumensaat et al., 2019). Furthermore, integrated management will play an increasing role in future infrastructure design, but models used for such processes are limited by the availability and quality of the data (Eggimann et al., 2017).
- 3 Application of data.** Simply assuring a high data reliability, however, does not release the biggest potential the sector is currently facing; the overall value created by the collected data from digital technologies is often too unclear (Sarni et al., 2019). In general, proof is still lacking as to whether digitalisation really leads to long-term savings and increased performance; this is often because it is unclear which exact challenges can and cannot be solved by mining the data (Blumensaat et al., 2019; Boyle et al., 2013; Cominola et al., 2015; Eggimann et al., 2017; Sarni et al., 2019).

I envision these three steps as an interlinked, cyclic process, whereby new insights from unsuccessful and successful data applications in step 3 may lead to a better understanding of data collection requirements in step 1. Furthermore,

Papers III, IV and VI highlighted that it is first the application of data that clarifies valid and invalid data, thus improving step 2.

It should be noted that many additional important challenges and topics related to digitalisation, such as privacy and cyber-security concerns or a changing customer–utility relationship raise important questions. These were, however, deemed outside the scope of this thesis.

1.4 Objectives

Currently, digitalisation is one of *the* main topics of the urban water sector. In this thesis, I will assess obstacles and opportunities as well as develop methods related to 1) data collection, 2) validation and reconstruction of data and 3) application of the data in water distribution systems and external water-related systems. This is done to guide utilities towards improving the operation and management of their systems. This thesis addresses topics 1–3 by answering the following questions.

- 1** Define adequate sampling resolution from smart meters.
 - What is the effect of changed sampling resolutions from smart meters in hydraulic simulations?
 - What is the effect of filling gaps between measurements with different methods?
- 2** Improve the reliability of collected data in urban water systems.
 - What are the common rates and types of anomalies in water network data?
 - How can we improve future data collection procedures?
- 3** Evaluate new ways to use collected data in water distribution systems.
 - Which benefits can be obtained by analysing temperature data collected in the water distribution network and from smart meters?
 - What is the impact on consumers and the distributed water quality when installing heat pumps in water distribution networks?
 - How can smart meter data be used to estimate wastewater flows?

A major focus of this thesis is the application of ‘real-world’ data to answer the above stated questions. In short, ‘real-world’ case studies from five different utilities are used to outline the current state of digitalisation.

1.5 Thesis structure

The thesis structure follows the data flows in a water utility (Figure 3).

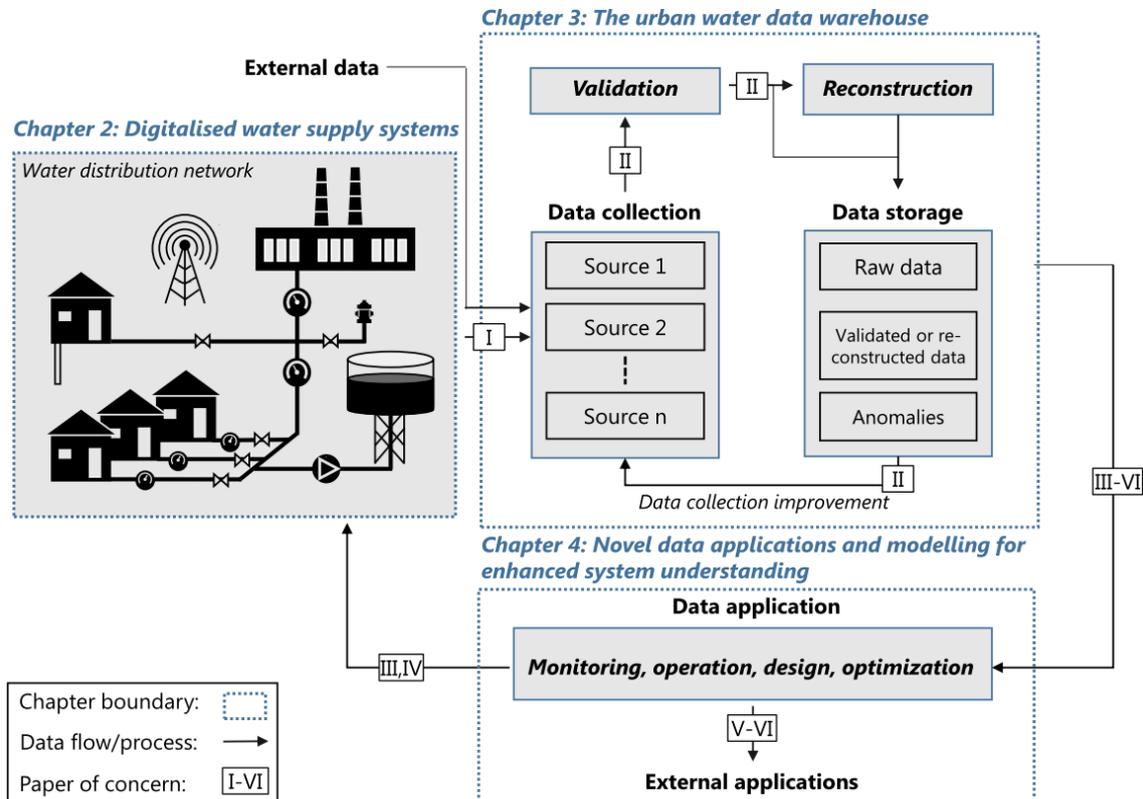


Figure 3. Conceptual overview of data flows in a water utility. The focus of each chapter is delineated by blue boxes.

The collected data needs to undergo various processes (indicated by boxes and arrows in Figure 3) prior to successful application, guiding the remainder of this thesis.

- Chapter 2 (Digitalised water supply systems) introduces commonly installed data-collection devices in drinking water distribution networks (WDNs), including the benefits and pitfalls of these technologies.
- Chapter 3 (The urban water data warehouse) highlights the need for utilities to rethink their data gathering (Step 1) and processing in forms of validation and reconstruction (Step 2) prior to successful application of the data (Papers I and II):
 - The first step includes data collection. Paper I highlights how sampling resolutions affect pressure and water age simulations as well as water loss assessments.

- Next, it is important to maintain a high level of reliability of the collected data. Paper II shows how the data can be validated and how errors in data collection can be identified. An example of reconstructing data is given in Section 3.2.4
- Chapter 4 (Novel data applications and modelling for enhanced system analysis) addresses new opportunities (Step 3) from data collected in WDNs for optimised water distribution system analysis (Papers III and IV) and external research fields (Papers V and VI).
 - WDN-related applications. Papers III and IV show how temperature data collected in the network can be used to improve the operation and management of WDNs.
 - External applications. Papers V and VI show that temperature and smart meter data are beneficial for management/operation of systems outside the ‘WDN bubble’, such as district heating and urban drainage.
- Chapter 5 (Conclusion) summarises the main results.
- Chapter 6 (Perspective) lists future research possibilities, risks and challenges that need to be addressed.

The applied real-world data across six case studies covers smart meters at the household level, meters at WDN level and manual sampling of flows, pressure and temperatures (Table 1).

Table 1. Overview of real flow (Q), pressure (P) and temperature (T) data collected by smart meters, sensors in the water distribution network and manual sampling.

Data source Parameter	Smart meter (household)		Water distribution network			Manual sampling	
	Q	T	Q	T	P	T	P
Paper I	X		X		X		
Paper II			X		X		
Paper III			X	X	X	X	X
Paper IV	X	X	X	X	X		
Paper V						X	
Paper VI	X		X				

2 Digitalised water supply systems

Currently, the amount of data collected from devices deployed in urban water systems is increasing notably. Examples include sensors monitoring various forms of parameters, such as flow and pressure meters, pumps and valves collecting data about their state (e.g. on or off) and acoustic loggers recording the level of noise at selected locations. The data collected from such devices is useful for monitoring, control, design and planning purposes. Kulkarni and Farnham (2016) gave an overview of (smart) ICT deployed in water distribution systems, subdividing the field of water monitoring into pressure and flow management, consumption monitoring, water loss management and water quality monitoring. I used these categories to classify commonly implemented data-collection devices according to their major objectives in WDNs (Table 2). Here, actuators (pumps and valves), as well as non-data-driven drivers (e.g. customer satisfaction) have been deemed outside the scope of this thesis. The following section addresses each of the devices listed in Table 2.

Table 2. Major data-collection devices deployed in water distribution networks and their application in water distribution system analysis, based on my categorisation.

Drivers Devices	Pressure & flow management	Consumption monitoring	Water loss management	Water quality monitoring
Flow & pressure meters (distribution level)	X	X	X	
Water quality sensors				X
Noise loggers			X	
Smart meters (consumer level)	X	X	X	X

2.1 Flow and pressure meters (at distribution level)

Pressure & flow management. Flow and pressure sensors installed at DMA level or other critical locations (e.g. tanks or reservoirs) have been an integral part of many utilities for decades (even though Puust et al. (2010) described the installation of flow meters as “a recent trend”). For example, in combination with hydraulic models, the collected data may be used to simulate the flow and pressure throughout the network and highlight areas with pressure deficits or long retention times. Pressure monitoring and management are also directly connected to water loss management, as higher pressure may lead to higher leakage rates (see below). Furthermore, pressure and flow data play a major role in real-time control of WDNs, as the data is used to remotely control the states of pumps and valves (e.g. Creaco et al., 2019). Some of these pumps and valves are termed ‘intelligent’ as they are capable of taking decisions independent of remotely controlled set-points, such as in the case of blackouts or connectivity issues (e.g. AVK, 2018).

Consumption monitoring. In terms of consumption monitoring, the data can be used to gather information about the consumers in the area, e.g. based on DMA-level measurements useful for data reconstruction purposes, demand forecasting and design of networks (e.g. Kirstein, 2016; Kirstein et al., 2016). Application of forecast models that rely on demand data from flow meters, is expected to increase, as already seen for optimal pump control in Dutch water supply systems (Bakker et al., 2014).

Water loss management. DMA data can be used for leakage detection. One of the most common approaches in using flow data for leakage detection is minimum night flow (MNF) analysis (for example, Puust et al., 2010). In MNF analysis, the water inflow into an area is measured (e.g. between 02:00 and 05:00) and legitimate uses are subtracted. High deviations in MNF can then be used to detect leakages and bursts. Good knowledge of the downstream area is important when installing sensors, as otherwise it is possible that the installed sensors will not be sensitive enough to detect the MNF, as seen in the applied data in Paper II. In addition, more sophisticated approaches can be applied, such as data-driven methods for automatic burst detection in WDNs (Wu and Liu, 2017). Here, the quality of the data constitutes the main uncertainty, with flow measurements being more reliable (though also more expensive) than pressure measurements (Wu and Liu, 2017). Most sensors are deployed at DMA level, being capable of detecting bursts in the range of 1.5–50% of average DMA inflow, but the success of the reviewed methods were difficult to compare owing to the varying nature of the case studies (Wu and Liu, 2017). Wu and Liu (2017) also state that finer sampling rates and data communication frequencies (e.g. < 5 minutes) could reduce detection times.

2.2 Water quality sensors

Water quality monitoring. In general, the monitoring of water quality has been deemed to be one of the most difficult parameters to “monitor remotely and reliably” (Makropoulos and Savić, 2019) owing, among other things, to the high number of possible contaminants and entry points (Adedoja et al., 2018; Eggimann et al., 2017). Moreover, Makropoulos and Savić (2019) state that additional work is needed on novel water quality sensors. An overview of microbial sensors from 2015 can be found in Tatari et al. (2016), showing that the response time of many sensors varies (examples in Tatari et al. (2016) between 10 minutes and 18 hours), depending on which parameters are measured. Tang and Albrechtsen (2019) list commercially available technologies for real-

time or near-real-time monitoring of water quality (albeit from the food industry), of which the most common parameters included pH, turbidity and conductivity. Currently, technologies behind sensors measuring microbial water quality as well as physiochemical properties of the water are being rapidly developed (Tang and Albrechtsen, 2019). Thus, it is believed that the number of water quality sensors will increase notably in the coming years. Potentially, the number of water quality parameters monitored can also lead to the detection of leakages, among other things because turbidity can increase during pipe burst events (Puust et al., 2010).

Temperature measurements in the WDN can be used as a simple water quality indicator and guide utilities to locations of water quality audits, among other things because temperature affects the growth rates of biofilm-forming bacteria (Liu et al., 2016) and because temperature measurements are an indicator of the water's residence time in the WDN (Papers III and IV). Thus, temperature can be selected as a parameter useful for the identification of relevant monitoring sites (for example, Larsen et al., 2017).

2.3 Noise loggers

Water loss management. Typically, acoustic loggers attached to pipe fittings record the level and spread of noise in pipes to detect leakages by statistical analysis (Puust et al., 2010). Leakage correlators use the signal between two adjacent noise loggers to correlate and narrow down the area of leakages (see, for example, Li et al. (2015) for further information on acoustic detection methods).

2.4 The smart meter revolution

Here, smart meters are understood to be meters at the household level. Smart meters exemplify the digital transformation in the water sector more than any other device, as they affect both utility and consumers and can play a major role in integrated urban water management. The meters are termed 'smart' because water consumption and other parameters are measured on a less than daily basis and are collected remotely, opening up various possibilities within data analytics with benefits for the utility and consumers (Boyle et al., 2013; Cominola et al., 2015). Also, technological advances have made it possible to implement additional monitoring capabilities in smart meters, no longer limiting meters to only monitoring the demand.

Advances in sensing have made meters (e.g. ultrasonic smart meters) more accurate and less fragile compared to mechanical meters, and the implementation

of smart meters has increased notably over the last few decades (Boyle et al., 2013; Cominola et al., 2015; Kamstrup, 2019a). For example, based on a survey of 55–60 participating utilities, the number of remotely read meters at household level in Denmark increased from 15% to 46% between 2013 and 2017 (DANVA, 2018). However, the literature also points out that on a global scale, the majority of digital metering rollouts were conducted on smaller scale trials. Furthermore, the rollout was slower than expected, partly because many of the benefits are difficult to monetize or utilities struggle to see the benefits of the data (Monks et al., 2019; Stewart et al., 2018; Paper IV).

Based on a comprehensive literature search and survey among experts, Monks et al. (2019) revealed in total 75 benefits for utilities and customers when implementing digital metering. These included, among other things, enhanced interaction between consumers and the utility, greenhouse gas reductions due to reduced driving, and more efficient billing (Monks et al., 2019). Smart meter deployment reduces one of the major limiting factors of water distribution system analysis: the unknown demand. However, as long as utilities do not apply this data in analyses, such as in near-real-time water loss assessments, in hydraulic online models or in demand forecasting, the benefits of smart meter deployment are less obvious for utilities with a basic level of digital adaptation (Figure 2). In the following, data-driven benefits of smart meter deployment will be discussed concerning the categories listed in Table 2.

Pressure & flow management. The increased data quantity and quality from smart meters can be a major driver for improved hydraulic model accuracy in terms of flow and pressure simulations. Naturally, the resemblance of the DMA inflow with smart meter consumption data increases when using daily or less-than-daily consumption data, compared to coarser audit data (e.g. quarterly) (Paper I). Smart meters may also be used to remotely disconnect users from water, e.g. during maintenance or when bills have not been paid (e.g. Blokker, 2019).

Consumption monitoring. One major driver for installing smart meters is water demand management, as smart meters can help to increase awareness about water usage and thus lead to reduced consumption (Boyle et al., 2013; Nguyen et al., 2018). Here, disaggregation of the consumption data into specific usages can increase the awareness of users about their consumption (e.g. the specific water usage of certain appliances) and help the utility with managing demand peaks (Cole and Stewart, 2013; Cominola et al., 2015; Nguyen et al., 2018;

Stewart et al., 2018). Digital metering can also be used to improve infrastructure planning and reduce network augmentation, which is particularly evident for consumers with little information prior to data collection (Gurung et al., 2016, 2014; Monks et al., 2019).

Water loss management. Smart meter data can help to establish temporary or permanent DMAs and enable detailed water balances and post-meter leakage detection, e.g. based on leak alarms (e.g. Monks et al., 2019). Furthermore, more accurate estimates of water loss components, such as background leakage, can be achieved (Loureiro et al., 2014). Newer smart meters may even be equipped with inbuilt noise loggers intended to improve post-meter and network leak detection (Kamstrup, 2019b). Moreover, Bragalli et al. (2019) analysed the impact of an increasing number of missing smart meters from apartment blocks and single houses on the error of the estimated water loss. The results of estimated water loss worsened notably for missing large-scale consumers (i.e. apartment blocks) compared to single houses. This is important for utilities, as the deployment of smart meters at large-scale consumers may be more troublesome, for example, because a different type of smart meter needs to be installed and a larger number of people/businesses are affected.

Water quality monitoring. Monks et al. (2019) list only few water quality related benefits of smart meters, but state that the implementation of water quality testing at customer meter level might help to reduce the number of required audits. Yet, according to Blokker (2019), no microbiological parameters (or parameters related to water quality, such as pH) are monitored with smart meters. Temperature is an often sampled parameter as it may be collected as spin-off from ultrasonic smart meters; however, utilities struggle to know how to use the data (Blokker, 2019; Paper IV). Additional potential benefits of (smart meter) temperature data based on Papers III and IV are listed in Sections 4.1 and 4.2.

Also, smart meter data can be beneficial as input to enhanced demand knowledge from a water-energy-nexus point of view (Stewart et al., 2018). Smart meter temperature and consumption data can help to identify certain pipe and soil characteristics and estimate the heat transfer potential of drinking water in WDNs, e.g. as a low-temperature source in district heating systems (Papers IV and V). Furthermore, smart meter data can be an integral part of estimating the wastewater flow and as a validation tool of in-sewer flow observations in urban drainage management (Paper VI).

3 The urban water data warehouse

Gartner’s IT glossary (Gartner, 2019) defines a data warehouse as a “storage architecture designed to hold data extracted from transaction systems, operational data stores and external sources”. I envision the *urban water data warehouse* as the location where data from multiple sources, such as sensors and asset data, is collected as well as made available for a variety of applications through extract, transform and load processes. As utilities become increasingly data-driven, it is important that they secure highly reliable data in the warehouse. However, prior to the rollout of new ICT, such as smart meters and flow or pressure sensors throughout their network, a utility should ask itself what its general goal is in installing the new technology. Is it to facilitate the billing of customers, to provide near-real time leakage monitoring and demand forecasts or to provide water quality estimates for use in the case of contamination?

If data-driven goals are of significance for a utility, the above-mentioned questions are important because the actual use of the data in the utility drives the maintenance of a high reliability of the collected data in WDNs (Papers I–IV and VI).

In other words, I speculate that data used proactively by utilities show fewer anomalies and these, if detected, are given higher priority. This is not the case for sparsely used data, where errors prevail for longer periods (Papers II and III). This may in some cases lead to a wrong understanding of the data quality and false trust in data.

The applicability of data depends, among other things, on where the data is collected. However, identifying optimal sensor locations – such as of water quality sensors for contaminant detection (e.g. Adedaja et al., 2018), pressure meters for burst detection (Wu and Liu, 2017), or which consumers should be favoured during smart meter enrolment (e.g. Bragalli et al., 2019) – were deemed outside the scope of this thesis.

As with the location of data collecting devices, the sampling resolution (i.e. the time step between measured data points) has a significant impact on the achievement of the ‘digitalisation goals’ of a utility, as described below.

3.1 Sampling resolution and gap filling

ICT deployments have a trade-off between transmission costs and applicability (amount of collected information). Among other things, transmission costs depend on the selected network, message sizes, and costs of data collection

(Kulkarni and Farnham, 2016; Mekki et al., 2019).

Cominola et al. (2018) and Nguyen et al. (2018) showed the benefits of having very fine sampling resolutions of smart meter demand observations (< 1 minute) resulting in successful end-use disaggregation. Furthermore, Cominola et al. (2018) showed that there can be up to 62% difference in the magnitude of peak demands based on a 10-second compared to a 24-hour sampling resolution. Gurung et al. (2014) showed that this is important because detailed knowledge about peak demands based on a fine sampling resolution can provide better demand patterns and subsequently reduce network augmentation. In terms of generating reliable water quality simulations, the spatial aggregation and sampling resolution of demands are of importance (Blokker et al., 2008). Depending on the set-up, a 1-hour time step can be sufficient for water quality simulations of large demand aggregations (e.g. transportation system models), but sampling resolutions below 5 minutes are required to generate reliable results for smaller networks (Blokker et al., 2010, 2008). Also, Creaco et al. (2017) showed that pressure simulations are largely affected by the sampling resolution and demand modelling approach. Whereas a top-down approach (applying a demand multiplier on coarse readings) first leads to acceptable simulations at sampling resolutions of demands and modelling time steps > 1 hour, a bottom-up stochastic approach (generating unique demand profiles for each consumer) can generate reliable pressure simulations at time steps > 2 minutes, which is important when running near-real-time simulations (Creaco et al., 2017).

Smart meter data represent the actual demand from each household. Incorporating such data into hydraulic models can enable the monitoring of WDNs in near real-time. Compared to stochastic bottom-up demand allocation approaches (e.g. Blokker et al. (2010) and Creaco et al. (2017)), Paper I applied the actual consumption from each individual household's smart meter as input to a hydraulic model of a case study DMA, to assess the effect of common smart meter sampling resolutions on pressure, water age and total consumption simulations. Accumulated volume readings from each smart meter were available with a sampling resolution averaging 30 minutes/sample. To align the random nature of each sample's timestamp, gaps between samples were filled by linear interpolation and a demand-pattern-based approach. Hereby, aligned and unique consumption data time series were generated for each smart meter, making it possible to aggregate the consumption from different consumers in the hydraulic model's nodes. The model was ran in 5-minute time steps to compare the model's results with the 5-minute sampling resolution of the DMA

inflow. This approach was repeated with gradually coarser sampling resolutions. Resolutions commonly sampled by Danish utilities were thus established, ranging between 30 minutes to 24 hours/sample (Paper I).

For pressure simulations, the results showed that there was little difference in the root mean square error (RMSE) between assessed sampling resolutions, mainly because of the overall low pressure head loss through the network. Even combining a representative demand pattern with very coarse demand readings (e.g. quarterly) turned out to be sufficient compared to finer sampling resolutions when simulating pressure (RMSE < 0.1 m) (Paper I).

In terms of water age, however, there is a clear benefit from collecting finer sampling resolutions. Thus, if utilities are concerned about retention times and interested in water quality simulations, finer sampling resolutions should be favoured (Paper I). This is also shown in Paper IV, where the sampling resolution of around 30 minutes/sample is the major restraining factor of more detailed temperature analyses (see also Section 4.2).

When very coarse readings (e.g. quarterly) are available in combination with reliable demand patterns, the RMSE of total DMA consumption was only around 0.5 m³/hour higher (though approximately equal to the DMA's MNF) than when using a fine sampling resolution (Paper I). However, such a good demand pattern might only be available from smart meter data which has a fine resolution.

When having sampling resolutions finer or equal to 2 hours, there was little difference between the two gap-filling methods; at coarser resolutions, however, the demand-pattern-based gap-filling method outperformed linear interpolation.

Thus, the following could be concluded on sampling resolutions and gap filling from smart meters.

- Fine sampling resolutions (< 1 hour) should be used when utilities are interested in a high level of detail about the consumption in their systems. This knowledge improves the validity of water quality simulations relying on accurate simulation of the water's retention time in the WDN.
- Coarse sampling resolutions (1–24 hours) are sufficient when the overall goal includes daily water loss assessments and when assessing the pressure in areas with overall low pressure head loss.

- At coarse sampling resolutions (> 2 hours), a demand pattern based approach to gap filling between samples is preferred over linear interpolation.

3.2 Validating and reconstructing data

Data cleansing, including validation and reconstruction processes, is a major concern in almost all fields of research and applications related to environmental water management (Sun and Scanlon, 2019). Owing the increasing amount of data collected, the quality control, validation and easier access to data are “expected to be at the heart of the next steps in hydroinformatics” (Makropoulos and Savić, 2019). Without such steps, the value of the collected data decreases notably and might lead to false acts on incorrect data (e.g. Blumensaat et al., 2019). The first statement in the Introduction could thus be reformulated:

“Water and wastewater utilities embracing digital solutions have to include data validation and reconstruction processes. There is really no alternative.”

Data validation has been part of the urban water sector for decades, covering fields such as urban hydrology (e.g. Branisavljević et al., 2011; Mourad and Bertrand-Krajewski, 2002) and water distribution systems (e.g. Cugueró-Escofet et al., 2016; García et al., 2017; Quevedo et al., 2010). Unfortunately, the implementation and focus on automatic data validation has not kept up with the interest in data collection and application. This needs to change as utilities become increasingly data-driven. The steps of validation and reconstruction may seem very obvious, but real-world examples (Papers II, III and VI) reveal that this is tricky and often only practised to a limited extent.

3.2.1 Need for automatic data cleansing

An example of why automatic validation and reconstruction of water network data is of such crucial importance is shown for a real-pressure time series collected at the exit of a DMA in Denmark (Figure 4). Visualization of the raw time series (Figure 4a) highlights data points that can be deemed unfeasible ($1.5 \cdot 10^9$ m). Removing these anomalies from further analysis results in Figure 4b, indicating that additional data points exist that are highly doubtful in the area (recurring pressure measurements). When omitting these anomalies, the time series looks more trustworthy, as shown in January 2014 (Figure 4c).

However, even here, anomalies are still visible in the form of two flatline segments. After a more exhaustive identification and potential reconstruction of invalid data points, the data quality may be considered as ‘good’. For example, Figure 4d shows a day of plausible measurements in 2016. Nevertheless, a control measurement was conducted on the same day by the utility, highlighting a pressure difference of 3 m. Additional control measurements conducted at earlier and later stages confirmed similar deviations. This shows that, even though the data looks ‘good’, there might be essential flaws in it, as the pressure sensor might have incorrect time settings, drifted over many years, or wrong location coordinates in the utility’s asset database.

Being one example out of thousands of sensors deployed means an obvious need for automatic validation and reconstruction of data. Even more important, this example shows that even though data looks ‘good’, users should retain some scepticism regarding the truth of the data, as this otherwise may lead to incorrect decisions, e.g. when this data is used in a real-time control setup.

3.2.2 Anomaly types

As seen in Figure 4, many types of anomalies exist in the data collected in WDNs, owing to transmission errors, meter malfunctions, system operations, etc. (for example, Loureiro et al., 2016; Quevedo et al., 2017). In Paper II,

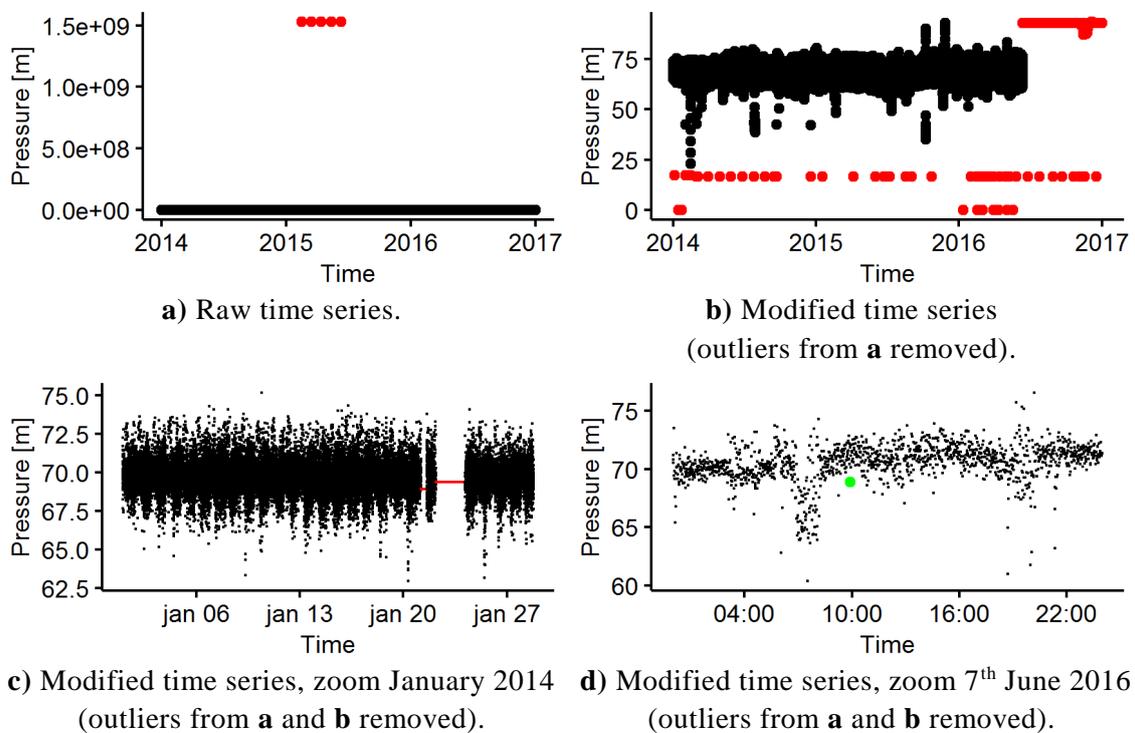


Figure 4. Example of pressure time series (● measurement ● anomaly ● control) collected at a district metered area outlet in Denmark.

anomaly types were subdivided into three different categories (Types 1–3), described in detail in Figure 5. However, Papers III, IV and VI revealed a more ‘general’ anomaly type that is difficult to quantify, but needs to be addressed: if the understanding of system attributes is misconceived (Type 0, Figure 5) the application of data will be troublesome.

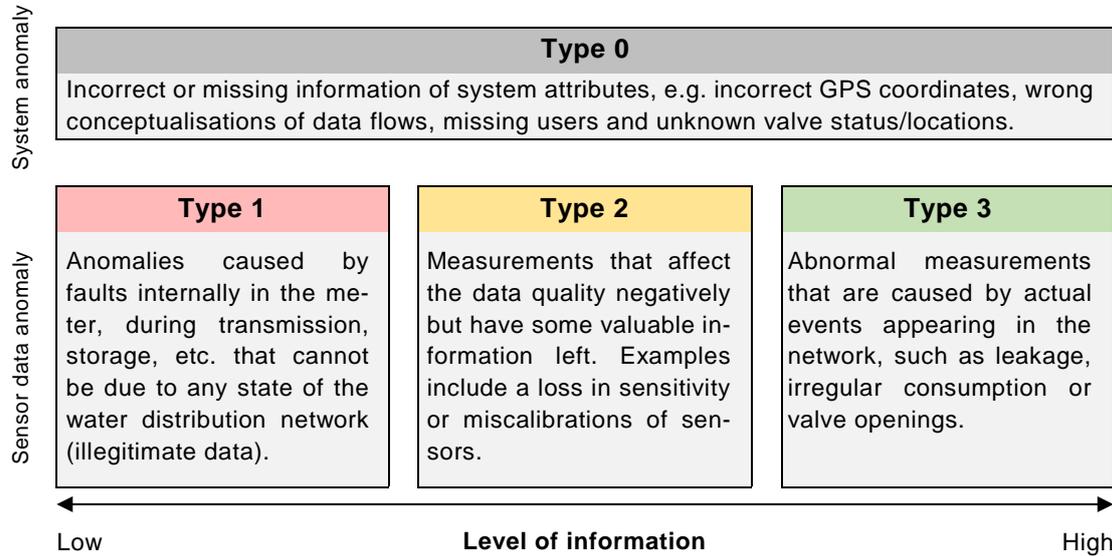


Figure 5. Classification of anomalous data collected in water distribution networks (Paper II) with additional system anomaly description of Type 0.

3.2.3 Identification of anomalies

Seven tests were proposed in Paper II to detect anomalies of Types 1 and 2, i.e. errors that do not reflect the true state of the water distribution systems: duplicate timestamp test; illegitimate format test; range test; rate of change test; flatline test; timestamp inconsistency test; and a timestamp drift test. An online tool was generated during the PhD research, whereby the seven tests can be run: <https://leakagemanagement.net/meter-validate/>.

The seven tests were run on flow and pressure data sets from three Danish utilities, covering on average 32 months. The results showed that a varying proportion of anomalies were found in all utilities’ data sets, averaging between 3% and 35%. Even though only a small proportion of data points was identified by the timestamp inconsistency test, the analysis showed that these anomalies covered around 10% of the time in the data sets of the three utilities (i.e., on average, 10% of data was missing). This clearly highlights the strong need for validation and reconstruction of the collected data in Danish utilities. Furthermore, the tool’s practicability for reducing the amount of anomalies prevailing for long periods is exemplified in a monitoring or operational setup (Figure 6). Figure 6a shows the incoming data for ten flow (Q) and ten pressure

(P) sensors over time, with invalidated data coloured in red from the seven tests. The figure shows 1) that some anomalies occur at the same time, and 2) that there are specific sensors with either no data or a higher number of anomalies. Such a visualisation can help the utility to discover overall issues within data collection and react quickly to such data quality and collection issues. Figure 6b shows the percentage of flatline anomalies (being the major anomaly contributor) for all 22 pressure meters in one of the utilities from Paper II. Such a visualisation can be used to prioritise which meters should be calibrated first and to rank the reliability of the meters.

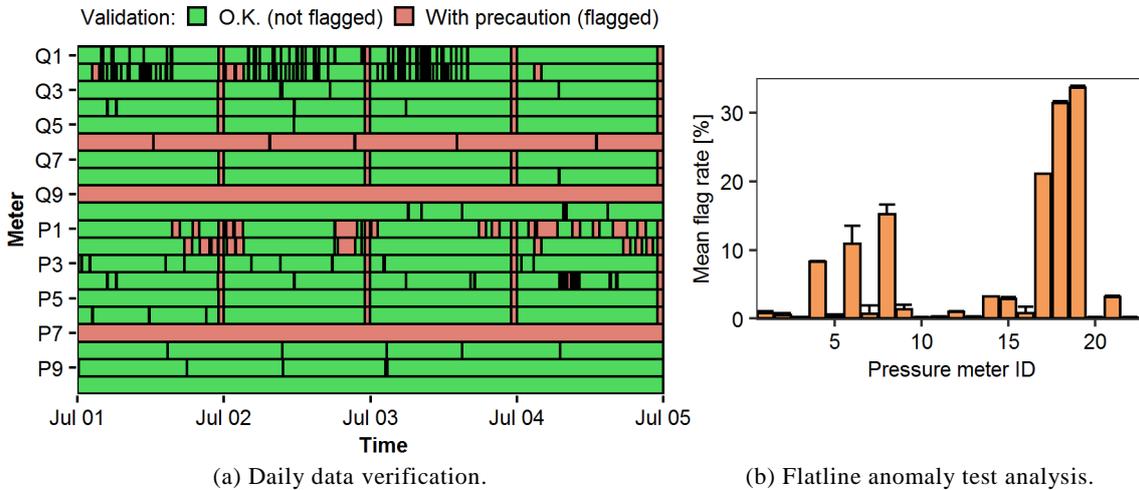


Figure 6. Example of anomaly visualisation for operational use. (a) Raw meter data validation from ten flow (Q) and ten pressure (P) meters in July 2015. (b) Mean flag rate based on all raw data points for the flatline test in pressure meters. Whiskers display the total flag rate based on all anomaly tests in the individual pressure meter, highlighting that flatline anomalies are the major anomaly source. From supporting information, Paper II.

3.2.4 The value of anomalies

Analysis of anomalies can help to identify general errors in data collection and transmission. For example, Figure 6a shows that data from multiple sensors were repeatedly invalidated for one hour around midnight. Whereas these meters were not physically attached to the same part of the WDN, the data was collected in the same database and the error is therefore likely to have originated from the utility’s database setup. Therefore, anomalies were stored as an amendment to the operational database in a ‘malfunction indicator database’ in Paper II. For example, the Jaccard Index (e.g. Tan et al., 2006) can then be applied to identify similarities between anomalies. Application of this similarity measure revealed multiple issues, potentially related to transmission and connectivity problems (Paper II).

3.2.5 Higher-level validation of anomalies and reconstruction

As stated in Paper II, various methods exist for reconstructing invalidated data. These include time-series analysis, physically based models or machine learning approaches such as artificial neural networks (ANNs) (e.g. Branisavljević et al., 2011; García et al., 2017; Mounce et al., 2010; Quevedo et al., 2010). Figure 7 shows an example of training multiple ANNs to reconstruct missing or invalidated pressure and flow data, respectively, and to validate Type 3 anomalies. The ANN training was solely based on validated data from other sensors with similar dynamics deployed in the utilities' networks (see Appendix B for further information).

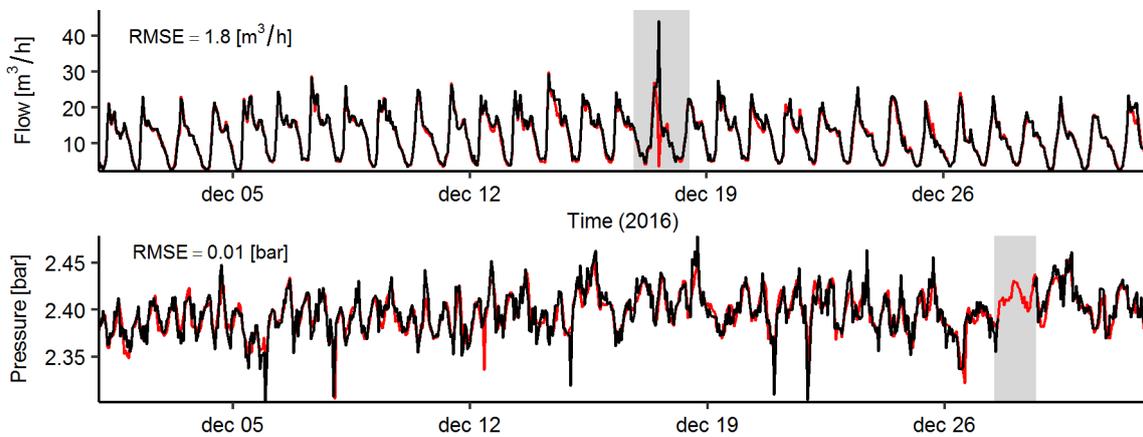


Figure 7. Measured (—) and reconstructed (—) pressure and flow time series from two Danish utilities. See Appendix B for more information. RMSE = root mean square error.

In both time series, the reconstructed data follow the real measurements adequately (Figure 7). The upper panel shows the potential of reconstructing data for validation purposes and identification of Type 3 anomalies. On 18 December a much higher flow rate was measured than predicted. As the reconstructed data is solely based on other sensors in the network, this might indicate that other sensors in the network were affected simultaneously, indicating that this event was a Type 3 anomaly caused by a legitimate measured network operation affecting larger parts of the network. The lower panel indicates that missing data around 28 December can be filled by ANNs to generate uniform data streams.

Thus, the following can be concluded based on the work from Paper II.

- Pressure and flow data sets from three Danish utilities revealed a high need for validation, as a large proportion of anomalies prevailed for long periods in the utilities' data sets, with one data set having 35% invalid data (Types 1 and 2).
- The proposed methodology helps utilities to monitor their data collection and can be used as an operational tool to quickly act on and reduce Type 1 and Type 2 anomalies.
- On average, 10% of the period covered by the utilities' data sets was missing. This highlights the need for reconstruction processes, as data applications require uniform and reliable data streams. Reconstruction can also be used to further validate Type 1, 2 and 3 anomalies.

4 Novel data applications and modelling for enhanced system analysis

After collection and quality control, the data can be applied for the monitoring, operation, control and design of WDNs. Research has begun into digital ‘multi-utilities’, whereby data from other sectors, such as gas and electricity, are coupled with water data to produce enhanced demand management strategies and the opportunity for new services and businesses (Stewart et al., 2018). However, silo thinking may be a limiting factor for sharing of data from different sources both within and between institutions, because utilities are complex organisations with multiple departments and differing objectives (Kulkarni and Farnham, 2016; Sarni et al., 2019). Thus, the first stage of data collection and processing (Figure 3) should not be seen as a process covering only data from the WDN, but also include data from other sectors and sources (that needs to be easily accessible (e.g. Makropoulos and Savić, 2019; Sarni et al., 2019)). Likewise, the application of data is not restricted solely to the field of WDN analysis (Figure 3).

Even if data sharing between institutions is successful, one major problem of the increased data collection is the missing expertise, experience and examples of possible applications of the collected data streams. In the following, I will demonstrate four examples of ‘surplus value’ from the collected data, with a major focus on temperature data, as utilities often do not know what to do with this data. The first two examples concern mainly optimised water distribution system analysis, whereas the latter two display the external fields of district heating and urban drainage.

Temperature simulations in water distribution networks

Besides the simple water quality monitoring benefits discussed in Chapter 2, I believe that temperature measurements from both smart meters and the WDN have an important function that is not yet fully exploited: As the water is heated up or cooled down throughout the network, temperature measurements store to some extent the history about the path of the water. It is this effect that is explored in Papers III–V.

Heat transfer models describe the change in water temperature over time in a WDN (Blokker and Pieterse-Quirijns, 2013; De Pasquale et al., 2017; Paper V), and include four particularly important parameters that need to be considered: 1) the undisturbed soil temperature, determining whether the drinking

water is heated up or cooled down throughout the network; 2) the inlet (initial) water temperature; 3) the time that the water has spent in the network; and 4) the heat transfer coefficient, taking into account thermal resistances of soil and pipe materials, among other things. The latter is subject to variations in the literature, and Papers III–V apply the Hubeck-Graudal implementation (Paper V).

Figure 8 illustrates the concept and effect of these four parameters on a simple network setup (Figure 8a) with two parallel running pipes with different insulating materials. Varying demand and closing of either pipe will change the time spent by the water in the WDN. Figure 8b shows the effect of varying demand on the retention time of the water in the WDN and thus the simulated end temperature, depending on whether the unlined cast iron or polyethylene pipe are closed or both pipes are kept open.

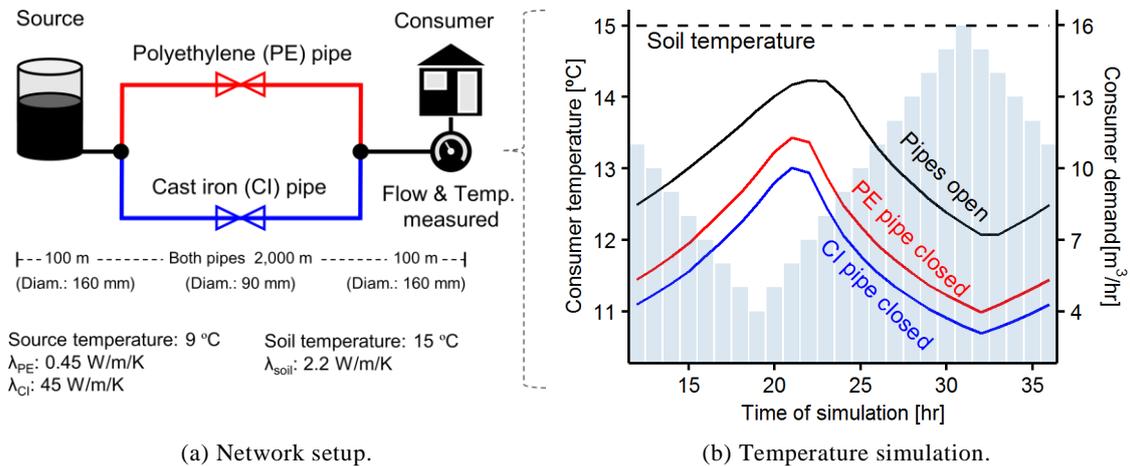


Figure 8. Modelled temperature variation when closing a polyethylene (PE) or unlined cast iron (CI) pipe, or having both pipes open. Based on the temperature model described in Paper IV.

In the following, two examples are described showing that temperature data from the distribution system, as well as from smart meters, can be used to update utilities’ understanding of their systems.

4.1 Improved system understanding through temperature modelling

As shown in Figure 8, the state of valves in particular can have a major effect on the temperature fluctuations in the WDN. To show whether the application of temperature data can help in identifying major system anomalies (Type 0, Figure 5) otherwise overseen in WDNs, the effect of open and closed valves on temperature simulations was analysed in Paper III.

Knowing the location and correct status of valves (i.e. open or closed) is of paramount importance for utilities. The list of possible negative consequences is long and includes: areas incorrectly narrowed down during contamination events; improperly working fire hydrants; and a greater proportion of affected customers without service during construction works (Deb et al., 2012; Delgado and Lansey, 2009; Wilson, 2011). Moreover, anecdotal evidence from Denmark indicates that DMA water balances that do not add up are often a result of unknown valve status. To mitigate these problems, utilities often maintain databases on the location and status of valves (Walski et al., 2003). Even if utilities are unaware about the actual status of a given number of valves, the state of the valves is often assumed known in WDN model analyses (e.g. Sophocleous et al., 2017; Wu et al., 2012). Including such possible incorrect valve settings in a hydraulic model generates unreliable results, particularly when simulating water quality (Savic et al., 2009).

In the literature, examples exist where the correct location and status of valves were identified by combining a hydraulic model with pressure and flow measurements (e.g. Delgado and Lansey, 2009; Do et al., 2018; Sophocleous et al., 2017; Walski et al., 2014; Wu et al., 2012). Yet none of these studies has incorporated temperature measurements as an additional parameter. One major issue of combining models with flow and pressure data to determine the valve status is the lack of significant head loss throughout the WDN. For example, there was as little as an average of 0.2 m head loss throughout the DMA in Paper I. The head loss can be increased, e.g. by opening hydrants, but it is a rather labour-intensive approach. Thus, it was investigated whether temperature data could simplify this approach and increase chances of correct valve status identification and location.

In Paper III, a semi-synthetic and a real case study were presented and a genetic algorithm (GA) was applied to identify the location and status of valves, including temperature data. This type of optimisation algorithm was run to identify the best combination of open and closed valves in the WDN by minimizing a fitness function. As the search space of open and closed valve combinations can be large, the number of assessed valves was reduced where possible. For example, valves that disconnected consumers entirely from the network when closed were not considered in the search space, since this would be reported by consumers otherwise. The aim of the semi-synthetic case study was to show whether temperature measurements alone, e.g. available from smart meters, could be utilised to identify correct valve status. In the second case study, real

temperature, flow and pressure measurements from a transportation network were applied and the valve settings were manually tested in the field.

4.1.1 Semi-synthetic case study

The semi-synthetic case study included real temperature measurements at the inlet of a DMA. First, the synthetic ‘true’ nodal temperatures were generated by closing five out of the 379 valves included in the search space. The remaining valves were left open. Next, starting with the assumption of all valves being open, the GA was applied to identify the five closed valves based solely on the ‘true’ nodal temperatures. These would in a real case system stem from smart meters; however, Paper IV indicated that in reality the collection of this ‘true’ data set is difficult (see also Section 4.2).

A total of 48 different GA setups were run, each identifying a best set of closed valves termed ‘best fit’. In these runs, the perfect solution was identified 16 times. In total, the ‘best fits’ identified 183 out of 240 possible ‘truly’ closed valves and 21 incorrectly closed valves. However, these incorrectly closed valves were often in close proximity to the actual closed valves, and thus the exercise could potentially still guide utilities to areas of concern. In some runs, the GA was stuck in local minima, among other things, because of small temperature differences occurring when closing selected valves, the relative large search space and the specified GA parameters.

The GA runs showed that filling the initial GA population with a weighted approach (incorporating some prior knowledge based on simulations) and a larger initial population had the highest impact on the outcome. This is of importance, as it can improve future modelling approaches. Where possible, the search space should be reduced by incorporating the knowledge of preliminary model runs and operators: “This valve should be closed in the model, and be excluded from optimisation, as it was tested last week”, increasing the overall success rate of the GA.

4.1.2 Real transportation network case study

The second case study analysed valve settings, a transportation WDN model and a week’s measured flow and pressure measurements at ten DMA inlets, one tank and two waterworks. As in the synthetic case, a GA was used to identify valves status. Even though some of these were marked as ‘closed’ by the utility, this information was not used in the modelling procedure owing to the high likelihood that it was incorrect. In the analysis, temperature measurements

at two DMA inlets in combination with all pressure measurements were incorporated into a fitness function. The GA was then run with no prior knowledge (i.e. all valves open) about the valves' actual status. Among other things, different weighting scenarios between temperature and pressure measurements were applied in the fitness function to highlight the individual and combined effect of the two parameters on valve status identification.

Identification of 'system anomalies' (Type 0) prior to application

The interest in temperature data revealed various Type 0 anomalies (Figure 5). First, two temperature meters had to be excluded from analysis as they measured doubtful values, ranging between < -5 and > 40 °C. Moreover, one meter with reliable measurements was thought to be located at the inlet of a tank, but preliminary temperature simulations resulted in a contradiction between the meter's temperature measurements and another meter's data set. Discussions with the utility revealed two different SCADA setup diagrams, where the meter on the first diagram was located at the tank inlet (incorrect) and at the outlet (correct) on the second diagram.

All pressure measurement sites showed reliable dynamics, but test samples taken at the sites revealed pressure offsets between -1.3 and 3 m. Consequently, the fitness function had to be modified by subtracting the simulated and measured median pressure to account for unknown drifts but still incorporating the dynamics in the data.

Identification of 'system anomalies' (Type 0) after application

The best fit of 30 GA runs resulted in the identification of 41 different out of a total of 106 possible closed valves. This number was reduced to nine valves by only considering those reappearing in at least 50% of the weighting scenarios. Two of these valves were identified only by the weighting scenario runs that solely included temperature measurements. Going to the field to check the actual settings of the valves revealed the following about the nine valves:

- Two valves were locked in and reported faulty (thus, these valves' status were unknown).
- Three valves were inaccessible, e.g. owing to heavy vegetation.
- One valve, marked as closed by the utility and scoring the highest score in all GA runs (closed in 24 out of 30 runs) was possibly open (unfortunately, the technician was in doubt).

- Three valves no longer existed, and these and adjacent parallel pipes should be removed from the utility's asset and hydraulic WDN model. The WDN model did thus not reflect the real system.

Overall, application of the temperature model and field tests casts doubt on the validity of the hydraulic model. Moreover, at one critical location, temperature measurements indicated clearly that the water was distributed and mixed differently than anticipated by the utility. Pressure measurements and simulations indicated that the water originated from a nearby waterworks, whereas temperature measurements indicated that the water originated from a nearby tank. The latter could partially be validated during a period where the waterworks was running on low capacity. Only the weighting scenarios considering temperature pointed to this critical location and a 'valve station' was discovered with locked-in and rusty valves and bypasses that were not part of the asset database.

The application of the methodology in Paper III revealed that poor data quantity and quality are major restricting factors when using hydraulic and temperature models for determining valve status. Additional validation tests of the temperature measurements and a more thorough analysis of the real status of the valves are required. At present, the method can be used to identify errors in the utility's asset database (i.e. Type 0 anomalies, Figure 5) and of Types 1 and 2, but it is first possible to pinpoint valves that really are closed if a high data reliability is also assured. Thus, whereas the valves might be closed as indicated by the models, the high uncertainty in the input data and unclear valve tests make it difficult to trust the modelling results at present.

The following can be concluded on the applied methodology.

- The evaluation of temperature data in combination with hydraulic data, prior to application of the methodology, revealed various errors otherwise overseen by the utility: invalid temperature and pressure observations and incorrect conceptualizations of where data is collected.
- The applied methodology resulted in the discovery of forgotten and faulty valves and other errors originating from the utility's asset database. However, as long as the high uncertainty in the input data is not reduced it is not possible to apply advanced methods such as the one presented for valve status detection.
- At one location, temperature data showed that the water was mixed differently than anticipated by the utility which, among other things, based its knowledge on (perhaps faulty) pressure measurements. This highlighted

the value of using temperature data for an improved understanding of the WDN.

4.2 Smart meter temperature data

One of the restraining factors of the applied methodology in Paper III was the reduced number of available temperature measurements as well as their limited spatial distribution. Thus, Paper IV looked into the question of whether smart meter temperature data can be used to overcome this limit. Moreover, Paper IV was applied to identify whether smart meter temperature data can be used to validate the temperature model also applied in Paper III (Section 4.1).

In the main, smart meter data is collected for improved billing purposes or enhanced water loss assessments; the temperature data, however, is collected as a ‘spin-off’ product and not optimised for usage. Whereas the temperature data applied in Paper III could be compared directly with temperature observations from the WDN, temperature data from smart meters does not directly represent the WDN temperature. This is because the water may spend a substantial amount of time in service pipes and (in-) house connections before being metered. Based on varying residence times in such connections, the smart meter temperature data can represent nearly the WDN temperature, the soil temperature, a mixture in between the two, or a temperature affected by other external heat sources such as the indoor air temperature (Figure 9). The applied real smart meter data in Paper IV thus required some classification and filtration prior to application.

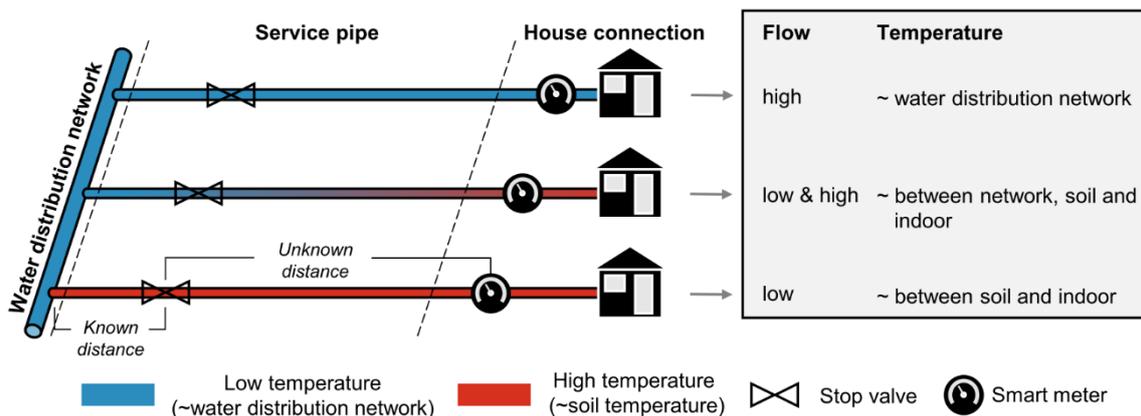


Figure 9. Variation in demand affects the residence time of the water in service pipes and house connections, leading to differing smart meter temperature measurements at the consumers’ homes. Example from summer, with higher soil than water source temperatures.

4.2.1 Filtration and classification of smart meter temperature data

In the case study DMA, two weeks of temperature and consumed volume smart

meter data were available at all homes. The data covered a period during August, with higher soil than DMA inlet temperatures. For most consumers, information about service pipes from the asset database were available, only covering the distance between WDN mains and stop valves on the consumers' properties (Figure 9). However, no information about the pipe stretch from stop valves to the smart meter locations existed and the smart meter location was unknown, as the obtained GPS coordinates only represented the coordinates of the consumers' properties. Based on this scarce information, the service pipes were prolonged linearly from each stop valve to the GPS coordinates of each property.

Identification of 'system anomalies' (Type 0) through soil temperature estimates

The soil temperature was used as a boundary condition in the heat transfer model, but the actual soil temperature in the DMA was unknown. To mitigate this, the soil temperature was estimated using the smart meter temperature data. Low or no consumption lead to stagnant or slow-flowing water in service pipes, with the water temperature gradually going into equilibrium with the soil temperature (Figure 9). Sudden demand pushing this 'stagnant volume' into consumers' homes can result in temperature samples which, to some extent, represent the soil temperature. Owing the uncertainty of service pipe lengths (and their diameters) between stop valves and smart meter locations and the time of actual consumption, a conservative filtering was applied to identify samples representing the soil temperature. This was done by only accepting samples representing: 1) retention times > 3 hours in the service pipe; 2) a consumption equal to 20–80% of the service pipe volume; and 3) temperature samples that were at most 15 minutes old (otherwise they may potentially have been affected by, and represented, indoor temperatures).

Samples passing this conservative filtration revealed a slowly decreasing soil temperature over two weeks with some clear outliers ($> \pm 3$ °C from mean). Manual inspection of all smart meter locations in a geographic information system revealed that these outlying temperatures were found in the only location with two-storey houses (each storey had its own smart meter). Thus, the actual service pipe and house connections were much longer than anticipated. Furthermore, outliers of soil temperature, being colder than the mean temperature, revealed locations where no stop valve information was available; thus, the model-building process connected service pipes to the nearest main. The too-cold soil temperature estimates indicated that the service pipe locations,

diameter or length at these locations were incorrect. Thus, analysing the temperature data based on knowledge of service pipes revealed incorrect and missing information about consumers in the utility's asset database (i.e. Type 0 anomalies, Figure 5). Real soil temperatures are currently being sampled to validate the ones estimated from smart meters.

Water distribution network temperatures

The heat transfer model is validated by comparing the simulated network temperature with 'measured' smart meter temperatures. Except for the DMA inlet, however, the actual temperature throughout the DMA was unknown, as no temperature sensors were installed directly on the WDN mains. Like the soil temperature, the 'measured' WDN temperature was therefore estimated using smart meter temperature data. WDN temperatures were estimated using smart meter samples where at least 400% of the service pipe volume was consumed within 15 minutes. This conservative filtering was applied to overcome issues based on too coarse network skeletonisation, and to reduce uncertainties owing to unknown service pipe and home connection lengths and diameters. Some samples passing the filtering were above the soil temperature, indicating uncertainty in the applied filtering process. This uncertainty is expected to stem from external heating sources, incorrect pipe characteristics estimations or higher local soil temperatures than expected. The remaining temperature samples showed expected variations and were used to assess the temperature model's validity.

4.2.2 Temperature model analysis with smart meter data

The temperature model, with the soil temperature estimated from smart meter measurements and the measured inlet DMA temperature used as boundary conditions, was run and compared to the estimated 'measured' network temperatures for validation purposes. Samples representing the WDN temperature showed good match with the simulated temperature in all but two nodes having a mean RMSE of 1 °C (Figure 10). Two nodes were removed from the analysis, as their RMSE stood out (> 6 °C, Figure 10a), improving the mean RMSE to 0.9 °C. For one node, the reason was a too coarse skeletonisation during the building process of the hydraulic model. In the hydraulic model, smart meter temperatures were used to estimate the network temperature in the nearest node; in the real WDN, however, the temperature samples represented the temperature in a side-branch, explaining the high temperature deviation between model and samples. In the other case, discussions with the utility revealed that a consumer had incorrect GPS coordinates and was not part of the DMA. Thus, running a temperature model showed that the model could validate data and

identify larger flaws such as incorrect consumers that would otherwise not be detected (i.e. Type 0 anomalies, Figure 5).

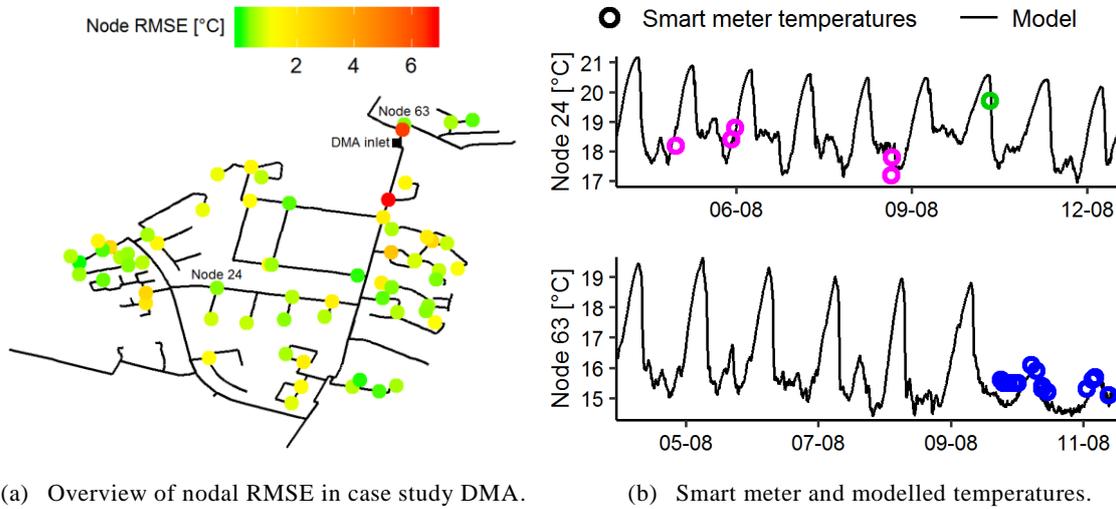


Figure 10. Examples of temperature variation throughout the case study water distribution network (Paper IV). (a) Nodal root mean square error variation (RMSE) based on temperature model and smart meter temperatures. (b) Example of two nodes with an RMSE of 0.36 °C (upper panel) and 0.32 °C (lower panel) where distinct colours represent individual homes’ smart meters.

I believe that this small analysis revealed only a small fraction of smart meter temperature potentials. Future applications include, among other things, the following (see Paper IV for additional potentials).

- Leakage detection, as higher flow rates lead to different temperature gradients than anticipated. Reliable temperature simulations can then be used to pinpoint locations with large differences between model output and temperature observations.
- Valve status identification (similar to Paper III), as the water is distributed differently than anticipated, changing the temperature profile throughout the WDN.
- Consumer alerts, e.g. when temperatures are too high or too low over a certain period.

Hydraulic network setup and sampling resolution

In particular, the way smart meter temperature data was sampled posed a challenge to the application of the data. Smart meters in the case study DMA sampled a large proportion of data, where the temperatures represented a mixture between service pipe and household connection temperatures. In Paper IV, less than 0.2% of all smart meter temperature data passed the conservative filtration

for estimates of soil and network temperatures. A more intelligent metering than sampling every hour, or quasi-random as in the case study, would be preferable. For example, sampling each time a certain volume has passed the meter could increase the percentage of applicable data. The average sampling resolution in the case study of around 30 minutes was insufficient. Thus, a finer sampling resolution (at least during periods of consumption) would improve the applicability, as the uncertainty of when the water actually has been used is then reduced. The second challenge includes the level of skeletonisation of the hydraulic model, which revealed a high impact on temperature simulations. A higher level of detail of the hydraulic model is required in future applications (for example, not bundling consumers in nodes), making it possible to represent temperatures at the actual location in the WDN, and also improving water quality simulations.

Currently the following can be concluded on smart meter temperature data.

- The temperatures throughout the DMA were simulated to a satisfying degree (average RMSE of 0.9 °C). This made it possible to use the temperature model and compare the output with smart meter temperature data to highlight incorrect information about service pipes and consumers in the asset database, not found easily without these temperature data.
- Based on the applied filtering, only a small proportion (< 0.2%) of the collected data could be used to represent WDN and soil temperatures, revealing that smart meter temperature data is not sampled in a preferable manner. An improvement would include finer sampling resolutions (< 30 minutes) or additional sampling during times of high consumption, which would increase the possibility of representing WDN temperatures.
- The list of potential valuable applications of smart meter temperature data needing further research is long, including improved leakage detection and valve status identification.

4.3 Drinking water as a low-temperature source in district heating systems

When electricity prices are low, the operation of electrically driven heat pumps, extracting energy from WDNs and transferring it to district heating systems, seems to be a perfect match. This is firstly because the heat provided from the WDN to the district heating system is potentially replacing less sustainable sources of energy, such as fuels used in combined heat and power

plants. Secondly, it is because WDNs have stable temperatures and are available in many places around the world. Finally, low drinking water temperatures may reduce the risk of biofilm growth (Liu et al., 2016) and increase the share of water complying with the recommended upper temperature limit of 12 °C at the tap in Denmark (Ministry of Environment and Food of Denmark, 2018b). Real-case examples of heat pumps installed in WDNs already exist, as exemplified by a smaller utility in Northern Jutland, Denmark, supplying around 15% of the annual required energy for district heating (Cronborg A/S, 2014). However, there is a potential socioeconomic downside: does the lower water temperature lead to more energy being required to heat the water at consumers' homes?

To analyse this effect, Paper V simulated the temperature throughout Copenhagen's WDN and included the effect of eight heat pumps, deployed at favourable locations in the model on the WDN mains. Figure 11 exemplifies the linkage between the WDN and a district heating system with a heat pump. The drinking water is cooled significantly at the heat pump location and, in this example, due to the soil being warmer than the water, some of the extracted energy is regained from the ground over time. Hereby, the water temperature slowly re-approaches the soil temperature downstream of the heat pump (Figure 11).

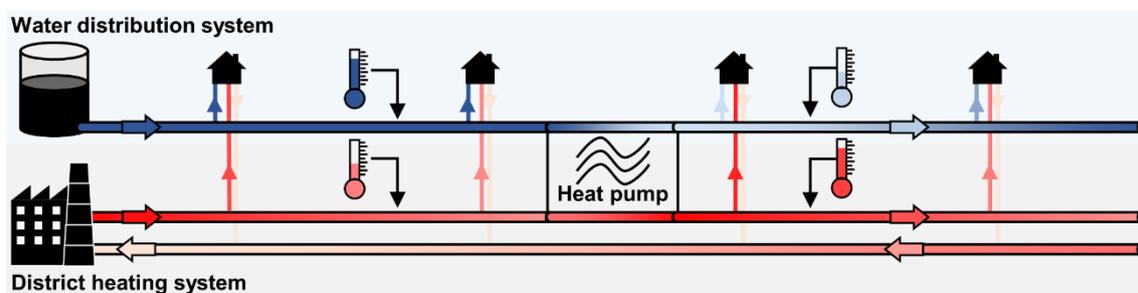


Figure 11. Linkage between a water distribution system and district heating system by a heat pump cooling the drinking water below soil temperature and transferring heat to the district heating system.

Case studies exist in the literature that assess the benefits and consequences of heat pump installations in WDNs (Blokker et al., 2013; De Pasquale et al., 2017; Paper V). However, owing to differing WDN characteristics (such as residence times) and differences in the temperature model setup (for example, whether a term describing the thermal resistance of the soil is included and whether soil temperatures were estimated or measured), results are not easily comparable and depend on the system setup.

4.3.1 Case study: Greater Copenhagen Utility, Denmark

The Hubeck-Graudal heat transfer model (Paper V) was validated against weekly sampled temperature measurements covering two years at 15 sampling locations. The simulated results showed a good match with the measurements.

The case study assessed the effect of reducing the drinking water temperature by 5 °C at eight favourable heat pump locations. The analysis revealed a total heat extraction potential of 29.2 MW for Copenhagen's WDN. When including the work supplied by the heat pumps (35.9 MW), this amounts to around 2% of Copenhagen's peak heat demand during January and is 50% higher than previously estimated by Bach et al. (2016). This large difference was mainly explained by not deploying heat pumps to reservoir locations only. The case study simulations revealed that 38% (11.1 MW) of the extracted heat will be returned by the soil (i.e. water that has been cooled through heat pumps will regain heat downstream). Moreover, 33% (9.6 MW) of the energy has, based on water demand assumptions (Rygaard et al., 2013), no effect on the district heating demand of consumers, resulting in a net heat potential of 20.7 MW from the source. In other words, a heat utilisation degree of 71% was computed for Copenhagen's WDN.

The coefficient of performance (COP) is a term describing the ratio between the amounts of energy provided and the work (i.e. electricity) required to run heat pumps. Usually, this value exceeds unity, showing that more heat is extracted than work required to run the heat pump. The COP ratio can be applied to heat pumps individually and to the entire system. When applied as system COP, the additional electricity and district heating demand required for heating at consumers' homes, because of cooler water temperatures downstream of heat pumps, are taken into consideration. The application of a conservative heat pump COP of 2.9 resulted in a system COP of 1.7. This value described a relatively low ratio of heat provided over the total work required by the system, when compared to more generic integrations of heat pumps in district heating systems (having a system COP of up to 5 (e.g. Ommen et al., 2014)). Even a less conservative heat pump COP of 4 would only increase the system COP by up to 1.9. Thus, a more thorough (economic) analysis is required to assess whether heat pumps extracting energy from WDNs are more competitive when other solutions, including wastewater, seawater and fresh water sources, are not available (e.g. Elías-Maxil et al., 2014).

The heat utilisation analysis of the WDN revealed that the thermal resistance

from pipe materials in general was insignificant compared to the thermal resistance from the soil material. As the exact thermal conductivity of the soil was unknown, a sensitivity analysis was conducted modifying the parameter by $\pm 30\%$. Whereas the analysis showed that the degree of heat utilisation of the soil changed by 11–15%, the overall degree of heat utilisation in this specific case study changed by only 2–4 percentage points.

Even though the energy benefits were debatable, the application of heat pumps had a significant positive impact on the end-use temperature, as the share of water delivered to users complying with the upper limit of 12 °C increased from 41% to 81% during August.

In Paper V, the degree of heat utilisation of the soil increased from 34% to 38% when including the heat recovery occurring in service lines. This additional heat recovery was, however, based on scarce information about service lines, such as average service pipe lengths and diameters and limited demand information stored in the utility's hydraulic model. The application of smart meter temperature data in Paper IV showed that the water in service pipes can be stagnant or slow-flowing for long periods, indicating that the heat transfer in service lines may be underestimated significantly. Moreover, the applied WDN model in Paper V did not contain up-to-date demand patterns, nor did the model reflect seasonal demand variations. Smart meter data, as available in Paper IV, could be useful to improve the demand dynamics in the system and thus be used to validate the heat transfer model even further, such as during periods where the WDN model's demand pattern very likely does not resemble the real demand in Copenhagen's WDN.

In Paper V, the heat transfer model was a valuable tool for assessing the effect of installing heat pumps and cooling the water in a WDN.

- Whereas a net potential of 20.7 MW can be extracted from Copenhagen's WDN, the system COP was as low as 1.7, indicating that other sources of energy should be considered prior to heat pump deployment in Copenhagen's WDN.
- The share of water at consumers' nodes complying with a statutory recommendation (< 12 °C) increased from 41% to 81% during August when installing heat pumps, thus improving the quality of the distributed water. This indicated a clear benefit of heat pump deployment.

- Future investigations should incorporate additional information, such as smart meter temperature data, to further validate the temperature model and assess the effect of heat recovery occurring in service lines.

4.4 Smart meter data as an add-on to urban drainage management

In Copenhagen, at least 83% of the consumed drinking water is expected to end up in the urban drainage system (Rygaard et al., 2013). This highlights the potential of using and linking smart meter consumption data with urban drainage models to estimate, for example, wastewater flows and constituents, as mentioned in the literature (Cole and Stewart, 2013; Monks et al., 2019; Nguyen et al., 2018). Comparing these estimates with in-sewer flow observations may furthermore be used to estimate infiltration/exfiltration in the urban drainage system. In Paper VI, this potential was tested on a real case study. Smart meter consumption data was bundled according to the meters' wastewater catchments (the 'DMA-equivalent' in urban drainage management) and compared to in-sewer observations from a combined sewer system. Hereby, the potential of using smart meter data to simulate the wastewater component of the dry weather flow in an urban drainage system was established, and the methodology could furthermore be used as a data anomaly and system validation tool.

4.4.1 Case study: Utility Elsinore, Denmark

In Paper VI, data from five in-sewer sensors covering three monitoring periods without rain was available for testing. Figure 12 displays the concept behind the applied method in Paper VI, where in-sewer observations were available from five catchments.

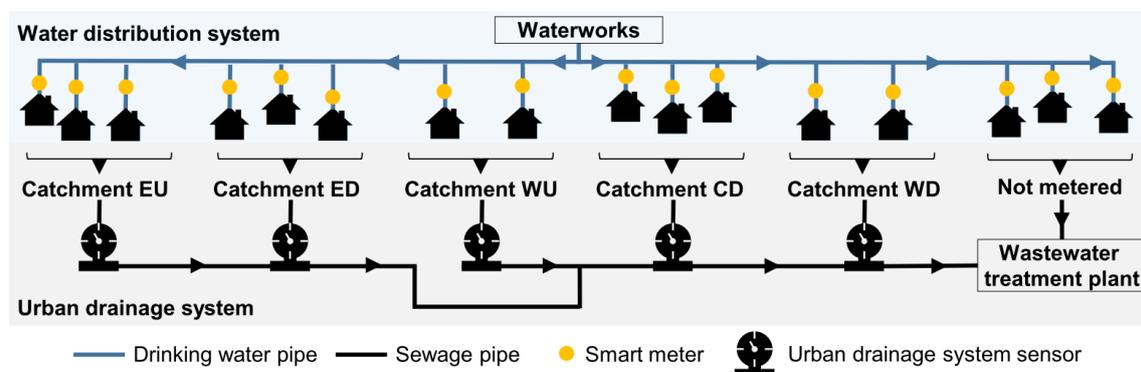


Figure 12. Linkage between a water distribution system with smart meters deployed and an urban drainage system consisting of five catchments. In-sewer sensors were installed downstream of these five catchments. Utility Elsinore, Denmark (Paper VI).

Moreover, data on the outflow from the waterworks and the inflow to wastewater treatment plants was available. Using this data, two approaches were tested:

- 1 Summing the smart meter data, based on upstream catchments.
- 2 Linking the smart meter data to an urban drainage model.

The first approach, being a more simple process, did not include the routing time through the urban drainage system.

Prior to the application, two typical problems related to digitalisation were identified. First, the evaluation of smart meter data from Utility Elsinore revealed Type 0 anomalies (Figure 5). These included a missing sub-catchment belonging to catchment 'EU' (Figure 12) in the wastewater plans of the utility as well as consumers without smart meters (only accounting for < 2% of the demand in the analysed area). Second, a substantial effort was put into gaining access to data from the utility and its partners. The smart meter and in-sewer sensor data, the WDN model and the urban drainage system model were managed by four individual contractors, and the wastewater treatment plant and waterworks data was obtained from various utility employees.

The summed smart meter data was compared with outflow measurements from the utility's waterworks. The results showed a waterworks outflow around 15% higher than the summed data in all three periods. This deficit might have been caused by, among other things, leakage, unknown consumers or sensor errors. Next, the wastewater treatment plant inflow was compared to the simulated inflow based on smart meters upstream of the five in-sewer sensors and additional catchments (Figure 12). Here, the simulated inflow was 50% higher than the observed one during two monitoring periods, but showed a high match during the third monitoring period. The smart meter data, however, showed a high match with audit data prior to smart meter deployment and did not markedly change its dynamics over the three monitoring periods. Thus, the smart meter data was deemed more valid than wastewater treatment plant observations. Most likely, deviations were caused by erroneous sensors at the wastewater treatment plant or other factors, such as exfiltration, occurring in the urban drainage system.

Looking at the five in-sewer observations only (Figure 12), the wastewater flow did not increase as expected in all downstream catchments, indicating possible exfiltration and infiltration of water into the urban drainage system, or in-sewer sensor errors. Moreover, residuals between observed and simulated

wastewater flow, varied greatly in all measuring periods and catchments, ranging from around -20 L/s (observation < simulation) to 70 L/s (observation > simulation). Different sources of error were discussed extensively in Paper VI, including consumed water not discharged to the sewer system or construction dewatering unintentionally released to the urban drainage system; however, no consistent positive or negative trend was seen in the mass balances of the catchments. It was thus concluded that the most likely reason for the observed discrepancies was erroneous in-sewer sensors. For example, the measurements in the most downstream catchment were at times 25 L/s higher than simulated, being an unrealistic additional volume of water ending up in the wastewater treatment plant (where differences between smart meter simulation and observed values were already too high). It is therefore believed that the degree of uncertainty of the in-sewer observations even exceeded expected uncertainties of up to 20%.

For each monitoring period, summed and simulated smart meter data flows upstream of each in-sewer sensor increased the further downstream the in-sewer sensors were placed (Figure 12). When compared with in-sewer observations (whose time dynamics are assumed trustworthy despite the uncertain flow magnitude), the included routing time by the simulated smart meter data showed a better timing of peaks and low points. This could be potentially useful for estimating infiltration of water into the urban drainage system and when used in a real-time control setup.

As the in-sewer observations were deemed the least trustworthy, the smart meter data posed a more reliable source for estimating the dry weather flow in the urban drainage system, even though it only represents the wastewater component. The case study showed that the application of smart meter data can also be used to highlight system anomalies, such as identifying unreliable data collected in the urban drainage system. Owing to the high uncertainty in the applied datasets, independent data sources are needed to verify that the in-sewer observations are indeed erroneous and to further estimate the other dry weather flow components, such as exfiltration and infiltration.

Thus, the application of smart meter data for comparison with in-sewer observations showed the following.

- Coupling smart meter data with urban drainage models increases the understanding of urban drainage systems and can be used as a tool to identify anomalies, including erroneous in-sewer observations.

- Owing to the large deviations between simulated wastewater flow and in-sewer observations, and because of various possible sources of error in the in-sewer data, smart meter data was deemed a more valid source for estimating the dry weather flow in urban drainage systems.
- Whereas linking the data from WDNs and urban drainage data was relatively simple, the access to data and information was complicated by being distributed between four contractors and several employees internally in the utility, which is an obstacle for further digitalisation of the water sector.

5 Conclusions

The water sector conceives digitalisation as a way to solve many of the challenges currently faced. Using six real-world case studies and data from five utilities, this thesis developed a range of novel methods to identify and address challenges and potentials for utilities to become increasingly data-driven.

Being data-driven requires action within at least three fields: 1) data collection, 2) data validation and reconstruction, and 3) application of the data. The investigated case studies revealed that these actions are highly interlinked and will enhance and complement each other.

Data collection. One typical example of digitalisation is the enrolment of smart meters, a process familiar to all five case study utilities. The sampling resolution of smart meters can have a major impact on reaching different data-driven goals. In terms of water loss assessment, commonly implemented sampling resolutions of consumption data ranging between 1 and 24 hours proved to be sufficient if representative demand patterns are available and used to interpolate in between adjacent data points. It is especially relevant to use such weighted demand-pattern-based approaches over linear interpolation for sampling resolutions > 2 hours. When it comes to water age simulations, however, sampling resolutions finer than 1 hour are needed to increase the validity of the simulations significantly. Moreover, application of smart meter temperature data showed that this data was not sampled optimally in regards to using the data for analysing the temperatures in the WDN. Only a small fraction of the smart meter temperature data ($< 0.2\%$ over 2 weeks) represented WDN and soil temperatures, which were essential inputs to simulate the temperatures in WDNs. Thus, additional sampling points or a finer sampling resolution than available in the case study utility (< 30 minutes) is likely needed.

Data validation and reconstruction. The case studies highlighted the need for an increased focus on data validation. This PhD project developed a systematic approach by categorizing anomaly types into four groups. Type 0 anomalies described ‘static’ data, representing missing and incorrect system attributes (e.g. wrong pipe diameters or consumer affiliations). Type 1 and 2 anomalies described sensor data that did not reflect the real state (e.g. illegitimate data) and misrepresented the actual state (e.g. drifted data from a miscalibrated sensor) of the WDN, respectively. Type 3 anomalies were used to classify abnormal measurements representing actual events (e.g. bursts) in the WDN.

The project implemented seven validation tests to identify Type 1 and 2 anomalies. Application of these tests revealed that many errors prevailed for long periods, with 3–35% invalidated and 10% missing data in the data sets from three utilities. Running such tests in an operational setting as well as analysing similarities in the occurrence of anomalies may help identify major flaws in the data collection schemes and help utilities to react in a fast manner, reducing the overall number of anomalies in the utilities' data sets.

High numbers of missing and invalidated data points also lead to a need for reconstruction processes. The thesis demonstrated data reconstruction based on artificial neural networks. Reconstructed consistent data streams can serve as an important input in, for example, online models. The reconstructed data can also help to identify Type 3 anomalies.

Application. In the literature, numerous studies highlight many potential data-driven applications of WDN analyses, but their success depends on the quality of the data. In this PhD thesis, the utilisation of real-world sensor data and application of advanced methods were particularly useful in highlighting Type 0 anomalies, i.e. flawed registrations in the utilities' asset databases. A large number of Type 0 anomalies limits the success of more advanced analyses.

A novel approach combined WDN temperature data, temperature simulations and a hydraulic WDN model, with the intention of identifying valves with unknown status. However, the analysis led to identification of errors in the utility's asset and sensor data instead, ultimately casting doubt on the validity of the hydraulic model. The combined temperature and hydraulic modelling thus turned out to have another benefit than originally intended. Moreover, WDN temperature data proved to have a clear value for the case study utility, as it indicated the water's origin and path taken in the WDN, which is otherwise not easily understood from pressure and flow data alone.

A low number of WDN temperature measurements can be overcome by using smart meter temperature data, which is potentially available from each household. In another case study, comparing smart meter temperatures with simulated temperature values throughout a DMA highlighted the applicability of the data to represent WDN temperatures (mean RMSE of 0.9 °C). Well-sampled data and detailed WDN models having distinct nodes for all consumers may provide temperature simulations suitable for leakage localisation or valve status identification.

Another use of a WDN temperature model was identified by analysing the effect of eight heat pumps deployed in Copenhagen between the drinking water and district heating systems. The analysis revealed a positive net heat extraction potential of 21 MW; this may be an important supplement, especially during peak heat demands. In addition, the analysis revealed that the installation of heat pumps may improve summer water temperatures markedly by increasing the share of consumers receiving water under 12 °C from 41% to 81% during the month of August.

In another case study, smart meter data from the WDN was coupled with an urban drainage model. Access to the data was complicated by being distributed between different contractors and staff members, representing a typical obstacle of digitalisation. Large deviations were observed between simulated wastewater flows and the in-sewer observations, and the smart meter data was deemed as a more valid source for estimating the dry weather flow in urban drainage systems.

This thesis revealed that there is a great potential hidden in the often unused data from water distribution systems. For example, it was demonstrated that smart meter data, and in particular temperature data (which is often overseen in optimisation and modelling tasks), add great value to the understanding and management of water distribution systems and beyond. Another important outcome of this thesis is the realization that it may be more challenging for utilities than expected to become increasingly data-driven. This is, among other things, because too many errors and uncertainties still exist in the utilities' asset databases and sensor data. In the years to come, utilities, technology providers and researchers need to collaborate to identify and reduce these uncertainties; thereby allowing the water sector to reach higher levels of digital maturity.

6 Perspective

The application of real-world data revealed various potentials, but also challenges that need to be addressed by utilities, researchers and technology providers, before utilities may become more data-driven.

Digitalised utilities. Even though the case study utilities have invested in monitoring capabilities in the form of SCADA systems and have started full smart meter rollouts, the use of the collected data is very limited. Many reasons may exist for this, including limited resources for analyses, missing analytical tools, or lack of funding for managing data and devices. This limited usage also means that utilities do not know the actual quality of their data. However, this thesis revealed that a generally low data reliability was a major issue in all case study utilities. Until reliability is improved, this will limit the potential of data-driven applications. I suspect that if utilities want to climb from ‘basic’ or ‘opportunistic’ digital adaptation to ‘systematic’ and later become ‘transformational’ (Figure 2), the following three steps are of importance.

- 1** High-quality asset data and a thorough system understanding need to be established. For example, are pipes, valves and consumers located as stated in the utility’s database? To the best of my knowledge, fully automated procedures for identifying such discrepancies are missing, but the application and model-aided evaluation of data is a first step to identifying and reducing such errors.
- 2** It is of paramount importance to secure a high data reliability. Otherwise, there might be no reason for deploying additional sensors. Applications depending on the data may run flawlessly in the beginning, but the number of errors normally increases over time owing to limited resources for re-calibrating sensors and keeping system changes up-to-date in databases. Thus, utilities need to allocate resources for a proper validation to maintain a high reliability of the data, and the use of data validation tools should be part of the daily tasks in a utility.
- 3** When the reliability of the data is assured and the number of major system anomalies is reduced, utilities need to increase their usage of the data. This usage will likely highlight additional minor ‘system anomalies’ and improve the understanding of the system. This process is expected to result in further improvements in data collection and reliability. This will also require involvement from the research community and technology providers, who

should help utilities to identify and provide optimal data collection possibilities and analytical tools as well as demonstrate the application of these in real-world case studies.

Research opportunities. Use of temperature data and simulation of the temperature throughout the WDN revealed various benefits. To further highlight the true potential of the temperature simulations for valve status identification, additional studies need to be conducted. These include in-depth sensitivity analyses discussing when the method can be applied successfully and when not (for example, identifying which temperature gradients are needed throughout the WDN). Moreover, the applied genetic algorithm and temperature modelling software (EPANET-MSX) turned out to have limitations in terms of speed and convergence success at the applied fine sampling resolutions and are therefore currently not suitable for daily operations; thus, there is a demand for additional research in computational optimisation.

The true value of using smart meter temperature data was inhibited by how the data was sampled and categorized. Whereas the sampling resolution could be increased by the technology provider, additional procedures for categorizing the temperature data into soil and WDN temperatures will be useful. Owing to the uncertain nature of service line data, approaches that are independent of the physical properties of service lines should be investigated. These include, for example, unsupervised clustering algorithms, identifying smart meter data that represent soil and WDN temperatures. Next, it should be investigated to what extent the data can be used for more advanced methods such as valve status detection or leakage localisation. Smart meter temperature data could also be used to improve the knowledge of utilities about heat recovery in service lines. If there is indeed a higher recovery potential, the benefit of heat pumps deployed in WDNs may be greater than estimated here.

7 References

- Adedoja, O.S., Hamam, Y., Khalaf, B., Sadiku, R., 2018. A state-of-the-art review of an optimal sensor placement for contaminant warning system in a water distribution network. *Urban Water J.* 15, 985–1000. doi:10.1080/1573062X.2019.1597378
- ASCE, 2011. *Failure to Act: The Economic Impact of Current Investment Trends in Water and Wastewater Treatment Infrastructure*. American Society of Civil Engineers, VA.
- AVK, 2018. Smart pressure management. Available at <https://www.avkvalves.com/-/media/websites/avk-global/files/downloads/avk-pressure-management-brochure> (accessed November 12, 2019).
- Bach, B., Werling, J., Ommen, T., Münster, M., Morales, J.M., Elmegaard, B., 2016. Integration of large-scale heat pumps in the district heating systems of Greater Copenhagen. *Energy* 107, 321–334. doi:10.1016/j.energy.2016.04.029
- Bakker, M., van Duist, H., van Schagen, K., Vreeburg, J., Rietveld, L., 2014. Improving the Performance of Water Demand Forecasting Models by Using Weather Input. *Procedia Eng.* 70, 93–102. doi:10.1016/j.proeng.2014.02.012
- Blokker, E.J.M., Pieterse-Quirijns, E.J., 2013. Modeling temperature in the drinking water distribution system. *J. Am. Water Work. Assoc.* 105, E19–E28. doi:10.5942/jawwa.2013.105.0011. Erratum (2018) in *J. Am. Water Work. Assoc.* 110, 98–98. doi:10.1002/awwa.1178
- Blokker, E.J.M., van Osch, A.M., Hogeveen, R., Mudde, C., 2013. Thermal energy from drinking water and cost benefit analysis for an entire city. *J. Water Clim. Chang.* 4, 11–16. doi:10.2166/wcc.2013.010
- Blokker, E.J.M., Vreeburg, J.H.G., Beverloo, H., Klein Arfman, M., van Dijk, J.C., 2010. A bottom-up approach of stochastic demand allocation in water quality modelling. *Drink. Water Eng. Sci.* 3, 43–51. doi:10.5194/dwes-3-43-2010
- Blokker, E.J.M., Vreeburg, J.H.G., Buchberger, S.G., van Dijk, J.C., 2008. Importance of demand modelling in network water quality models: a review. *Drink. Water Eng. Sci.* 1, 27–38. doi:10.5194/dwes-1-27-2008
- Blokker, M., 2019. Making water meters truly smart by adding additional features. Poster presented at the 17th International Computing & Control for the Water Industry Conference (CCWI). University of Exeter, UK.
- Blumensaat, F., Leitão, J.P., Ort, C., Rieckermann, J., Scheidegger, A., Vanrolleghem, P.A., Villez, K., 2019. How Urban Storm- and Wastewater Management Prepares for Emerging Opportunities and Threats: Digital Transformation, Ubiquitous Sensing, New Data Sources, and Beyond - A Horizon Scan. *Environ. Sci. Technol.* 53, 8488–8498. doi:10.1021/acs.est.8b06481
- Boyle, T., Giurco, D., Mukheibir, P., Liu, A., Moy, C., White, S., Stewart, R., 2013. Intelligent Metering for Urban Water: A Review. *Water* 5, 1052–1081. doi:10.3390/w5031052
- Bragalli, C., Neri, M., Toth, E., 2019. Effectiveness of smart meter-based urban water loss assessment in a real network with synchronous and incomplete readings. *Environ. Model. Softw.* 112, 128–142. doi:10.1016/j.envsoft.2018.10.010
- Branisavljević, N., Kapelan, Z., Prodanović, D., 2011. Improved real-time data anomaly

- detection using context classification. *J. Hydroinformatics* 13, 307–323. doi:10.2166/hydro.2011.042
- Cole, G., Stewart, R.A., 2013. Smart meter enabled disaggregation of urban peak water demand: precursor to effective urban water planning. *Urban Water J.* 10, 174–194. doi:10.1080/1573062X.2012.716446
- Cominola, A., Giuliani, M., Castelletti, A., Rosenberg, D.E., Abdallah, A.M., 2018. Implications of data sampling resolution on water use simulation, end-use disaggregation, and demand management. *Environ. Model. Softw.* 102, 199–212. doi:10.1016/j.envsoft.2017.11.022
- Cominola, A., Giuliani, M., Piga, D., Castelletti, A., Rizzoli, A.E., 2015. Benefits and challenges of using smart meters for advancing residential water demand modeling and management: A review. *Environ. Model. Softw.* 72, 198–214. doi:10.1016/j.envsoft.2015.07.012
- Creaco, E., Campisano, A., Fontana, N., Marini, G., Page, P.R., Walski, T., 2019. Real time control of water distribution networks: A state-of-the-art review. *Water Res.* 161, 517–530. doi:10.1016/j.watres.2019.06.025
- Creaco, E., Pezzinga, G., Savic, D., 2017. On the choice of the demand and hydraulic modeling approach to WDN real-time simulation. *Water Resour. Res.* 53, 6159–6177. doi:10.1002/2016WR020104
- Cronborg A/S, 2014. Energi fra drikkevand bruges til fjernvarme. Available at <https://csr.dk/energi-fra-drikkevand-bruges-til-fjernvarme> (accessed November 12, 2019).
- Cugueró-Escofet, M.À., García, D., Quevedo, J., Puig, V., Espin, S., Roquet, J., 2016. A methodology and a software tool for sensor data validation/reconstruction: Application to the Catalonia regional water network. *Control Eng. Pract.* 49, 159–172. doi:10.1016/j.conengprac.2015.11.005
- DANVA, 2018. Vand i tal 2018 (in Danish). Danish Water and Wastewater Association, Skanderborg. Available at https://www.danva.dk/media/5002/2018_vand-i-tal.pdf (accessed November 12, 2019).
- De Pasquale, A.M., Giostri, A., Romano, M.C., Chiesa, P., Demeco, T., Tani, S., 2017. District heating by drinking water heat pump: Modelling and energy analysis of a case study in the city of Milan. *Energy* 118, 246–263. doi:10.1016/j.energy.2016.12.014
- Deb, A.K., Snyder, J.K., Grayman, W.M., 2012. Management of valves to improve performance reliability of distribution systems. *J. Am. Water Works Assoc.* 104, E370–E382. doi:10.5942/jawwa.2012.104.0087
- Delgado, D.M., Lansley, K.E., 2009. Detection of Closed Valves in Water Distribution Systems, in: *Water Distribution Systems Analysis 2008*. American Society of Civil Engineers, Reston, VA, pp. 1–7. doi:10.1061/41024(340)82
- Do, N.C., Simpson, A.R., Deuerlein, J.W., Piller, O., 2018. Locating Inadvertently Partially Closed Valves in Water Distribution Systems. *J. Water Resour. Plan. Manag.* 144, 04018039. doi:10.1061/(ASCE)WR.1943-5452.0000958
- Eggimann, S., Mutzner, L., Wani, O., Schneider, M.Y., Spuhler, D., Moy de Vitry, M., Beutler, P., Maurer, M., 2017. The Potential of Knowing More: A Review of Data-Driven Urban Water Management. *Environ. Sci. Technol.* 51, 2538–2553.

doi:10.1021/acs.est.6b04267

- Elías-Maxil, J.A., van der Hoek, J.P., Hofman, J., Rietveld, L., 2014. Energy in the urban water cycle: Actions to reduce the total expenditure of fossil fuels with emphasis on heat reclamation from urban water. *Renew. Sustain. Energy Rev.* 30, 808–820. doi:10.1016/j.rser.2013.10.007
- García, D., Creus, R., Minoves, M., Pardo, X., Quevedo, J., Puig, V., 2017. Data analytics methodology for monitoring quality sensors and events in the Barcelona drinking water network. *J. Hydroinformatics* 19, 123–137. doi:10.2166/hydro.2016.048
- Gartner, 2019. IT Glossary. Available at <https://www.gartner.com/it-glossary/> (accessed November 12, 2019).
- German Environment Agency, 2019. Trockenheit in Deutschland - Fragen und Antworten (in German). Available at <https://www.umweltbundesamt.de/themen/trockenheit-in-deutschland-fragen-antworten> (accessed November 12, 2019).
- Gurung, T.R., Stewart, R.A., Beal, C.D., Sharma, A.K., 2016. Smart meter enabled informatics for economically efficient diversified water supply infrastructure planning. *J. Clean. Prod.* 135, 1023–1033. doi:10.1016/j.jclepro.2016.07.017
- Gurung, T.R., Stewart, R.A., Sharma, A.K., Beal, C.D., 2014. Smart meters for enhanced water supply network modelling and infrastructure planning. *Resour. Conserv. Recycl.* 90, 34–50. doi:10.1016/j.resconrec.2014.06.005
- Kamstrup, 2019a. Mechanical vs. Ultrasonic. Available at <https://www.kamstrup.com/en-en/water-solutions/themes/mechanical-vs-ultrasonic> (accessed November 12, 2019).
- Kamstrup, 2019b. flowIQ 2200 - Advanced residential water meter with acoustic leak detection. Available at <https://www.kamstrup.com/en-en/water-solutions/smart-water-meters/flowiq-2200> (accessed November 12, 2019).
- Kirstein, J.K., 2016. Using high-frequency data to optimize water supply operation - Case Study: Halsnæs Water Supply, Denmark. MSc Thesis. Technical University of Denmark, DTU Environment, Kgs. Lyngby.
- Kirstein, J.K., Borup, M., Rygaard, M., Høgh, K., 2016. Målerdata kan gemme på gratis informationer for forsyningerne (in Danish). *Danskvand* 84, 50–51.
- Kulkarni, P., Farnham, T., 2016. Smart City Wireless Connectivity Considerations and Cost Analysis: Lessons Learnt From Smart Water Case Studies. *IEEE Access* 4, 660–672. doi:10.1109/ACCESS.2016.2525041
- Larsen, S.L., Christensen, S.C.B., Albrechtsen, H.-J., Rygaard, M., 2017. GISMOWA: Geospatial Risk-Based Analysis Identifying Water Quality Monitoring Sites in Distribution Systems. *J. Water Resour. Plan. Manag.* 143, 04017018. doi:10.1061/(ASCE)WR.1943-5452.0000754
- LEAKman, 2018. 2,300 smart meters installed in demonstration areas. Available at <https://leakagemanagement.net/demo-facilities/2300-smart-meters-installed-in-demonstration-areas/> (accessed November 12, 2019).
- Li, R., Huang, H., Xin, K., Tao, T., 2015. A review of methods for burst/leakage detection and location in water distribution systems. *Water Sci. Technol. Water Supply* 15, 429–441. doi:10.2166/ws.2014.131
- Liu, S., Gunawan, C., Barraud, N., Rice, S.A., Harry, E.J., Amal, R., 2016. Understanding,

- Monitoring, and Controlling Biofilm Growth in Drinking Water Distribution Systems. *Environ. Sci. Technol.* 50, 8954–8976. doi:10.1021/acs.est.6b00835
- Loureiro, D., Alegre, H., Coelho, S.T., Martins, A., Mamade, A., 2014. A new approach to improve water loss control using smart metering data. *Water Sci. Technol. Water Supply* 14, 618–625. doi:10.2166/ws.2014.016
- Loureiro, D., Amado, C., Martins, A., Vitorino, D., Mamade, A., Coelho, S.T., 2016. Water distribution systems flow monitoring and anomalous event detection: A practical approach. *Urban Water J.* 13, 242–252. doi:10.1080/1573062X.2014.988733
- Makropoulos, Savić, 2019. Urban Hydroinformatics: Past, Present and Future. *Water* 11, 1959. doi:10.3390/w11101959
- Mekki, K., Bajic, E., Chaxel, F., Meyer, F., 2019. A comparative study of LPWAN technologies for large-scale IoT deployment. *ICT Express* 5, 1–7. doi:10.1016/j.icte.2017.12.005
- Ministry of Environment and Food of Denmark, 2018a. Miljø- og fødevarerministerens endelige besvarelse af spørgsmål nr. 775 (MOF alm. del) stillet 8. juni efter ønske fra Julie Skovsby (S) (in Danish). Available at <https://www.ft.dk/samling/20171/almudel/mof/spm/775/svar/1510482/1938569.pdf> (accessed November 12, 2019).
- Ministry of Environment and Food of Denmark, 2018b. Bekendtgørelse om vandkvalitet og tilsyn med vandforsyningsanlæg (statutory order no. 1068, in Danish). Available at <https://www.retsinformation.dk/Forms/R0710.aspx?id=202768> (accessed November 12, 2019).
- Monks, Stewart, Sahin, Keller, 2019. Revealing Unreported Benefits of Digital Water Metering: Literature Review and Expert Opinions. *Water* 11, 838. doi:10.3390/w11040838
- Mounce, S.R., Boxall, J.B., Machell, J., 2010. Development and Verification of an Online Artificial Intelligence System for Detection of Bursts and Other Abnormal Flows. *J. Water Resour. Plan. Manag.* 136, 309–318. doi:10.1061/(ASCE)WR.1943-5452.0000030
- Mourad, M., Bertrand-Krajewski, J.-L., 2002. A method for automatic validation of long time series of data in urban hydrology. *Water Sci. Technol.* 45, 263–270. doi:10.2166/wst.2002.0601
- Moy de Vitry, M., Schneider, M.Y., Wani, O., Manny, L., Leitão, J.P., Eggimann, S., 2019. Smart urban water systems: what could possibly go wrong? *Environ. Res. Lett.* 14, 081001. doi:10.1088/1748-9326/ab3761
- Nguyen, K.A., Stewart, R.A., Zhang, H., Sahin, O., Siriwardene, N., 2018. Re-engineering traditional urban water management practices with smart metering and informatics. *Environ. Model. Softw.* 101, 256–267. doi:10.1016/j.envsoft.2017.12.015
- Ommen, T., Markussen, W.B., Elmegaard, B., 2014. Heat pumps in combined heat and power systems. *Energy* 76, 989–1000. doi:10.1016/j.energy.2014.09.016
- Puust, R., Kapelan, Z., Savic, D.A., Koppel, T., 2010. A review of methods for leakage management in pipe networks. *Urban Water J.* 7, 25–45. doi:10.1080/15730621003610878
- Quevedo, J., Garcia, D., Puig, V., Saludes, J., Cugueró, M.A., Espin, S., Roquet, J., Valero,

- F., 2017. Real-time Monitoring and Operational Control of Drinking-Water Systems, Real-Time Monitoring and Operational Control of Drinking- Water Systems, *Advances in Industrial Control*. Springer International Publishing, Cham. doi:10.1007/978-3-319-50751-4
- Quevedo, J., Puig, V., Cembrano, G., Blanch, J., Aguilar, J., Saporta, D., Benito, G., Hedo, M., Molina, A., 2010. Validation and reconstruction of flow meter data in the Barcelona water distribution network. *Control Eng. Pract.* 18, 640–651. doi:10.1016/j.conengprac.2010.03.003
- Rygaard, M., Alsbjørn, L., Ejsing, M., Godskesen, B., Hansen, R., Hoffmann, B., Jørgensen, C., Ledgaard, K., Friis Møller, H.-M., Poulsen, M.-B.B., Schmidt, M., Tarp-Johansen, L., Zambrano, S.V.K., 2013. Sekundavand i Nordhavn: En forundersøgelse til strategi for alternativ vandleverance. Technical University of Denmark, DTU Environment, Kgs. Lyngby.
- Rygaard, M., Binning, P.J., Albrechtsen, H.-J., 2011. Increasing urban water self-sufficiency: New era, new challenges. *J. Environ. Manage.* 92, 185–194. doi:10.1016/j.jenvman.2010.09.009
- Sarni, W., White, C., Webb, R., Cross, K., Glotzbach, R., 2019. Digital Water - Industry leaders chart the transformation journey. International Water Association. London. Available at https://iwa-network.org/wp-content/uploads/2019/06/IWA_2019_Digital_Water_Report.pdf (accessed November 12, 2019).
- Savic, D.A., Kapelan, Z.S., Jonkergouw, P.M.R., 2009. Quo vadis water distribution model calibration? *Urban Water J.* 6, 3–22. doi:10.1080/15730620802613380
- Sophocleous, S., Savić, D.A., Kapelan, Z., Giustolisi, O., 2017. A Two-stage Calibration for Detection of Leakage Hotspots in a Real Water Distribution Network. *Procedia Eng.* 186, 168–176. doi:10.1016/j.proeng.2017.03.223
- Stewart, R.A., Nguyen, K., Beal, C., Zhang, H., Sahin, O., Bertone, E., Vieira, A.S., Castelletti, A., Cominola, A., Giuliani, M., Giurco, D., Blumenstein, M., Turner, A., Liu, A., Kenway, S., Savić, D.A., Makropoulos, C., Kossieris, P., 2018. Integrated intelligent water-energy metering systems and informatics: Visioning a digital multi-utility service provider. *Environ. Model. Softw.* 105, 94–117. doi:10.1016/j.envsoft.2018.03.006
- Sun, A.Y., Scanlon, B.R., 2019. How can Big Data and machine learning benefit environment and water management: a survey of methods, applications, and future directions. *Environ. Res. Lett.* 14, 073001. doi:10.1088/1748-9326/ab1b7d
- Tan, P.N., Steinbach, M., Kumar, V., 2006. *Introduction to Data Mining*, Pearson International Edition. Pearson Addison Wesley, Boston.
- Tang, C., Albrechtsen, H.-J., 2019. Water quality monitoring in the food and beverage industry: Sensors for monitoring physical, chemical and microbial water quality. Technical University of Denmark, DTU Environment, Kgs. Lyngby. [in press].
- Tatari, K., Corfitzen, C.B., Albrechtsen, H.-J., Christensen, S.C.B., 2016. Sensors for microbial drinking water quality. Technical University of Denmark, DTU Environment, Kgs. Lyngby.
- United Nations, 2019. *The Sustainable Development Goals Report 2019*. United Nations Publications, New York. Available at <https://unstats.un.org/sdgs/report/2019/The->

Sustainable-Development-Goals-Report-2019.pdf (accessed November 12, 2019).

- Walski, T., Sage, P., Wu, Z., 2014. What Does it Take to Make Automated Calibration Find Closed Valves and Leaks?, in: World Environmental and Water Resources Congress 2014. American Society of Civil Engineers, Reston, VA, pp. 555–565. doi:10.1061/9780784413548.059
- Walski, T.M., Chase, D. V, Savic, D.A., Grayman, W., Beckwith, S., Koelle, E., 2003. Advanced water distribution modeling and management. Haestad press, Waterbury, CT.
- Watertech A/S, 2005. Vandforbrug og forbrugsvariationer (in Danish). Miljøstyrelsen (Danish EPA). Available at <https://www2.mst.dk/Udgiv/publikationer/2005/87-7614-592-1/html/default.htm> (accessed November 12, 2019).
- Wilson, C., 2011. Sustainable distribution system asset management: Risk mitigation through control point management. J. Am. Water Works Assoc. 103, 100–103. doi:10.1002/j.1551-8833.2011.tb11456.x
- Wu, Y., Liu, S., 2017. A review of data-driven approaches for burst detection in water distribution systems. Urban Water J. 14, 972–983. doi:10.1080/1573062X.2017.1279191
- Wu, Z.Y., Song, Y., Syed, J.L., 2012. Optimization Model for Identifying Unknown Valve Statuses and Settings, in: WDSA 2012: 14th Water Distribution Systems Analysis Conference, 24-27 September 2012 in Adelaide, South Australia. pp. 975–983.
- WWAP, 2019. The United Nations World Water Development Report 2019: Leaving No One Behind. UNESCO World Water Assessment Programme, Paris. Available at <https://unesdoc.unesco.org/ark:/48223/pf0000367306> (accessed November 12, 2019).

8 Appendix

A Scopus search query

For hits related to publications that cover digitalisation in the water sector, the following query was used on <https://www.scopus.com>:

```
TITLE-ABS-KEY(("digitalisation" OR "digitalization" OR "digital transformation" OR "data-driven" OR "digital water" OR "water 4.0" OR "internet of water" OR "smart water" OR "intelligent meter" OR "smart meter") AND "water") AND (EXCLUDE(SUBJAREA,"MEDI" ) OR EXCLUDE ( SUBJAREA,"IMMU") OR EXCLUDE(SUBJAREA,"HEAL" ) OR EXCLUDE(SUBJAREA,"PSYC") OR EXCLUDE(SUBJAREA, "PHAR") OR EXCLUDE(SUBJAREA, "NURS") OR EXCLUDE(SUBJAREA, "VETE") OR EXCLUDE(SUBJAREA, "DENT") OR EXCLUDE(SUBJAREA, "ARTS") OR EXCLUDE(SUBJAREA, "NEUR"))
```

To account for general water-related publications, ("digitalisation" OR "digitalization" OR "digital transformation" OR "data-driven" OR "digital water" OR "water 4.0" OR "internet of water" OR "smart water" OR "intelligent meter" OR "smart meter") was removed from the upper query. The hit ratio was then computed between both outputs.

B Data reconstruction

During anomaly testing of the raw data in Paper II, missing data periods were identified and dubious observations flagged; however, most post-applications of the collected meter data will benefit from a continuous and flawless input data stream. Missing and flagged data points (as well as valid data points for visualization purposes) were reconstructed based on estimates from feed-forward artificial neural networks (ANNs) (Tan et al., 2006). The ANN models were trained based on the resilient back-propagation algorithm with weight backtracking (Riedmiller and Braun, 1994)¹. To reconstruct meter observations, it was decided to construct simple ANNs that consisted of a maximum of 10 input neurons, 1 hidden layer including a maximum of 10 hidden neurons, and 1 output neuron. The hourly observations of a meter were predicted based on an ANN model that was trained on the hourly or specific lagged hourly

¹ Riedmiller, M., Braun, H., 1994. A direct adaptive method for faster backpropagation learning: the RPROP algorithm, in: IEEE International Conference on Neural Networks. IEEE, pp. 586–591. doi:10.1109/ICNN.1993.298623

values from other meters in a utilities network. The selection of other meters favoured highly correlated data sets. To ensure uniform time intervals, the data was aggregated to hourly data similar to test VII (see Paper II). Only pre-validated data of tests I–VII were used for creating the hourly averages. This improved the reliability of the model results, as neural networks are sensitive to the presence of noise in the training data (Tan et al., 2006).

Applied training of artificial neural networks

ANNs comprise three types of layer: one input, multiple hidden and one output layer. Neurons connect the different layers by specific weights and an activation function, which in our case is the hyperbolic tangent function (\tanh):

$$\tanh(z) = \frac{e^z - e^{-z}}{e^z + e^{-z}}, \quad (1)$$

where z represents a linear combination of weighted inputs from neurons. The complete selection process of input and hidden neurons and training of ANN models for a meter data set is summarized in the following four steps (Figure B-1):

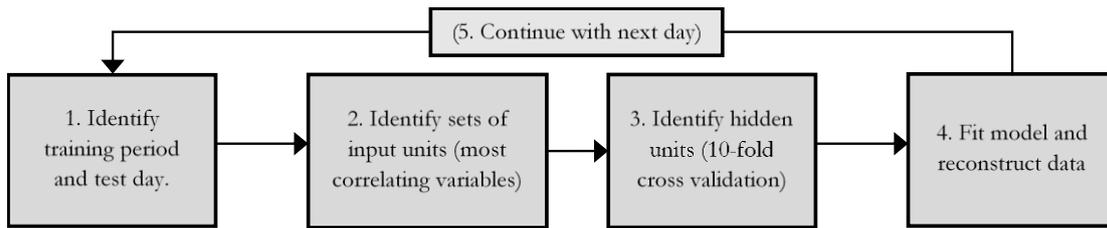


Figure B-1. Daily meter data reconstruction process based on artificial neural network models.

- 1 A meter data set was selected for reconstruction. As utilities may change their network set-up frequently, significant changes may be observed in the statistical properties of pressure and flow time series. Therefore, each daily model is only trained on the available data sets of the previous 60 days. Thus, a test day and the previous 60 days were used as test and training periods, respectively. If more than one third of the data was missing from the test meter in the training period, the test day was skipped for data reconstruction.
- 2 The correlation between the meter to be reconstructed and all available data sets was computed based on the training period. The five meter data sets with the highest absolute correlation were selected as ANN model input units. Also, the hourly lags 1, 2, 6, 12 and 24 of the most correlating meter

data set were included as input neurons. The response variable itself was not used as an input variable; this was done to overcome issues that may arise when only the analytical redundancy of the meter is used for model building. For example, time series models based on one meter might fail to predict measurements correctly that include sudden network operations (e.g. Quevedo et al., 2010).

- 3 For each set of input units, the number of neurons in the hidden layer was determined by 10-fold cross validation (CV) (Tan et al., 2006) applied on the training data set to avoid model overfitting. In short, the training set was divided into 10 partitions, where one partition was equal to the ‘CV test set’ in each fold. An ANN model was trained on the remaining nine parts and the sum of squared error between the model prediction and each ‘CV test set’ was computed. This procedure was repeated ten times. The number of neurons having the lowest mean squared sum of errors was selected for model training in the next step.
- 4 Having determined a final number of input units and number of neurons in the hidden layer, the model setup was trained on the entire training set available. The test period was then applied to the model(s) and the daily data was reconstructed.

Several conditions could occur that require the training of additional models or the discarding of model training and thus data reconstruction. Relating to step 2 of the ANN construction process illustrated in Figure B-1, the following conditions had to be met regarding the training and testing of data sets with missing data:

- If any input unit in combination with the reference meter had more than one third missing data in the training period, the input unit was excluded and another, based on its correlation with the reference meter, was selected.
- If the set of input units had more than one third missing timestamps in the training period, the input unit with the most missing timestamps was removed and a new input unit, based on its correlation with the reference meter, was selected.
- If an input unit had one or multiple missing timestamps in the training and/or test data set, additional input-unit sets and thus models covering these timestamps had to be constructed. The input units not covering the period were excluded from the new models.

- If no input-unit set could be determined that met all of the above-stated criteria for certain timestamps, the number of most correlating variables and hence input units was reduced by one. In this case, the entire procedure of step two was repeated until it is possible to reconstruct all timestamps. If it was not possible to find a set of input units, selected timestamps were skipped for reconstruction.

9 Papers

- I Kirstein, J.K., Høgh, K., Rygaard, M. & Borup, M. (2019).** Effect of data sampling resolution of smart meter readings in water distribution network simulations. *Manuscript in preparation.*
- II Kirstein, J.K., Høgh, K., Rygaard, M. & Borup, M. (2019).** A semi-automated approach to validation and error diagnostics of water network data. *Urban Water Journal*, **16**(1), 1–10, doi:10.1080/1573062X.2019.1611884
- III Kirstein, J.K., Liu, S., Høgh, K., Borup, M. & Rygaard, M. (2019).** Valve status identification by temperature modelling in water distribution networks. *Manuscript in preparation.*
- IV Kirstein, J.K., Høgh, K., Rygaard, M. & Borup, M. (2019).** Using smart meter temperature and consumption data for water distribution system analysis. *Manuscript in preparation.*
- V Hubeck-Graudal, H., Kirstein, J.K., Ommen, T., Rygaard, M. & Elmegaard, B. (2019).** Drinking water supply as low-temperature source in the district heating system: a case study for the city of Copenhagen. *Submitted.*
- VI Lund, N.S.V., Kirstein, J.K., Mikkelsen, P.S., Madsen, H., Mark, O. & Borup, M. (2019).** Using smart meter water consumption data and in-sewer flow observations for model based sewer system analysis. *Submitted.*

I

Effect of data sampling resolution of smart meter readings in water distribution network simulations

J.K. Kirstein, K. Høgh, M. Rygaard and M. Borup

Effect of data sampling resolution of smart meter readings in water distribution network simulations

Jonas Kjeld Kirstein^{1,*}, Klavs Høgh², Martin Rygaard¹ and Morten Borup¹

*corresponding author

¹DTU Environment, Technical University of Denmark, Kgs. Lyngby, Denmark. jkir@env.dtu.dk; mryg@env.dtu.dk; morb@env.dtu.dk.

²NIRAS A/S, Allerød, Denmark; kvh@niras.dk.

Abstract Water usage data collected from smart meters at the end user can improve the accuracy and applicability of water distribution network models. Collecting and storing large amounts of data across hundreds or more smart meters is costly which makes it important to consider what makes a sufficient sampling resolution. This paper explores the effect of data sampling resolution in smart meter data on model performance in regard to flow, pressure and water age. Furthermore, the paper investigates the effect of using either linear interpolation or demand patterns to fill gaps when data is sampled coarsely. The study was based on real data from 526 smart meters in a district metered area in Denmark. The results showed that smart meter data can greatly improve modelling results, and if the temporal data resolution is coarser than 2 hours, then demand pattern based gap filling should be used. Furthermore, the study showed that a reduction in sampling resolution from 24 to 6 hours is of small benefit compared to refining the resolution even further.

1 Introduction

In recent years, utilities around the world have installed smart (water) meters at an increasing rate. This implementation entails a wide range of benefits for utilities and customers, including reduced operational costs, enhanced demand management, efficient pipe network infrastructure planning, accurate billing, improved leakage detection and higher customer satisfaction (Stewart et al. 2018; Monks et al. 2019). In model based water distribution systems analysis, the quantity and quality of the available data is often one of the most restraining factors (Savic et al. 2009) and the expected widespread use of smart meter data has the potential to improve the applicability and accuracy of the models (Gurung et al. 2014, 2016).

The sampling resolution, or sampling rate, of the data describes the temporal distance between a device's water meter readings (WMRs) and has a major influence on the usefulness of the data. Cominola et al. (2018) showed that a fine sampling resolution (ranging from minutes to seconds) increased significantly the accuracy of end-use disaggregation. Furthermore, they showed that the magnitude and timing of the peak demand can vary up to 62% and more than 15 hours, respectively, when changing the sampling resolutions from 10 s to one day. Also, Gurung et al. (2014) showed how enhanced demand patterns could be generated based on a very fine sampling resolution (5 s) in smart meter data sets, reducing measured peak demand levels and updating the time of peak occurrence. This is crucial information, as demand patterns and associated peaks are used to design and optimise water distribution systems (Gurung et al. 2014; Cominola et al. 2018). In terms of water quality modelling, Blokker et al. (2008) showed that the spatial aggregation and sampling resolution

play a significant role. For larger transportation networks water quality models including advection-reactions can be sufficient at 1 hr sampling resolutions; but at finer spatial aggregations, water quality models should also include dispersion reactions, with sampling resolutions at least below 5 min. This should be done, for example, when simulating water age in parts of the networks where laminar flows occur (Blokker et al., 2010). However, in Blokker et al. (2011), a hydraulic modelling time step of 15 min was determined to be sufficient for accurate residence time computations using a bottom-up demand allocation approach (i.e. unique demand patterns for each household). Creaco et al. (2017) showed that when using the conventional top-down approach (i.e. allocating strongly correlated demand patterns to nodes) pressure head simulations require larger time steps (here, ≥ 1 hour) before being deemed reliable. Also, Creaco et al. (2017) showed that the bottom-up approach combined with extended period simulations are capable of generating accurate pressure and flow simulations at time steps larger than 2 min.

Even though case studies have shown the benefits of finer sampling resolutions, the increased data volumes, originating from implementing advanced metering infrastructure, also pose challenges for utilities. One example includes data management, as utilities may struggle with identifying the best type and frequencies of data needed for tasks related to operation and planning of water distribution systems (Boyle et al. 2013). With a finer sampling resolution, the required data storage per installed meter grows, and there is an important trade-off between sampling resolution and battery life of battery-driven smart meters. One possible solution to mitigate the latter challenges is the emergence of low power wide-area networks (Stewart et al. 2018) that reduce the required energy during transmission, but at the cost of coverage/range, payload length (bytes per message) and maximum number of messages send per day per device (Mekki et al. 2019). Common for most large-scale rollouts of smart meters is that the smart meters use their own transmission channels with more focus on energy efficiency than reliability, as opposed to traditional system sensors connected to the SCADA system. This makes smart meter data conceptually different from other types of data from water distribution systems. Where data from traditional sensors are transmitted and stored at regular configured intervals, the smart meter readings can arrive at more irregular intervals depending on the system setup.

Our study aims to provide information that support utilities in selecting the sampling resolution best suited for their needs. This is done by resampling collected fine-resolution smart meter data to coarser resolutions and investigate its impact on modelling results. We also investigate how different approaches to fill data gaps in data sets with random or uniformly sampled time step intervals affect hydraulic simulation. This is done by comparing gap filling by linear interpolation with interpolation based on expected demand pattern data. We compare simulation results of water consumption, pressure head and water age based on data set scenarios with varying sampling resolution and gap filling methods with a reference simulation based on the finest sampling resolution available.

2 Methodology

The impact of sampling methods and sampling resolutions on total district metered area (DMA) consumption, pressure head and water age simulations were assessed for a number of data scenarios

(Table 1). The DMA consumption represents the summed water usage of all users in the DMA over time.

Table 1. Various data set scenarios applied to water distribution network simulation.

Sampling resolution [hr]	Finest available; 0.5; 1; 2; 3; 4; 6; 12; 24
Sampling method	Uniform; random
Gap filling methods	Linear interpolation; demand pattern
Data basis (for demand pattern generation)	Full; as sampling resolution

Sampling resolution

The total sampling resolution for the entire data set of $N_{devices}$ smart meters is calculated as:

$$\text{sampling resolution} = \frac{\text{Covered sampling period} * N_{devices}}{\sum_{i=1}^{N_{devices}} n(i)} \quad (1)$$

where $n(i)$ is the number of WMRs for the i -th device. The calculated sampling resolution defines the average temporal distance between WMRs in the data set.

Sampling method

The sampling interval between a device's individual WMRs can be uniform (e.g. every full hour) or pseudorandom for similar sampling resolutions. Pseudorandom data may occur when data transfer depends on water consumption triggering the transmission of messages, unstable connections to data collection devices, or similar. Scenarios with pseudorandom sampling intervals were created by deleting data points at random from the raw data set, until the required sampling resolution according to (1) was achieved. Scenarios with uniform sampling intervals required the generation of WMR sampling points at uniform time step intervals. Uniform intervals were generated by linear interpolation of gaps between WMRs in the raw data set.

Gap filling method

The WMRs come in irregular time intervals and need to be aligned for hydraulic simulations. We ran simulations at a resolution of 5 minutes in order to compare the model results with the DMA inflow observations and not to lose too much information in periods with small temporal distances between WMRs. Therefore, the WMR data need to be resampled at 5 min resolution which implies estimating values in between the sometimes much coarser WMR data. In the following this process is called gap filling. Gaps were filled either by distributing the consumption in-between two WMR readings according to a representative demand pattern or by linear interpolation between the same WMR data. Figure 1a shows an example 24-hour demand pattern used to fill gaps between four random WMRs with accumulated volume readings from a smart meter (Figure 1b). Even though the total consumption over the shown period was the same, both estimation methods resulted in different flow rates (Figure 1c).

Demand pattern generation

Demand patterns for the total consumption in the DMA were generated based on the available WMRs of all smart meters in a given period that was independent of the period used for model simulation. First, all WMRs were linearly interpolated to the nearest uniform timestamp t (with an interval of 5 min). Next, the mean consumption of all consumers \bar{x}_t , for each time step t available on a working

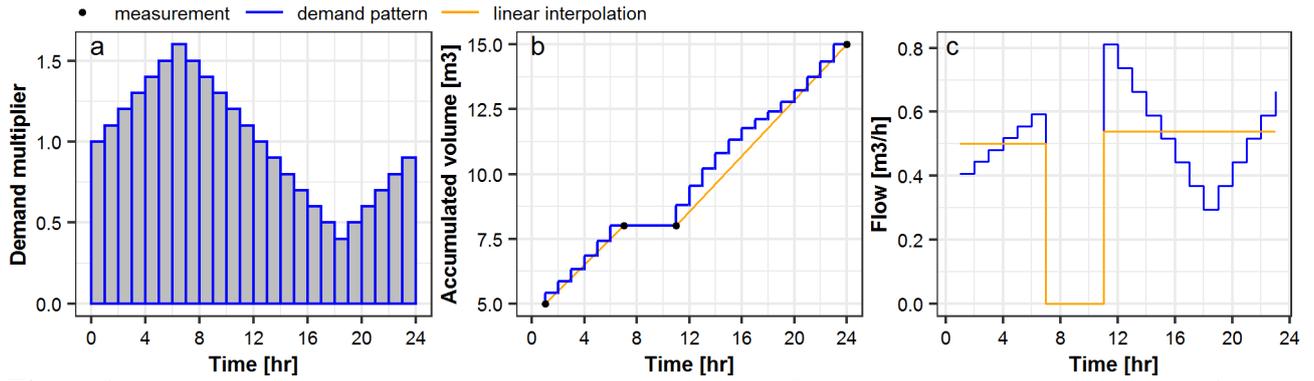


Figure 1. A demand pattern (a) or linear interpolation are used to fill data gaps between a smart meter's accumulated volume measurements (b), resulting in different flow rates over time (c).

day and a weekend day was computed, respectively. The demand patterns for a working day and weekend day were then found by:

$$\text{Demand multiplier}(t) = \frac{\bar{x}_t * n(\bar{x}_t)}{\sum_{t=1}^{\text{timestamps}} \bar{x}_t} \quad (2)$$

where $n(\bar{x}_t)$ is the total number of time steps per 24 hours (each day = 288 time steps).

Data basis for demand pattern generation

The data basis describes whether the period used for generating the demand patterns was using the finest sampling resolution available (termed 'full' availability) or whether the data to generate demand patterns was equal to the applied sampling resolution scenario listed in Table 1.

Error assessment

The root mean square error (RMSE) was computed for 1) the difference between measured inflow to the DMA and total water consumption of various data set scenarios (Table 1) and 2) between results of pressure head and water age simulations (only at nodes with consumers assigned) of a reference model and the models based on the varying data set scenarios. The reference model was based on the finest sampling resolution available, and linear interpolation was used as data gap filling method to generate uniform time steps.

Standard model

A 'standard model' and its simulation results were included, representing today's best practise. In this model, consumers were assigned a demand average based on, e.g., quarterly meter readings, in combination with a demand pattern derived for the given type of consumer.

3 Case study

The methodology was applied to smart meter data collected in a DMA in Brønderslev, Denmark (Table 2).

Table 2. Information about a district metering area and the hydraulic network in Brønderslev, Denmark. WMR = Water meter reading.

Smart meters	Total WMRs (01 st -15 th 08/18 and 01 st -15 th 01/19)	Sampling resolution [min] (01 st -15 th 08/18 and 01 st -15 th 01/19)	Network junctions (with smart meter assigned)	Network pipes (average length [m])
526	368,840/412,016	29/26	128 (103)	133 (72)

All of the installed smart meters are of type Diehl's water meter, HYDRUS, based on ultrasonic technology (Diehl Stiftung & Co. KG 2019). The analysed data included accumulated volume readings and respective timestamps. The meters transmit WMRs every 20 s to a system of antenna towers that collects data from the surrounding smart meters and transmit the WMR closest to the nearest full hour to the utility. If more than one tower picks up a signal from a given meter within an hour, a record from each of these towers will be stored in the utility's database. This implies that the sampling frequency can seem rather irregular or random (Figure 2).

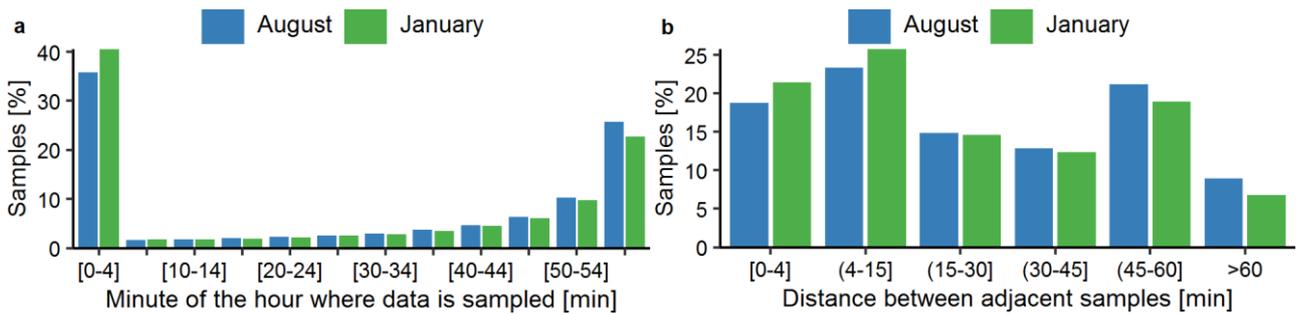


Figure 2. Distribution of the minute of the hour for all water meter readings (a) and the distribution of temporal distances between adjacent samples for individual meters (b).

Demand patterns were generated based on data from the period 1st to 15th August 2018 (Table 2). All remaining analyses were performed using data from the period between 1st and 15th of January 2019. All hydraulic simulations were run in EPANET (Rossman 2000). The hydraulic model consists of 133 pipes and 128 nodes and smart meter data were assigned to the nearest node, averaging 4 smart meters per network node. Measured pressure with a sampling interval of 5 min at the single inlet to the DMA was used as boundary condition, and the model was accordingly ran in 5-minute time steps.

The raw WMR data set (Table 2) represents the reference case resolution with average sampling intervals of 29 and 26 min/sample/device during 2018 and 2019, respectively. Simulation results based on the first 84 hours were not included when computing the RMSE to avoid impact from the initial water age conditions in the model.

In the standard model, the averaged demand of each consumer, based on the difference between the first and last available WMR in the analysed period, was assigned to the associated nodes in the

model. Furthermore, the finest demand pattern generated from the 2018 data set was linked to this model.

4 Results

Demand patterns

The generated demand patterns for a weekend and a working day in August 2018 are shown in Figure 3. Unsurprisingly, the level of detail of the demand patterns decreases gradually with a coarser sampling resolution. In general, a decreasing peak demand and changes in the timing of the peaks appear in connection with a coarser sampling resolution. The demand patterns show less variation for sampling resolutions of > 6 hours and do no longer clearly follow the daily pattern of the finest sampling resolution. The uniform sampling method results in discretized demand multipliers (Figure 3b and Figure 3d). For example, a uniform sampling resolution of 12 hours, starting at 00:00 hrs, only includes two WMRs per day per device which makes it difficult to assess the water consumption behaviour.

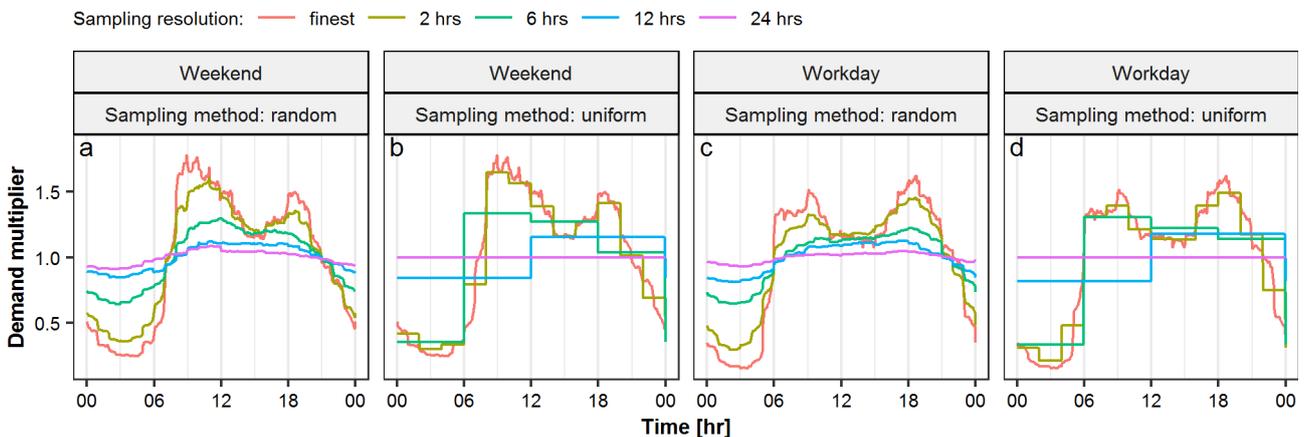


Figure 3. Generated DMA demand patterns, based on varying sampling methods and resolutions during 01st – 15th of August 2018.

Measured vs. simulated consumption

Figure 4 illustrates how the total water consumption varies for the different data set scenarios (Table 1) in comparison to the measured flow into the DMA over two selected days. At the finest sampling resolution, the total consumption resembles the measured inflow into the DMA to a high degree. Thus, it can be assumed that the amount of unaccounted water consumption is negligible in this DMA. Furthermore, it is evident from the Figures 3 and 4 that the choice of gap filling method makes only little difference when the sampling resolution is 1 hour or less. At coarser resolutions, however, the demand pattern based gap filling outperforms linear interpolation. By comparing the demand pattern lines on Figure 5, it is evident that the gap filling works best when based on fine resolution data (i.e. ‘full data set of 2018’) – which especially is the case when the sampling resolution is coarse.

Figure 5 clarifies that the use of a good demand pattern resulted in a simulated total consumption closer to the measured inflow than when using linear interpolation at a sampling resolution coarser than 1 hour. However, when the sampling method was uniform and the information to generate demand patterns was equal to the sampling interval, the difference between linear interpolation and

demand pattern gap filling is no longer visible (Figure 5b). The standard model outperformed linear interpolation at sampling resolutions coarser than 2 hours. At a sampling resolution coarser or equal to 12 hours, the standard model achieved results comparable to using demand pattern gap filling based on the full data set. For all gap filling methods the modelling results continue to improve as the resolutions become finer.

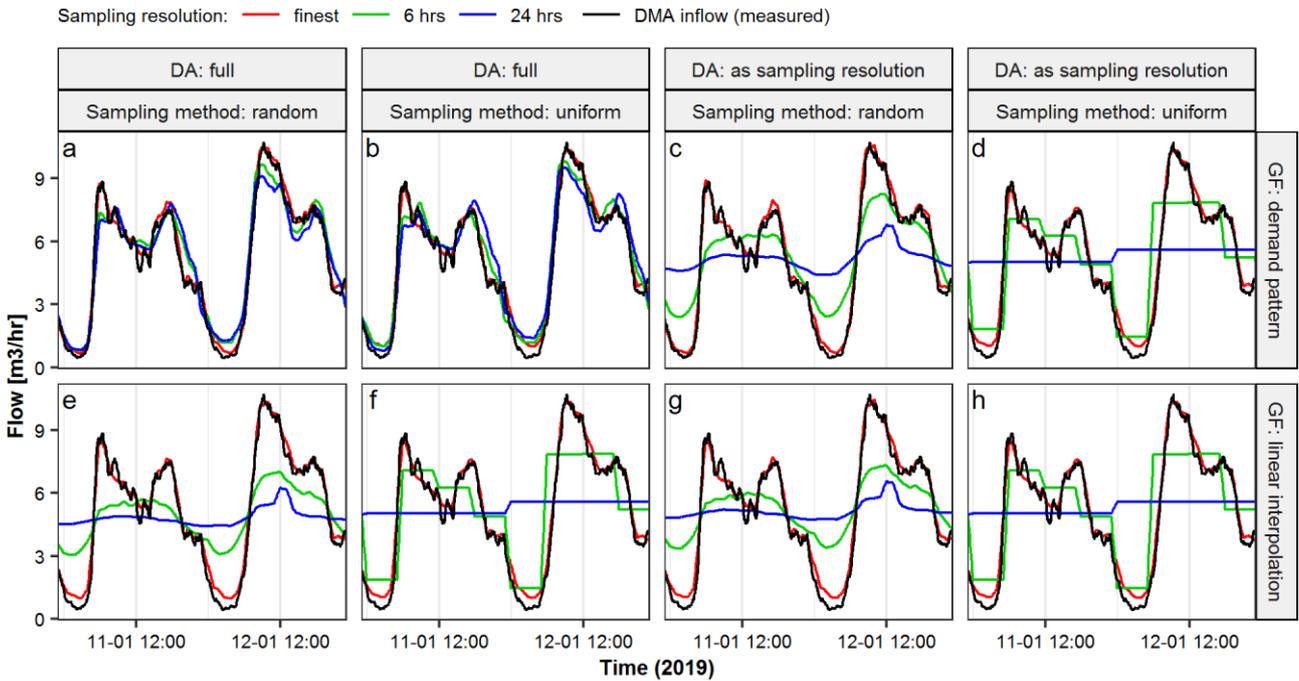


Figure 4. Example of simulated total consumption over two days in 2019 in DMA with varying sampling resolution, sampling methods and gap filling methods. For smooth illustration, a rolling mean window of 1 hour was applied to the data set. DA = data availability; GF = gap filling.

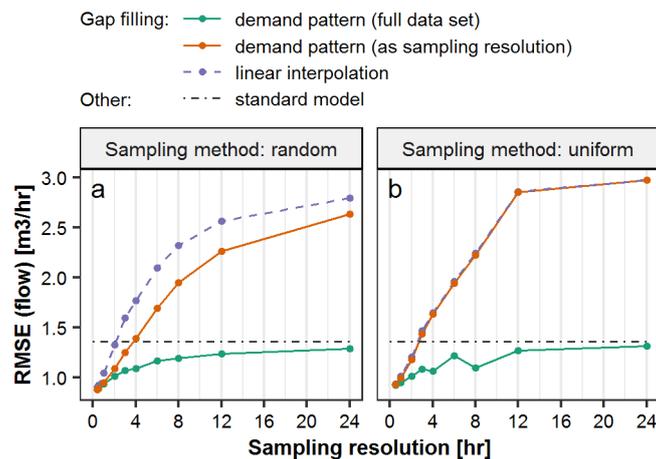


Figure 5. Root mean square error (RMSE) between measured inflow in DMA and accumulated demand in water distribution network simulations at various sampling resolutions and gap filling methods. The standard model represents a top-down approach where the same demand pattern was assigned to each node based on the consumption over 5 months.

Simulation of pressure head and water age

The average pressure head loss between the inlet and network nodes with smart meters is 0.16 m in the reference model. This low head loss is due to over-dimensioned pipes, owing fire safety regulations, which is typical for Danish water distribution networks. Consequently, also the average RMSE of change in pressure over a coarser sampling resolution is low (Figure 6a–b) and reaches mean RMSE values of only approximately 0.15 m at the coarsest sampling resolution of 24 hours. However, Figure 6 still illustrates differences in sampling methods and sampling resolutions. For sampling resolutions coarser than ~ 2 hours the gap filling methods have a large impact on the results, with clearly the best results achieved when using demand pattern based gap filling with a demand pattern based on the finest resolution available from 2018. However, when the available information to generate demand patterns is reduced, also the difference between linear interpolation and demand pattern gap filling decreases (Figure 6b). The standard model resulted in better pressure results than linear interpolation at sampling resolutions coarser than 4 hours, as well as to comparable results as demand pattern gap filling based on the full data set for sampling resolutions of around 12 hours or more.

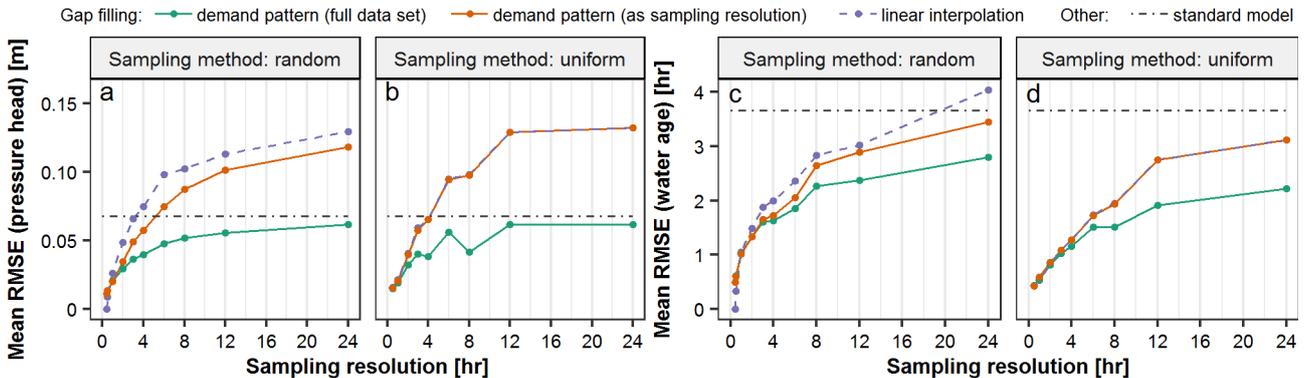


Figure 6. Mean root mean square error (RMSE) of pressure head (a–b) and water age (c–d) simulations with varying sampling methods and sampling resolutions based on all nodes with smart meters installed.

The average water age in the reference model was 17.7 hours. Depending on sampling resolution and method the average RMSE varied between 0 and 4 hours (Figure 6c–d). The choice of gap filling method had smaller impact on the water age than on pressure and flow, which especially is noticeable for sampling resolutions of 4 hours or less. This can be explained by the fact that the age of the water is a result of many hours of consumption so the age is less dependent on short term variations. As seen with the pressure head change, the demand pattern based gap filling still resulted in better overall results, however. Moreover, the water age simulations performed better with uniform sampling. This can be explained by the fact that the random sampling method can have smart meters without any or only few WMRs, whereas the uniform sampling method still required data from all meters at a given sampling resolution. Thus, if certain areas in the network had no WMRs for long periods in the random sampling method, the error in water quality simulation at these locations increased. The greatest difference between finer sampling resolutions and the standard model were seen in the water age simulations (Figure 6d). Demand pattern gap filling outperformed the standard model at all sampling resolutions and only at a sampling resolution of 24 hours the standard model was better than linear interpolation. This shows that smart meter data improve significantly water quality simulations,

while there is little difference in pressure and total demand simulations when using sampling resolutions greater or equal to four hours.

5 Conclusions

By using household smart meter data to specify consumers demand in a water distribution model and by varying the data sampling resolution of the smart meter data we found that:

- If there is a limit under which a finer resolution will not improve modelling results any further, then it is below 30 minutes.
- There is not much difference between a 6 hour and 24 hour resolution – the biggest improvements occur when the resolution becomes finer than 6 hours.
- It is much better to use representative demand patterns than linear interpolation to fill gaps between observations. A data resolution of 4 hours led to comparable or worse results in terms of pressure and flow than 24 hour data with the demand pattern based gap filling method.
- Water age simulations improve greatly by the use of smart meter data almost regardless of how the data is treated, while simulations of flow and pressure are affected more by choice in gap filling method.

Overall, it can be concluded that smart meter data can greatly improve modelling results, and if the temporal data resolution is coarser than 2 hours, then demand pattern based gap filling should be used.

Acknowledgment

We thank Brønderslev Forsyning A/S, in particular Henrik Horsholt Christensen, Per Grønvald and Thorkil Bartholdy Neergaard, for providing data and answering questions related to their water supply. This project was partly funded by the Danish Eco-Innovation Program (MST-141-01277/NST-404-00378).

References

- Blokker, E.J.M., Beverloo, H., Vogelaar, A.J., Vreeburg, J.H.G., van Dijk, J.C., 2011. A bottom-up approach of stochastic demand allocation in a hydraulic network model: a sensitivity study of model parameters. *J. Hydroinformatics* 13, 714–728. doi:10.2166/hydro.2011.067
- Blokker, E.J.M., Vreeburg, J.H.G., Beverloo, H., Klein Arfman, M., van Dijk, J.C., 2010. A bottom-up approach of stochastic demand allocation in water quality modelling. *Drink. Water Eng. Sci.* 3, 43–51. doi:10.5194/dwes-3-43-2010
- Blokker, E.J.M., Vreeburg, J.H.G., Buchberger, S.G., van Dijk, J.C., 2008. Importance of demand modelling in network water quality models: a review. *Drink. Water Eng. Sci.* 1, 27–38. doi:10.5194/dwes-1-27-2008
- Boyle, T., Giurco, D., Mukheibir, P., Liu, A., Moy, C., White, S., Stewart, R., 2013. Intelligent Metering for Urban Water: A Review. *Water* 5, 1052–1081. doi:10.3390/w5031052
- Cominola, A., Giuliani, M., Castelletti, A., Rosenberg, D.E., Abdallah, A.M., 2018. Implications of data sampling resolution on water use simulation, end-use disaggregation, and demand management.

Environ. Model. Softw. 102, 199–212. doi:10.1016/j.envsoft.2017.11.022

- Creaco, E., Pezzinga, G., Savic, D., 2017. On the choice of the demand and hydraulic modeling approach to WDN real-time simulation. *Water Resour. Res.* 53, 6159–6177. doi:10.1002/2016WR020104
- Diehl Stiftung & Co. KG, 2019. Hydrus. Available at <https://www.diehl.com/metering/en/portfolio/water-metering/water-metering-product/hydrus/63697/> (accessed July 20, 2019).
- Gurung, T.R., Stewart, R.A., Beal, C.D., Sharma, A.K., 2016. Smart meter enabled informatics for economically efficient diversified water supply infrastructure planning. *J. Clean. Prod.* 135, 1023–1033. doi:10.1016/j.jclepro.2016.07.017
- Gurung, T.R., Stewart, R.A., Sharma, A.K., Beal, C.D., 2014. Smart meters for enhanced water supply network modelling and infrastructure planning. *Resour. Conserv. Recycl.* 90, 34–50. doi:10.1016/j.resconrec.2014.06.005
- Mekki, K., Bajic, E., Chaxel, F., Meyer, F., 2019. A comparative study of LPWAN technologies for large-scale IoT deployment. *ICT Express* 5, 1–7. doi:10.1016/j.ict.2017.12.005
- Monks, Stewart, Sahin, Keller, 2019. Revealing Unreported Benefits of Digital Water Metering: Literature Review and Expert Opinions. *Water* 11, 838. doi:10.3390/w11040838
- Rossman, L.A., 2000. EPANET 2 Users Manual. Cincinnati, OH.
- Savic, D.A., Kapelan, Z.S., Jonkergouw, P.M.R., 2009. Quo vadis water distribution model calibration? *Urban Water J.* 6, 3–22. doi:10.1080/15730620802613380
- Stewart, R.A., Nguyen, K., Beal, C., Zhang, H., Sahin, O., Bertone, E., Vieira, A.S., Castelletti, A., Cominola, A., Giuliani, M., Giurco, D., Blumenstein, M., Turner, A., Liu, A., Kenway, S., Savić, D.A., Makropoulos, C., Kossieris, P., 2018. Integrated intelligent water-energy metering systems and informatics: Visioning a digital multi-utility service provider. *Environ. Model. Softw.* 105, 94–117. doi:10.1016/j.envsoft.2018.03.006

II

A semi-automated approach to validation and error diagnostics of water network data

J.K. Kirstein, K. Høgh, M. Rygaard and M. Borup

RESEARCH ARTICLE



A semi-automated approach to validation and error diagnostics of water network data

Jonas Kjeld Kirstein ^a, Klavs Høgh^b, Martin Rygaard ^a and Morten Borup ^a

^aUrban Water Systems, Department of Environmental Engineering, Technical University of Denmark, Kgs. Lyngby, Denmark; ^bNIRAS A/S, Allerød, Denmark

ABSTRACT

We propose a method for quality assurance of raw data from water distribution networks in near real-time. Well-known and novel data analysis methods, including a timestamp drift test, are combined to produce a malfunction indicator database for diagnosing anomalies within data acquisition practices. The method was applied to 112 flow and 111 pressure data sets, covering on average 32 months, located throughout the distribution networks of three Danish utilities. Around 10% of measurements in the utilities' meter data sets were absent and 3–35% were categorized as dubious or erroneous. The most common types of anomalies for flow and pressure data were flatline and time stamp inconsistencies. Time drifts were identified in all three utilities and a similarity analysis revealed a simultaneous occurrence of many anomalies. These high rates could have been avoided if the proposed method had been implemented to automatically highlight meter errors and system-wide problems in data collection.

ARTICLE HISTORY

Received 30 April 2018
Accepted 22 April 2019

KEYWORDS

Data validation; error diagnostics; water supply

1. Introduction

Erroneous meter data from water distribution networks can lead to incorrect conclusions during data fusion and data analysis in the water supply sector. Reliable data is a fundamental prerequisite for leakage detection, water quality monitoring, hydraulic modelling and network optimization. To secure a high data integrity it is therefore essential to validate the collected data (Quevedo et al. 2017) and to maintain high data accuracy and data reliability, as promoted by the International Water Association (IWA) (Alegre et al. 2006). This is of growing importance, as utilities are drawing inferences at an increasing rate from data-driven applications and the collected data itself as sensors and data transmission become cheaper. Combining all data sources can be used proactively to improve network operations in a utility (Machell et al. 2014). Andrews et al. (2017) found that it is common for many utilities to struggle with the validation of their own data. However, they also found that, once validated, the data builds the foundation for advanced applications, such as a successful water loss control program. Data validation and anomaly detection methods are not new to urban water engineering. Examples include rainfall data (Jørgensen et al. 1998; Estévez, Gavilán, and Giráldez 2011), urban hydrology (Mourad and Bertrand-Krajewski 2002; Branisavljević, Kapelan, and Prodanović 2011), wastewater treatment (Rosen, Röttorp, and Jeppsson 2003; Puig et al. 2008) and the collection of data on water quality (García et al. 2017) and flow (Quevedo et al. 2010; Cugueró-Escofet et al. 2016) in water distribution networks. Typically, validation and anomaly detection are performed in a combination of 'low-level' methods, based on simple heuristics and limited statistical knowledge, and 'high-level' model-based approaches. High-level approaches use the spatial and temporal redundancy in the available data sets to

flag suspicious and erroneous data (Quevedo et al. 2017). Even though data validation does not represent a new research field, it is still not a priority issue at many water supply utilities. Whereas larger Danish water utilities are benchmarked on an annual basis to compare their performance in terms of operational costs, unregistered water losses, etc. (DANVA 2016), no comparison or benchmark exists in their data acquisition and management practices. In the American Water Works Association's Free Water Audit Software (WLCC 2014) utility practices are graded; the software, however, does not focus on anomalies in raw meter data. In this paper, we make a call for a renewed interest in data quality and improved management of errors in logged data.

Defining dubious data points as outliers (in this paper referred to as anomalies) is a subjective concept, as classification varies among practitioners, researchers and the applied methodologies (Helsel and Hirsch 2002; Rosen, Röttorp, and Jeppsson 2003; Branisavljević, Kapelan, and Prodanović 2011). The cause for erroneous or abnormal data is a combination of many factors, such as meter malfunctions, problems occurring in the data transmission and storage processes, changes in network system operations or burst pipes (Loureiro et al. 2016; Quevedo et al. 2017). We suggest dividing anomalous data into three categories:

- Type 1: Anomalies caused by faults internally in the meter, during transmission, storage, etc. that cannot be due to any state of the water distribution network (illegitimate data).
- Type 2: Measurements that affect the data quality negatively but have some valuable information left; examples include a loss of sensitivity and miscalibration of sensors.

- Type 3: Abnormal measurements that are caused by actual events appearing in the network, such as leakage, irregular consumption or valve opening.

Type 1 contains data with no relevance for the operation of the system, because the integrity of the data is damaged or missing. This type of anomaly has also been described as ‘dirty data’ (Mounce, Boxall, and Machell 2010; Machell et al. 2014). Examples include missing or illegitimate observations, flatline segments, erroneous timestamps or duplicates. Applying time series data that includes type 1 data in hydraulic, leakage or water quality models can lead to incorrect conclusions. Since type 1 data contains no valuable information, it is preferable to remove type 1 anomalies from the time series before further analyses are made. Compared to type 1, anomalies of types 2 and 3 have a higher level of information and the major difference lies in the integrity. Type 2 includes erroneous data that can be corrected at times. These include, for example, errors due to wrong time settings, miscalibration of instruments or loss of sensitivity as seen in water quality sensors (García et al. 2017). Anomalous data of type 3 represent valuable information from correct physical measurements but comprises all abnormal observations that are caused by actual events in the network, such as burst pipes, pump trips, irregular consumption or valve operations. Research on abnormal consumption, leakage and/or burst detection methodologies is dependent on anomalies of type 3, while types 1 and 2 reduce the reliability of such applications (Wu and Liu 2017). It is our aim to develop a structured and improved identification of type 1 and 2 anomalies without misclassifying type 3 anomalies.

The existing literature focuses on batch validation of already collected data to prepare the data for further applications. The current work focuses on near real-time validation of the data as it is collected, while producing diagnostic plots that help the operators to detect errors in the data collection. Proper visualization of anomalies is just as important as their detection, since this enables the operator to correct for errors on a daily basis. If this is not to some extent automated in near real-time, experience shows that errors can prevail for months or years. Also, benchmarking, in terms of the number of anomalies registered over a selected period, can help utilities evaluate their data collection and acquisition performance.

To produce maximum benefits from the collected information, a well-maintained operational database containing validated and easily accessible data should be constructed (Cugueró-Escofet et al. 2016; Loureiro et al. 2016). In such a database, invalid and missing observations can be represented by coexisting qualified estimates based on, for example, time series analysis, physically based models or machine learning approaches such as artificial neural networks (Mounce, Boxall, and Machell 2010; Quevedo et al. 2010; Branisavljević, Kapelan, and Prodanović 2011; Cugueró-Escofet et al. 2016; García et al. 2017). Such reconstruction can benefit from the prior identification of anomalies, but we consider reconstruction of data outside the scope of the article.

Here we present a data processing framework that identifies anomalies in the collected data and store these in a malfunction indicator database (MAID) that is used for further analysis and visualization of anomalies. The method can validate the collected data without expert assistance and without sensor-specific parameterization, which will be useful when applied in future systems that include thousands of online sensors. If needed, the test parameters can be tuned to change the sensitivity of anomaly detection. To identify patterns in error occurrence, a diagnostic tool based on the Jaccard coefficient is proposed. The tool gives operators an idea of where to improve future data collection procedures. The method is demonstrated on 223 meter data sets from three Danish utilities where no data validation procedures are currently in place.

2. Methodology

2.1. Data processing framework

A conceptual scheme of a utility’s data collection, processing and analysis system is shown in Figure 1. The system starts with the collection and storage of data from various sources, such as water quality sensors, pressure and flow meters, in a raw meter database. The next step is to categorize every raw data point as either an anomaly or a valid data point, based on a series of anomaly testing. The results from the tests are binary flags (true or false) stored in the MAID as an amendment to the operational database. The MAID highlights anomalies logged by sensors and facilitates the investigation of possible patterns in errors occurring in the data series. Since

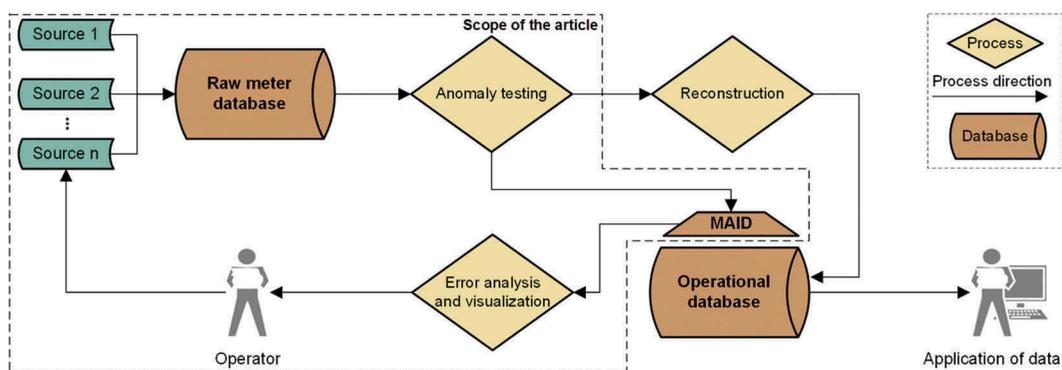


Figure 1. Data processing framework illustrating the transformation of data collected in a raw meter database to an operational database. The malfunction indicator database (MAID) amendment stores flags from the anomaly testing process, then uses this for the analysis and visualization of errors, ultimately suggesting improvements in current data collection procedures.

the test flags are binary values, the storage demand for the MAID entries is negligible compared to, for example, timestamp entries in the raw and operational database.

In the following, the entries in the databases are referred to as matrices and vectors:

- A raw data set from one meter consists of n observations $\mathbf{x} = [x_1, \dots, x_n]$ with respective timestamps $\mathbf{t} = [t_1, \dots, t_n]$.
- The MAID data for each data series is an $n \times m$ matrix \mathbf{M} of binary values, where m is the number of tests included in the anomaly testing phase. The value of $\mathbf{M}_{k,i}$ is true only if the k -th test for the i -th time step detects an anomaly.

Flagged and missing data should be reconstructed and stored together with the validated data in an operational database. Such data are often stored in a uniform manner to account for differences in timestamp intervals between various data streams. This application of the data is outside the scope of this article (Figure 1). However, our error analysis and visualization step can use the flags stored in the MAID to provide both short- and long-term diagnostics as well as day-to-day visualizations of errors for use in the daily operation. Having detected anomalies, operators can use this information to investigate whether the data are, in fact, erroneous and to improve future data collection processes (Figure 1).

2.2. Anomaly testing of raw data

Seven tests (I–VII) form the anomaly testing process of the raw data that aim to identify type 1 and type 2 anomalies.

A short summary explaining the occurrence of selected anomalies and justifying the necessity of the tests is given in the supporting information (SI) A. Tests I and II are run in a sequential manner to provide regular data streams to the subsequent tests III–VI that can be run in parallel. Finally, all data points flagged by tests I–VI are excluded in test VII. Whereas parts of the described tests are included in standard extract, transform, load (ETL) processes, their implementation may vary among practitioners. Moreover, it is important to note that this list of anomaly tests is not exhaustive, as it is based on the most common anomalies that appeared in the analysed data sets. Figure 2 exemplifies periods with validated and flagged data for each of the seven tests in a simplified manner. Based on time, a data point is seen as valid (marked green) until a new data point, pattern or inconsistency violating the tests' constraints has been registered (marked red).

1. Duplicate timestamp test

The duplicate timestamp test (Figure 2(a)) flags any t_i that is not unique in \mathbf{M}_1 . Similarly, the raw data is checked for a regular data stream at the 'communications level' in Cugueró-Escofet et al. (2016). If duplicate timestamps can be replaced by known values, such as by an operator, it is possible to adjust the data set without compromising future applications. Otherwise, all duplicates represent a type 1 anomaly.

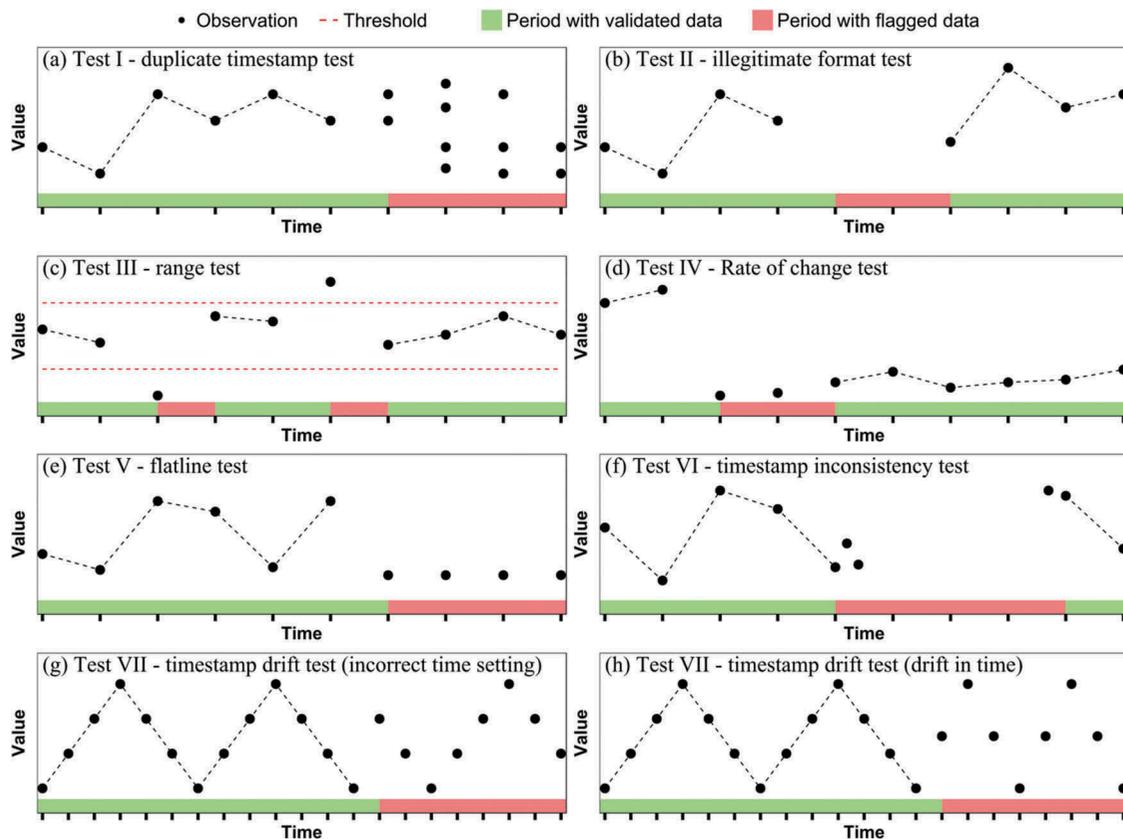


Figure 2. Overview of anomaly tests. Green and red distinguish valid and invalid data. Succeeding validated data are connected with a dashed line.

II. Illegitimate format test

In the illegitimate format test (Figure 2(b)), all timestamps t_i that are linked with non-numerical observations x_i are flagged in \mathbf{M}_2 .

III. Range test

The range test (Figure 2(c)) identifies all x_i below or above a minimum or maximum threshold, i.e. physically unfeasible or dubious compared to the distribution of the data series. Inspired by the locally realistic range test (Mourad and Bertrand-Krajewski 2002), we flag values in \mathbf{M}_3 based on the distance between the median and the 2.5th and 97.5th percentiles of the data series, multiplied by a factor of α and β , respectively.

$$\mathbf{M}_{3,i} = \text{true} \begin{cases} < P_{50\%}(x) + [P_{2.5\%}(x) - P_{50\%}(x)] * \alpha, \alpha > 1 \\ > P_{50\%}(x) + [P_{97.5\%}(x) - P_{50\%}(x)] * \beta, \beta > 1 \end{cases} \quad (1)$$

Being similar to the interquartile range, this test is not as sensitive to outliers in data series as standard score-based methods (i.e. using distance to the mean relative to the standard deviation). This is particularly important, as anomalies in the data series should not be allowed to obstruct the identification of future anomalies. There is always a risk of type 3 being flagged, since the test is based on the historical distribution of the data (e.g. if the data set has not yet covered legitimate seasonal changes). This risk will decrease if α and β are increased but at the cost of sensitivity. Note that α and β values of 1 will result in 5% of the data being flagged as anomalies, while larger values for these parameters rapidly decrease this fraction, such that α and β values of 1.5 for normal distributed data will result in just 0.3% being flagged.

IV. Rate of change test

A rate of change test (Figure 2(d)) identifies unusual changes in observations over time based on the concepts of the step test (Estévez, Gavilán, and Giráldez 2011), signal gradient test (Mourad and Bertrand-Krajewski 2002) and trend level test (Cugueró-Escofet et al. 2016). A threshold θ for a likely rate of change is defined by computing the 97.5th percentiles of all absolute rates of change in a meter data set multiplied by a factor of λ :

$$\theta \geq P_{97.5\%} \left(\frac{|\Delta x|}{\Delta t} \right) * \lambda \quad (2)$$

Here, the rate of change $\Delta x/\Delta t$ is the ratio between the difference of two subsequent observations and their respective differences in time. In the case that the rate of change between two observations is above θ , the measured value is flagged in \mathbf{M}_4 . Having flagged a value for two consecutive observations, the rate of change for the next observation is computed based on the last valid measurement and is tested for compliance with the threshold of θ . Flagging is stopped if a measurement is below θ . As in test III, the factor θ is based on the historical distribution of the data and might lead to false positives. Additional information about test IV is listed in SI B.

V. Flatline test

The flatline test identifies successive identical observations (Figure 2(e)). Here, flatline segments are flagged in \mathbf{M}_5 if successive identical observations more or equal to s cover a period of more than p minutes:

$$\mathbf{M}_{5,[i,\dots,i+s-1]} = \text{true} \{x_i = x_{i+1} \dots = x_{i+s-1} \text{ and } t_{i+s-1} - t_i > p\} \quad (3)$$

VI. Timestamp inconsistency test

The timestamp inconsistency test (Figure 2(f)) identifies irregularities in timestamp intervals. This is done to check that data is recorded and stored at a constant interval between measurements. Varying timestamp intervals indicate errors between measuring devices and data collection. Data are flagged in \mathbf{M}_6 if 1) the gap between measured data is more than q minutes; or 2) there are identical subsequent timestamp intervals less or equal to r . The latter check is needed to avoid data being marked as flawed because of a persistent change in logging frequency:

$$\mathbf{M}_{6,i} = \text{true} \{t_{i+1} - t_i > q \text{ or } t_{i+1} - t_i \neq (t_i - t_{i-1}, \dots, t_{i-r+2} - t_{i-r+1})\} \quad (4)$$

VII. Timestamp drift test

The timestamp drift test is used to identify meters with drifting or changing time settings. Figure 2(g-h) shows an example of a meter suddenly having a wrong time setting and a meter drifting in time. The test works by assuming similar patterns in the data over a selected period of days and weeks, which allows for the identification of time shifts by comparing a given week with previous weeks. This implies that the test only works if the overall patterns in the data do not change with time. For the test, data excluding flagged values from tests I–VI is aggregated to hourly data to provide uniform time series without changing time intervals, and to reduce the computational time. Changing time settings are identified by the following method for a data set of length j days.

First, a reference day rd and a test day td , with a period of w weeks between, are selected, i.e. the test and reference day represent the same weekday. Next, the reference day (including the lags $l = \pm 12$ h) is cross-correlated against the test day and the peak correlation value $C_{daily}(d)$ and the corresponding hourly time lag $P_{daily}(d)$ are determined:

$$C_{daily}(d) = \arg \max \left\{ \text{corr}[rd(l), td] = \frac{S_{rd(l)td}}{S_{rd(l)}S_{td}}, l = 0, \pm 1, \dots, \pm 12 \right\}, \quad d = 1, 2, \dots, n \quad (5)$$

$$P_{daily}(d) = \text{lag } l \text{ of } C_{daily}(d) \quad (6)$$

where $S_{rd(l)td}$ is the covariance of $rd(l)$ and td , and $S_{rd(l)}$ and S_{td} the standard deviation of $rd(l)$ and td , respectively. $C_{daily}(d)$ and $P_{daily}(d)$ are invalidated if more than one third of the data in the reference or test period is unavailable. Next, $C_{daily}(d)$ is used to compute a weekly correlation measure, $C_{weekly}(d)$:

$$C_{weekly}(d) = \frac{\sum_{l=0}^6 C_{daily}(d-l)}{\text{Valid test days in } C_{daily}[d : (d-6)]}, d = 7, 8, \dots, j \quad (7)$$

$C_{weekly}(d)$ is invalidated if more than three days of data are missing. Finally, a weekly measure of the hourly drift, $P_{weekly}(d)$, is determined:

$$P_{weekly}(d) = \frac{\sum_{l=0}^6 C_{daily}(d-l) * P_{daily}(d-l)}{\sum_{l=0}^6 C_{daily}(d-l)}, d = 7, 8, \dots, j \quad (8)$$

$P_{weekly}(d)$ is not computed in the case of previously invalidated values of $C_{weekly}(d)$.

To avoid a single week with a changed consumption pattern (such as a vacation week) resulting in time drift flags, steps 1–3 are run with three different reference periods and w is set to a 1-, 4- and 8-week interval. A drift is identified in the case that $P_{weekly}(d)$ exceeds a threshold of ± 2 hours, allowing the time to fluctuate slightly before a flag is set. Additionally, a flag is only raised in \mathbf{M}_7 if the following two conditions are met:

- (1) $C_{weekly}(d) \geq 0.8$, for at least two reference weeks w .
- (2) $|P_{weekly}(d)| > 2$, for more than two subsequent test days d .

These conservative conditions ensure that the test is run only on data sets with a regular daily and weekly pattern, thus reducing the risk of a false alarm.

Parameter selection. It is unavoidable that type 3 anomalies are sometimes misclassified as type 1 or 2. This is one reason why dubious data points should not be deleted or replaced but rather flagged. Most tests presented have parameters that can change the sensitivity of the test and thereby also result in a trade-off between the misclassification of type 1 and 2 anomalies versus type 3 anomalies. To what extent such misclassification is a problem is dependent on the usage of the data. Thus, favourable parameter settings may differ between specific data sets. In Sec. 4.1, we propose a default parameter setting that works well with all data from the three utilities in the current study.

When a high proportion and variety of type I and II anomalies are present in raw data sets, the distribution of the data, independent of meter and data type, differs significantly (SI C). In such cases, outlier cut-off values based on historic averages and standard deviations may have an unwanted influence on the detection of anomalies. For tests III and IV, it was decided to base anomaly detection on the distance between median and low/high percentiles in the raw data sets, which will provide a more robust approach in case of skewed distributions.

Finally, utilities may assign varying quality levels to flagged anomalies that depend on the severity of a test's outcome, e.g. for prioritization during operational troubleshooting. In this study, no quality criterion was included.

2.3. Similarity analysis of anomalies

Similarities between flagged attributes in the MAID are analysed based on the Jaccard coefficient J , suitable for sparse

and asymmetric binary data (Tan, Steinbach, and Kumar 2006). The similarity between a single binary feature \mathbf{M}_k in the MAID of two selected meters φ and ω are assessed by:

$$J(\varphi, \omega) = \frac{f_{11}}{f_{01} + f_{10} + f_{11}}, \text{ where } t(\varphi) \cap t(\omega) \quad (9)$$

Here, f_{11} is the number of attributes where φ and ω are equal to one (i.e. flagged attributes), f_{01} is the number of attributes where only φ is equal to one, f_{10} is the number of attributes where only ω is equal to one, and $t(\varphi)$ and $t(\omega)$ represent the meters' timestamps. The Jaccard coefficient ranges between no similarity (0) and a perfect similarity (1) and is only computed for matching timestamps. Here, we consider data entries with identical minute stamps to be matching.

A Jaccard coefficient close to unity means that flags occur simultaneously for two meters, which can be used to interpret the source and nature of the anomalies. If, for example, multiple sensors in an area are flagged as 'out of range' during the same periods of time, it is most likely due to actual physical conditions in the system (type 3 anomalies) rather than sensor errors. However, if sensors are located far from each other but share a high Jaccard coefficient, it might be due to database or transmission errors.

3. Case studies

Data from three Danish water utilities of varying size were analysed (Table 1). The three utilities represent typical water network data acquisition practices in Danish utilities, where mainly flow and pressure data are collected without a systematic analysis and tracking of errors and anomalies. Data is currently used for estimating minimum night flows and producing daily averages. The data sets from utilities A and C are based on pressure and flow measurements at storage locations (waterworks and tanks), DMA inlets and outlets, and other critical locations of the WDN such as pumping stations. In utility B, flow and pressure data is only collected in pipes that are highly important to the utility. Most of the installed devices in utilities A and C are Siemens MAG6000 and MAG8000 flow meters, whereas utility B measures flow with Primayer's PrimeProbe3. On average, a meter data set covered a period of 32 months. Additional information about the analysed utilities and an exhaustive description of the analysed raw meter data sets, including examples of analysed time series, can be found in SI C.

4. Results and discussion

All results listed in the following are a part of the error analysis and visualization process (Figure 1).

4.1. Anomaly testing of tests I–VI

We investigated the impact of various test parameters by manually looking through time series from the utilities, hereby identifying to what extent we agreed with the test results. Eventually, we selected a single set of parameters that works well for all three utilities (Table 2). It is this parameter set that is used for all results presented in the following, and we envisage that this

Table 1. Summary of analysed data sets from three utilities in Denmark. Consumption data from DANVA (2016). Meter types: P = pressure; Q = flow. See also SI C for additional information.

Utility	A	B	C
Name	Halsnæs Vand A/S	HOFOR (Greater Copenhagen Utility)	Nordvand A/S
Water supplied [Mm ³ /yr]	0.6	50.5	7.1
Consumers [10 ³]	10.4	585.8	142.7
Utility mains [km]	169	1085	529
Consumer/utility mains [km]	61.5	539.9	269.8
Waterworks	3	7	3
District Metered Areas (DMAs)	29	-	22
Number of meters/data sets [type]	68 (30P & 38Q)	44 (22P & 22Q)	111 (59P & 52Q)
Data period (avg. duration)	11/13–01/17 (Ø ≈ 25 months)	01/14–07/17 (Ø ≈ 33 months)	01/14–05/17 (Ø ≈ 35 months)
Total data points (P:Q)	33,487,833 (46:54)	5,691,133 (50:50)	157,876,298 (55:45)
Q – Average (min – max) ⁽¹⁾ of medians [m ³ /h]	5.7 (0–60.8)	2.9 (–3.4–12.1)	61 (–12.6–463.1)
P – Average (min – max) ⁽¹⁾ of medians [bar]	2.1 (0.3–4.3)	2.9 (1.4–5.1)	6.3 (4.8–7.4)
Q – Timestamp intervals [%]	91% in 1 min, 9% in other	54% in 5 min, 46% in 15 min	98% in 1 min, 2% in other
P – Timestamp intervals [%]	92% in 1 min, 8% in other	54% in 5 min, 46% in 15 min	97% in 1 min, 3% in other

Notes: 1) average, min and max of medians based on the median values of a utility's individual raw meter data sets.

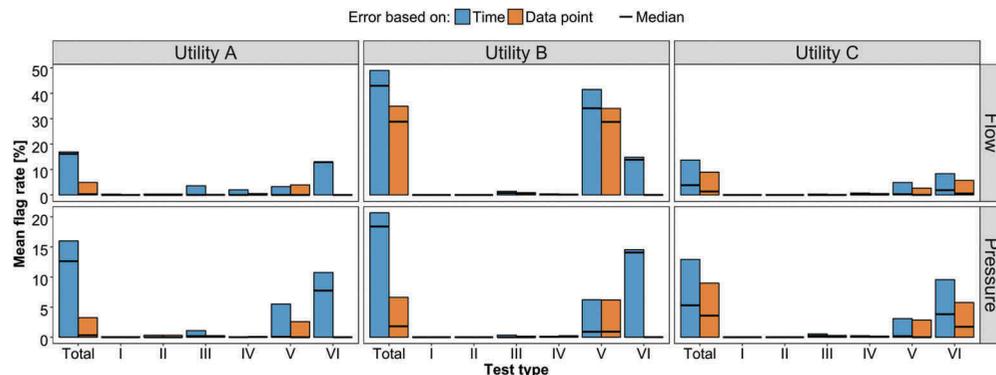
Table 2. Default set of parameters applied to the anomaly tests III–VI.

Test	Parameters
III	$\alpha, \beta = 2.5$
IV	$\lambda = 2.5$
V	$s = 3, p = 60$ min
VI	$q = 60$ min, $r = 3$

parameter set could be used as a default parameter set for other utilities as well. This default parameter set can be seen as a qualified estimate that provides a good starting point; however, each utility company can alter the parameters according to its knowledge about its system and individual sensors. The distribution of raw and validated data for utilities A–C on different time scales is shown in SI C. In SI D we discuss how a variety of different parameter selections affects the overall test results, and an example of data validation and analysis for operational use is shown in SI E. In a near real-time implementation of tests III and IV, the test parameters that are based on the historical distribution should be updated at a user-defined interval. In this study, fixed parameters based on the entire data set available were selected.

In the following, we focus on the results stored in the MAID. Two mean error rates for tests I–VI were computed based on each meter's individual data set (Figure 3). The

'data point error rate' is the average rate of errors for each sensor in the utility based on the number of flagged measurements, while the 'time error rate' is the corresponding value based on the time affected by flagged measurements. Independent of utility and meter type, the time error rate varied between 12% and 49%. Thus, it is common for all three utilities to struggle to maintain a consistent and anomaly free data stream. Only utility C has a median value of the time error rate below 5% of data entering the system in constant intervals (test VI), illustrating a better data acquisition, maintenance and handling procedure than the other utilities. For utility A, the time error rate was above 15% for both flow and pressure measurements. A similar high rate for test VI indicates that this is mainly caused by missing or inconsistent data, as there is a large discrepancy between the mean flag rate of time and data points. In utility C, however, the data and time error rate are more alike. This is due to the setup of the utility's database, as only one data point is stored if measurements do not change over 15 minutes. The highest time error rates were found in utility B, with up to 49% and 20% of flow and pressure measurements, respectively. Missing data at such high rates not only indicates clear systematic problems within data acquisition practices but also challenges the

**Figure 3.** Mean meter flag rate of six different anomaly tests for the three utilities. The line across each bar represents the median flag rate. The *data point* bars show the number of flagged data points relative to the total number of data points, while the *time* bars show the period of time represented by the flagged data points relative to the total operating time of individual meters. Test types: I – duplicate timestamp, II – illegitimate timestamp, III – range, IV – rate of change, V – flatline and VI – timestamp inconsistency.

use and reliability of the data for later applications. For utilities A and B, the timestamp inconsistency bar reveals almost no error rate based on data points but a high rate based on time, averaging 10% for all three utilities. Thus, long periods of no data collection are not unusual. On average, 35% of the flow measurements were flagged as flatlines in utility B. The lowest average error rate based on data points was found in pressure meters of utility A (3%).

Parameter selection and implications

The selected parameters have some implications for the generated results. For example, there is a high risk that the parameter selection of the flatline test (Table 2) has led to neglect of anomalies at lower timestamp intervals. A future implementation could include a varying resolution of the number of significant digits in the flatline test or minimum rate of change test to identify meters measuring data with only minor changes, such as highlighting failing pressure sensors. In the case of data sets with true flatline segments (e.g. an emergency pump flow meter) utilities should have the possibility to suppress the generated anomalies or to skip certain tests from the overall analysis. It is possible that physical boundaries in the range and rate of change tests are known, and these should be applied instead of the formulated methods. Moreover, water supply data can exhibit strong seasonality on daily, weekly and yearly scales, and ideally, the period used for determining the various test parameters should be large enough to cover all of these scales. When this is not the case, the operator should be aware that the reliability of the tests is reduced. In general, it would be beneficial to compare anomalies identified by utility personnel or experts with the flags being raised by the method, as conducted by Branisavljević, Kapelan, and Prodanović (2011).

4.2. Anomaly testing of timestamp drift test (test VII)

It is commonly assumed that the internal clock of the data sources (i.e. meter) has an adequate accuracy and only needs

few adjustments over its lifetime. To check if this is true, we applied the timestamp drift test. We applied various combinations of the parameters defining the conditions raising a flag and discussed their outcome in SI D. In the following, the test was applied according to the parameters and conditions stated in test VII (Section 2.2).

Figure 4 illustrates an example of a flow meter where a time drift is visualized. The upper row in Figure 4(a) displays the measured flow over time, and it is difficult to identify any problems from the time series itself except for missing data around the 8th of April. However, a clear drift is visible, when illustrating the time series by a week-by-week profile sectioned into the applied test and reference weeks (Figure 4(b)). The second row of Figure 4(a) shows the computed correlation value $[C_{weekly}(d)]$ for the illustrated period, based on the timestamp drift test with a 1-, 4- or 8-week interval. A correlation value above the threshold of 0.8 is observed for most of the period, i.e. a clear daily/weekly recurring pattern is present in the meter time series. The third row shows at which hour the highest correlation value $[P_{weekly}(d)]$ was determined. It can be seen that there is an upwards trend, illustrating the drift in time. The drift started at the end of March and was identified by all three reference intervals. It appears that the logger time is corrected at the beginning of April. However, whereas the 1-week interval slowly decreases towards a peak correlation hour of ± 2 (approx. 15th of April), the 4- and 8-week intervals indicate that the new time setting is approximately 6 hours different to the setting before the drift started. Thus, the drift test is capable not only of identifying drifts in time but also of showing when a time setting is different (potentially erroneous) from earlier settings.

Test VII is summarized for all utilities in Figure 5. The first test condition can be interpreted as the percentage of data where the test could be run. For example, column 1 of utility A shows that 58% of the flow data passed test condition 1. In other words, 58% of the data included a regular weekly flow pattern as required by the test. Column 2 represents the number of meters meeting test condition 2, i.e. meters where a time drift was found at least once. This was the case for 10 meters in utility A.

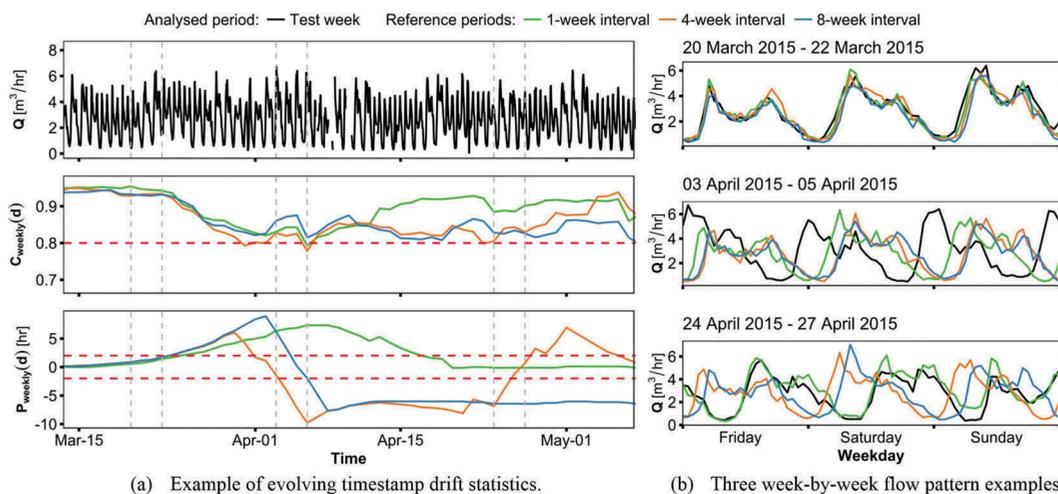


Figure 4. Example of a drifting flow meter in utility A. (a) The upper row displays the measured flow over time. The second and third row illustrate the computed correlation values $[C_{weekly}(d)]$ and $[P_{weekly}(d)]$ based on the applied test and three reference weeks needed in test VII. The drift is visible in (b) where the reference weeks are plotted against the test week at three selected periods.

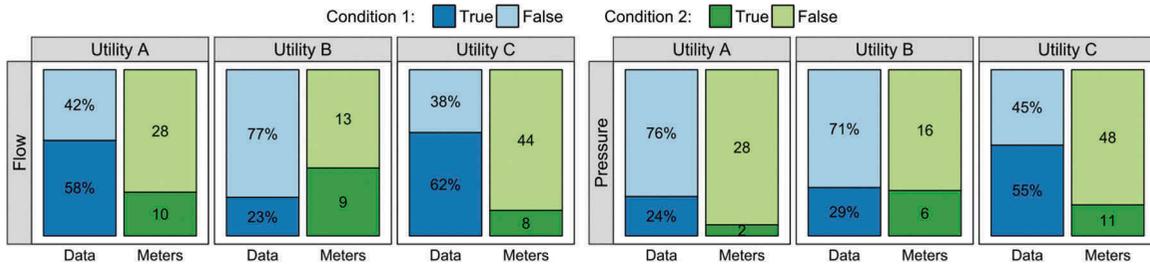


Figure 5. Summary of timestamp drift test. The data column accounts for the percentage of data where test condition 1 was met ($C_{weekly}(d) \geq 0.8$, for at least two reference weeks w). The meter column accounts for the number of meters meeting condition 2 at least once over the course of the measured period ($|P_{weekly}(d)| > 2$, for more than two subsequent test days d).

The test was conducted only on periods excluding flawed data (Figure 3). Consequently, the lowest percentage of data feasible for drift checking was found in utility B, with only 23% of flow and 29% of pressure data (Figure 5). However, 15 meters out of 44 drifting at some point is still a very high proportion of meters with a drift error. The highest share of data passing condition 1 was found in utility C, with up to 62% of flow and 55% of pressure data. Utility C confirmed that it observed drifting time in many of its battery-driven flow meters. To some extent, this validates the applied method. However, utility C had the lowest number of flow meters with a drift. A likely explanation includes the more proactive approach taken by the utility to correct meter errors. If the utility identifies and corrects a drift at a sufficiently early stage, the test will no longer raise an alarm owing to the conditions stated. In general, the test depends on a regular recurring pattern in the data and thus will not be applicable to all data series. Also, it is possible that the test might identify when a time setting is very different (potentially erroneous) compared to earlier settings but does not represent a drift. Nevertheless, our example illustrates the usefulness of testing for drifts.

4.3. Similarity analysis of anomalies

The Jaccard coefficient determines whether there are similarities or convergences in anomalies throughout the network, which can be an effective indicator of where to look for errors. Figure 6 illustrates the Jaccard coefficient for selected tests from utilities

A and C. A Jaccard coefficient between 0 (marked green) and 1 (marked red) illustrates the degree of similarity in the occurrence of anomalies. In the figure, each flow and pressure meter was assigned an ID according to the total number of meter data sets and the Jaccard coefficient was evaluated for each set. Duplicate (M_1) and illegitimate datum flags (M_2) were evaluated as one category. These anomalies, as well as timestamp inconsistencies, occur in most cases simultaneously in meters ID 1–8 and ID 38–44 in utility A, evident from a Jaccard coefficient close to unity for these meters (Figure 6). These meters were not connected to the same supply areas nor did they use the same data transmission, but their data was collected in the same database, different from the remaining meters. This observation can be used by the utility to revise and improve its database setup. Furthermore, in the timestamp inconsistency test (test VI), a large number of flow meters and pressure meters have a high similarity in sets of two, visualized in the form of an almost straight line between flow and pressure sensor IDs in the Jaccard coefficient plot. As most pressure meters were installed to send data through a flow meter, any lack of data transmission would affect both flow and pressure observations at the same time. In the case of utility C, flagged values within test VI share a high similarity within a group of pressure meters and these meters are all located in the same supply area. Interestingly, measurements below the minute step appear simultaneously at the highlighted meters, indicating unknown meter or database settings or malfunctions. Jaccard coefficients for all remaining tests I–VI in utility A–C can be found in SI F.

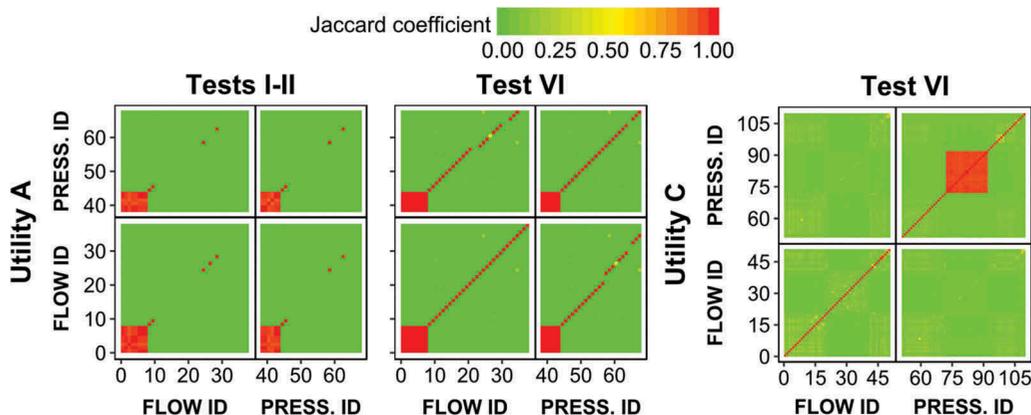


Figure 6. Jaccard coefficient computed for different tests from the anomaly testing covering flow and pressure meters. Results from the illegitimate and duplicate timestamp test were merged and illustrated in one column. Test types: I – duplicate timestamp, II – illegitimate timestamp and VI – timestamp inconsistency.

The occurrence of systematic errors (Figure 6) and large numbers of anomalies in the various tests (Figure 3) emphasises that there is a need for systematic data validation. It is our hope that the proposed methodology can assist a utility in systematically analysing errors and thereby markedly increase data quality in the future.

4.4. Future work

The next logical step includes a reconstruction process, generating qualified estimates of data being flagged as anomalies. The reconstruction method can utilize the MAID database in the training/calibration process to avoid being influenced negatively by anomalies of types 1 and 2. Additional data sets that can help to signal false misclassification of anomalies and improve the reconstruction process should be incorporated where possible. This includes data sets signalling the status of pumps, valves, external temperature or battery charge in selected devices. During reconstruction, the spatial and temporal redundancy between meters can be exploited (see, e.g. Cugueró-Escofet et al. (2016)); this, among other things, can make it possible to actively identify and flag type 3 anomalies. For example, due to a sudden drop in pressure, it is possible that a 'rate of change' flag was stored in the MAID. If the reconstructed value based on nearby sensors was close to the observed flagged value, the anomaly is unlikely to be due to a sensor error and can thus be categorized as type 3 anomaly.

5. Conclusion

This study revealed the need for the validation of meter data collected in water distribution networks. We have proposed a semi-automatic method to verify the collected data and highlight anomalies to deal with the increasing volume of collected data. The implementation of this method in a near real-time version, including a proper visualization of the error flags, will make it easy for the operators, on a day-to-day basis, to see whether sensors and data collection systems are working properly and to identify and correct errors when they arise. Having applied the method on 223 meter data sets from three Danish utilities, we found that, on average, at least 10% of the time that a meter collected data it was flagged as anomalous. For one utility, the collected flow data contained anomalous characteristics for an average of more than 35% of the time. Certain anomalies occur simultaneously throughout the network. Highlighting the occurrence of such similarities can help to improve future data collection and thus reduce the number of stored anomalies. The high rates of errors/anomalies could have been avoided if the proposed method had been implemented to automatically highlight meter errors in near real-time.

Acknowledgements

We thank all participating partners of the LEAKman project, Halsnæs Forsyning and Martin Brandt-Ewon from Nordvand A/S for providing data and answering questions related to data management practices.

Disclosure statement

No potential conflict of interest was reported by the authors.

Funding

This work was supported by the LEAKman project and partners under the Danish Eco-Innovation Program (MST-141-01277/NST-404-00378).

ORCID

Jonas Kjeld Kirstein  <http://orcid.org/0000-0002-0752-3061>
 Martin Rygaard  <http://orcid.org/0000-0001-8578-8842>
 Morten Borup  <http://orcid.org/0000-0002-6531-9464>

References

- Alegre, H., J. M. Baptista, C. E. Jr, F. Cubillo, P. Duarte, W. Hirner, W. Merkel, and R. Parena. 2006. *Performance Indicators for Water Supply Services*. London: IWA Publishing.
- Andrews, L., K. Gasner, R. Sturm, W. Jernigan, S. Cavanaugh, and G. Kunkel. 2017. *Utility Water Audit Validation: Principles and Programs*. Denver: Water Research Foundation. <http://www.waterrf.org/PublicReportLibrary/4639b.pdf>.
- Branisavljević, N., Z. Kapelan, and D. Prodanović. 2011. "Improved Real-Time Data Anomaly Detection Using Context Classification." *Journal of Hydroinformatics* 13 (3): 307–323. doi:10.2166/hydro.2011.042.
- Cugueró-Escofet, M. À., D. García, J. Quevedo, V. Puig, S. Espin, and J. Roquet. 2016. "A Methodology and A Software Tool for Sensor Data Validation/Reconstruction: Application to the Catalonia Regional Water Network." *Control Engineering Practice* 49 (Apr.): 159–172. doi:10.1016/j.conengprac.2015.11.005.
- DANVA (Danish Water and Wastewater Association). 2016. *Water in Figures 2016*. Skanderborg: DANVA. <https://www.danva.dk/publikationer/benchmarking-og-statistik/water-in-figures-pdf/water-in-figures-2016/>.
- Estévez, J., P. Gavilán, and J. V. Giráldez. 2011. "Guidelines on Validation Procedures for Meteorological Data from Automatic Weather Stations." *Journal of Hydrology* 402 (1–2): 144–154. doi:10.1016/j.jhydrol.2011.02.031.
- García, D., R. Creus, M. Minoves, X. Pardo, J. Quevedo, and V. Puig. 2017. "Data Analytics Methodology for Monitoring Quality Sensors and Events in the Barcelona Drinking Water Network." *Journal of Hydroinformatics* 19 (1): 123–137. doi:10.2166/hydro.2016.048.
- Helsel, D. R., and R. M. Hirsch. 2002. *Statistical Methods in Water Resources Techniques of Water Resources Investigations. Book 4, chapter A3*. Reston, VA: U.S. Geological Survey. <https://pubs.usgs.gov/twri/twri4a3/>.
- Jørgensen, H. K., S. Rosenørn, H. Madsen, and P. S. Mikkelsen. 1998. "Quality Control of Rain Data Used for Urban Runoff Systems." *Water Science and Technology* 37 (11): 113–120. doi:10.1016/S0273-1223(98)00323-0.
- Loureiro, D., C. Amado, A. Martins, D. Vitorino, A. Mamade, and S. T. Coelho. 2016. "Water Distribution Systems Flow Monitoring and Anomalous Event Detection: A Practical Approach." *Urban Water Journal* 13 (3): 242–252. doi:10.1080/1573062X.2014.988733.
- Machell, J., S. R. Mounce, B. Farley, and J. B. Boxall. 2014. "Online Data Processing for Proactive UK Water Distribution Network Operation." *Drinking Water Engineering and Science* 7 (1): 23–33. doi:10.5194/dwes-7-23-2014.
- Mounce, S. R., J. B. Boxall, and J. Machell. 2010. "Development and Verification of an Online Artificial Intelligence System for Detection of Bursts and Other Abnormal Flows." *Journal of Water Resources Planning and Management* 136 (3): 309–318. doi:10.1061/(ASCE)WR.1943-5452.0000030.
- Mourad, M., and J.-L. Bertrand-Krajewski. 2002. "A Method for Automatic Validation of Long Time Series of Data in Urban Hydrology." *Water Science and Technology* 45 (4–5): 263–270. doi:10.2166/wst.2002.0601.
- Puig, S., M. C. M. van Loosdrecht, J. Colprim, and S. C. F. Meijer. 2008. "Data Evaluation of Full-Scale Wastewater Treatment Plants by Mass Balance." *Water Research* 42 (18): Elsevier Ltd: 4645–4655. doi:10.1016/j.watres.2008.08.009.

- Quevedo, J., D. Garcia, V. Puig, J. Saludes, M. A. Cugueró, S. Espin, J. Roquet, and F. Valero. 2017. "Sensor Data Validation and Reconstruction." In *Real-Time Monitoring and Operational Control of Drinking-Water Systems*, edited by V. Puig, C. Ocampo-Martínez, R. Pérez, G. Cembrano, J. Quevedo, and T. Escobet, 175–193. Cham: Springer International Publishing.
- Quevedo, J., V. Puig, G. Cembrano, J. Blanch, J. Aguilar, D. Saporta, G. Benito, M. Hedo, and A. Molina. 2010. "Validation and Reconstruction of Flow Meter Data in the Barcelona Water Distribution Network." *Control Engineering Practice* 18 (6): 640–651. doi:10.1016/j.conengprac.2010.03.003.
- Rosen, C., J. Röttorp, and U. Jeppsson. 2003. "Multivariate on-Line Monitoring: Challenges and Solutions for Modern Wastewater Treatment Operation." *Water Science and Technology* 47 (2): 171–179. doi:10.2166/wst.2003.0113.
- Tan, P.-N., M. Steinbach, and V. Kumar. 2006. *Introduction to Data Mining*. Boston: Pearson Addison Wesley.
- WLCC (Water Loss Control Committee). 2014. "AWWA Water Audit Software (Version 5)." Microsoft Excel file. Denver: American Water Works Association.
- Wu, Y., and S. Liu. 2017. "A Review of Data-Driven Approaches for Burst Detection in Water Distribution Systems." *Urban Water Journal* 14 (9): 972–983. doi:10.1080/1573062X.2017.1279191.

A semi-automated approach to validation and error diagnostics of water network data

Supporting Information (SI)

Jonas Kjeld Kirstein¹, Klavs Høgh, Martin Rygaard and Morten Borup

The aim of the SI is to offer additional insight into selected methods and results in the paper ‘*A semi-automated approach to validation and error diagnostics of water network data*’.

We refer to the corresponding section of the main paper at the beginning of each section of the SI.

Contents: Supporting information

A	Anomaly sources	2
B	Anomaly testing - additional information	2
C	Data set description	3
C.1	Time series examples	3
C.2	Comparison of raw and validated data.....	4
D	Sensitivity of parameters applied in the anomaly testing	11
D.1	Tests III–VI: Parameter sensitivity	11
D.2	Test VII: Sensitivity of conditions.....	13
D.3	Test VII: Additional notes	14
E	Example of data validation and analysis for operational use	15
F	Extended Jaccard coefficient analysis	16
G	References	17

¹ E-mail: jkir@env.dtu.dk; Technical University of Denmark, Department of Environmental Engineering, Bygningstorvet, Building 115, 2800 Kgs. Lyngby, Denmark.

A. Anomaly sources

Section 2.1 of the main paper

Table A-1 lists major anomaly sources that contribute to errors and irregularities captured in the raw meter data sets. In addition, it is stated which of the applied tests in the anomaly testing were applied and designed to identify mainly a certain type of anomaly.

Table A-1. Potential anomaly sources that contribute to errors and irregularities and the tests applied to detect their main occurrence in raw meter data. Anomaly sources based on actual events in the network are not considered.

Anomaly source	Anomaly background	Main tests
 Operator	<p>Wrong time settings (incl. ‘drifting’ clocks) are an effect of erroneous synchronization or manual manipulation of the internal clock, incorrect adjustment for daylight saving time, etc. These settings can bias data applications such as DMA balances and reduce the overall value of smart online control of the water network.</p> <p>Many data applications depend on complete data streams. Errors within transmission and storage procedures, power outages, etc. can lead to illegitimate and missing data periods and highlight systematic problems of data acquisition.</p> <p>The miscalibration and loss in sensitivity of sensors, mechanical errors, and observations above or below cut-off values/deadbands decrease the overall reliability of the collected data streams. Among other things, these anomalies typically highlight meter malfunctions.</p>	I, VII
 Sensor		
 Transmission		
 Storage/application		II, VI
		III–V

B. Anomaly testing - additional information

Section 2.2 of the main paper

TEST IV) A description of the flagging procedure in the rate of change test (test IV) is summarized in the following:

- 1) Identify all $u \in t_n$ where $|\Delta x|/\Delta t \geq \theta$.
- 2) Split u into v sets of consecutive numbers.
- 3) Flag all first entries within v in \mathbf{M}_4 .
- 4) Return to step 1 and leave the flagged values out of the analysed data set. Stop, if no further u are found.

C. Data set description

Section 3 and 4 of the main paper

C.1 Time series examples

To provide insight on the analysed data series from the three utilities, we grouped measurements from selected meter data sets into a week's 168 hours. Then, to illustrate the variation in demand/flow and pressure, the percentile distribution was computed for each hour. The outcome for flow and pressure time series is shown in Figure C-1 and C-2, respectively.

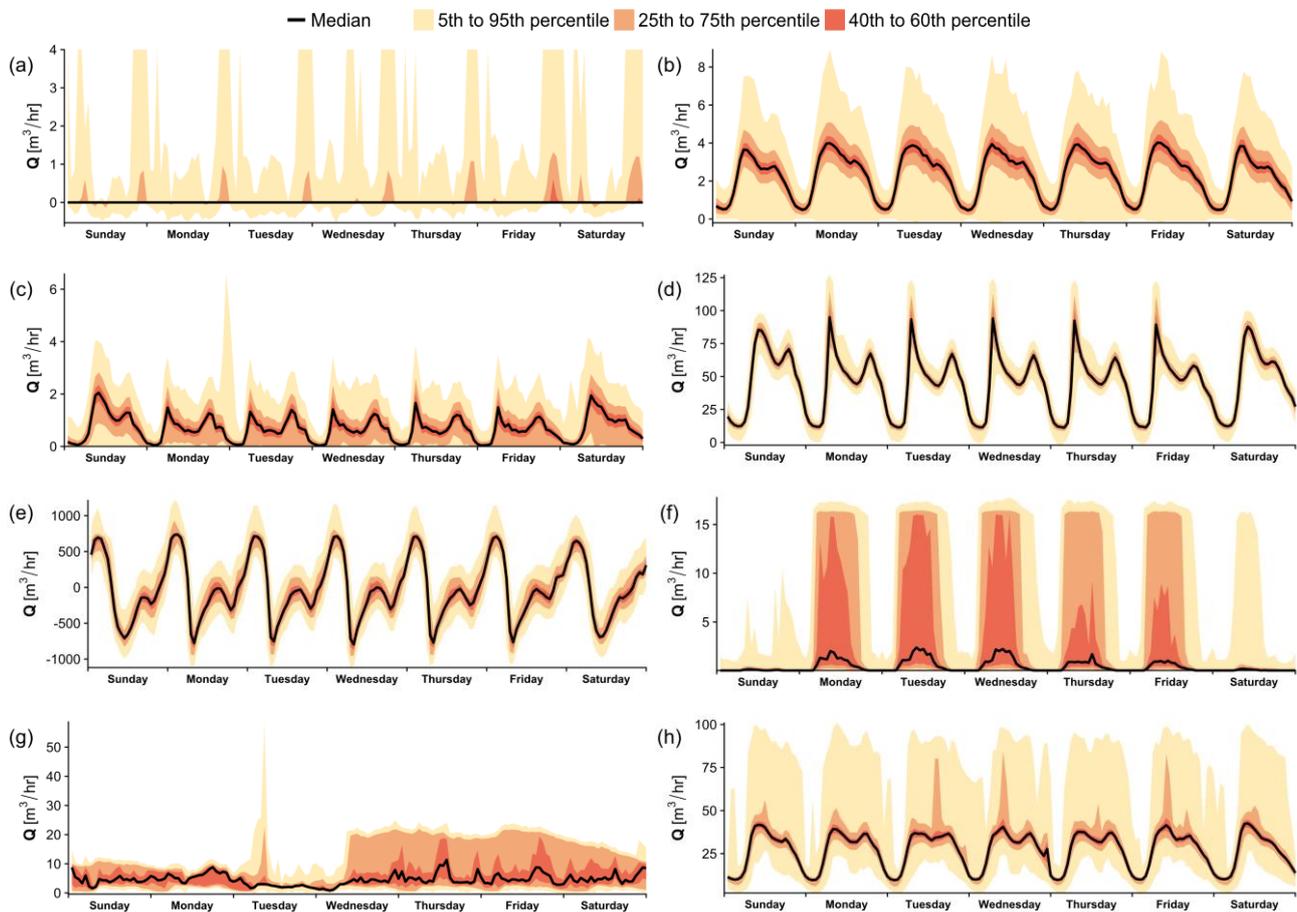


Figure C-1. Example of eight flow time series applied in the case study. The percentile distribution is based on having grouped all measurements into a week's 168 hours.

Figure C-1 shows that there is a great variety in the analysed flow time series. For example, the median flow and lower percentiles in Figure C-1a reflect that no flow is usually measured at all; however, there is a tendency towards higher flow rates during the hours before midnight, as shown by the 25th to 75th percentiles. Other time

series show a distinct daily pattern, as in Figure C-1c; this is the result of changing flow rates between summer and winter months. Figure C-1d shows an observation point with minor seasonal variations. Whereas these time series have a residential pattern, Figure C-1f shows high deviations in flow only during working days. As with flow, in selected pressure series a distinct daily pattern is visible, pressure being higher during the night than over the course of the day (e.g. Figure C-2d and C-2g). The pressure series in Figure C-2b and Figure C-2h indicate a higher share of anomalies, as a large proportion of measurements is around or even below 0 bar.

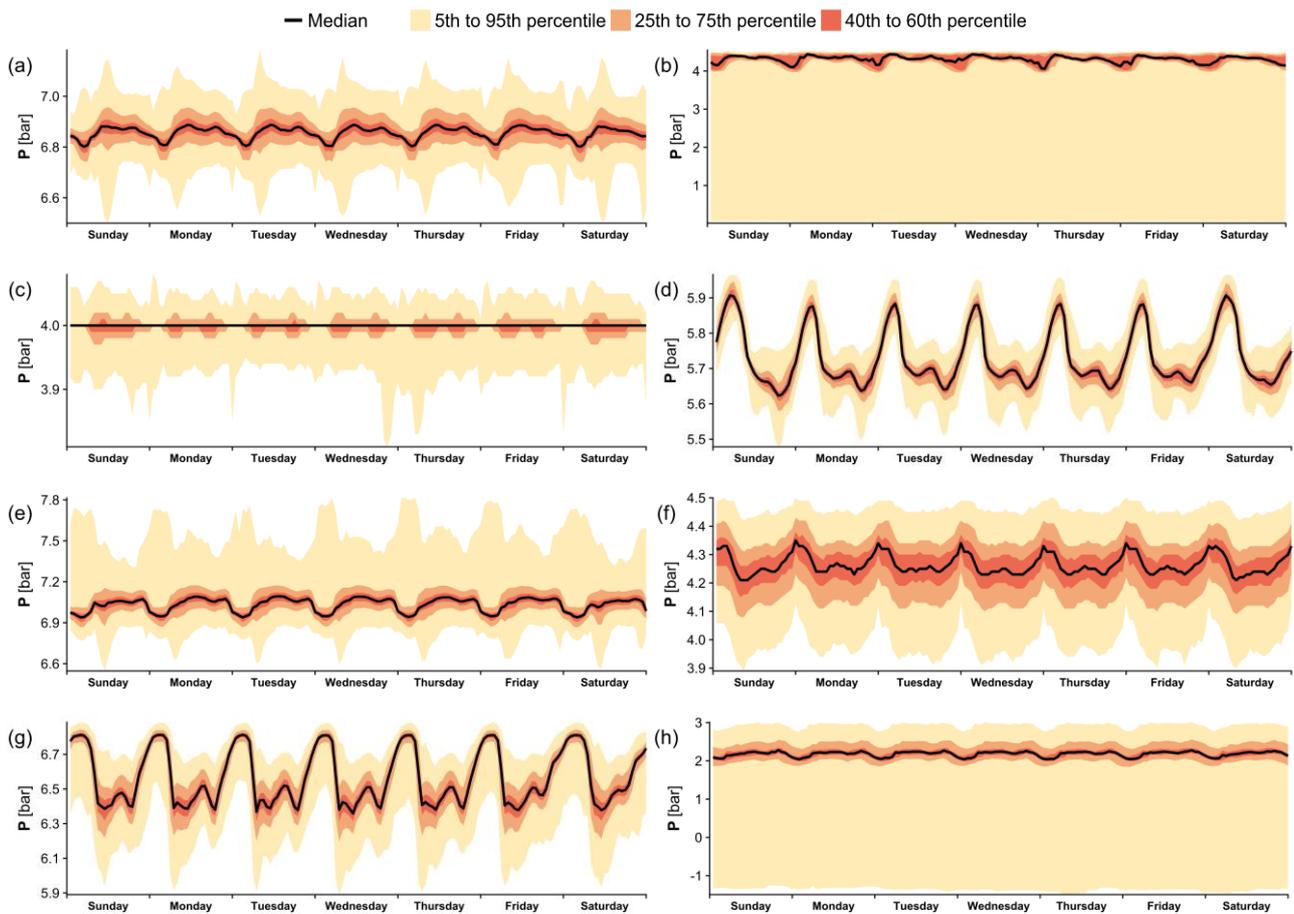


Figure C-2. Example of eight pressure time series applied in the case study. The percentile distribution is based on having grouped all measurements into a week's 168 hours.

C.2 Comparison of raw and validated data

All validated data in this section is based on the same parameter selection as summarized in Table 2. Table C-1 summarizes some key features of the applied data sets before and after data validation. It can be seen that the largest data sets were collected by utility C. Most of the data sets of utilities A and C are collected in 1-min timestamp intervals, whereas the highest share of data from utility B is stored at 5- or 15-min intervals.

Our data validation found the highest number of invalid data points (29.88%) in flow measurements of utility B.

Table C-1. Total number of data points in flow (Q) and pressure (P) meter data sets from utility A-C before and after data validation including timestamp interval distribution.

Utility	Type	Figure	Total data points	Timestamp interval distribution [%]					
				1 [min]	5 [min]	10 [min]	15 [min]	60 [min]	Other [min]
<i>Raw data</i>									
A	Q	Figure C-2(a)	18,150,889	91.15	0.37	6.11	0.46	0.93	0.99
	P		15,336,944	92.45	<0.0	6.06	0.41	0.22	0.85
B	Q	Figure C-3(a)	2,842,734	<0.0	53.85	<0.0	46.12	<0.0	0.03
	P		2,848,399	<0.0	53.78	<0.0	46.19	<0.0	0.03
C	Q	Figure C-4(a)	71,249,575	97.80	0.07	0.02	0.29	<0.0	1.82
	P		86,626,723	96.50	0.03	<0.0	0.02	<0.0	3.45
<i>Validated data (excluding anomalies flagged by test I-VI)</i>									
A	Q	Figure C-2(b)	(-8.09%) 16,682,587	91.40	0.30	6.60	0.49	0.97	0.23
	P		(-4.85%) 14,592,458	93.38	<0.0	5.96	0.34	0.19	0.12
B	Q	Figure C-3(b)	(-29.88%) 1,993,408	<0.0	67.46	0.02	30.93	0.03	1.57
	P		(-7.05%) 2,647,703	<0.0	54.14	0.06	45.46	<0.0	0.34
C	Q	Figure C-4(b)	(-4.52%) 68,030,580	99.21	0.04	0.01	0.01	<0.0	0.72
	P		(-8.75%) 79,043,261	97.83	0.11	0.04	0.02	<0.0	2.00

Figure C-3 exemplifies the possible distribution of raw data collected in selected flow and pressure meters for each of the three utilities before and after application of data validation (tests I–VI). Each of the six histograms represents a unique distribution pattern and it is difficult to identify clear similarities in the underlying distributions; however, some features are shared between selected meters, such as a lower bound of zero in the flow meters of utilities A and C. Moreover, all raw data histograms, except the pressure meter in utility C, indicate that a large proportion of type I and II anomalies have been stored in the data. For example, in the shown (raw data) pressure meter of utility B, more than 10,000 data points have a pressure value below 0 bar and a smaller proportion includes values around 30 bar. Observations within this range are highly unlikely, considering the utility’s pressure meters are located on distribution mains inside the city. After data validation, a large number of anomalies have been removed from the histograms. In terms of pressure data, all 0-bar data was deemed invalid in utility A. In utility B, a reduced amount of infeasible pressure measurements is still kept in the data, demanding a stricter choice of parameters. In the case of utility C, no values were flagged. Similarly, the number of ‘clear’ anomalies decreased in the three flow meters.

The histograms with validated data show the problems behind a global parameter selection of the different tests when applied to utilities with varying network setups. As in the case of pressure measurements, a simple test of min/max measurements above or below a certain threshold could have removed the remaining infeasible measurements. Thus, including specific system knowledge could increase the value of the validation method. In some cases, this knowledge is not available and a stricter choice of parameters could be a solution, although this could also increase the detection of false positives. In the case of the shown pressure meter in utility C, the selection of parameters has not negatively affected the anomaly detection visually, i.e. there was no flagging of values that appear correct.

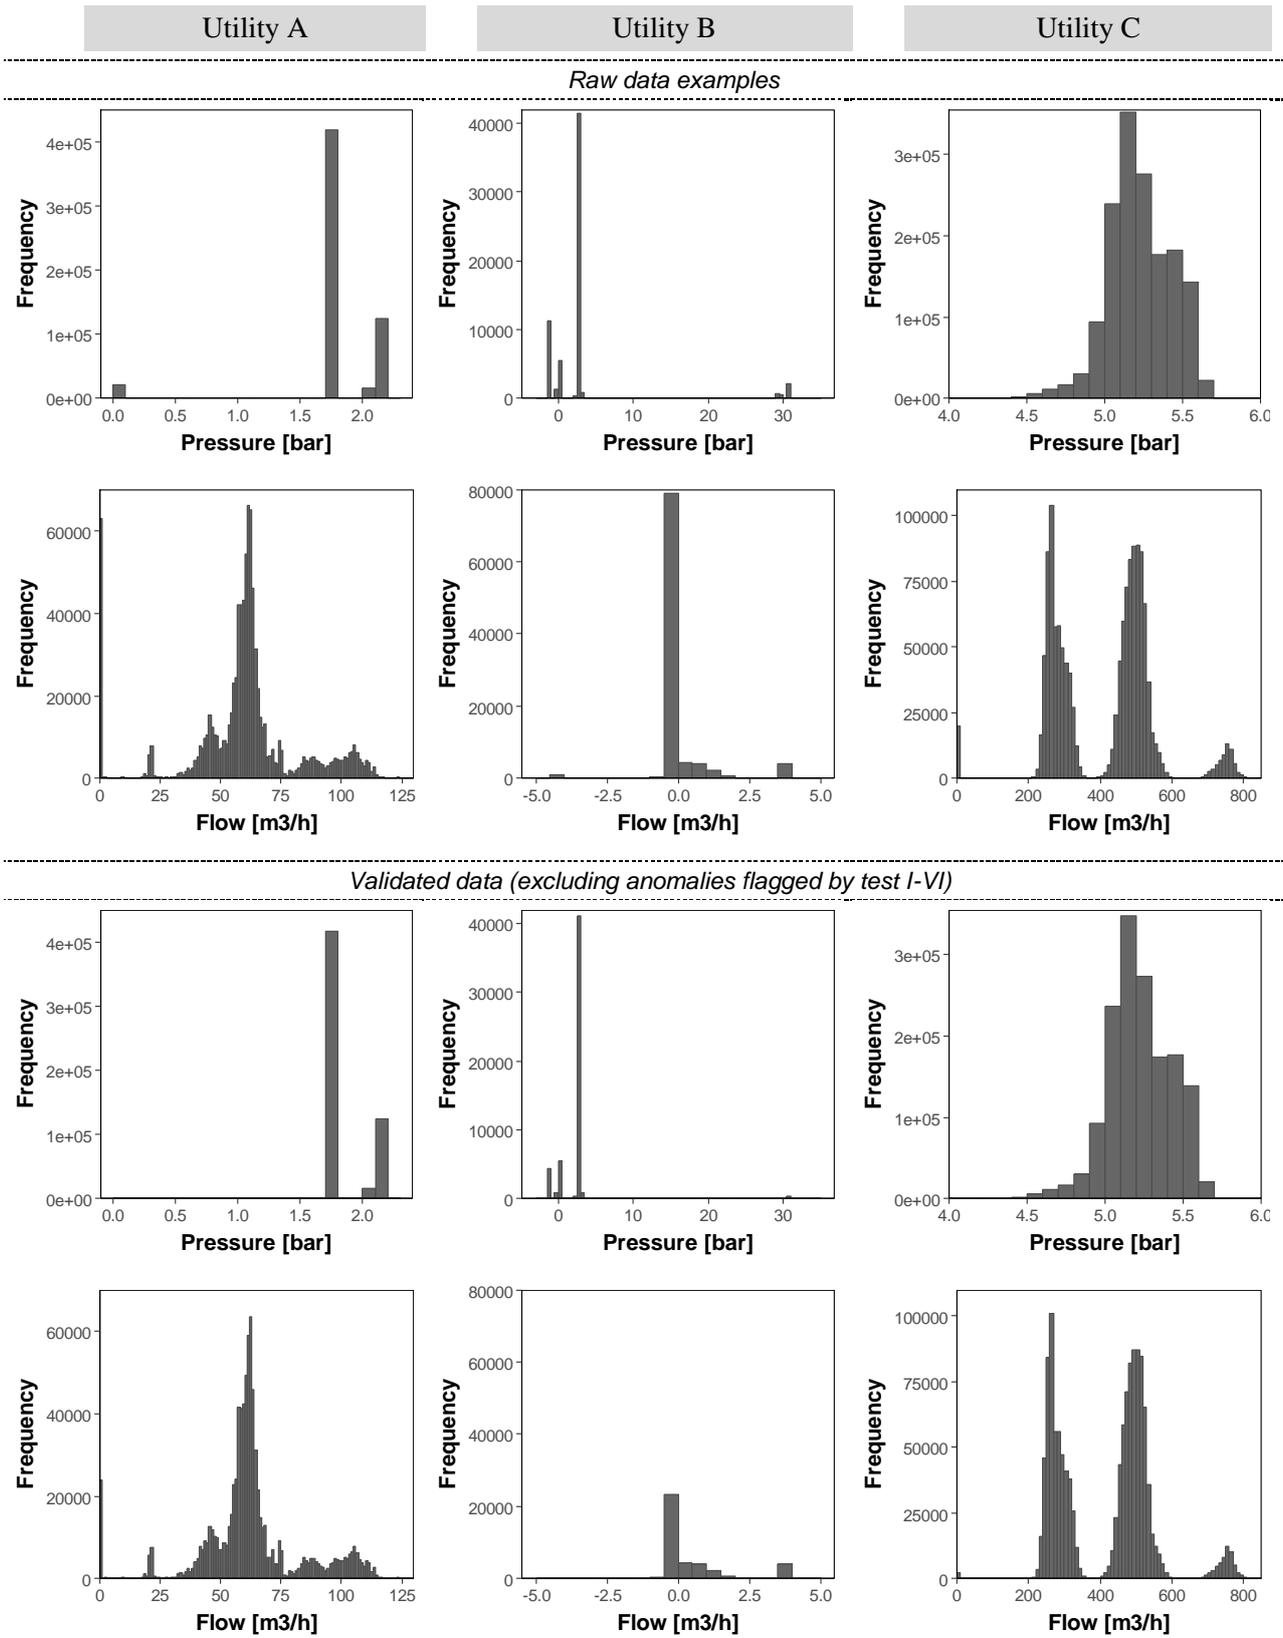


Figure C-3. Flow and pressure data histograms from six selected meters in utility A-C before and after data validation. Owing to the selected histogram bin size and margins, a minor share of data is not shown at maximum 0.06% (in this case raw pressure data in utility C).

Figure C-4 to Figure C-6 show the raw and validated data sets (tests I–VI) of utilities A–C distributed on daily, weekly and monthly time scales. Even though around 8% of the flow data (Table C-1) was flagged as anomalous in utility A (Figure C-4a) the data validation has no clear effect on the distribution of the utilities' flow data (Figure C-4b). Also, a daily pattern and some seasonality are visible in terms of higher summer consumption. The raw pressure data show a clear deviation between the 50th and 75th percentiles in November and December, and in the 23th hour compared to the remaining time steps. This can be explained by the fact that parts of the data collection system stopped working properly at the end of October 2016 and first continued functioning at the beginning of 2017. Moreover, in utility A, often at 23.00 hours each day, invalid (non-numeric) data was collected. Having applied the data validation tests, the 75th percentile decreased notably in February, in October and in the first hour of the day. This can partly be explained by around 25% of the pressure data being metered around 4 bar and the remaining 75% around 2 bar. Even slight deviations in the number of flagged values can move the 75th percentile. In February, for example, multiple timestamp duplicates (test I) were flagged in the high pressure meters, moving the 75th percentile towards a lower pressure. Among other things, a higher number of flatline anomalies (test V) had the same effect in the first hour.

The raw data of utility B is displayed in Figure C-5a. With regards to flow data, low flow hours (20.00 to 04.00 hours), certain days (Saturday and Sunday) and certain months (December–May) are particularly prone to having a median flow of around 0 m³/h. Having applied the data validation procedure (Figure C-5b), the median flow increased on all time scales. This is mainly because of the high percentage of flatline anomalies (test V) flagged in utility B that occur during low flow periods. The raw and validated flow data display a daily flow pattern as seen in utility A. However, the median has increased significantly in the validated data sets, indicating a good performance of the selected test parameters. Also, a very high pressure measured in April, May and June (Figure C-5a) was no longer visible after data validation (Figure C-5b). Possible explanations include anomalies flagged by tests III and IV. Finally, the data validation has not flagged a certain percentage of negative pressure values. This can be explained by the fact that the majority of pressure values in selected meter data sets, erroneously, are negative. As in utility A, there is no general difference in the observed pressure on different time scales.

In the case of utility C, there are only minor differences between the raw data set (Figure C-6a) and the validated data set (Figure C-6b). This can partly be explained by a more proactive approach in utility C to repair and solve issues within the data acquisition in a reasonable amount of time, reducing the overall number of anomalies in the raw data sets. As in utility B, bidirectional flows are also captured in the data sets, as can be seen in the hourly flow plot. The greatest difference between raw and validated data sets can be seen in the 2.5th pressure percentile on a monthly scale. In the validated data set, the lower hinge is more constant around 50 mWC. This stabilization was most likely caused by anomalies flagged by test V. A daily flow and pressure pattern is visible in the raw and validated data sets, but cannot be seen on a larger time scale.

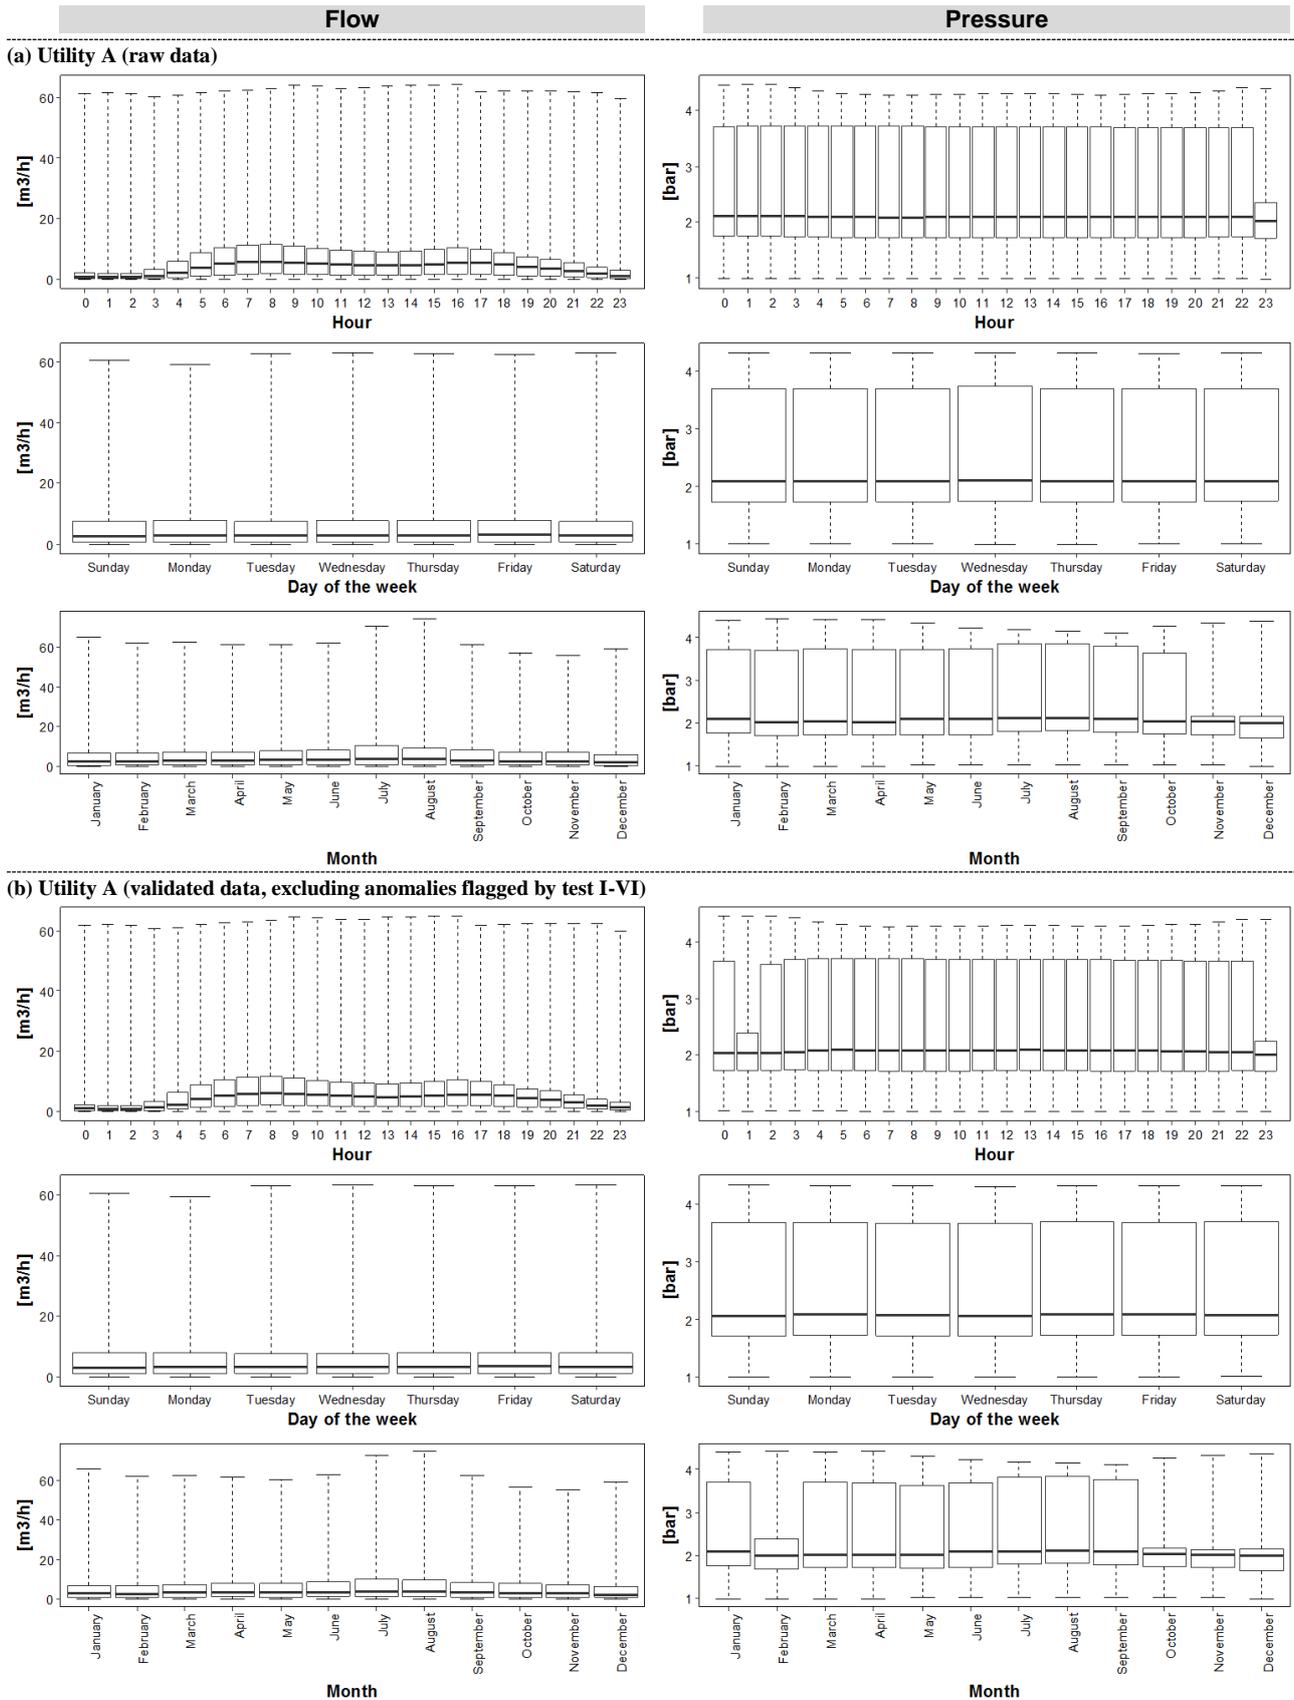


Figure C-4. Summary of raw (a) and validated (b) data sets in Utility A. Upper and lower hinge and whiskers represent 25th, 75th, 2.5th and 97.5th percentiles. The line across the box displays the median. Except non-numeric values, no outliers have been removed from the raw data analysis. Values below the 2.5th and above the 97.5th percentiles are not displayed.

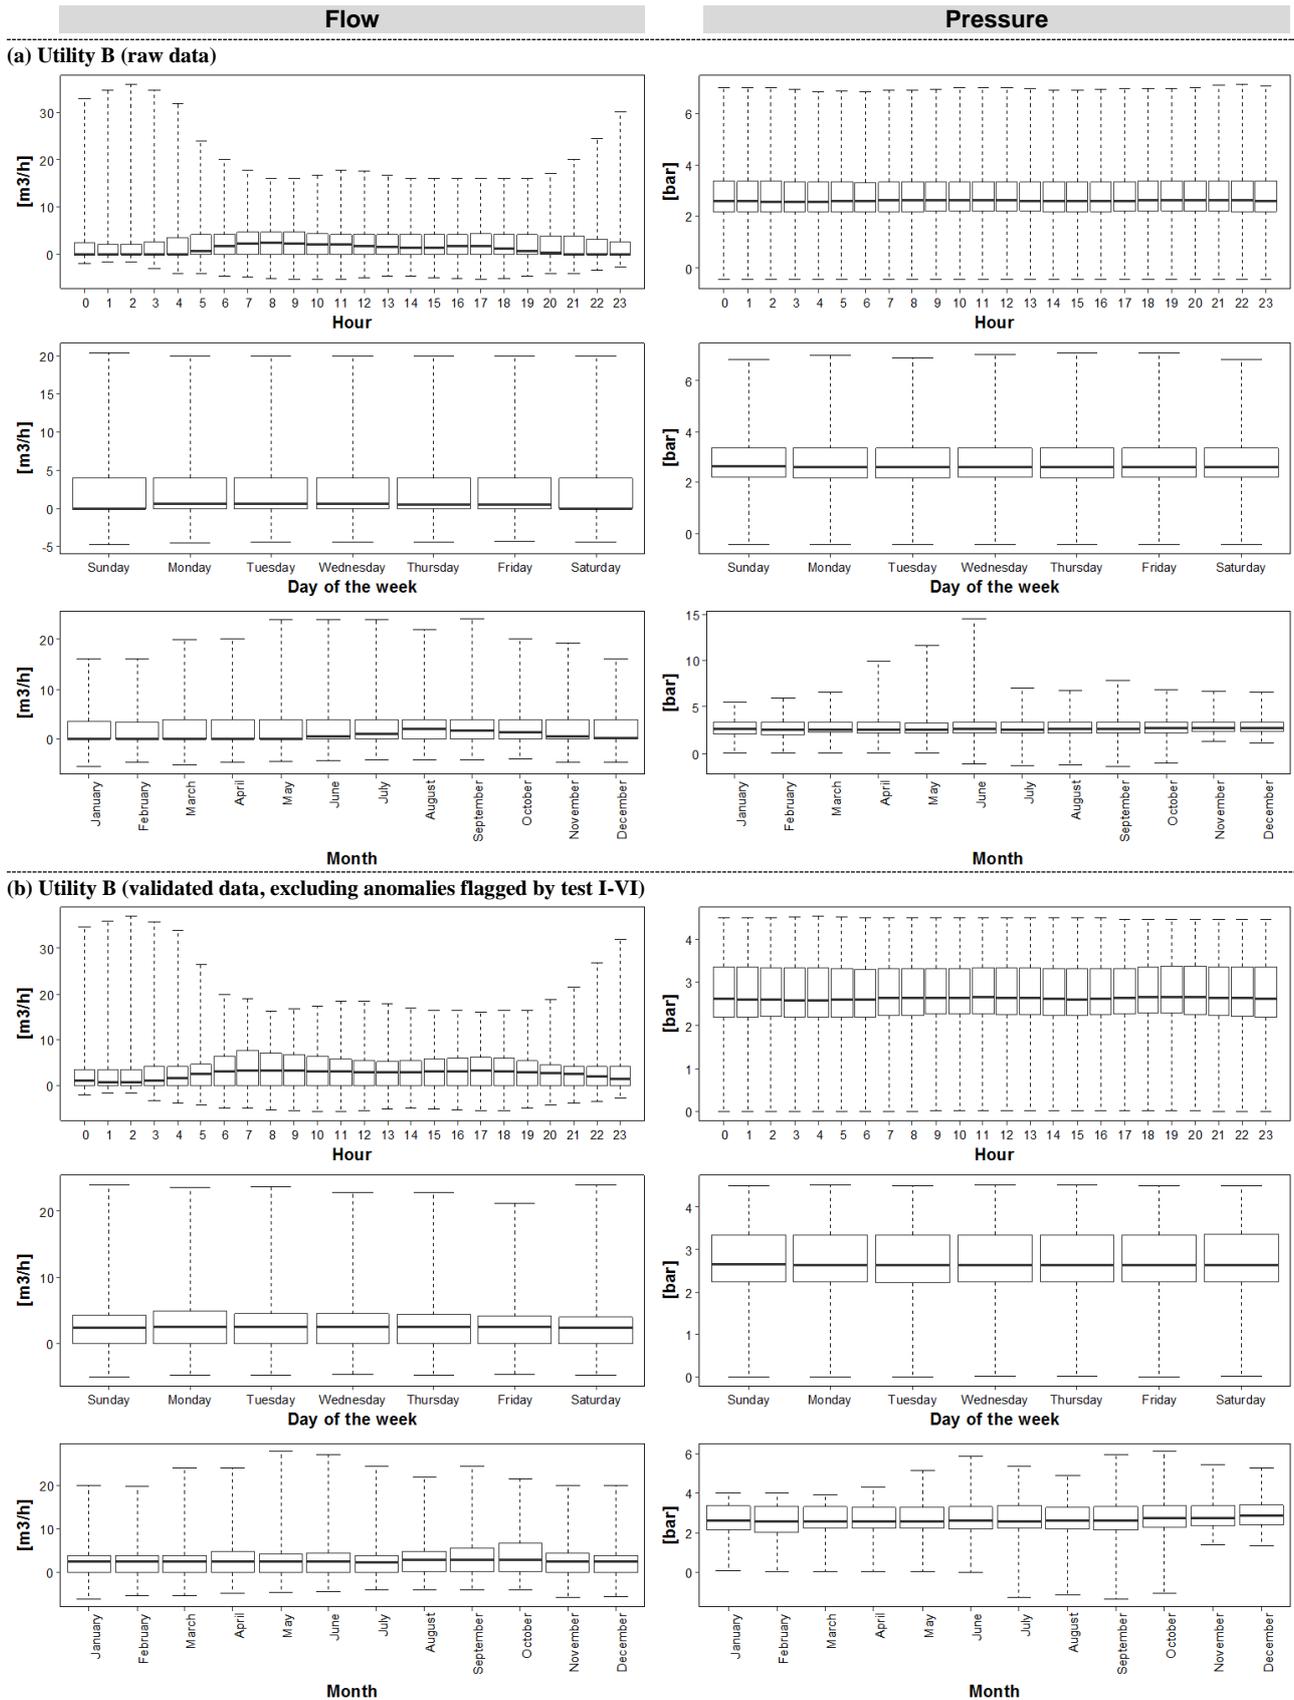


Figure C-5. Summary of raw (a) and validated (b) data sets in Utility B. Upper and lower hinge and whiskers represent the 25th, 75th, 2.5th and 97.5th percentiles. The line across the box displays the median. Except non-numeric values, no outliers have been removed from the raw data analysis. Values below the 2.5th and above the 97.5th percentiles are not displayed.

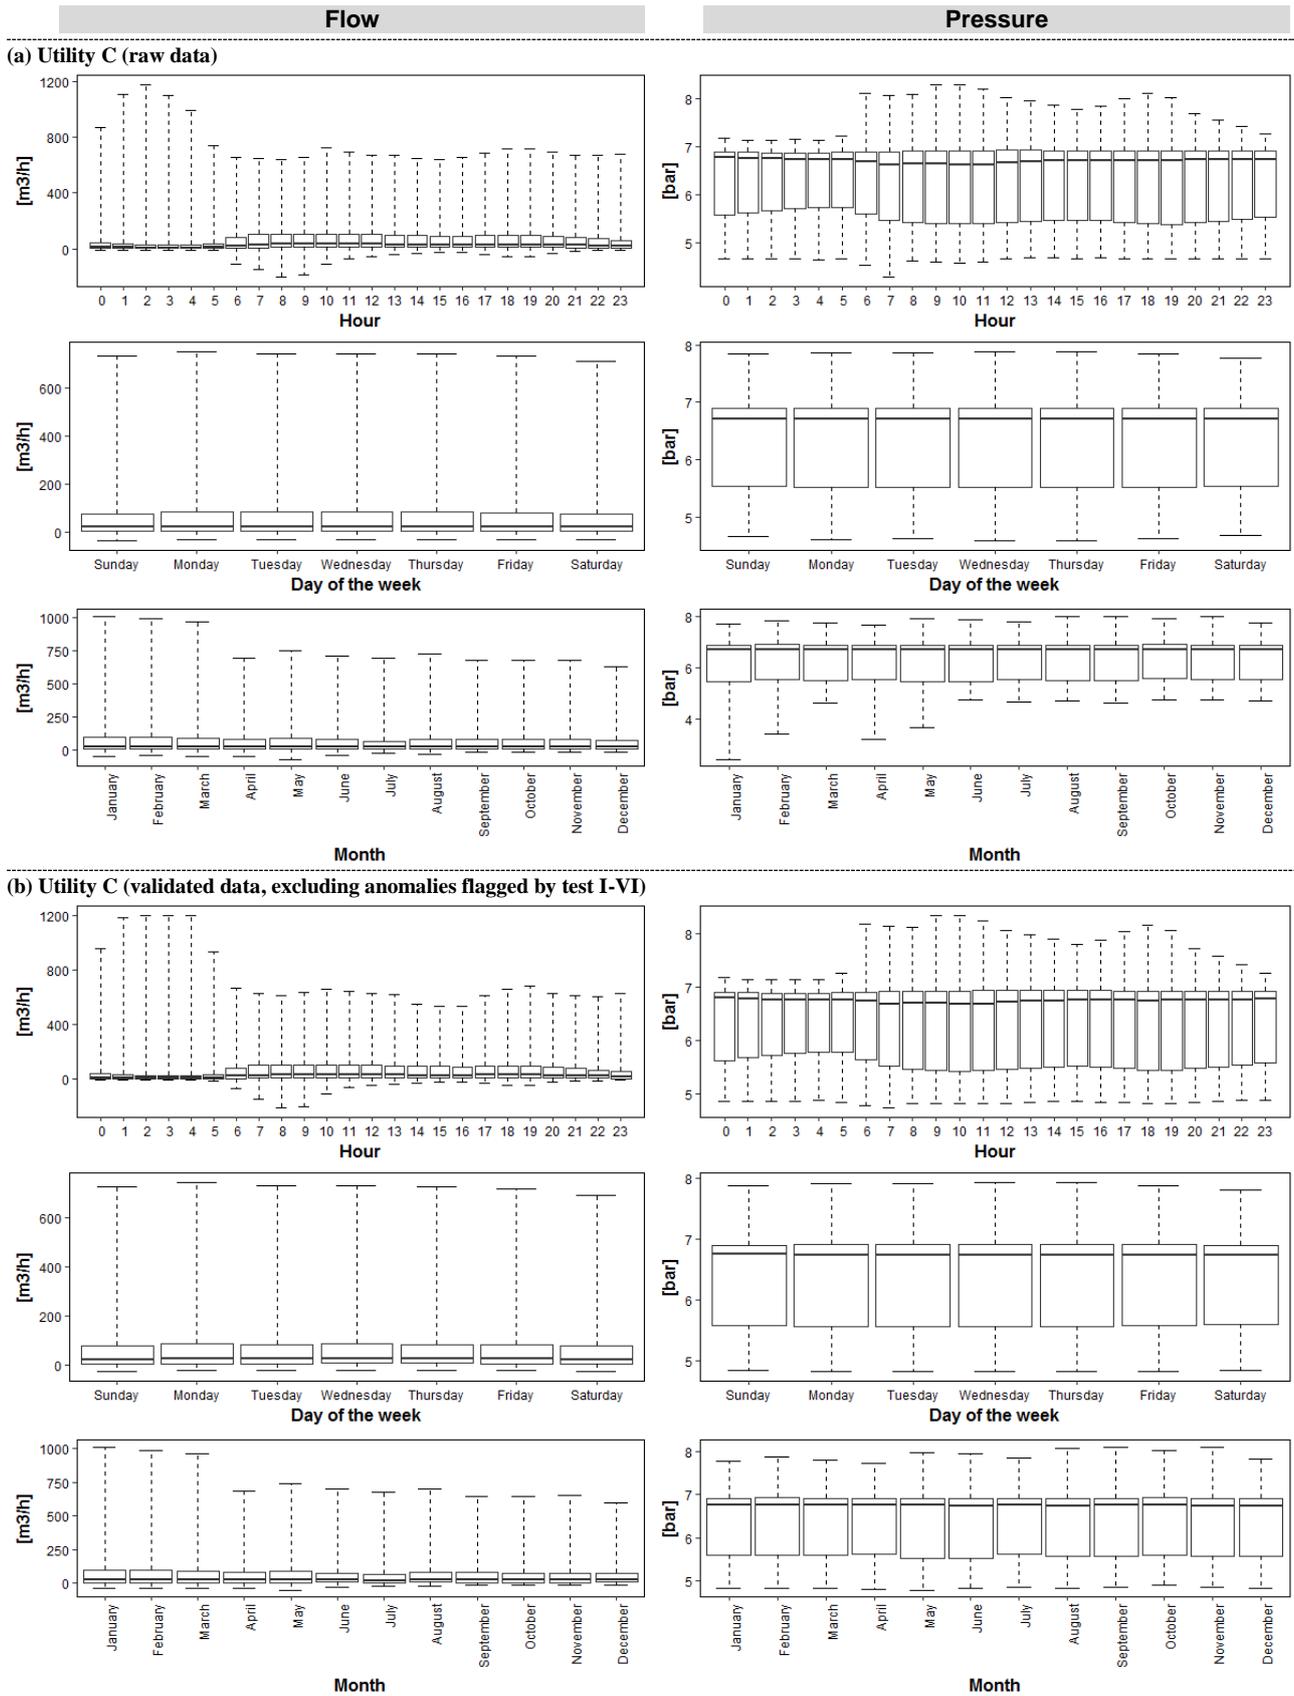


Figure C-6. Summary of raw (a) and validated (b) data sets in Utility C. Upper and lower hinge and whiskers represent the 25th, 75th, 2.5th and 97.5th percentiles. The line across the box displays the median. Except non-numeric values, no outliers have been removed from the raw data analysis. All values below the 2.5th and above the 97.5th percentiles are not displayed.

D. Sensitivity of parameters applied in the anomaly testing

Section 2.2, 4.1 and 4.2 of the main paper

It is almost inevitable that erroneously flagged data will be contained in the malfunction indicator database (MAID). In general, incorrect flagging of meter data occurs when test parameters are chosen that are too sensitive. Moreover, the range (test III) and rate of change test (test IV) depend on the historical distribution of the data. In such a case, it is difficult to avoid flagging of type 3 anomalies if the analysed data set covered periods where the system behaviour changed drastically. Thus, to reduce false alarm rates and flagging of type 3 anomalies, parameters had to be chosen carefully in selected tests of the anomaly testing framework. First, the results of different parameter combinations for tests III–VI are shown and discussed, following which examples are given of parameter combinations affecting the applied conditions of test VII.

D.1 Tests III–VI: Parameter sensitivity

Figure D-1 and Figure D-2 illustrate the sensitivity of the anomaly tests III–VI to the variation of parameters within the selected tests for all flow and pressure meters, respectively. Each column illustrates how the average percentage of flagged data points changed by a given parameter set. In the range test (test III) and change in rate test (test IV), higher percentile rates in combination with increasing the values of α/β and λ led as intended to a lower number of flagged values.

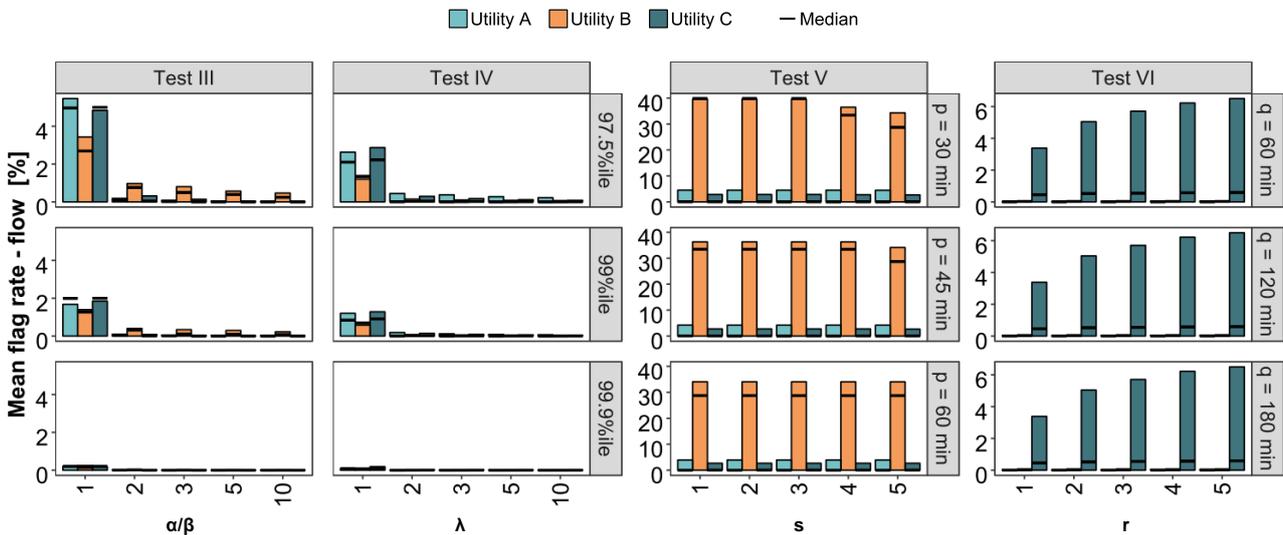


Figure D-1. Effect on error rates of all flow meters by varying test parameters on the range test (test III), change in rate test (test IV), flatline test (test V), and timestamp inconsistency test (test VI), based on the raw data from three utilities.

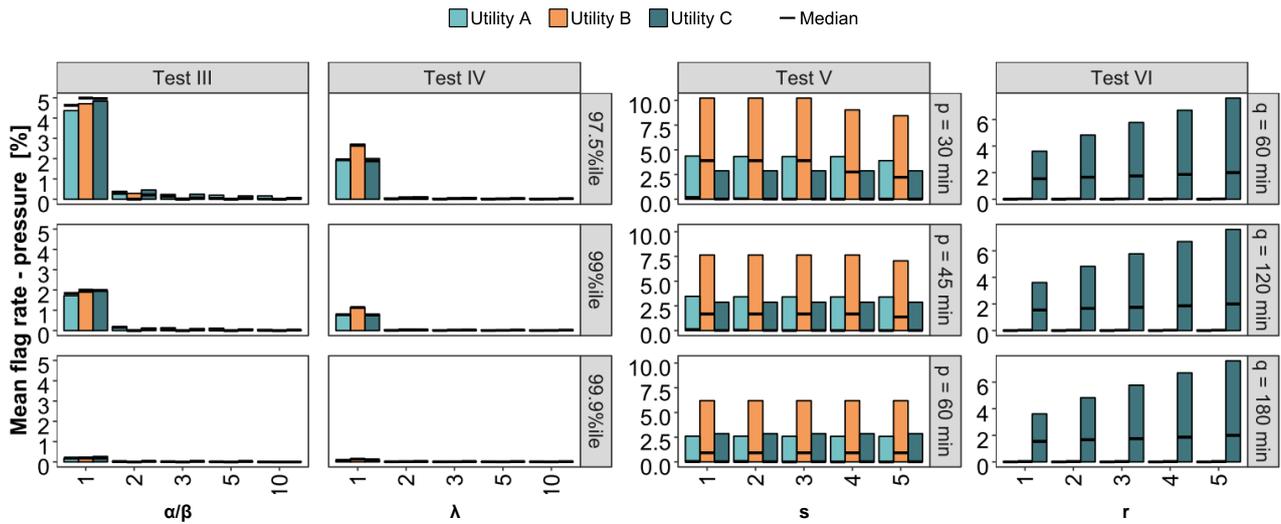


Figure D-2. Effect on error rates of all pressure meters by varying test parameters on the range test (test III), change in rate test (test IV), flatline test (test V), and timestamp inconsistency test (test VI), based on the raw data from three utilities.

In general, the figures illustrate that certain types of anomaly vary between the utilities. In the case of near real-time applications, however, it is important for the utility to identify a set of parameters that does not generate too many false alarms while correctly identifying anomalies. As the variation of parameters has a clear impact on the total number of data points, the parameters need to be fine-tuned. A utility should not only fine-tune the test parameters based on the measured parameters, but also change the test parameters independently for certain meter groups. For example, meters with varying sampling intervals or different objectives (e.g. DMA inlet or emergency pump monitoring) will probably need different optimal parameter settings. The flatline test column (test V) shows that varying the length of time and number of consecutive steps having an identical value has less influence on the mean flag rate, meaning that the flatline segments captured by these parameter sets in general occur over a longer period. The figure also illustrates that utility B has a notably higher frequency of flatline segments than the other two. It is likely that decreasing p would increase the error rate in utilities A and C notably, as a larger share of their data was measured at a higher sampling frequency. In this study, the originally stored number of significant digits in each meter set was included in the flatline test. Likely, more flatline segments would have been detected, when a reduced number of significant digits was included in the test. Moreover, the timestamp inconsistency test (test VI) shows that, for utility C only, varying parameter r has a visible influence on the mean rate of flagged values. This can be explained by the fact that the data collection system in utility C deletes equal measurements of less than 15 minutes duration.

D.2 Test VII: Sensitivity of conditions

In the timestamp drift test (test VII), the number of meters where a drift had been identified was initially determined by two conditions (Sec. 2.2). The sensitivity of test VII was assessed by varying the four parameters used to raise a flag in the conditions, namely $C_{weekly}(d)$, the number of reference weeks w included, $P_{weekly}(d)$ and the number of subsequent days d . Figure D-3 and Figure D-4 illustrate the effect of changing these parameters on the total number of meters with a drift for all flow and pressure measurements, respectively. As with the sensitivity of tests III–VI, an increase in the individual parameters led to a decrease in the number of identified anomalies, i.e. meters with drift. Also, this test needs to be fine-tuned by sensitivity; certain patterns in a utility might occur that would raise the false alarm rate. It seems that an increase in $P_{weekly}(d)$ had the largest impact on the overall number of meters where a drift was identified. A similar effect was seen for $C_{weekly}(d)$, because an increasing threshold also reduces the amount of data available in the test. The application of rather loose conditions shows that most drifts were identified in utilities A and C, which is linked to the total number of meters installed. Interestingly, increasing the parameters has only a small effect on the overall number of meters with a drift in utility B. A clear change in the number of meters is first seen when $P_{weekly}(d) > 3$ hours.

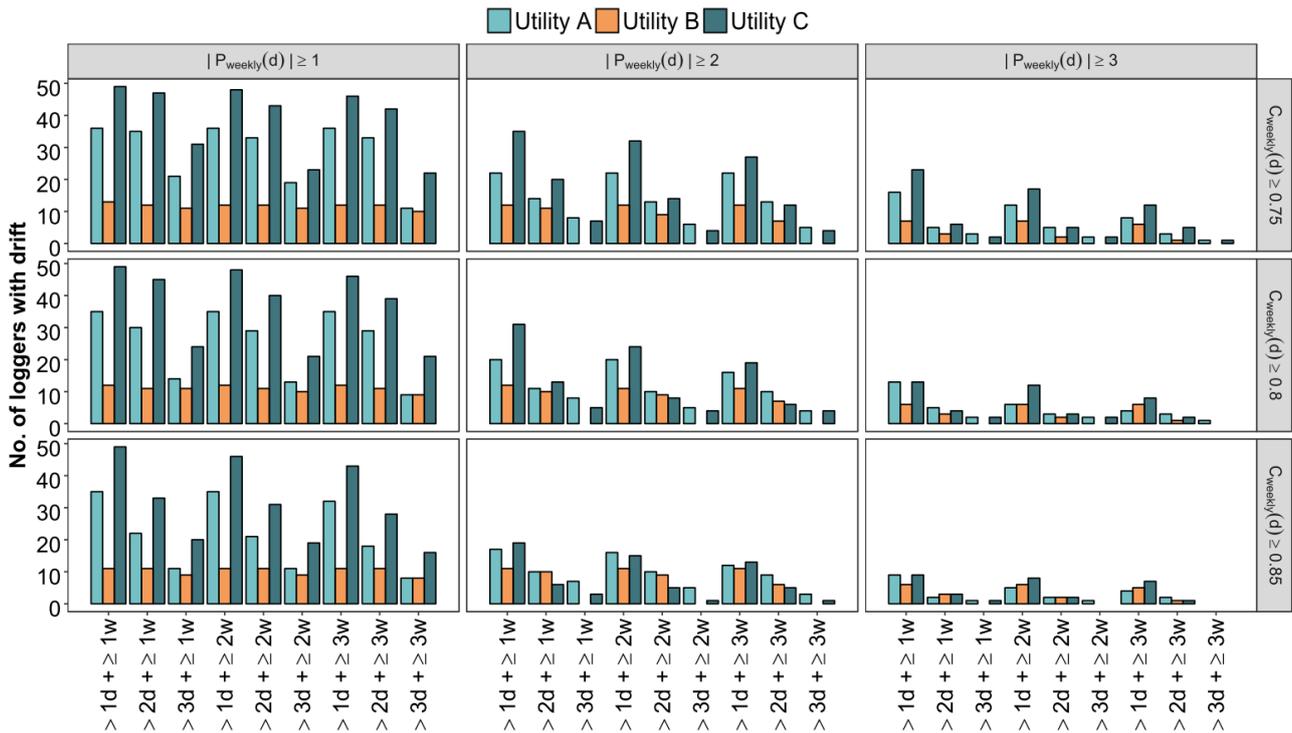


Figure D-3. Sensitivity analysis of the timestamp drift test (test VII) for all flow meters. Four parameters were varied: 1) the weekly correlation value [$C_{weekly}(d)$], before the data was accepted in the test; 2) the threshold hour before a test identifies a logger as drifting [$P_{weekly}(d)$]; 3) the number of consecutive days d in which the test has to identify a drift before a flag is raised ($1d-3d$); and 4) the number of reference weeks w ($1w-3w$) that have to agree on a drift before it is flagged.

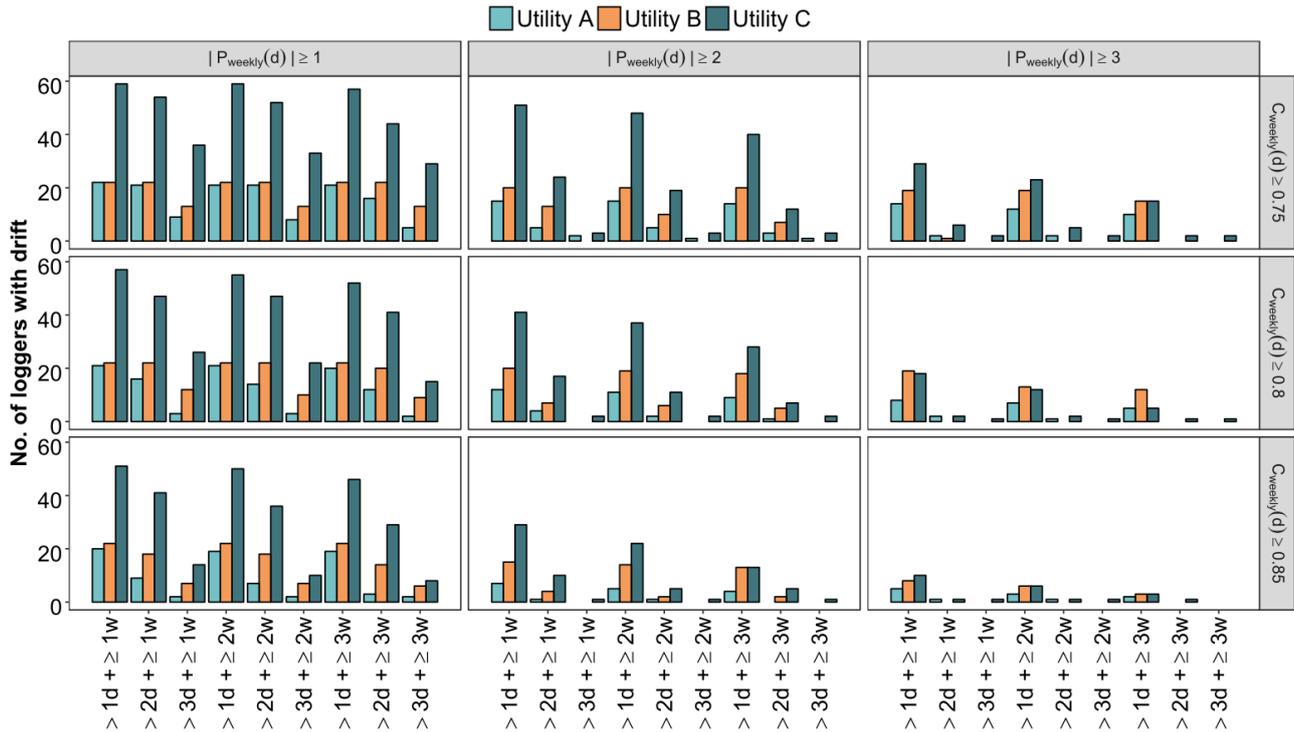


Figure D-4. Sensitivity analysis of the timestamp drift test (test VII) for all pressure meters. Four parameters were varied: 1) the weekly correlation value [$C_{weekly}(d)$], before the data was accepted in the test; 2) the threshold hour before a test identifies a logger as drifting [$P_{weekly}(d)$]; 3) the number of consecutive days d in which the test has to identify a drift before a flag is raised ($1d-3d$); and 4) the number of reference weeks w ($1w-3w$) that have to agree on a drift before it is flagged.

D.3 Test VII: Additional notes

The timestamp drift test is still under development and its reliability and feasibility need to be improved, as only a relatively low percentage of the collected data could be run by the test with the applied conditions. Another opportunity for improvement is where the test identifies drifts that, for example, reflect correct changes in the DMA set-up. Also, future applications should flag data at a lower scale than ± 2 hours. This would make it possible to detect incorrect daylight-saving time transitions on a daily operation. It would be interesting to analyse whether certain events, such as a sudden drop in pressure or fluctuation in flow, where the exact time is known, can be used to verify the occurrence of drifts identified by test VII. Furthermore, drifts could be verified by using the temporal and spatial redundancy between meters, e.g. by including a measure of similarity between similar time series. Since test VII is dependent on a regular recurring pattern in the data, it will never be applicable for all data series; the current implementation is therefore to be seen as a first draft that illustrates the usefulness of such a test, and we expect to be able to improve it.

E. Example of data validation and analysis for operational use

Section 4.1 of the main paper

Figure E-1 illustrates two examples of data validation and analysis for operational use. Figure E-1a provides an example of the ‘flag status’ based on all anomaly tests of 20 sensors in utility A over a period of four days. Flagged data was found from single (minute) data points for up to several days of consecutive flagged data (e.g. pressure meter P7). For example, short periods of invalidated data are visible at a higher rate in the flow meters Q1 and Q2. Q6, Q9 and P7 indicate ‘dirty data’, of which Q9 and P7 likely have no data available at all.

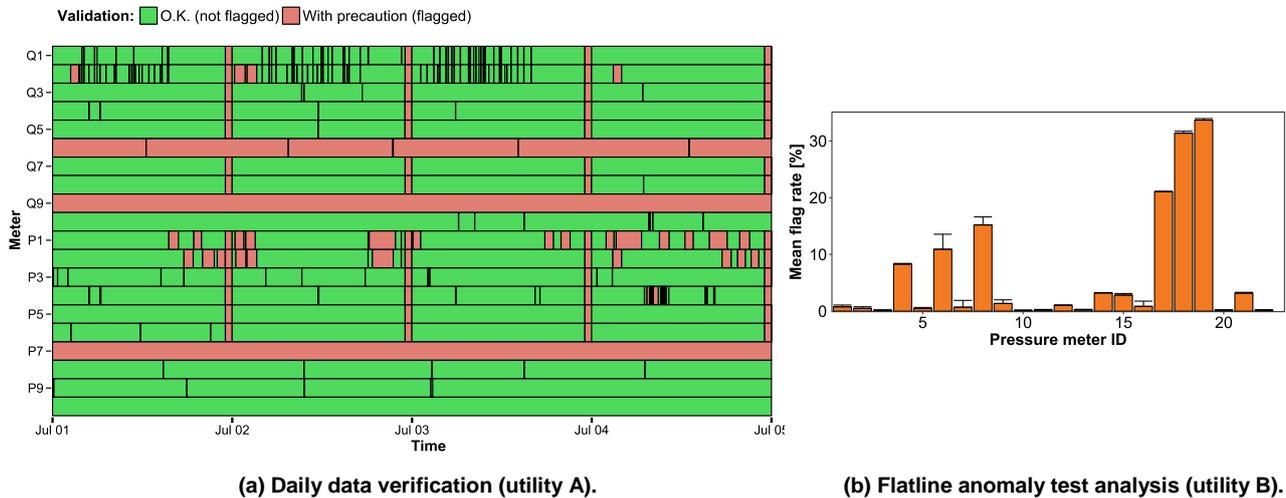


Figure E-1. Example of anomaly visualization for operational use. (a) Raw meter data validation from ten flow (Q) and pressure (P) meters in utility A in July 2015. (b) Mean flag rate based on all raw data points for the flatline test in pressure meters of utility B. Whiskers display the total flag rate based on all anomaly tests in the individual pressure meter.

It can be beneficial to focus on the individual error rates of the meters (Figure E-1b). For example, identifying differences in the individual error rates can be used to highlight meters not being sensitive enough to detect certain flow regimes at the installed location. In our case, the pressure meters P17–P19 tended to have a significant higher rate of flatlines than the remaining meters, which calls for inspection of these meters.

F. Extended Jaccard coefficient analysis

Section 4.3 of the main paper

Figure F-1 illustrates the Jaccard coefficients based on flags determined by tests I–VI for all utilities. We merged the anomalies of tests I and II into one group owing to the low number of flags in both categories. In the case that a meter contains flags in a certain anomaly process, the Jaccard coefficient by itself results in a high similarity of 1. If a meter data set contains no anomalies, no Jaccard coefficient, i.e. a value of zero, is computed. This results in the almost full red straight lines such as those seen in the range test of utilities A and B. In case of utility B, only a certain similarity pattern is visible for the flatline test of flow meters. This pattern is likely to be due to the high rate of flatline segments already shown in Figure 3 in the main paper. Also, Figure 3 shows that the number of missing data points (not based on time) was relatively low for utility B; however, the Jaccard coefficient of around 0.5 in Figure F-1, covering almost the entire timestamp inconsistency plot, indicates that large parts of the system lack data or went offline at the same time. High rate of change test and flatline test similarities are seen between six pressure meters (ID > 90, Figure F-1). All six pressure meters are positioned in an emergency pumping station. Thus, under normal circumstances all meters should log more or less constant values and change drastically if pumping starts. This might be reflected in terms of flags captured by the two tests. According to utility C, the meters are connected to the same programmable logical controller before being sent to the raw database, potentially being the source of the anomaly.

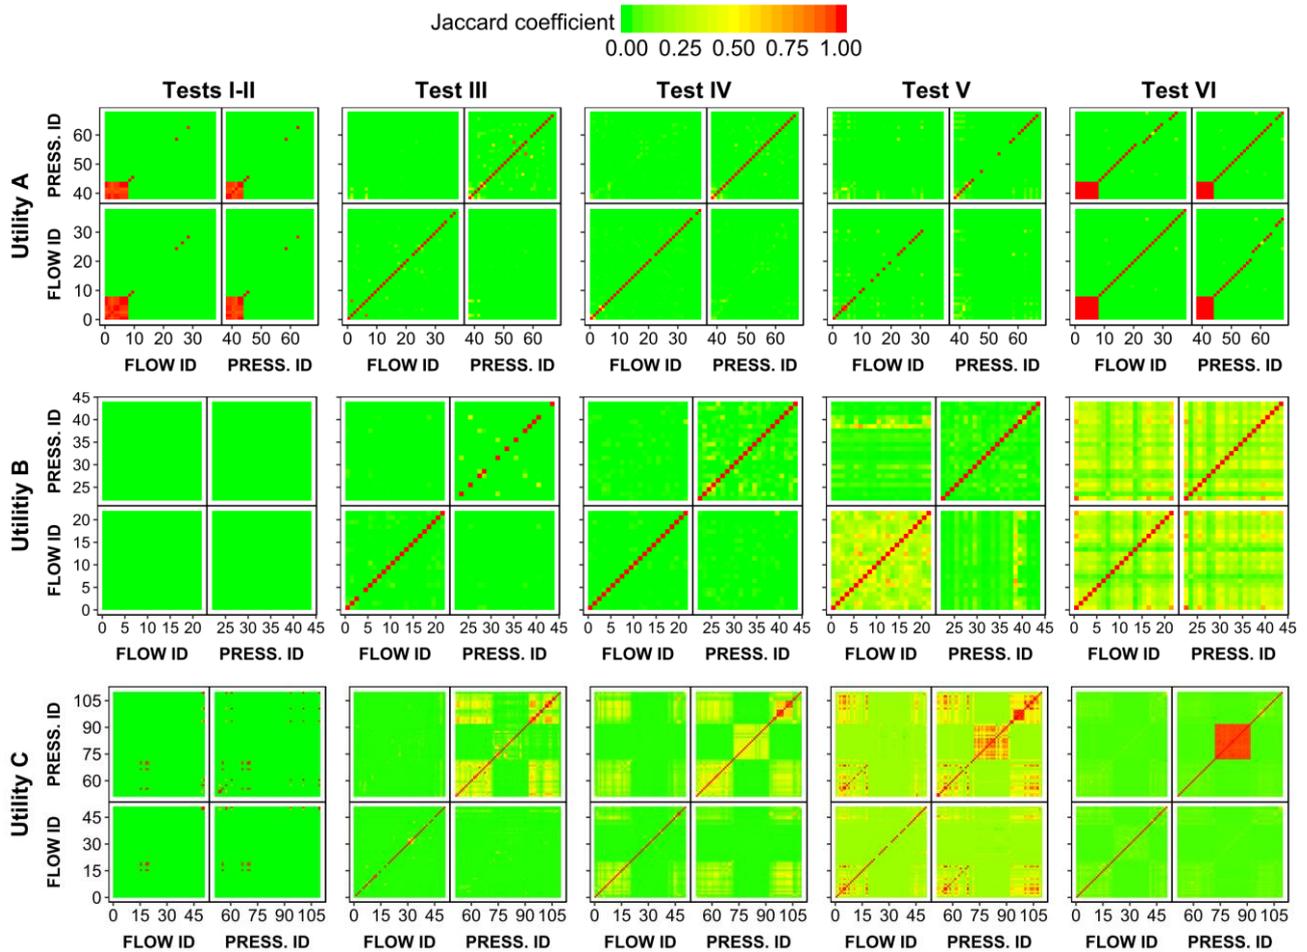


Figure F-1. The Jaccard coefficient computed for six different tests from the anomaly testing, based on the available flow and pressure meters. A Jaccard coefficient between zero and one describes no and very high similarity respectively for the occurrence of flags in the applied tests. Test types: I – duplicate timestamp, II – illegitimate timestamp, III – range, IV – rate of change, V – flatline and VI – timestamp inconsistency.

G. References

- DANVA (Danish Water and Wastewater Association). 2016. Water in Figures 2016. Skanderborg: DANVA.
<https://www.danva.dk/publikationer/benchmarking-og-statistik/water-in-figures-pdf/water-in-figures-2016/>



Valve status identification by temperature modelling in water distribution networks

J.K. Kirstein, S. Liu, K. Høgh, M. Borup and M. Rygaard

Valve status identification by temperature modelling in water distribution networks

Jonas Kjeld Kirstein^{1,*}, Shuming Liu², Klavs Høgh³, Morten Borup¹ and Martin Rygaard¹

*corresponding author

¹DTU Environment, Technical University of Denmark, Kgs. Lyngby, Denmark. jkir@env.dtu.dk; morb@env.dtu.dk; mryg@env.dtu.dk.

²School of Environment, Tsinghua University, Beijing, China. shumingliu@tsinghua.edu.cn.

³NIRAS A/S, Allerød, Denmark; kvh@niras.dk.

Abstract An up-to-date knowledge about the location and status of valves is of high priority for many utilities. In this study, we extend a water distribution network model with a temperature model in order to utilize temperature data for the model-based identification of valve status and location using a genetic algorithm (GA). A semi-synthetic and a real case study were run using temperature data from two district metering areas (DMAs) in Copenhagen, Denmark. First, a DMA model with real temperature data at the inlets and synthetic temperature data at the DMA's nodes was tested to identify five closed valves. A weighted initial population approach increased the GA's rate of successful convergence. In the second study, a transportation network model was used, including hydraulic and sample boundary temperature data, and the findings were checked in the field. The results of the GA search and subsequent testing revealed faulty valves and incorrect GIS information and has helped the utility to identify locations critical to their understanding of the system. The results show that the application of temperature data and models is useful for understanding the dynamics in water distribution networks.

Keywords: Hydraulic model; Temperature; Temperature modelling; Valve; Water Distribution Network.

1 Introduction

The unknown or doubtful status of valves in water distribution networks (WDN) is a common problem within the field of WDN operation and maintenance. Ultimately, this uncertainty slows down the operation and maintenance of the WDN and places in jeopardy the results obtained from hydraulic modelling. Aging infrastructures and the establishment of district metering areas (DMAs) lead to an increased focus on valves in the WDNs, as these control points are necessary when dealing with pipe repair and replacement (Walski et al., 2003; Sage, 2014; Wilson, 2011). During such processes, the ability to operate a utility's valves is important as it has an overall effect on the reliability of the WDN, such as the number of customers without service (Deb et al., 2012). For example, when a valve cannot be shut off, on average 1.7 other valves need to be closed, affecting larger service areas and resulting in higher collateral damage (Wilson, 2011). Moreover, valves with unknown status can pose a risk to the safety of consumers if a fire occurs (e.g. Delgado and Lansey, 2009) or lead to deterioration of the supplied water quality by unintentionally creating sections of the WDN with stagnant water. Incorrect settings of valves within hydraulic models have a crucial impact on the modelling results. For example, where real flows are not accurately represented by a hydraulic model because of incorrect valve settings, it is not possible for a water quality model to generate reliable

results (Savic et al., 2009). It is thus of major importance for utilities to have an up-to-date knowledge about the location and status of valves in the WDN.

Many utilities maintain valve books, i.e. databases with information about the status, operability, location and type of valves, and this information serves as a crucial input within asset management (Walski et al., 2003; Wilson, 2011). Based on this documentation, valve settings (i.e. open or closed) are often assumed to be known when calibrating WDN models (Sophocleous et al., 2017). However, with limited resources spent on the testing, mapping and maintenance of valves, it is questionable whether the available information on the WDN is accurate and up-to-date. Wu et al. (2012) state that it is common for utilities to be unaware about the correct status of a limited number of valves. Also, the experience from Danish utilities shows that the manual search for the correct valve status and location in the field is a time-consuming and often unsuccessful task. This can be due to many reasons. Examples include poor documentation, mechanical failure of valves, and uncertainty based on whether valves that have been closed during construction, testing or system maintenance have been reopened afterwards (Delgado and Lansey, 2009; Do et al., 2018).

In recent years, the number of data collection devices, such as smart meters installed at consumers' homes and flow and pressure meters in the WDN, has increased significantly. These advances open up several opportunities within the field of urban water research. One advance can be found within the area of WDN model calibration, where the lack of sufficient field data of high quality is often seen as one of the most restricting factors, eventually determining model accuracy (Savic et al., 2009; Walski et al., 2003). Likewise, the identification of closed valves and leaks depends on the quantity and quality of data available (Walski et al., 2014).

Typically, closed valves are identified, when there is a high level of discrepancy between the modelled and observed pressures throughout the network (Walski et al., 2014, 2003). It is often difficult to identify such deviations, as the head loss is low during normal flow conditions. Hydrant flow tests or opening of blow-off valves can increase the head loss through the system, but it is very time consuming (and thus costly) when applied in a WDN model building process (Ray et al., 2007). Delgado and Lansey (2009) identified a closed valve in a 'desktop study' by a hydraulic transient analysis, examining the reflection of simulated pressure transients. In the majority of valve status identification studies, genetic algorithms (GAs), combining results of pressure and flow simulations and real/synthetic measurements, have been applied (Do et al., 2018; Sophocleous et al., 2017; Walski et al., 2014, 2003; Wu et al., 2012). Moreover, most research has focused on fully closed or opened valves, whereas Do et al. (2018) and Wu et al. (2012) also looked for partially closed valves. The GA is used to explore the search space of possible solutions (i.e. combinations of closed and opened valves) by minimizing a selected fitness function.

Walski et al. (2014) list additional recommendations and requirements on how to apply GAs successfully for valve status and leak detection. These include applying the GA only to those parts of the WDN model where sufficient data are available and reducing the search space where possible, ultimately preventing the GA from guessing on selected valve statuses. Walski et al. (2014) recommend to look for a trend within the best solutions and use this information to guide the utility to a delineated area in the WDN when searching for leaks. We transfer this recommendation to the

valve status identification process, as combinations of valve statuses may lead to similar modelling results, eventually highlighting areas of concern in the WDN. The search space of the GA can be large when it is assumed that all valve locations and statuses in the WDN are doubtful. Wu et al. (2012) therefore assumed that only a small fixed number of valves have an incorrect setting, hereby reducing the search space. Do et al. (2018) and Sophocleous et al. (2017) simplified the search space by only looking at valves with a high sensitivity to changes.

Some studies have modelled the water temperature in the WDN. Examples include analysis of the water temperature in general (Blokker and Pieterse-Quirijns, 2013), and water quality in terms of disinfectant by-product generation (Eck et al., 2016) and as a heat source (Blokker et al., 2013; De Pasquale et al., 2017; Hubeck-Graudal et al., 2019). We hypothesize that temperature data can be a major benefit in addition to conventional flow and pressure data when the water temperature changes over time due to the surrounding soil temperature. Water heated up or cooled down throughout the WDN gives information about its path from the source to the place of measurement.

We present a method that combines hydraulic and temperature simulations and analyses differences with actual measurements from a real WDN to assess its potential to identify unknown and known valve statuses and locations. The method was evaluated using a semi-synthetic example as well as a real case study. The semi-synthetic example, with temperature measurements assumed known at all nodes in a DMA, was used to evaluate the pros and cons of various GA set-ups. The real case study applied the method to a transportation network with known input and output flows and pressures, as well as temperature at selected locations. The status and location of selected valves were tested in the field and checked against the obtained modelling results.

2 Methodology

The proposed method was divided into the following four steps (Figure 1):

- 1) Compilation of hydraulic and temperature input data, including information on the location and status of valves and the pipe network.
- 2) Calibration, based on data from 1) of a combined hydraulic and temperature model of the WDN.
- 3) Evaluation of discrepancies between real and simulated results.
- 4) Updating utility information about their WDN.

Hydraulic and water temperature models were simulated in EPANET (Rossman, 2000) using the multi-species extension EPANET-MSX (Shang et al., 2007). Preliminary network operations and input files were modified through the *EPANET-Matlab Toolkit* (Eliades et al., 2016). We performed most simulations on a high-performance computing cluster with a maximum of 120 MSX simulations running in parallel to reduce the computational time needed for the GA to converge. We used the MATLAB GA function based on mixed integer optimization (Deb, 2000; Deep et al., 2009).

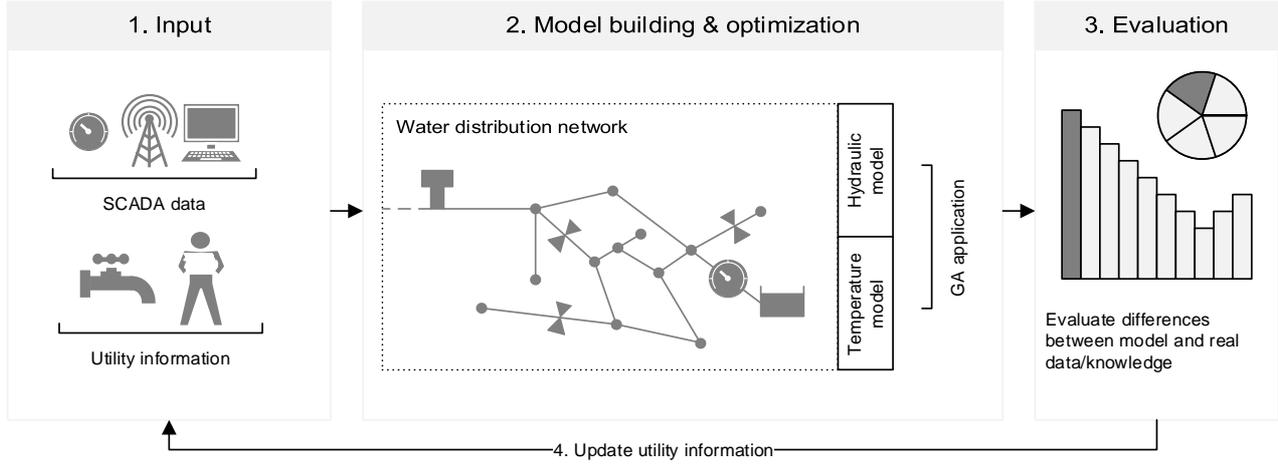


Figure 1. Overview of the four-step valve status identification process. GA = genetic algorithm; SCADA = Supervisory control and data acquisition.

2.1 Temperature modelling

In WDN temperature models, the change in the water temperature over time can be described by the following expression (Blokker and Pieterse-Quirijns, 2013):

$$\frac{dT_w}{dt} = \frac{2K_p}{c_p \rho r_i} (T_w - T_s), \quad (1)$$

where T_w is the water temperature [°C] and T_s the undisturbed soil temperature [°C], i.e. the location closest to pipe burying depth where the soil temperature is unaffected by the heat transfer between soil and pipe. c_p is the heat capacity of water [J/kg/K], ρ is the density of water [kg/m³] and r_i is the inner radius of the analysed pipe [m]. K_p [W/m²/K] is the heat transfer coefficient that defines the overall resistance in heat conduction between water and the surrounding soil. K_p is subject to the major differences in the applied temperature models in the literature (Blokker and Pieterse-Quirijns, 2013; De Pasquale et al., 2017; Hubeck-Graudal et al., 2019). Here we used the Hubeck-Graudal model (Hubeck-Graudal et al., 2019), where the thermal resistance of the surrounding soil (R_s) in the shape factor model (De Pasquale et al., 2017) was modified to account for a fictitious soil layer H :

$$K_p = (R_p + R_w + R_s)^{-1} = \left(\frac{\ln \left(\frac{r_o}{r_i} \right) r_i}{\lambda_{pipe}} + \frac{1}{a} + \frac{\ln \left(\frac{2H}{r_o} \right) * r_i}{\lambda_{soil}} \right)^{-1} \quad (2)$$

R_p and R_w account for the thermal resistance between the pipe wall and flowing water, r_o [m] is the outer radius of the analysed pipe, and λ_{pipe} and λ_{soil} are the thermal conductivity [W/m/K] of the pipe material and surrounding soil, respectively. Moreover, H [m] is defined as an additional soil layer, correcting for the convective boundary resistance at the soil surface (Kvisgaard and Hadvis, 1980):

$$H = z + \phi \lambda_{soil} \quad (3)$$

With ϕ describing a correction factor equal 0.07 [K·m²/W]. The soil resistance term R_s is based on the assumption that the height z [m] of the soil layer above the buried pipe (measured from the centre of the pipe) is larger than the inner diameter. a [W/m²/K] is the convective heat transfer coefficient

between the pipe wall and water. This convective heat transfer coefficient can be determined as (Cengel, 2003):

$$a = \frac{\lambda_{water} Nu}{2r_i}, \quad (4)$$

where λ_{water} [W/m/K] is the thermal conductivity of the water and the Nusselt number Nu [-]. Nu depends on the flow conditions in the WDN and is approximated by the Colburn equation (Cengel, 2003):

$$Nu = 0.023 Re^{0.8} Pr^{1/3}, \quad \left(\begin{array}{l} 0.7 \leq Pr \leq 160 \\ Re > 10,000 \end{array} \right), \quad (5)$$

where Re and Pr are the Reynolds Number and Prandtl Number, respectively. Additional information on the computation of Re , Pr and temperature models, i.e. undisturbed soil (Blokker and Pieterse-Quirijns, 2013), shape factor, De Pasquale (De Pasquale et al., 2017) and Hubeck-Graudal model can be found in the supporting information (SI) A. There, a parametric analysis, inspired by the work of De Pasquale et al. (2017), is used to discuss differences in choice of temperature model. In SI B, the effect of different selected parameters on the Hubeck-Graudal model is discussed.

2.2 Reducing the valve search space

Our study considered fully opened or closed isolation valves that control the pipe flow. In the WDN model, the number, location and status of all valves were assumed to be unknown, i.e. a valve was assigned to each pipe segment resulting in a total of n valves. Next, we identified k valves where closing the pipe segment had no impact on the WDN model (see next paragraph) or would disconnect users or parts of the network from a water source. This filtering reduced the total ‘valve search space’ to $n-k$ valves. We divided this valve search space reduction process into three categories: ‘dead ends’, ‘source disconnection’ and ‘valve clustering’ (Figure 2). It must be noted that, when a utility has a good knowledge of its valves, n should simply equal the known number of valve locations in the network.

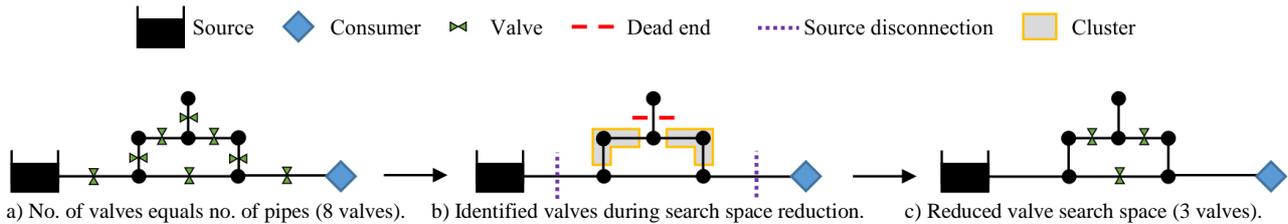


Figure 2. Example of reducing the search space for valve status identification in a water distribution network. The shown clusters represent an example of how a given proximity threshold leads to clustered valve segments.

2.2.1 Dead ends

Valves leading to ‘dead end’ segments of a WDN with no water flow have no impact on the surveyed WDN model parameters during calibration. Dead ends were removed from the valve search space (Figure 2) based on the topology of the WDN model and defined demands at consumers/sources.

2.2.2 Source disconnection

Valves that would potentially disconnect consumers from a water source were also filtered from the search space (Figure 2), since this would have been reported by the consumers.

2.2.3 Valve clustering and sensitivity

We clustered adjacent pipes of the WDN model, without any nodal demand in between, where the closing of each individual valve was assumed to have a little impact on the system in terms of pressure and temperature. This is based on a selected maximum proximity as a threshold (Figure 2). A pseudocode based on hierarchical clustering and excluding ‘dead end’ as well as ‘source disconnection’ segments can be found in SI C.

2.3 GA set-up

In the following only the reduced valve search space was taken into account. GAs have been widely used in water distribution network optimization (e.g. Savic and Walters (1997)). The GA used an initial population of size p , where each individual accounted for a WDN model set-up that stored information on whether a valve was closed (0) or open (1). New generations were constructed based on the individuals’ fitness values and the selected GA parameters *elite count* (EC), *cross-over fraction* (CF) and *mutation rate* (MR). The GA setup used to calibrate a WDN model was sectioned into the six steps I–VI (Figure 3).

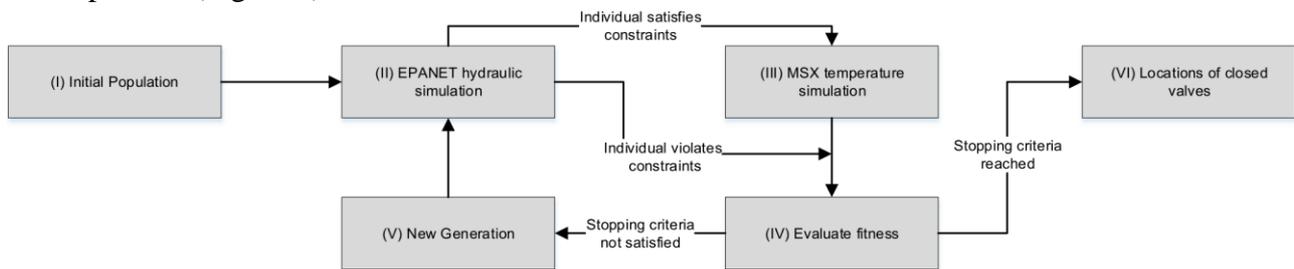


Figure 3. Concept for the genetic algorithm application to identify closed valves in a water distribution network model.

2.3.1 Constraints and initial population

The definition of an initial GA individual was subject to two constraints:

1. No consumers allowed disconnected from a source.
2. No negative pressure allowed in the hydraulic model simulation.

The initial population (step I, Figure 3) comprised only individuals that complied with these constraints, i.e. only ‘feasible’ solutions. The same two constraints were applied between steps II and III (Figure 3). If individuals were deemed infeasible during the GA run, they were assigned a penalty value greater than the maximum feasible individual. Two different initial population set-ups were generated. In the first set-up, each valve in a population was randomly drawn to be closed on an average of x -percent of the time. For example, with x being equal to 5% and an initial $p = 500$, each valve from the search space was closed on average 25 times in the initial population’s 500 WDN setups. We refer to this initial population scenario as *random*. In the second set-up, the initial

population was based on the similarity measures (SM) and an associated probability P , and this scenario is referred to as *weighted*. The SM was computed as follows:

$$SM(i) = \frac{\sum_{t=1}^N |\Delta T_i^{obs,pipe}(t) - \Delta T_i^{open,pipe}(t)|}{\sum_{t=1}^N |\Delta T_i^{obs,pipe}(t) - \Delta T_i^{closed,pipe}(t)| + \sum_{t=1}^N |\Delta T_i^{obs,pipe}(t) - \Delta T_i^{open,pipe}(t)|}, \quad i = 1, \dots, pipes \quad (6)$$

Here, $pipes$ refers to the number of pipes in the reduced valve search space and N describes the total number of time steps. ΔT_i^{pipe} accounts for the differences in temperature over time between the two nodes defining pipe i . Thus, $\Delta T_i^{open,pipe}$ defines the difference in simulated temperature between the starting and ending nodes of pipe i when all valves are open; $\Delta T_i^{closed,pipe}$ stores the simulated temperature differences over pipe i when only that particular pipe's valve has been closed in the network; and $\Delta T_i^{obs,pipe}$ covers the observed temperature differences over pipe i with unknown valve status. An SM close to unity represents a higher resemblance of the observed temperature pattern (i.e. with unknown valve status) to the one of the closed pipe segment, i.e. the valve is more likely closed. An SM value close to zero likely represents an open valve. Based on SM, a probability P for a valve being open or closed x -percent of the time in the initial population was computed:

$$P_x(i) = \frac{SM(i) * x * valves}{\sum_{i=1}^{pipes} SM(i)}, \quad i = 1, \dots, pipes \quad (7)$$

When generating the initial *weighted* population, P_x was used as a probability for a valve being closed or open.

Having determined an initial population or new generation, a hydraulic simulation was run for each individual (step II, Figure 3). If the individual's output was 'feasible', the individual's temperature results were simulated (step III, Figure 3).

2.3.2 Fitness functions

Goodness of fit (step IV, Figure 3) was evaluated by a weighted least squares fitness function (see e.g. Wu (2009) for additional functions) extended to include terms for the modelled and observed pressure and temperature:

$$F_1 = \sum_{t=0}^{N_t} \left[W_H \sum_{i=1}^{N_H} (H_i^{sim} - H_i^{obs})^2 + W_T \sum_{i=1}^{N_{T,node}} (T_i^{sim} - T_i^{obs})^2 \right], \quad (8)$$

where N_t is the number of timesteps, N_H and $N_{T,nodes}$ are the number of pressure and temperature measurement sites, respectively, H_i^{sim} and T_i^{sim} are the simulated head and temperature, H_i^{obs} and T_i^{obs} are respectively the observed head and temperature in the network, and W_H and W_T are the weighting factors for the head H and temperature T , respectively. We also propose an alternative fitness function that included the temperature difference over a pipe segment instead of the nodal temperatures:

$$F_2 = \sum_{t=0}^{N_t} \left[W_H \sum_{i=1}^{N_H} (H_i^{sim} - H_i^{obs})^2 + W_T \sum_{i=1}^{N_{T,pipes}} (\Delta T_i^{sim,pipe} - \Delta T_i^{obs,pipe})^2 \right], \quad (9)$$

where $N_{T,pipes}$ is the number of pipes with temperature measurements at both vertices. $\Delta T_i^{sim,pipe}$ represents the temperature differences during simulation between the starting and ending vertices of each pipe, similar to that used in equation (6). Closing one incorrect valve can result in high nodal temperature differences between simulated and measured values in a large area. In such a case, the application of F_1 may incorrectly guide the GA to search for closed valves inside the affected area. When applying F_2 , the error of the downstream area may be less affected by valve closure as the function is concerned with temperature differences over all pipe segments. However, the temperature deviation over the pipe segment with a closed valve should lead to a significantly higher error. F_2 supports the GA in prioritizing the selected pipe segment with doubtful temperature differences.

3 Case studies

The method was evaluated in two applications based on a DMA and a transportation WDN model of the Danish water utility Novafos, located in the northern Copenhagen area (Figure 4). The WDN models were created from a larger model provided by the utility.

The two model applications differed greatly in terms of data availability and network set-up (Table 1). In the transportation network (Figure 4a), the drinking water was distributed from two waterworks (WW) to ten DMAs and temporarily stored in one elevated tank. A large part of the transport network had parallel pipes, almost covering the area from WW1 to Tank, as shown by the dashed close-ups (Figure 4a) illustrating examples of the two parallel pipes, including transverse and bypass connections. Pressure and flow were measured on a minute-by-minute basis at all points where water was entering or exiting the transportation network. Temperature was measured per minute at two DMA entries, namely DMA3 and DMA5 of Søndersø DMA (Figure 4b).

Table 1. Key parameters of the applied water distribution network (WDN) models. P = Pressure; Q = Flow; T = Temperature; PE = Polyethylene.

WDN model name	Nordvand transportation network	DMA (Søndersø)
Nodes (junctions/waterworks/tanks)	219 (216/2/1)	413 (413/0/0)
Metered nodes (inputs/outputs/tanks)	13 (2/10/1)	4 (2/2/0)
Meter types	P, Q, (2 outputs T)	P, Q, (2 inputs T)
Metered input range [m ³ /hr]	405-510.7	38.5-113
Data period	11.-17. February 2019	4.-10. March 2016
Consumers (metered annually)	11	2328
Annually metered consumption [m ³ /hr]	0.2	48.9
Pipes	231	488
Σ Pipe length [km]	41.2	42.3
Material Distribution [%]	Cast Iron (19), Ductile Iron (1), PE (3), Concrete (74), Fibre Cement (3)	Cast Iron (57), Ductile Iron (4), PE (39)
Shut-off valves	84	298

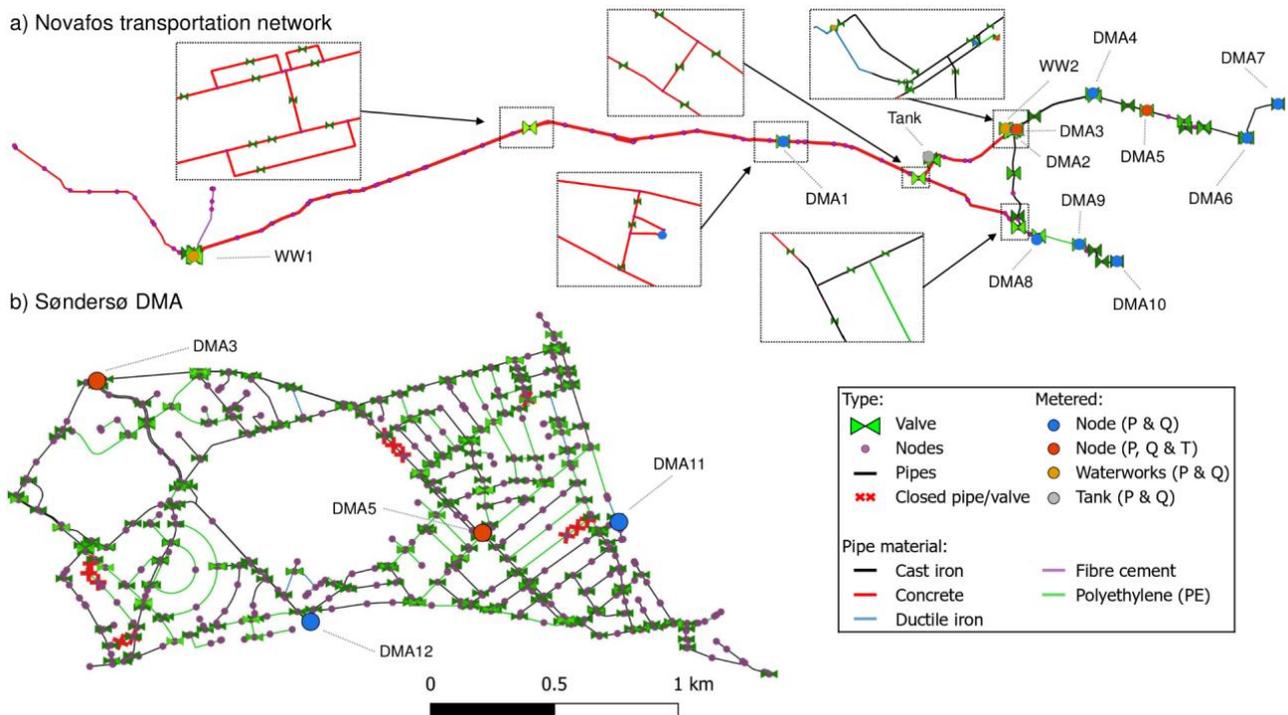


Figure 4. Novafos transportation water distribution network model. Lighter valve colours represent a higher number of overlapping valves. A close-up of the area around WW2 is illustrated in Figure 10b. WW = waterworks; P = pressure; Q = flow; T = temperature.

We cleaned the data for anomalies found in the raw data (e.g. illegitimate values) by the methodologies described in Kirstein et al. (2019) and removed temperature values below or equal to 0 °C. All data sets were averaged to uniform data streams with 5-minute intervals. Whereas real measurements were available at all major consumption nodes in the transportation network, the semi-synthetic DMA model used measured temperature for the inlets DMA3 and DMA5, pressure at DMA12 and flows at DMA3, DMA5 and DMA11. In the DMA model, five valves were closed (Figure 4b) and a synthetic temperature data set including temperatures at each node of the DMA was created to represent the ‘unknown’ valve status. The semi-synthetic data set represented the current development of increased temperature collection by smart meters at the end-user. In both case studies, the model started with all valves assumed open. The transportation network study compared the utility’s actual valve status with the valve status estimated by the GA. After calibration, the valves identified by the GA were checked in the field.

3.1 GA and hydraulic model parameters

In both case studies, a one week period was analysed (Table 1). In the DMA case study, only temperature data was included in the fitness function. Examples of how the temperature was affected in selected nodes over the course of one week by opening all valves vs. closing five valves can be found in SI D. The GA was run for the DMA and transportation network case study using a pre-set list of parameters (Table 2). We used suggested default EC and CF values for effective integer GA (MathWorks, 2018).

Table 2. GA model set-ups applied in the DMA and transportation network case study.

Model	DMA	Transportation network
Clustering threshold	200 m	500 m
Initial Population	Random; Weighted	Random
p	500; 1000	500
x [%]	1; 2.5; 5	1; 2.5; 5
EC [%]	10	10
CF:MR	9:1; 9.5:0.5	9:1; 9.5:0.5
Fitness Function	Eq. (8); Eq. (9)	Eq. (10)
F, subject to	T	T; H
Weights ($W_T:W_H$)	-	0:1; 1:3; 1:1; 3:1; 1:0
Total runs	48	30

The GA stopped for a ‘perfect fit’ (i.e. fitness value equal 0), after 50 generations, or when the mean fitness value was equal to the best fit.

3.1.1 Fitness function and constraints (transportation network case study)

The most recent field test of the pressure measurements from September 2018 revealed an offset in the range of -1.3 to +3 mWC. As a consequence of these field tests, F_1 was changed to account for the unknown temporal pressure offsets in the transport network case study by subtracting the median of the measured and simulated pressure values, making the fitness function reflect the variance of values rather than absolute values:

$$F_{1,modified} = \sum_{t=0}^{N_t} \left[W_H \sum_{i=1}^{N_H} ([H_i^{sim} - \text{median}(H_i^{sim})] - [H_i^{obs} - \text{median}(H_i^{obs})])^2 + W_T \sum_{i=1}^{N_{T,node}} (T_i^{sim} - T_i^{obs})^2 \right] \quad (10)$$

In the transportation network model, WW2 served as head boundary and its pressure data were raised by 0.7 mWC according to the most recently measured offset.

Owing the on-and-off status of very high pump rates at WW1, the model was very sensitive to generating negative pressures in certain set-ups. Consequently, solutions (including negative pressure) previously deemed ‘infeasible’ between steps II and III (Figure 3) were accepted.

3.2 Temperature model parameters

Soil temperatures measurements were based on measured temperatures from a water sampling location with long residence time at the periphery of the Copenhagen WDN (Figure 5). The high daily variations seen in the DMA inlet temperatures, up to 2.3 °C during winter and over 5.5 °C during summer, were a consequence of both the origin of the water and of the water’s retention time in the system. Most likely, a high proportion of the water at DMA3 originated from WW1 and Tank (Figure 5a). Long residence time led to cooling of the drinking water during winter and heating during

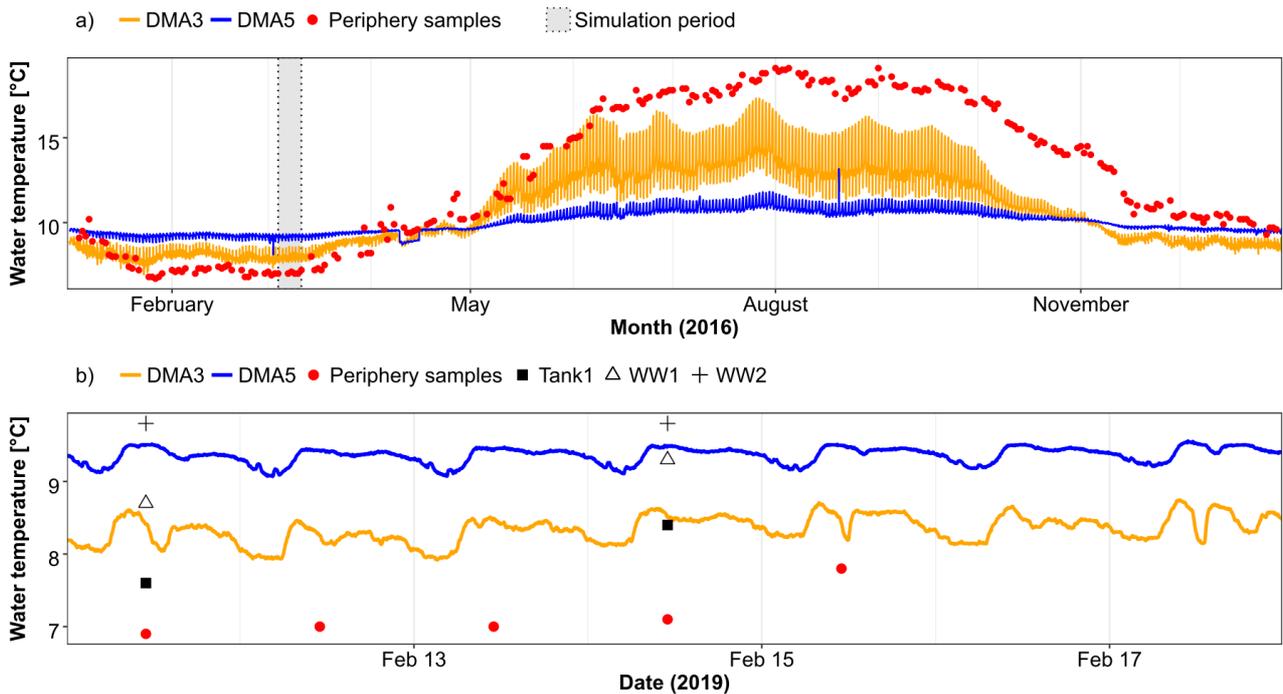


Figure 5. Measured water temperatures at DMA3 and DMA5 including samples taken at WW1, WW2 and Tank and samples taken at the periphery of Copenhagen water distribution network assumed representative for the soil temperature. The analysed period of March 2016 is illustrated in SI D.

summer before it reached DMA3. The water temperature at DMA5 showed less variation. It is expected that most of its water originated from WW2. During a week in April 2016, the temperature measured at DMA5 resembled that of DMA3 (Figure 5a), explained by a reduced production in WW2. It is likely that this change in operation caused water from WW1 and Tank to enter DMA5.

Samples taken in the analysed period (Figure 5b) suggested constant temperatures for all three water sources in the network. The temperature model used average source temperatures (Table 3). The temperatures at WW2 and the soil temperature in the transportation network case study were modified up to 0.3 °C (within the sampled temperature range) to improve the overall fit. This modification is within the stated thermometer uncertainty of ± 0.5 °C. λ_{soil} was set equal to 2.2 W/m/K based on Danish moraine clay measurements (Ditlefsen et al., 2011).

4 Results and discussion

4.1 DMA (semi-synthetic case study)

The valve search space of the DMA set-up was reduced from 488 to 379 pipes. Three out of five valves had a SM > 0.75 and were also unknowingly to the GA closed in the model set-up (Figure 6a). Furthermore, for a SM value close to zero, multiple valves indicated that they were likely open. Two closed valves had an SM of around 0.5, which indicated that, not considering any interactions of multiple closed valves, the status of the two valves had little or no impact on the temperature distribution over their respective pipes. It appears that the SM gave an initial indication of the correct valve status for selected pipe segments and serves to guide the GA towards an improved fit.

Table 3. Parameters applied for simulating water temperature. The mean sample drinking water temperatures were applied, except for ¹⁾ minimum sampled temperature; ²⁾ temperature was lowered by 0.2 °C (Section 4.2.1); ³⁾ max value omitted when computing the mean.

Parameter	Symbol	Value	Unit	Reference
Temperature at WW1 (Mar-16/Feb-19)	T_{WW1}	8.6/9	[°C]	Sampled
Temperature at WW2 (Mar-16/Feb-19)	T_{WW2}	9.4 ¹ /9.6 ²	[°C]	Sampled
Temperature at Tank1 (Mar-16/Feb-19)	T_{Tank1}	7.75/8	[°C]	Sampled
Undisturbed soil temperature (Mar-16/Feb-19)	T_{soil}	7.02/7 ³	[°C]	Sampled
Heat capacity of water [10°C]	c_p	4,188	[J/kg/K]	Dinçer and Zamfirescu (2016)
Density of water [10°C]	ρ	999.7	[kg/m ³]	Dinçer and Zamfirescu (2016)
Dynamic viscosity of water [10°C]	μ	1.31e-3	[Ns/m ²]	Dinçer and Zamfirescu (2016)
Thermal conductivity of cast iron pipes	$\lambda_{pipe,CI}$	45	[W/m/K]	Czichos (2000)
Thermal conductivity of cast ductile iron pipes	$\lambda_{pipe,DI}$	32.4	[W/m/K]	DIMG (1990)
Thermal conductivity of concrete pipes	$\lambda_{pipe,C}$	1.16	[W/m/K]	Lauritsen et al. (2012)
Thermal conductivity of polyethylene pipes	$\lambda_{pipe,PE}$	0.45	[W/m/K]	Czichos (2000)
Pipe wall thickness	d_{wall}	1/9· d_i	[m]	Hubeck-Graudal et al. (2019)
Thermal conductivity of the water	λ_{water}	0.57	[W/m/K]	Dinçer and Zamfirescu (2016)
Thermal conductivity of the soil	λ_{soil}	2.2	[W/m/K]	Ditlefsen et al. (2011)
Burying depth	z	1.4+2· d_{wall}	[m]	Hubeck-Graudal et al. (2019)

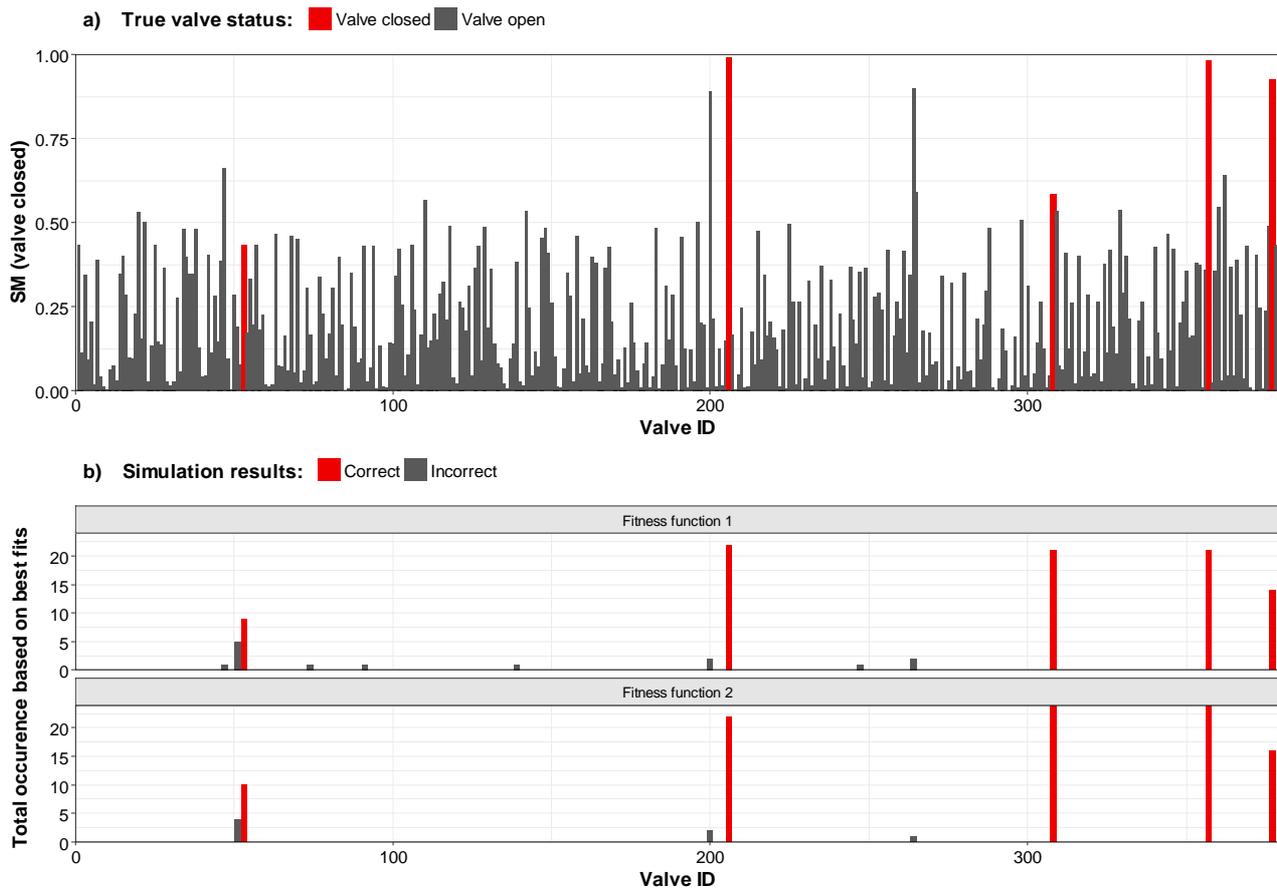


Figure 6. a) Similarity measure (SM) of a valve being closed (1) or open (0) used to compute an initial population. b) The total number of times a valve had been closed during the 24 GA runs of fitness functions F_1 and F_2 .

The development of individual GA runs over their generations can be found in SI E. Figure 6b illustrates the total number of times valves were closed based on the best fit in each individual GA run for F_1 and F_2 , respectively. Most GAs stopped when there was no more variation in the population, meaning that the best fit was equal to the mean fit. As the figure shows, only valve ID 53 was identified during less than 50% of the runs. Instead, another valve close to this valve was often found to be the best fit (grey bar, ID 51), likely causing the GA to become stuck in a local minimum. The remaining four closed valves were correctly identified in at least 58% of all runs, irrespective of selected fitness function. Interestingly, the valve with ID 308 was closed in most runs, even though it did not show a particularly high SM value (Figure 6a). This is an example of the effect of interacting valve statuses. Initially, the valve's SM value close to 0.5 indicated that it was 'unclear' for the GA whether the valve was open or closed. As the valve was closed in most simulations, it can be argued that the closure of this valve only led to significant better fitness values when another valve was closed simultaneously. In both fitness functions, only 8 out of 24 runs found the perfect solution. Overall, F_2 resulted in better fits. Whereas F_1 identified the status of 14 valves incorrectly and that of 87 valves correctly, F_2 identified 7 valves incorrectly and 96 correctly (Figure 6b). Valves identified incorrectly tended to have high SM values (Figure 6a). The tendency of better F_2 vs. F_1 results is also shown in Figure 7.

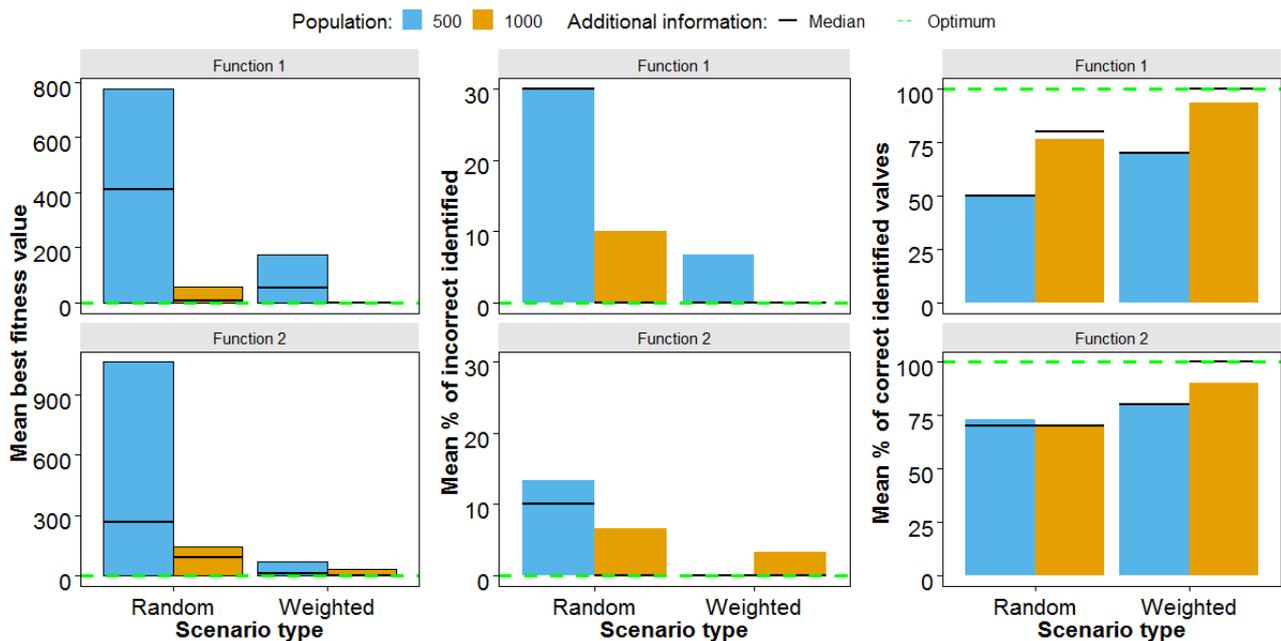


Figure 7. Averaged fitness function values, based on the best fit in various GA scenarios.

Here, the first row shows the mean best fitness value, based on the various GA runs, and the mean percentages of incorrectly and correctly identified valves with F_1 . The second row illustrates the results based on F_2 . Most importantly, however, Figure 7 indicates that a higher population size and an improved initial population (weighted scenario) had a greater influence on the final outcome than the selection of F_1 or F_2 . Irrespective of selected fitness function, a better mean final fitness value was reached using the weighted population. The mean best fitness value was reduced by 78% and 96% for populations of 500 and 1000, respectively with F_1 , and by 94% and 79%, respectively when selecting F_2 . In terms of wrongly identified valves, F_2 seemed to generate better results; however, the

median of both fitness functions indicated that there were certain outliers in the data set establishing the higher means. Also, a falsely identified valve is not immediately bad (e.g. valve ID 51, Figure 6a) because this might give the utility a good indication of a location for an unknown closed valve.

4.2 Transportation network (real case study)

4.2.1 Model validation

A tracer analysis conducted in the hydraulic model revealed that close to 100% of the water at DMA5 originated from WW2 (total distance approx. 2.4 km), which was in accordance with the utility’s understanding of the system; thus DMA5 was a good location to illustrate the simulated vs. observed temperature (Figure 8). First, the temperature was simulated including the mean sampled WW2 temperature (Figure 8a). Next, we re-simulated the temperature at DMA5 by having lowered the source temperature of WW2 by 0.2 °C, thus improving the overall fit (Figure 8b). This was assumed

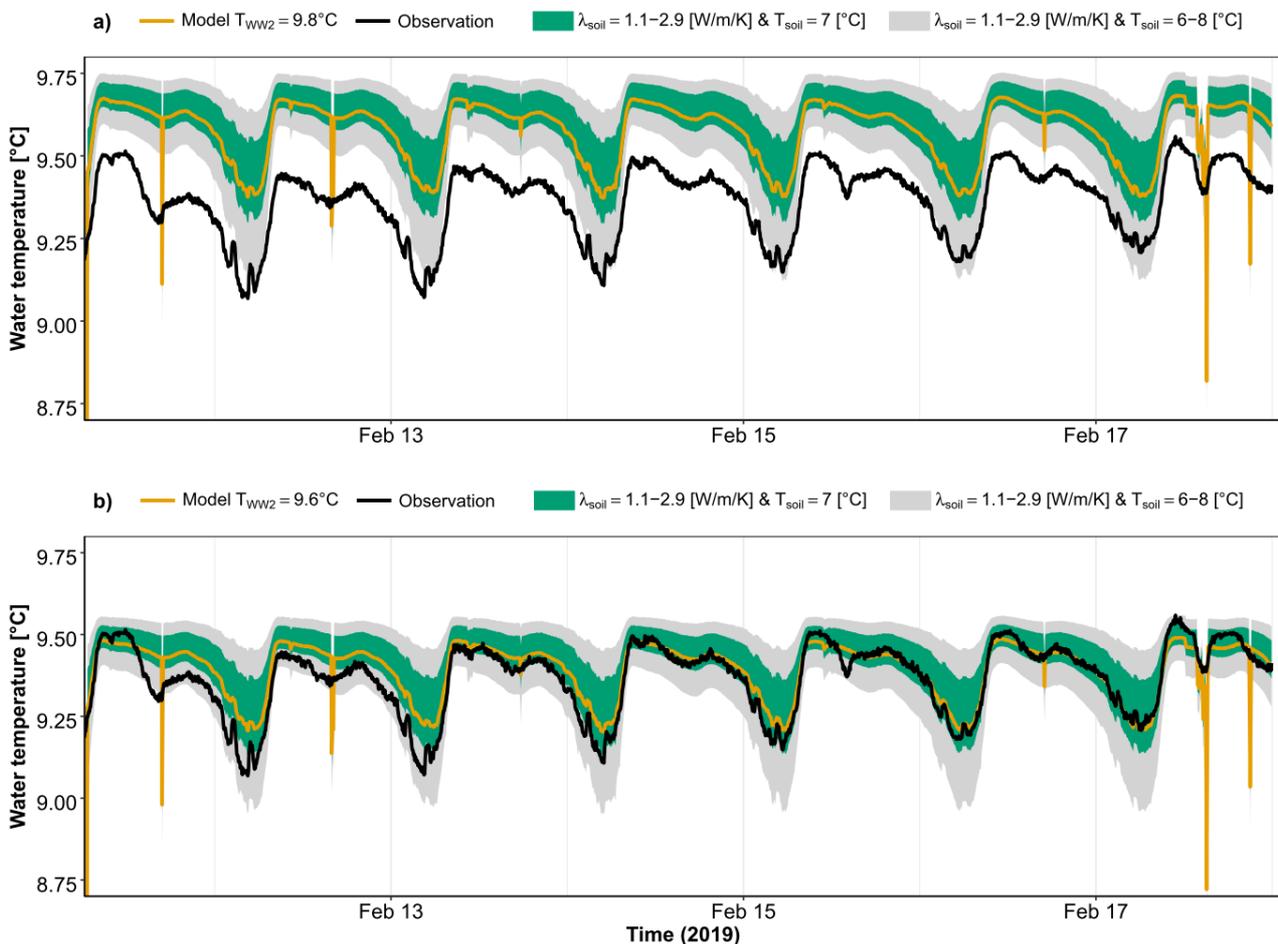


Figure 8. Measured vs. simulated drinking water temperature downstream at entry DMA5 based on varying λ_{soil} and T_{ww} . The figures show 1) the effect of only varying the source temperature at WW2; 2) the maximum temperature deviation by including the λ_{soil} range of 1.1–2.9 W/m/K (inner shaded area around model temperatures) based on measured variations in the λ_{soil} of moraine clay (Ditlefsen et al., 2011) including T_{soil} listed in Table 3; and 3) the maximum temperature deviation by varying the $T_{soil} \pm 1$ °C and the soil conductivities between 1.1–2.9 W/m/K (outer shaded area around model temperatures).

to be within the possible boundaries of the real water temperature at the waterworks, as the sampled source temperature had an uncertainty of ± 0.5 °C. The figure displays some of the uncertainty likely related to soil temperature and conductivity in the temperature model by showing temperature bands for ranges of $T_{soil} \pm 1$ °C and likely variations in the λ_{soil} of moraine clay between 1.1 and 2.9 (Ditlefsen et al., 2011).

Regardless of selected model parameters, the simulated temperature does follow the overall pattern of temperature change at DMA5 (Figure 8a).

4.2.2 GA results

The valve search space of the transport WDN model was reduced from 231 to a total of 106 valves. Additional information about this reduced search space can be found in SI F. The best fits of all individual GA runs revealed great variation and uncertainty in the valves identification (Table 4). A total of 41 different valves were closed in 30 different runs. Some valves, e.g. ID18, however, appeared in 4–6 best fits regardless of weighting scenario (i.e. $W_T:W_H$). A higher weighting of the temperature model also led to a higher number of closed valves. This could be explained by the fact that the temperature model was capable of identifying closed valves that cannot be identified when analysing the pressure alone. On the other hand, it is likely that the temperature model tried to improve its fit by closing selected valves that had a small but beneficial effect on the temperature simulation and thus fitness function, which was not seen when using a high pressure weighting factor.

When looking at the individual weighting scenarios, only small variations existed between the best fit value and identified valves. For example, in the scenario with a pressure weight equal to 100%, runs 2–6 achieved similar best fits and all five runs indicated that three valves should be closed (Table 4). Whereas valve ID 18 and 87 appeared in all five best fits, the third ID varied, only marginally improving the overall fit. An explanation could include a too loose clustering, including valves at close range whose change in status had a low but still improving impact on the fitness function. Future implementations should assess the sensitivity and refine the identified results, to highlight which valves have a clear effect on the overall fitness function and which could be excluded from the GA analysis, e.g. as done by Do et al. (2018).

The best fit for DMA3 (run 28), when using a weighting factor that only considered the temperature data (Table 4), is shown in Figure 9 and termed *calibrated*. Regardless of which valve settings were used, the simulated temperature for DMA3 did not seem to match the overall pattern of the measured temperature as well as in DMA5 (Figure 8). However, compared to the original settings by the utility (termed *utility*) or all valves being open (*all open*) the new temperature profile was improved (Figure 9). Up to now, it had been the understanding of the utility that all water at DMA3 originated from WW2, as their valve setting *utility* indicated. Our results indicated that most water at DMA3 originated from WW1 and Tank (Figure 5a), which was partly confirmed by the *calibrated* GA results, closing valves that connected DMA3 with WW2. Also, the head in the section no longer followed the same pattern as when using the standard valve settings (Figure 9).

Table 4. Individual GA calibration results of transport WDN model. The generation column displays the generation when the GA reached the best fit the first time before reaching the stopping criteria. Bold valve IDs highlight valves closed in the utility model. x = initial percentage closed valves; CF = cross-over fraction; MR = mutation rate; W_T = temperature weighting factor; W_H = pressure head weighting factor; F_T = best fit of temperature values; F_H = best fit of pressure values.

Run	x [%]	CF:MR [%]	$W_T:W_H$ [%]	Generation	No. closed valves (open:closed in utility model)	Closed valves [ID]	Best fit ($F_T:F_H$)
1	1	90:10	0:100	15	5 (0:5)	16, 28, 31, 56, 87	10004 (2655:10004)
2	1	95:05	0:100	6	3 (1:2)	18 , 79, 87	8846 (4886:8846)
3	2.5	90:10	0:100	7	3 (1:2)	18 , 79, 87	8846 (4886:8846)
4	2.5	95:05	0:100	4	3 (1:2)	18 , 35, 87	8846 (4886:8846)
5	5	90:10	0:100	8	3 (1:2)	18 , 77, 87	8846 (4886:8846)
6	5	95:05	0:100	7	3 (1:2)	18 , 51, 87	8846 (4886:8846)
7	1	90:10	25:75	12	5 (0:5)	16, 28, 31, 76, 87	8167 (2655:10004)
8	1	95:05	25:75	13	6 (1:5)	16, 28, 31, 60 , 70, 87	8167 (2655:10004)
9	2.5	90:10	25:75	12	5 (1:4)	18 , 23, 87, 102	7834 (4778:8853)
10	2.5	95:05	25:75	18	4 (1:3)	8, 18 , 87, 103	7839 (4479:8859)
11	5	90:10	25:75	13	4 (1:3)	18 , 67, 79, 87	7834 (4778:8853)
12	5	95:05	25:75	20	5 (2:3)	7 , 18 , 87, 102	7856 (4886:8847)
13	1	90:10	50:50	6	3 (0:3)	16, 31, 36	6234 (2274:10194)
14	1	95:05	50:50	8	4 (1:3)	16, 18 , 31, 45	6216 (2259:10171)
15	2.5	90:10	50:50	10	4 (1:3)	16, 31, 32 , 105	6234 (2274:10194)
16	2.5	95:05	50:50	13	4 (1:3)	16, 18 , 31, 76	6216 (2258:10174)
17	5	90:10	50:50	16	5 (1:4)	16, 18 , 31, 85, 86	6215 (2259:10172)
18	5	95:05	50:50	17	7 (3:4)	16, 18 , 31, 60 , 86, 91, 93	6215 (2259:10171)
19	1	90:10	75:25	21	4 (1:3)	16, 18 , 31, 76	4237 (2258:10174)
20	1	95:05	75:25	12	5 (2:3)	16, 18 , 31, 60 , 91	4237 (2258:10174)
21	2.5	90:10	75:25	7	4 (1:3)	16, 18 , 31, 70	4237 (2258:10174)
22	2.5	95:05	75:25	10	4 (1:3)	16, 18 , 31, 56	4237 (2258:10174)
23	5	90:10	75:25	17	6 (2:4)	16, 18 , 22, 31, 32 , 75	4237 (2259:10171)
24	5	95:05	75:25	19	6 (2:4)	16, 18 , 22, 31, 60 , 75	4237 (2259:10171)
25	1	90:10	100:0	13	6 (2:4)	16, 17, 18 , 34, 64, 106	294 (294:214850)
26	1	95:05	100:0	21	6 (1:5)	28, 35, 40, 75, 94 , 97	1682 (1682:28361)
27	2.5	90:10	100:0	18	7 (3:4)	15 , 16, 17, 18 , 24 , 63, 64	294 (294:214850)
28	2.5	95:05	100:0	20	8 (4:4)	15 , 16, 17, 18 , 24 , 34, 60 , 64	294 (294:214850)
29	5	90:10	100:0	26	8 (4:4)	15 , 16, 17, 18 , 24 , 32 , 64, 90	294 (294:214850)
30	5	95:05	100:0	15	7 (3:4)	16, 18 , 60 , 64, 70, 93 , 104	297 (297:214846)

Temperature measurements at DMA3 indicated that a large proportion of the water originated from Tank, as the temperature was slightly off the tank's boundary condition of 8 °C for long periods (Figure 9). Solely calibrating with pressure data did not lead to a correct simulation of the water origin. Modelling and use of temperature can add an extra dimension to the valve search process and help utilities to increase the understanding of their water distribution system. Future studies should explore the possible high variation in the tank's outlet temperature indicated by the two temperature samples from the tank (Figure 5b).

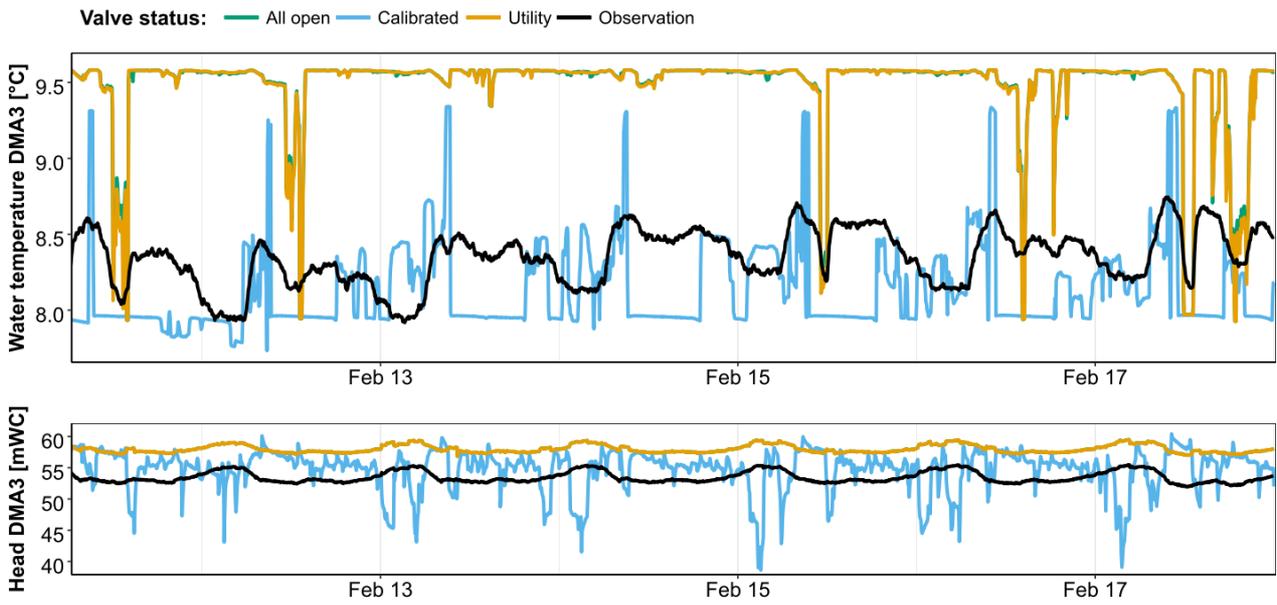
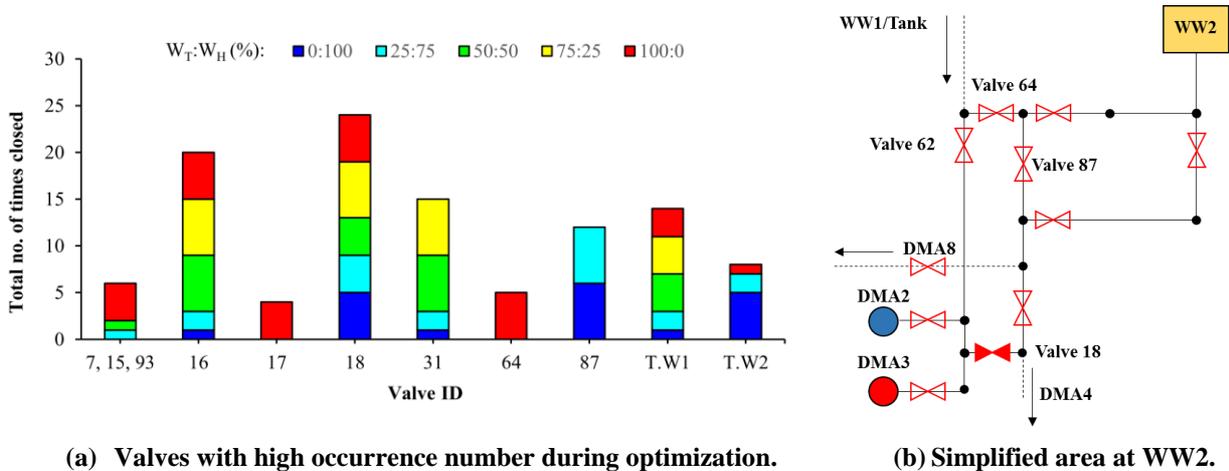


Figure 9. Temperature and head at DMA3 modeled for three valve scenarios. All valves open (green); a calibration with a weighting factor of 100% temperature measurements (blue); with the utilities knowledge of opened and closed valves (yellow); and the measured temperature at the location (black). Also, the measured and modelled head at the location are shown in the lower figure.

4.2.3 Testing in the field

Valves identified as closed in more than 50% of a weighting scenario (Table 4), i.e. more than 3 times, were considered for further analysis (Figure 10a). The three valves with ID 7, 15 and 93 were located in the same ‘valve station’ and achieved these conditions as a group. Sectioning the parallel pipes running between WW1 and Tank (Figure 4a) into six groups, depending on their orientation (west or east) and location between the transverse connections, led to two additional valve groups achieving the 50% closures, i.e. groups TW.1 and TW.2 (Figure 10a). In reality, there were no valves located in TW.1 and TW.2 and this could indicate that the model tried to compensate for unknown model uncertainties.



(a) Valves with high occurrence number during optimization.

(b) Simplified area at WW2.

Figure 10. (a) The total occurrence of valves that have been closed at least once in more than 50% of the time of a given weighting scenario. (b) A simplified version of the network around valve 18, 64 and 87.

The valve status identification process (Table 5) highlighted valves with uncertain status and guided the utility to locations that needed further maintenance and investigation and where the model had clear discrepancies in relation to measured data.

In particular, the field results from valves ID 18, 64 and 87 (Table 5) could explain some of the modelling results (Figure 10b). For example, when closing valve ID 64, the simulated temperature at valve DMA3 resembled, to a higher degree, the temperature related to a long residence time. This was also seen in the real measurements. However, with the GA evaluations only considering pressure optimization, valve 64 was not closed and most of the water at DMA3 was simulated to originate from WW2 with a simulated temperature incorrectly closer to the waterworks temperature than found by measurement. The valve status calibration process of the transport WDN might have suffered from the missing GIS information that was revealed during field testing of the valves. Further studies should explore additional factors and parameters not considered here, for example partially closed valves, wrong diameters, or missing pipe connections that might affect the results.

Table 5. Summary of valve field testing.

Valve ID	Closed	Open	Locked-in	Inaccessible	Information
7, 15, 93	(X)				All bypasses and valves, part of the utility's GIS system, no longer exist in reality.
16			X		Reported faulty.
17				X	This valve could not be tested, but additional valves not reported in the GIS system were visible around the area, being part of the parallel running pipes.
18		(X)			The technician was not sure whether the valve had been open or closed.
31				X	
64			X		At this location a 'valve station' was discovered, including two valves part of the GIS system (Valve ID 62 and 64, Figure 10b). Both valves contained a bypass with an additional valve not part of the GIS system. Valve ID 64 including the bypass valve were reported faulty.
87				X	Heavy vegetation.

4.3 Current and future implications of temperature modelling and valve testing

Analysis of the temperature data and linked calibration process revealed that temperature data can be used to increase understanding of the WDN. The temperature profile of Figure 5a reveals the origin of the water, which is not obvious from pressure and flow measurements only. The additional value of temperature data and modelling was particularly true in the transportation network case study, where pressure measurements had not been calibrated for a long period. As concluded in Walski et al. (2014), the quantity and in particular quality of the data determine the success of the GA for optimal valve status identification.

Danish WDNs tend to have large diameters, resulting in very small head losses throughout the network. Here, partially and fully closed valves, or leakage only have minor effects on the pressure.

Our study showed that when including a temperature model, the valve status identification process was improved, as additional critical locations were identified that otherwise were overlooked when solely using pressure data. Further studies should analyse to what extent temperature data can be used together with pressure and flow measurements to improve the detection of leaks as well as partially and fully closed valves.

Several uncertainties in the temperature model were not assessed in this study. For example, all boundary temperatures were assumed to be constant, but would have increased the validity of the model outcome if the temperature input was measured. Also, the soil temperature and the conductivity of the soil and pipe materials were likely to vary, e.g. depending on location and type of (road) surface, but were applied in a simplified manner. This parameter uncertainty is of high importance. Zlatanovic et al., (2017) conducted a sensitivity analysis by varying the input parameters of a domestic drinking water temperature model, similar to those of the WDN temperature models, by $\pm 10\%$. Their analysis revealed that parameters representing the outer pipe surface resistance and the ambient (soil in WDN) and inlet temperatures had the largest influence on the model output. Nevertheless, the temperature modelling in our study showed that it has potential in the fields of water distribution network operation, maintenance and design.

5 Conclusion

We extended a WDN model with a temperature model in order to utilize temperature data for model-based identification of valve status and location. A GA was applied to search for the optimal valve status by minimizing a fitness function dependent on the difference between real and simulated pressure and temperature data. The method was evaluated by two case studies. In the first semi-synthetic case study, five valves were closed and we assessed two fitness functions that considered the nodal temperature difference and the temperature difference over all pipes, respectively. Each function was assessed in 24 different set-ups and a best fit was found in 8 runs, independent of fitness function. The second fitness function resulted in better outcomes, where 96 out of 120 possible closed valves were identified. Moreover, it was found that a larger population size and a weighted, instead of random, initial population had a greater impact on the final outcome than modification of the fitness function. In the second real case study, a transportation network with temperatures measured at two locations revealed that the water is distributed or mixed differently than anticipated by the utility. Having included the temperature model during the calibration process revealed critical locations not found when solely using pressure data. Field testing of valves, identified by the calibration process, revealed various valves that were unknown to the utility, had doubtful status or were not accessible. The study revealed that the quality and quantity of data were the most restraining factors, challenging the interpretation of the results. The analysis of temperature data and the application of the temperature model showed great promise as an add-on in the field of water distribution network modelling, and as a tool for water utilities to increase their knowledge of their water distribution systems.

6 Acknowledgments

We thank all participating partners of the LEAKman project, in particular Mette Fogtmann, Novafos. Moreover, we would like to thank Bernd Dammann from DTU Compute and Hugo Maxwell Connery from DTU Environment.

7 References

- Blokker, E.J.M., Pieterse-Quirijns, E.J., 2013. Modeling temperature in the drinking water distribution system. *J. Am. Water Work. Assoc.* 105, E19–E28. doi:10.5942/jawwa.2013.105.0011. Erratum (2018) in *J. Am. Water Work. Assoc.* 110, 98–98. doi:10.1002/awwa.1178
- Blokker, E.J.M., van Osch, A.M., Hogeveen, R., Mudde, C., 2013. Thermal energy from drinking water and cost benefit analysis for an entire city. *J. Water Clim. Chang.* 4, 11–16. doi:10.2166/wcc.2013.010
- Cengel, Y.A., 2003. *Heat Transfer: A Practical Approach*. Mc Graw-Hill, Boston.
- Czichos, H., 2000. Werkstoffe (in German), in: *Die Grundlagen Der Ingenieurwissenschaften*. Springer Berlin Heidelberg, Berlin, Heidelberg, pp. 511–600. doi:10.1007/978-3-662-06652-2_5
- De Pasquale, A.M., Giotri, A., Romano, M.C., Chiesa, P., Demeco, T., Tani, S., 2017. District heating by drinking water heat pump: Modelling and energy analysis of a case study in the city of Milan. *Energy* 118, 246–263. doi:10.1016/j.energy.2016.12.014
- Deb, A.K., Snyder, J.K., Grayman, W.M., 2012. Management of valves to improve performance reliability of distribution systems. *J. Am. Water Works Assoc.* 104, E370–E382. doi:10.5942/jawwa.2012.104.0087
- Deb, K., 2000. An efficient constraint handling method for genetic algorithms. *Comput. Methods Appl. Mech. Eng.* 186, 311–338. doi:https://doi.org/10.1016/S0045-7825(99)00389-8
- Deep, K., Singh, K.P., Kansal, M.L., Mohan, C., 2009. A real coded genetic algorithm for solving integer and mixed integer optimization problems. *Appl. Math. Comput.* 212, 505–518. doi:10.1016/j.amc.2009.02.044
- Delgado, D.M., Lansley, K.E., 2009. Detection of Closed Valves in Water Distribution Systems, in: *Water Distribution Systems Analysis 2008*. American Society of Civil Engineers, Reston, VA, pp. 1–7. doi:10.1061/41024(340)82
- DIMG, 1990. *Ductile Iron Data for Design Engineers*. Rio Tinto Iron & Titanium Inc., Available at <https://www.ductile.org/didata/> (accessed October 20, 2018).
- Dinçer, İ., Zamfirescu, C., 2016. *Drying Phenomena, First Edit.* ed, *Drying Phenomena: Theory and Applications*. John Wiley & Sons, Ltd, Chichester, UK. doi:10.1002/9781118534892
- Ditlefsen, C., Vankilde-Pedersen, T., Højberg, A.L., 2011. D3 Geologi og jordvarmeboringer (in Danish). Available at http://geoenergi.org/xpdf/jordvarme-og-geologi-d3_2.pdf (accessed November 13, 2019).
- Do, N.C., Simpson, A.R., Deuerlein, J.W., Piller, O., 2018. Locating Inadvertently Partially Closed Valves in Water Distribution Systems. *J. Water Resour. Plan. Manag.* 144, 04018039. doi:10.1061/(ASCE)WR.1943-5452.0000958
- Eck, B.J., Saito, H., McKenna, S.A., 2016. Temperature dynamics and water quality in distribution systems. *IBM J. Res. Dev.* 60, 7:1-7:8. doi:10.1147/JRD.2016.2594128
- Eliades, D.G., Kyriakou, M., Vrachimis, S.G., Polycarpou, M.M., 2016. EPANET-MATLAB Toolkit : An Open-Source Software for Interfacing EPANET with MATLAB. 14th *Comput. Control Water Ind. Conf. CCWI 2016* 1–8. doi:10.5281/ZENODO.437751

- Hubeck-Graudal, H., Kirstein, J.K., Ommen, T., Rygaard, M., Elmegaard, B., 2019. Drinking water supply as low-temperature source in the district heating system: a case study for the city of Copenhagen (submitted).
- Kirstein, J.K., Høgh, K., Rygaard, M., Borup, M., 2019. A semi-automated approach to validation and error diagnostics of water network data. *Urban Water J.* 1–10. doi:10.1080/1573062X.2019.1611884
- Kvisgaard, B., Hadvis, S., 1980. Varmetab fra fjernvarmeledninger (in Danish). Technical Report. Technical University of Denmark, Kgs. Lyngby.
- Lauritsen, A.B., Andersen, N.B., Andreasen, B., Drivsholm, C., Hammer, F., Ulbjerg, F., Madsbøll, H., Poulsen, H., Tommerup, H., de Wit, J., 2012. *Varme ståbi* (in Danish), 6th ed. Nyt Teknisk Forlag, Kbh.
- MathWorks, 2018. Mixed Integer Optimization. Available at <https://se.mathworks.com/help/gads/mixed-integer-optimization.html> (accessed September 09, 2018).
- Ray, R., Walski, T., Gangemi, T., Gennone, M., Juergens, B., 2007. Lessons Learned in Calibrating an IDSE Model. *World Environ. Water Resour. Congr.* 2007 1–10. doi:10.1061/40927(243)491
- Rossman, L.A., 2000. *EPANET 2 Users Manual*. Cincinnati, OH.
- Sage, P., 2014. Practical Methods to Obtain Improved Outputs from Water Network Modelling Optimization. *Procedia Eng.* 70, 1450–1459. doi:10.1016/j.proeng.2014.02.160
- Savic, D.A., Kapelan, Z.S., Jonkergouw, P.M.R., 2009. Quo vadis water distribution model calibration? *Urban Water J.* 6, 3–22. doi:10.1080/15730620802613380
- Savic, D.A., Walters, G.A., 1997. Genetic Algorithms for Least-Cost Design of Water Distribution Networks. *J. Water Resour. Plan. Manag.* 123, 67–77. doi:10.1061/(ASCE)0733-9496(1997)123:2(67)
- Shang, F., Uber, J.G., Rossman, L.A., 2007. *EPANET Multi-Species Extension User's Manual* 115.
- Sophocleous, S., Savić, D.A., Kapelan, Z., Giustolisi, O., 2017. A Two-stage Calibration for Detection of Leakage Hotspots in a Real Water Distribution Network. *Procedia Eng.* 186, 168–176. doi:10.1016/j.proeng.2017.03.223
- Walski, T., Sage, P., Wu, Z., 2014. What Does it Take to Make Automated Calibration Find Closed Valves and Leaks?, in: *World Environmental and Water Resources Congress 2014*. American Society of Civil Engineers, Reston, VA, pp. 555–565. doi:10.1061/9780784413548.059
- Walski, T.M., Chase, D. V, Savic, D.A., Grayman, W., Beckwith, S., Koelle, E., 2003. *Advanced water distribution modeling and management*. Haestad press, Waterbury, CT.
- Wilson, C., 2011. Sustainable distribution system asset management: Risk mitigation through control point management. *J. Am. Water Works Assoc.* 103, 100–103. doi:10.1002/j.1551-8833.2011.tb11456.x
- Wu, Z.Y., Song, Y., Syed, J.L., 2012. Optimization Model for Identifying Unknown Valve Statuses and Settings, in: *WDSA 2012: 14th Water Distribution Systems Analysis Conference*, 24-27 September 2012 in Adelaide, South Australia. pp. 975–983.
- Zlatanovic, L., Moerman, A., van der Hoek, J.P., Vreeburg, J., Blokker, M., 2017. Development and validation of a drinking water temperature model in domestic drinking water supply systems. *Urban Water J.* 14, 1031–1037. doi:10.1080/1573062X.2017.1325501

Supporting Information (SI)

Jonas Kjeld Kirstein¹, Shuming Liu, Klavs Høgh, Martin Rygaard and Morten Borup

The aim of the SI is to offer additional insight into selected methods and results in the paper ‘*Valve status identification by temperature modelling in water distribution networks*’.

We refer to the corresponding section of the main paper at the beginning of each section of the SI.

Contents: Supporting information

A	Temperature model	23
A.1	Parametric analysis	24
B	Effect of individual parameters on Hubeck-Graudal model	26
C	Valve search space reduction.....	27
D	Temperature in selected nodes (open vs. closed).....	28
E	Genetic algorithm results	29
F	Example of valve search space reduction	31
H	References.....	32

¹ E-mail: jkir@env.dtu.dk; Technical University of Denmark, Department of Environmental Engineering, Bygningstorvet, Building 115, 2800 Kgs. Lyngby, Denmark.

A. Temperature models

Section 2.1 of the main paper

The *undisturbed soil model* (Blokker and Pieterse-Quirijns, 2013) takes the thermal resistance between the pipe wall (R_p) and the flowing water (R_w) into account:

$$K_{p,1} = (R_p + R_w)^{-1} = \left(\frac{d_{wall}}{\lambda_{pipe}} + \frac{1}{a} \right)^{-1}, \quad (1)$$

where λ_{pipe} [W/m/K] is the thermal conductivity of the pipe material, d_{wall} [m] the pipe wall thickness and a [W/m²/K] the convective heat transfer coefficient listed in the main paper. The convective heat transfer coefficient requires the computation of the Reynolds (Re) and Prandtl (Pr) numbers (Cengel, 2003):

$$Re = \frac{\rho v L_c}{\mu}, \quad (2)$$

where ρ [kg/m³] is the density of the water, v [m/s] the velocity of the water, μ [kg/m/s] the dynamic viscosity of the water and L_c [m] the characteristic length (hydraulic diameter) of the pipe. Moreover, Pr is computed as (Cengel, 2003):

$$Pr = \frac{\mu C_p}{\lambda_{water}}, \quad (3)$$

where C_p [J/kg/K] is the specific heat of the water and λ_{water} [W/m/K] the thermal conductivity of the water. De Pasquale et al. (2017) modified the *undisturbed soil model* by including a term for the thermal resistance of the surrounding soil, R_s , and changed the resistance of the pipe wall to account for a cylindrical shape, termed the *shape factor model*:

$$K_{p,2} = (R_p + R_w + R_s)^{-1} = \left(\frac{\ln\left(\frac{r_o}{r_i}\right)r_i}{\lambda_{pipe}} + \frac{1}{a} + \frac{2\pi r_i}{\lambda_{soil} S} \right)^{-1}, \quad (4)$$

where r_i [m] and r_o [m] are the inner and outer radii of the pipe, respectively, λ_{soil} [W/m/K] is the thermal conductivity of the surrounding soil, and S is the conduction shape factor of the pipe as described in Aziz (2003):

$$S = \frac{2\pi}{\cosh^{-1}\left(\frac{z}{r_o}\right)}, \quad (5)$$

where z [m] is the burying depth of the pipe. Also, De Pasquale et al. (2017) took into account the disturbance caused by the pipe to the surrounding soil (Krarti and Kreider, 1996), termed the *De Pasquale model*:

$$K_{p,3} = (R_p + R_w + R_s)^{-1} = \left(\frac{\ln\left(\frac{r_o}{r_i}\right)r_i}{\lambda_{pipe}} + \frac{1}{a} + \frac{\ln\left(\frac{l+r_i+d_{wall}}{r_i+d_{wall}}\right)r_i}{\lambda_{soil}} \right)^{-1}, \quad (6)$$

where d_{wall} [m] describes the pipe wall thickness and l [m] describes a fictitious soil layer thickness increasing the thermal resistance:

$$l = \sqrt{\frac{\left(\frac{\lambda}{\rho C_p}\right)_{soil}}{\omega}} \quad (7)$$

Here, $c_{p,soil}$ is the heat capacity of the soil [J/kg/K], ρ_{soil} the density of the soil [kg/m³] and ω the angular frequency [s⁻¹] of a sinusoidal function describing the soil temperature. In De Pasquale et al. (2017) this model was assumed to be superior to the *shape factor model*, as it was capable of incorporating the delay between a pipe's heat flux and the ambient temperature. Hubeck-Graudal et al. (2019), however, found that the *shape factor model* produced similar results to the *De Pasquale model*.

A.1 Parametric analysis

A parametric analysis was conducted to highlight major differences in the four temperature models: *undisturbed soil model*, *shape factor model*, *De Pasquale model* and *Hubeck-Graudal model*. In the analysis, a simple water distribution network (WDN) model was applied consisting of one pipe supplying water to a single consumer. Inspired by the work of De Pasquale et al. (2017), the hydraulic model parameters used in all four temperature models are summarized in Table A-1.

Table A-1. Model set-up parameters, inspired by De Pasquale et al. (2017). For examples of thermal conductivity of selected pipe materials and soils see Czichos (2000) and Blokker and Pieterse-Quirijns (2013). Experimental volumetric heat capacity based on Pagola et al. (2017).

Parameter	Symbol	Value	Unit
Inlet temperature	T_{inlet}	12	°C
Soil temperature	T_{soil}	5.71	°C
Flow/demand	Q	0.01	m ³ /s
Pipe length	L	1000	m
Pipe Diameter	D	0.1	m
Pipe burying depth	z	1.5	m
Thermal conductivity of cast iron	λ_{pipe}	45	W/m/K
Thermal conductivity	λ_{soil}	1.89	W/m/K
Volumetric heat capacity of soil	ρc_p	2.61	MJ/m ³ /K
Angular frequency	ω	1.99E-07	s ⁻¹

In the parametric analysis, the effect of changing a set of pipe materials, diameters and flow values vs. an increasing pipe length is analysed according to the values listed in Table A-2.

Table A-2. Parametric analysis of water temperature models based on varying pipe material, pipe diameters and flow, where $T_{inlet} = 12$ °C and $T_{soil} = 5.71$ °C. Inspired by De Pasquale et al. (2017). For assumptions on remaining model parameters see Table A-1.

Parameter set-up	Pipe material [W/m/K]	Diameter [mm]	Flow [L/s]	Soil material [W/m/K]
(a) Pipe material	Cast iron: 45; polyethylene: 0.4	100	10	1.89
(b) Pipe diameter	Cast iron: 45	200; 500	10	1.89
(c) Flow	Cast iron: 45	350	25; 75	1.89
(d) Soil material	Cast iron: 45	100	10	1;4

Figure A-1 and Figure A-2 display the simulated water temperature based on different pipe set-ups (Table A-2). The simulated temperature differs significantly only between the *undisturbed soil model* and the remaining three models. The *undisturbed soil model* reaches the soil temperature significantly faster, because a term for the thermal resistance of the surrounding soil is missing. The figures also show that it can be assumed that modelling differences obtained from the three remaining models can be assumed to be negligible.

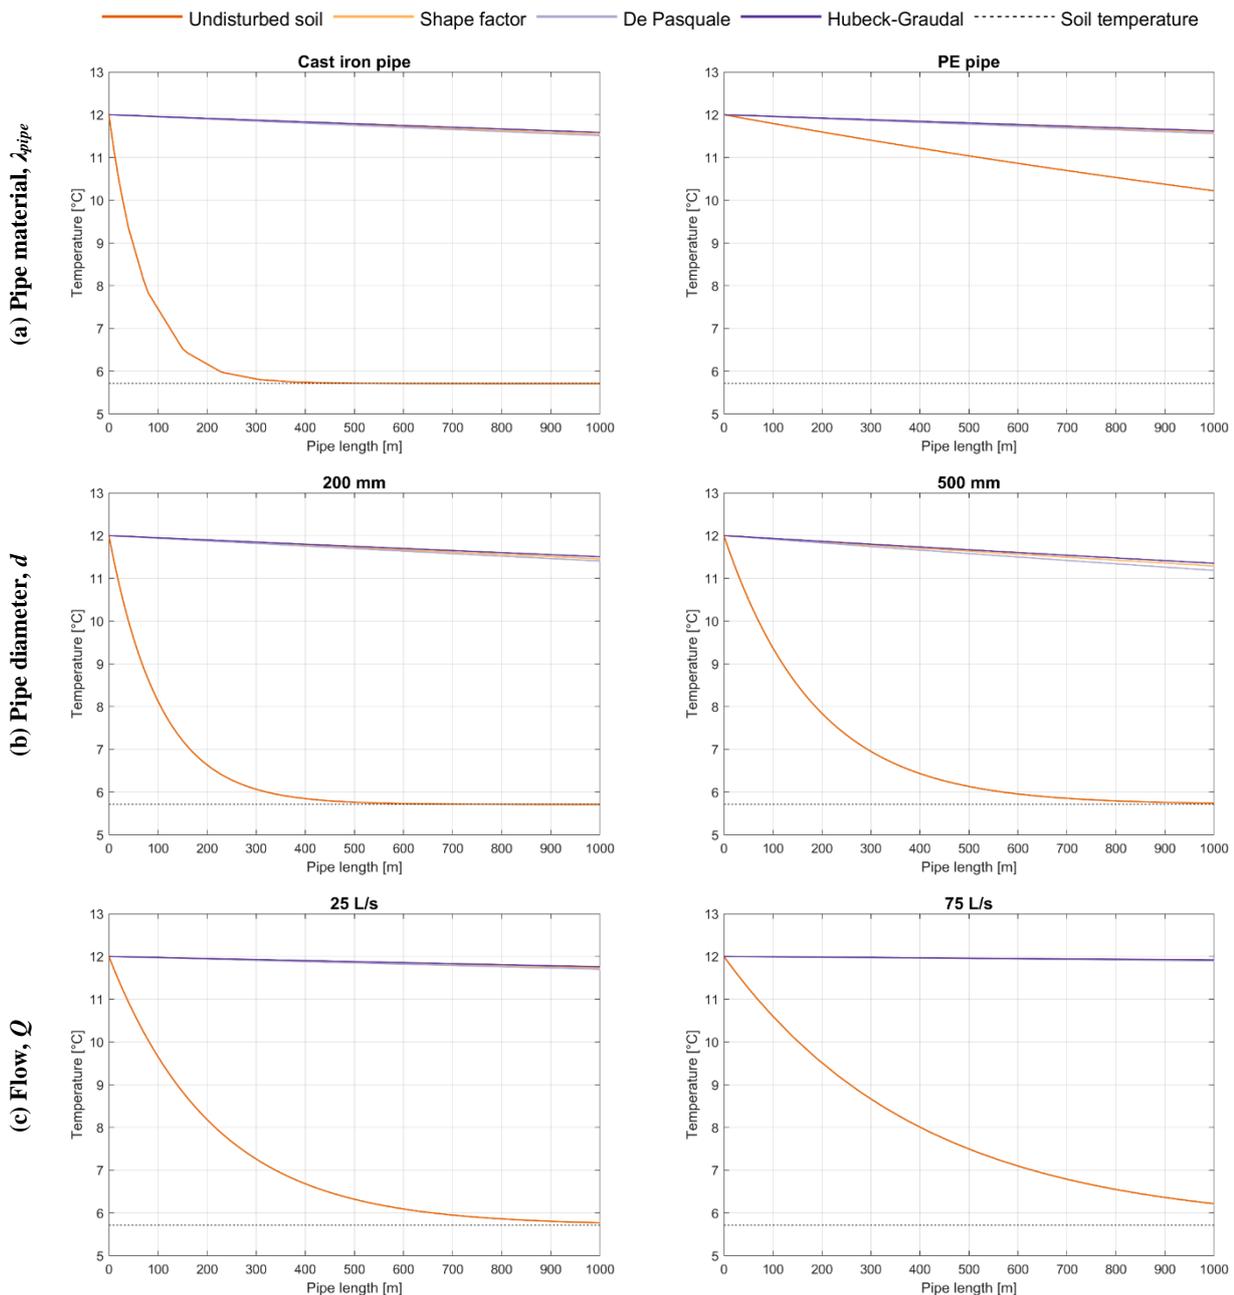


Figure A-1. Temperature profile of four different temperature models based on different set-ups, inspired by De Pasquale et al. (2017). For individual parameters of each set-up, see Table A-2.

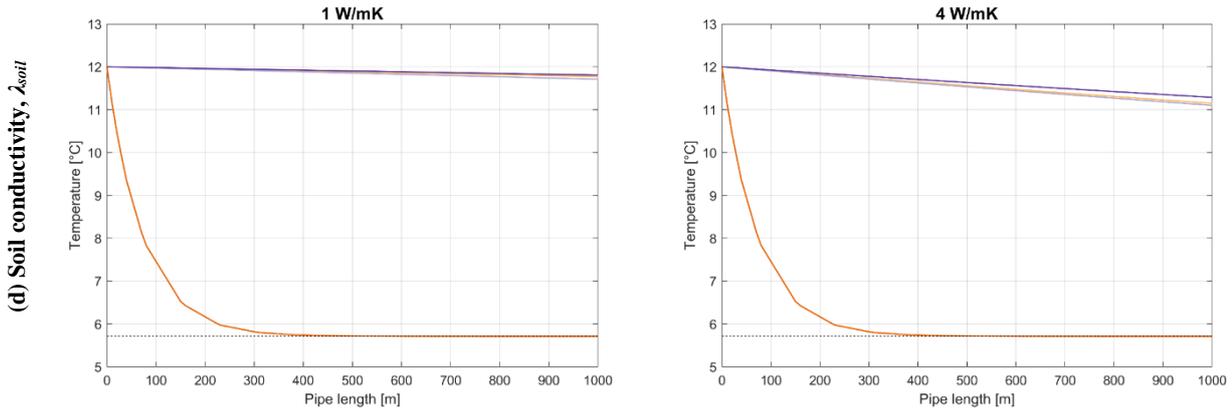


Figure A-2. Temperature profile of four different temperature models based on different set-ups, inspired by De Pasquale et al., (2017). For individual parameters of each set-up, see Table A-2. In (d), no λ_{soil} is included in the undisturbed soil model.

B. Effect of individual parameters on Hubeck-Graudal model

Section 2.1 of the main paper

In this paper, the *Hubeck-Graudal model* was used for simulating the WDN water temperature. Regardless of the selected temperature model, a substantial knowledge about the real WDN and the soil environment of the WDN is required to achieve an accurate WDN temperature model. Examples include data about the pipe characteristics, such as burying depth, material, inner and outer diameter, thermal conductivity and the temperature of the surrounding soil. We applied the *Hubeck-Graudal model* on a test set-up (Table A-1) with one pipe supplying water to a single consumer. Based on this set-up, the effect of varying three selected parameters (λ_{soil} , λ_{pipe} and d) vs. the soil temperature on the outlet water temperature is illustrated in Figure B-1. Similar to Zlatanovic et al. (2017) (who assessed a domestic water temperature model), the analysis revealed that the thermal resistance of the surrounding soil, R_s , represents between 80% and 99% of the overall thermal resistance, depending on whether PE (0.45 W/m/K) or cast iron pipes (45 W/m/K) were applied.

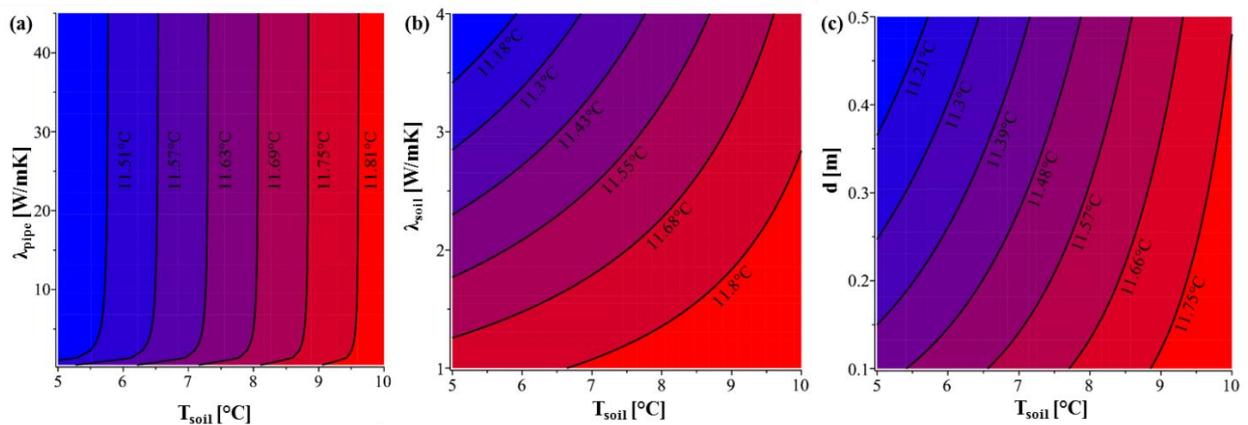


Figure B-1. Water temperature variation in the *Hubeck-Graudal model*, based on varying soil temperature and three parameters: soil conductivity (λ_{soil}), pipe conductivity (λ_{pipe}), and pipe diameter (d). Standard test set-up parameters are listed in Table A-1.

Only pipes having a low thermal conductivity (<5 W/m/K, such as PE or concrete pipes) have a major insulating effect and thus decreasing the rate of cooling of the water (Figure B-1a). Figure B-1b shows that an increase in soil conductivity, λ_{soil} , also leads to a lower water temperature. Figure B-1c shows that increasing diameters lead to a slight decrease in water temperature, as the retention time of the water increases.

C. Valve search space reduction

Section 2.2 of the main paper

Table C-1. Pseudo code/algorithm for ‘valve clustering’ segments.

Require: net – network model; V_d – the index of valves/pipe segments not closable as being deemed ‘dead end’; V_s – the number of valves/pipe segments not closable as being deemed in the category ‘source disconnection’; $thresh$ – proximity threshold for clustering.

Ensure: V_c – the number of valves/pipe segments not closed during calibration as being deemed in the category ‘valve clustering’

Based on net , establish the graphs $G = (V, E)$ with V and E representing the nodes and links of net , respectively.

Remove all edges e from G , being marked as not closable in V_s and V_d and all edges linked to vertices v with an assigned demand.

Store the $deg(v)$ of G in M .

Remove all edges e from G that are not connected to at least one vertex with $deg(v) = 2$

Separate G at each vertex v where $M(v) > 2$.

G consists now of multiple connected graph components cgc (only straight line segments).

for each cgc do

run hierarchical agglomerative clustering with the minimum proximity technique between vertices. Clusters are formed at heights $\leq thresh$

Remove the most central edge from the cluster (this edge is used for valve calibration)

Store the remaining edges from the cluster in V_c

end do

All valve search space reduction approaches introduced in the main paper utilize simple graph theory. In short, the network topology can be represented as an undirected graph $G = (V, E)$. G consists of a set V of vertices v (nodes) and set E of edges e (links). Each edge is attached to two vertices. The degree $deg(v)$ of a vertex defines the number of ingoing and outgoing connections from a vertex. For example, automatically, all edges e containing a vertex with $deg(v) = 1$ cannot be fully closed if the vertex represents a consumer, as it would otherwise disconnect the node from the network. In some cases, we applied the breadth-first search (BFS) algorithm (e.g. Cormen et al. (2003)) to traverse a graph from a certain vertex until a stop criterion was met. In the valve clustering process, we applied agglomerative hierarchical clustering. For additional information see, for example, Tan, Steinbach, and Kumar (2006). The pseudo code of our implementation of ‘valve clustering’ segments is listed in Table C-1.

D. Temperature in selected nodes (open vs. closed)

Section 3.1 of the main paper

Figure D-1 illustrates the simulated temperature in six nodes scattered around the network when all valves are open (yellow) or five valves are closed (blue) in the district metering area (DMA) case study.

The nodes represent a broad overview of temperature differences observed throughout the DMA network over the course of a week. Some nodes have higher deviations over the entire week, with respect to all valves being open and five being closed. This is the case for node A, where a high difference in the simulated temperatures with up to 1 degree difference is seen. Other nodes, such as nodes C and E, have the largest differences on selected days around and after the 8th of March, respectively. In some nodes, such as node D, valve closing has no effect on the temperature profile.

Figure D-2a shows the measured temperature over the course of 2016 at DMA3 and DMA5. Moreover, the temperature measured in a week in March 2016 can be seen in Figure D-2b. Here, the soil temperature is more or less constant around 7 °C, used in the DMA case study.

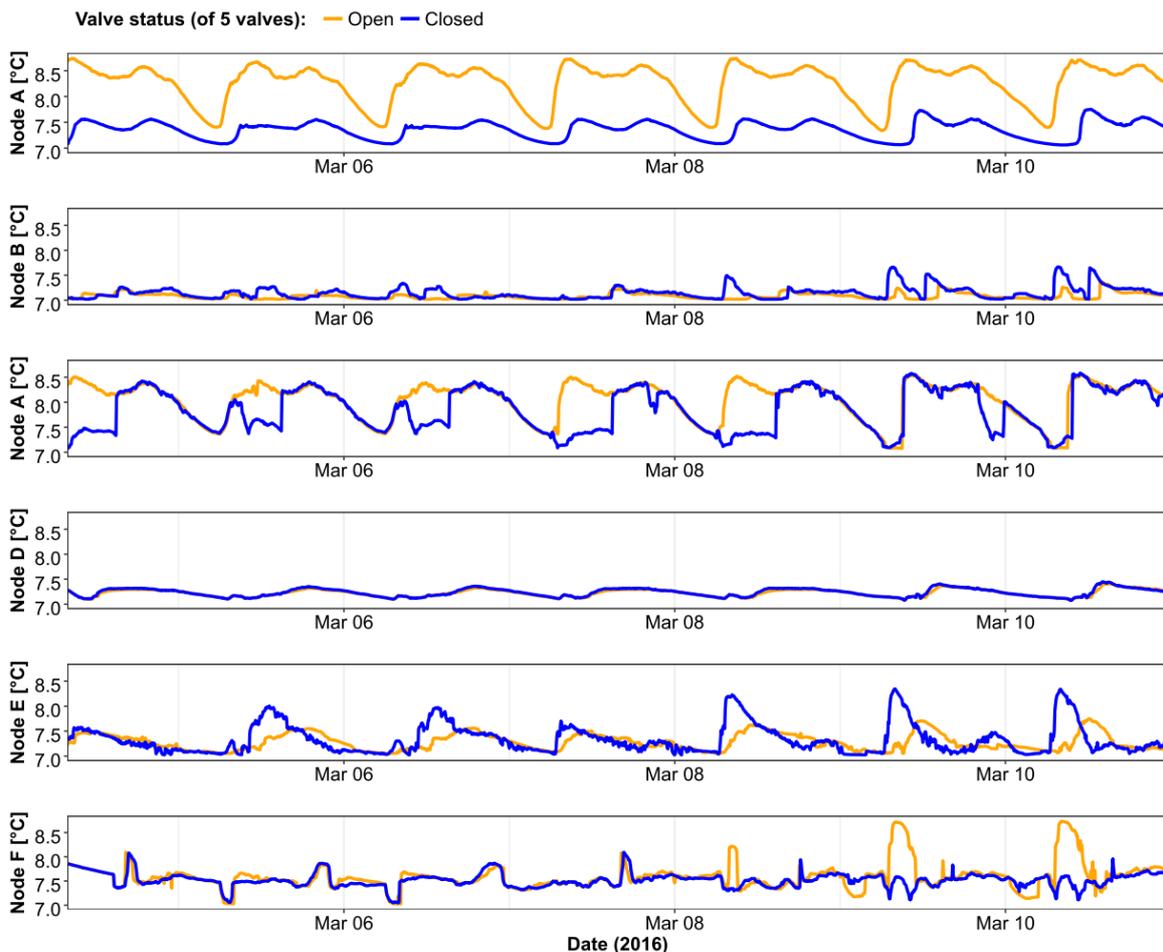


Figure D-1. Simulated water temperature in selected nodes of the DMA model with all valves open (orange) and five valves closed (blue).

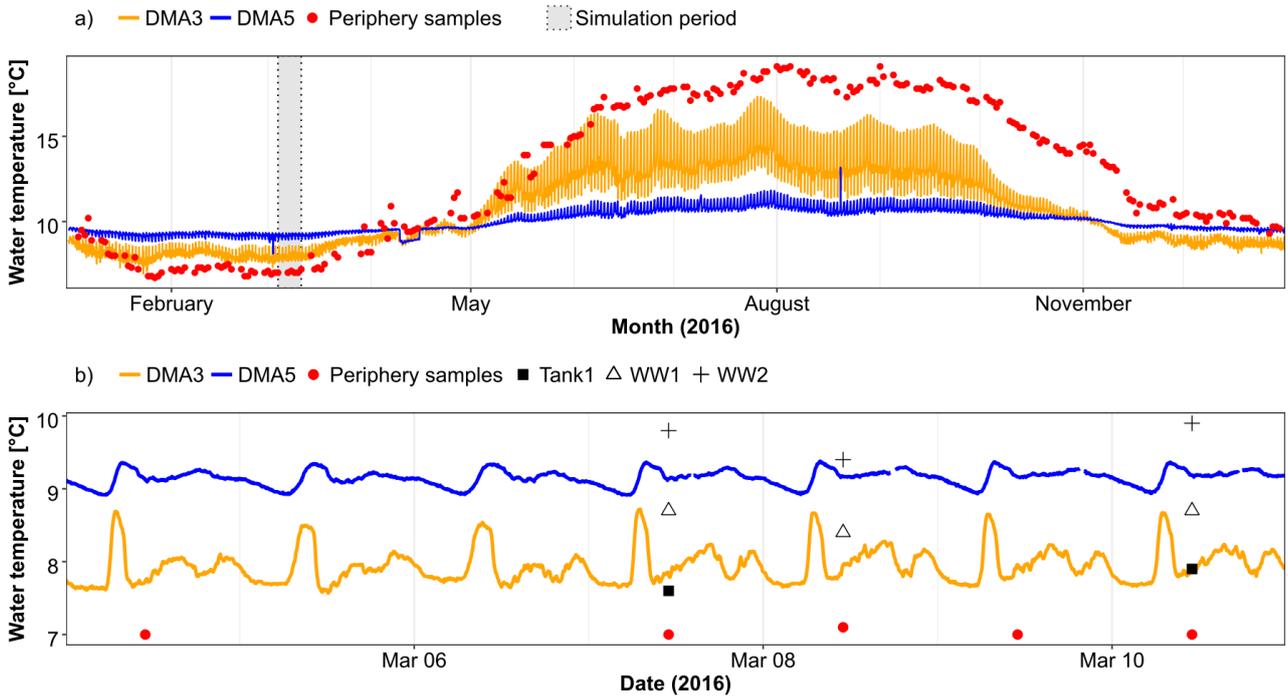


Figure D-2. Measured water temperatures at DMA3 and DMA5 including samples taken at WW1, WW2 and Tank and samples taken at the periphery of Copenhagen water distribution network assumed representative for the soil temperature.

E. Genetic algorithm results

Section 4.1 of the main paper

The mean and median best fitness values based on various initial population scenarios are summarized in Table E-1.

Table E-1. Average best fitness function value (median in parentheses) based on different initial genetic algorithm population and two different cross over rates.

Fitness function	1% <i>Random</i>	1% <i>Weighted</i>	2.5% <i>Random</i>	2.5% <i>Weighted</i>	5% <i>Random</i>	5% <i>Weighted</i>
F ₁	698.97 (6.94)	128.46 (50.98)	105.18 (44.61)	132.51 (10.47)	446.97 (367.7)	0.00 (0)
F ₂	1368.62 (94.79)	51.30 (5.94)	44.40 (5.94)	89.75 (82.85)	397.70 (434.4)	5.94 (5.94)

Even though the 5% *weighted* scenario led to the best results in both fitness functions, no clear trend in selecting the most suitable initial percentage of closed valves is visible from the scenarios listed in Table E-1. The median value indicates that there is a high variability in the results found, depending on the cross-over fraction to mutation rate (CF:MR) ratio and population size. Moreover, except for the 2.5% *random* scenario, the *weighted* scenario tends to have led to better fitness values.

The best fit development of each individual genetic algorithm (GA) run is summarized in Figure E-1. Figure E-1a shows the development over each generation based on fitness function F₁, while Figure E-1b shows that for fitness function F₂. Within the first generations, the *weighted* scenarios converge faster towards a lower fitness value than the *random* scenarios. This can partially be explained by the

fact that the weighted scenario has a better initial best fit. Moreover, it should be kept in mind that these better starting conditions are based on the computation of 380 simulations required to compute the SM values that were needed to generate the initial populations. If time is a constraint, selecting the *weighted* scenario seems to be a better selection than selecting the *random* approach.

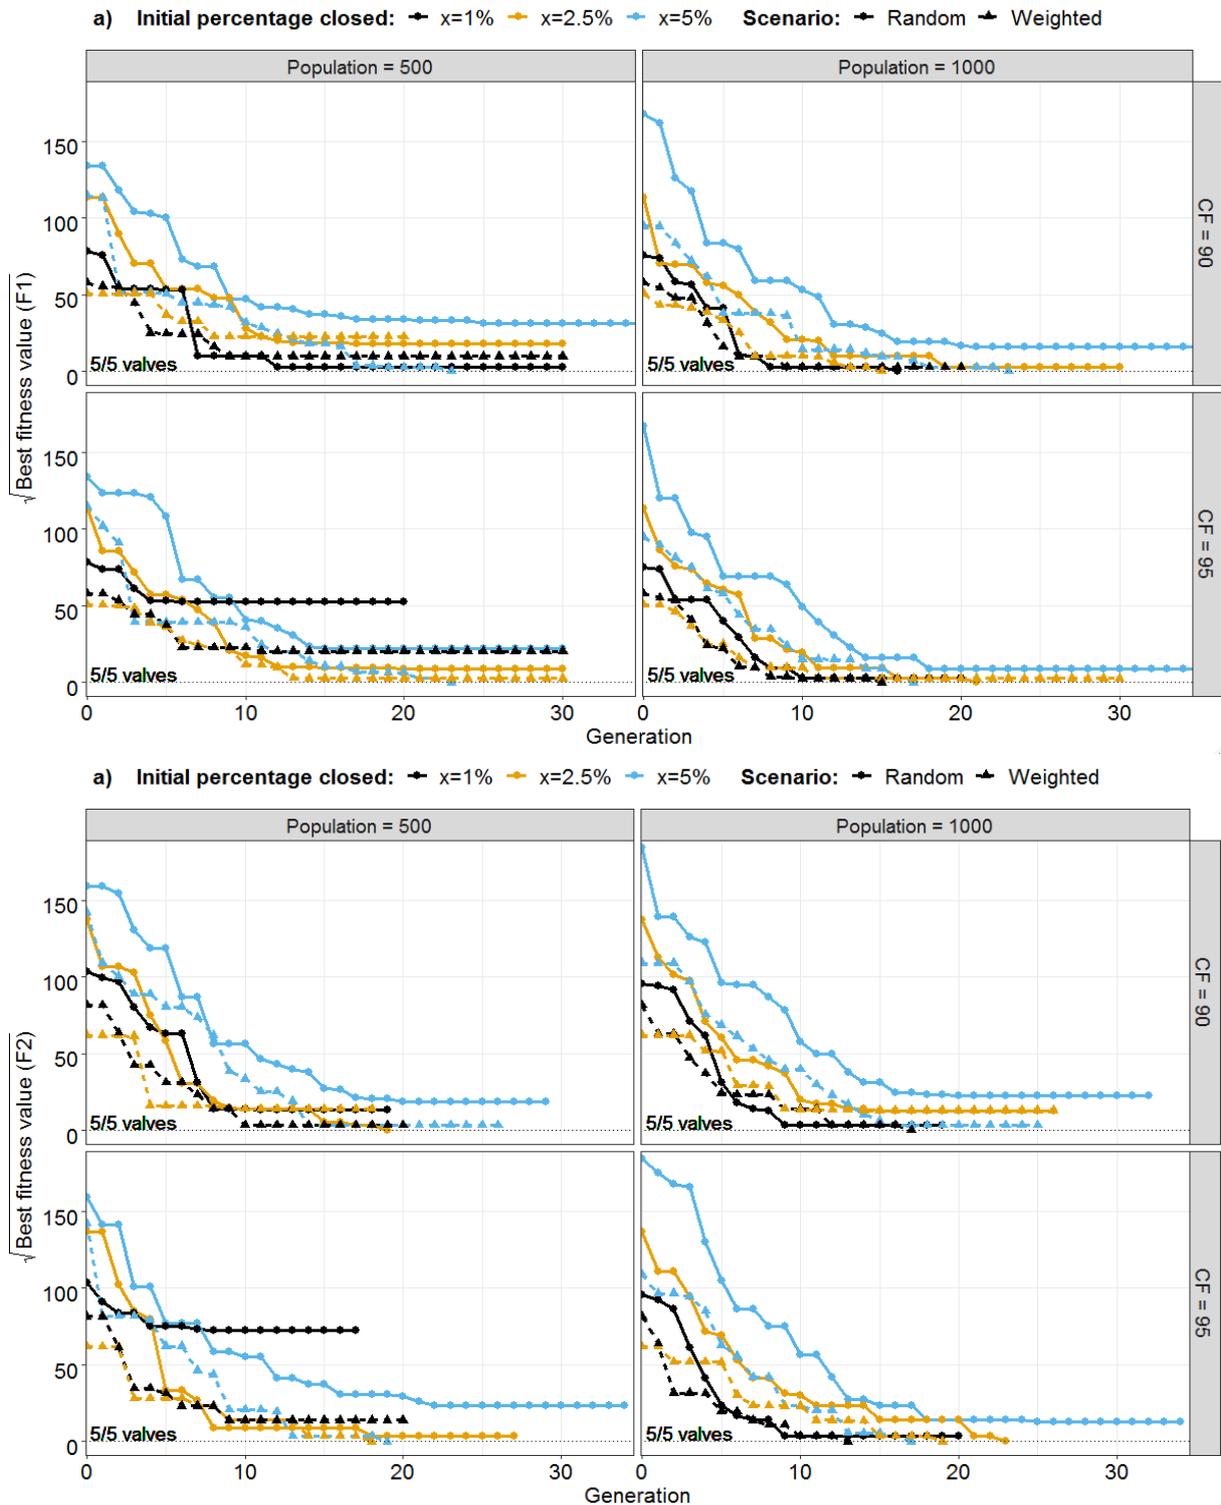


Figure E-1. Best fit of DMA model with various genetic algorithm (GA) set-ups; a) shows the GA runs with fitness function F_1 ; and b) the GA runs with fitness function F_2 . CF = Cross-over fraction. MR = mutation rate.

F. Example of valve search space reduction

Section 4.2.3 of the main paper

The reduced valve search space for the transportation case study is summarized in Table F-1 and illustrated in Figure F-1. Large parts connected upstream of WW1 have no consumption and can thus be removed from the calibration process (i.e. ‘dead ends’). Also, large parts downstream of WW2 are branched and can thus not be closed as they will disconnect consumers from a water source. In particular, pipe segments of the parallel pipes running between WW1 and Tank1 have been clustered (Figure F-1).

Table F-1. Valve search space reduction process.

Definition	Number of valves
#0 All pipes can be closed	231
#1 ‘Dead end’ stretches	-35
#2 No source disconnection	-59
#3 Cluster no demand stretches	-31
Total analysed valves in model	106

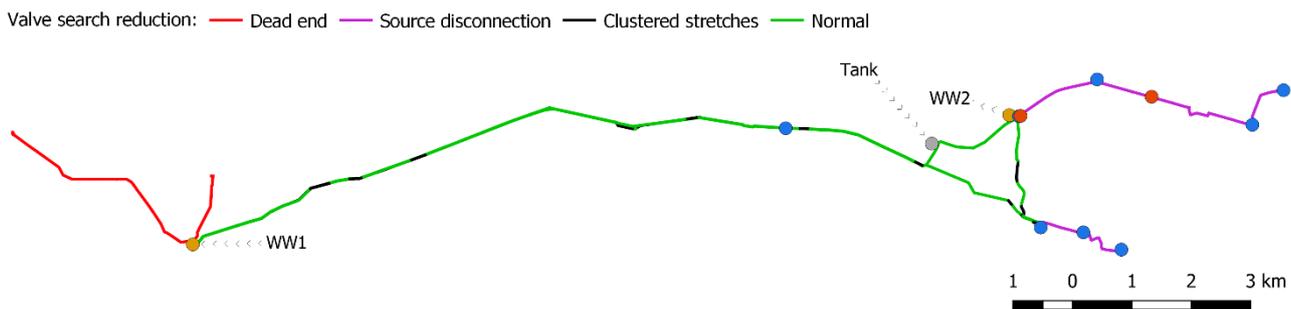


Figure F-1. Pipe segments not used for valve status identification, i.e. segments with no consumption and flow (dead ends), disconnecting consumers from a source and stretches clustered based on their proximity of 500m.

G. Additional results

Section 4.2.3 of the main paper

Figure G-1 shows the simulated temperature results at DMA5, when all valves are open (green), based on the best fit generated during GA calibration (blue) with a weighting of $W_T = 100\%$, the utility’s network setup (yellow) and the observation (black).

Only the valve setting of ‘all valves open’ has a minor effect on the simulated water temperature at DMA5 compared to the other two model settings. In terms of the simulated head, all three model setups follow the same pattern.

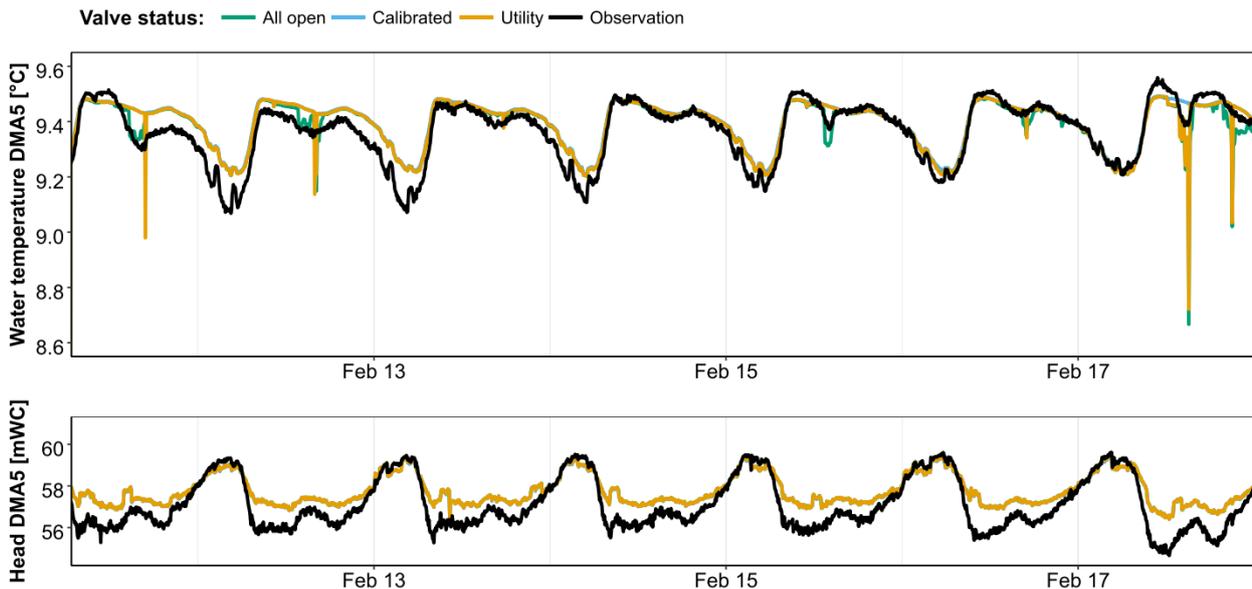


Figure G-1. Temperature and head at DMA5 based on various valve scenarios. All valves open (green); the model has only been calibrated with a weighting factor of 100% temperature measurements (blue); with the utilities knowledge of opened and closed valves (yellow); and the measured temperature at the location (black). Also, the measured and modelled head at the location are shown in the lower figure.

H. References

- Aziz, A., 2003. Conduction Heat Transfer, in: Bejan, A., Kraus, A.D. (Eds.), Heat Transfer Handbook. John Wiley & Sons, Inc., Hoboken, New Jersey.
- Blokker, E.J.M., Pieterse-Quirijns, E.J., 2013. Modeling temperature in the drinking water distribution system. *J. Am. Water Work. Assoc.* 105, E19–E28. doi:10.5942/jawwa.2013.105.0011. Erratum (2018) in *J. Am. Water Work. Assoc.* 110, 98-98. doi:10.1002/awwa.1178
- Cengel, Y.A., 2003. Heat Transfer: A Practical Approach. Mc Graw-Hill, Boston.
- Cormen, T.H., Leiserson, C.E., Rivest, R.L., Stein, C., 2003. Introduction to algorithms. MIT Press, New York.
- Czichos, H., 2000. Werkstoffe, in: Die Grundlagen Der Ingenieurwissenschaften. Springer Berlin Heidelberg, Berlin, Heidelberg, pp. 511–600. doi:10.1007/978-3-662-06652-2_5
- De Pasquale, A.M., Giostri, A., Romano, M.C., Chiesa, P., Demeco, T., Tani, S., 2017. District heating by drinking water heat pump: Modelling and energy analysis of a case study in the city of Milan. *Energy* 118, 246–263. doi:10.1016/j.energy.2016.12.014
- Hubeck-Graudal, H., Kirstein, J.K., Ommen, T., Rygaard, M., Elmegaard, B., 2019. Drinking water supply as low-temperature source in the district heating system: a case study for the city of Copenhagen (submitted).
- Krarti, M., Kreider, J.F., 1996. Analytical model for heat transfer in an underground air tunnel. *Energy Convers. Manag.* 37, 1561–1574. doi:10.1016/0196-8904(95)00208-1
- Pagola, M.A., Jensen, R.L., Madsen, S., Poulsen, S.E., 2017. Measurement of thermal properties of soil and concrete sample (No. 235), DCE Technical Report. Aalborg. Available at [https://vbn.aau.dk/ws/portalfiles/portal/266378485/Measurement_of_thermal_properties_of_soil_and_concrete_s_amples.pdf](https://vbn.aau.dk/ws/portalfiles/portal/266378485/Measurement_of_thermal_properties_of_soil_and_concrete_samples.pdf) (accessed November 12, 2019).
- Tan, P.N., Steinbach, M., Kumar, V., 2006. Introduction to Data Mining, Pearson International Edition. Pearson Addison Wesley, Boston.
- Zlatanovic, L., Moerman, A., van der Hoek, J.P., Vreeburg, J., Blokker, M., 2017. Development and validation of a drinking water temperature model in domestic drinking water supply systems. *Urban Water J.* 14, 1031–1037. doi:10.1080/1573062X.2017.1325501

IV

Using smart meter temperature and consumption data for water distribution system analysis

J.K. Kirstein, K. Høgh, M. Rygaard and M. Borup

Using smart meter temperature and consumption data for water distribution system analysis

Jonas Kjeld Kirstein^{1,*}, Klavs Høgh², Martin Rygaard¹ and Morten Borup¹

*corresponding author

¹DTU Environment, Technical University of Denmark, Kgs. Lyngby, Denmark. jkir@env.dtu.dk; mryg@env.dtu.dk; morb@env.dtu.dk.

²NIRAS A/S, Allerød, Denmark; kvh@niras.dk.

Abstract Smart meter temperature data are increasingly collected in water supply systems, but the utilities are unaware of potential successful applications. In this study, smart meter consumption and temperature data from a district metered area (DMA) in Denmark was analysed for various potentials. First, application of a consumption-based filtration method categorized the data into estimates of soil and water distribution network temperatures. Soil temperatures deviating highly from the mean revealed smart meters with incorrect service line information stored in the utility's asset database. Simulating the temperature throughout the DMA and comparing it with the smart meter temperatures showed a good resemblance. One node, however, with a high error, revealed a consumer that was incorrectly affiliated to the DMA, which is otherwise not easily identified from common hydraulic data or the billing system. Future applications with finer sampling resolutions are expected to reveal further potentials of smart meter temperature, e.g. useful within the fields of leakage detection, pipe connectivity and valve status identification.

Keywords: District metered area; Hydraulic modelling; Smart meter; Temperature; Water distribution network

1 Introduction

Smart meters measuring the consumers' water consumption and automatically transmitting the data to the utility are currently being implemented many places around the world. Temperature data may become a valuable 'by-product' from smart meters. The temperature measurements are used in the computation of volume observations and are thus available for other purposes. Smart meter manufacturers advertise the use of temperature data for alarming consumers about the risk of frost, and the possibility of being used as a tool for identifying areas with stagnant water owing to higher water temperatures (e.g. Rokkjær, 2018). In practice, however, smart meter temperatures are not mentioned when identifying benefits of digital metering (Boyle et al., 2013; Monks et al., 2019) and utilities are unsure what to do with the collected temperature data (e.g. Blokker, 2019). In Kirstein et al. (2019) it was shown that temperature data directly collected from the water distribution network (WDN) can have a value in identifying system anomalies, such as incorrect information stored in a utility's asset database, as well as in increasing the understanding of how the water is distributed. However, one of the limitations in the study was the limited amount of temperature sensors and low spatial distribution of the available data throughout the WDN. Here, smart meter data can play a substantial role in closing this gap, as temperature data may be available from each household.

Unfortunately, smart meter temperatures ($T_{smart\ meter}$) do not directly represent WDN temperatures, as the water temperature is influenced by various factors on its way to the consumer (Figure 1).

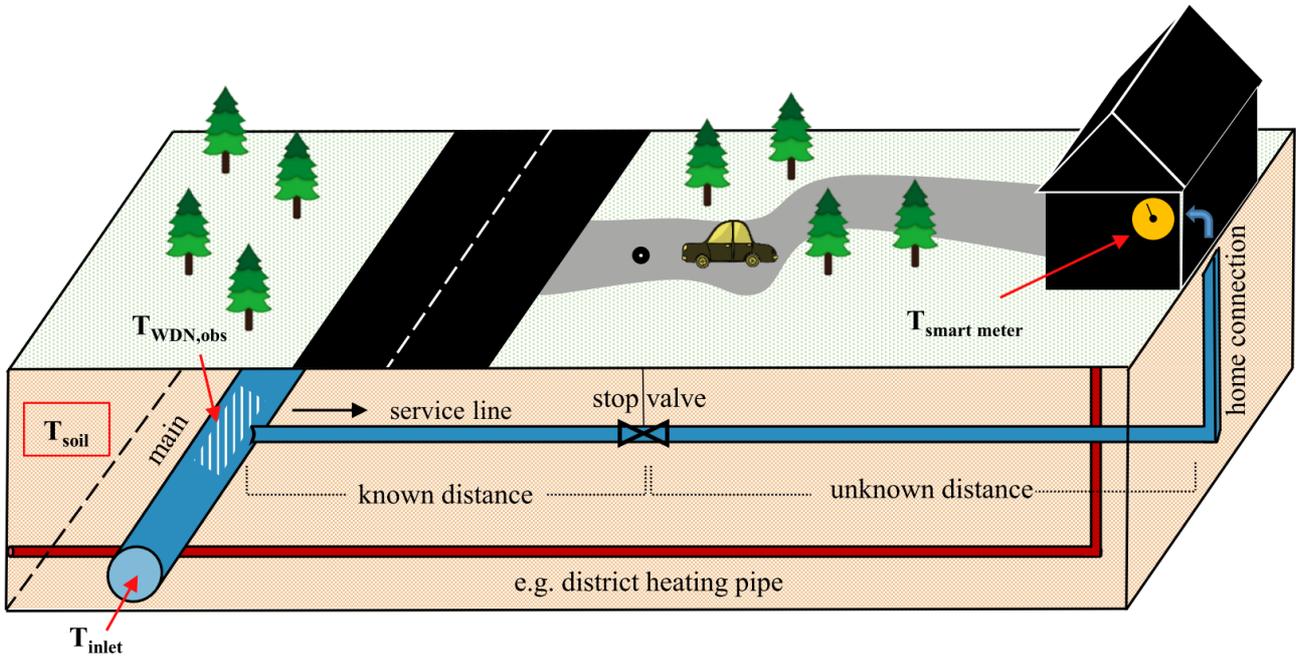


Figure 1. Cross-sectional overview of a water distribution network main connecting a household via a service line. External factors (e.g. vegetation, pavement or district heating pipes) affect the local soil temperature. WDN = water distribution network.

Water stagnates in the flowmeter during times of no consumption and approximates the ambient (e.g. indoor) temperature. At the same time, the water inside the service line approaches the surrounding soil temperature (T_{soil}), which may further be influenced by local features in the area such as pavement, vegetation, district heating and urban drainage pipes (Figure 1). However, if the consumption is high enough, the water has only spent a brief period inside service lines and home connections and could thus be used as a proxy for the WDN temperature ($T_{WDN,obs}$, Figure 1). A common service line of a single-family house has an inner diameter of 33 mm (e.g. Wavin, 2019) and an average length of 17 m (Table 1), thus holding a total volume of around 15 L. According to a Danish utility (Vejen Forsyning, 2014), a 5 min shower with a water saving shower head uses between 30–40 L. Thus, the water inside the service line is replaced more than twice during such an event and a smart meter temperature sample should be able to represent $T_{water,obs}$ (Figure 1) if sampled at the right time. Not only the simulation, but also the analysis of the temperature data in itself has many potential benefits. For example, during times of no consumption the water temperature in service lines approximates the surrounding soil temperature. A sudden demand may push this water volume to the meter, and the measured temperature at this point in time may be a good estimate of the soil temperature.

1.1 Potentials of smart meter temperature data

We suggest to explore the potentials for analysing WDNs based on a heat transfer model by comparing measured smart meter temperature samples with simulated temperature values. The

analysis and simulation of temperature data from smart meters may have multiple benefits, summarized in the following.

Water quality. Owing the large number of possible entry points for contaminants, it is challenging to monitor water quality of WDNs reliably (Eggimann et al., 2017; Makropoulos and Savić, 2019). Temperatures affect bacterial growth in the WDN (Liu et al., 2016) and are of importance when selecting preferable locations for water quality audits, as they may indicate areas of stagnant or slow-flowing water (e.g. Larsen et al., 2017). Moreover, temperatures can affect the generation of disinfection by-products (e.g. Kyriakou et al., 2016). Thus, the spatial availability of smart meters and temperature samples from each household may improve water quality estimates throughout the WDN.

WDN analysis & asset management. Water temperatures can indicate the origin of the water, otherwise not easily seen from flow and pressure data (Kirstein et al., 2019) or requiring exhaustive tracer studies. Moreover, temperature data can be used as an additional parameter for calibration of WDNs, such as valve status identification where valve closures affect the temperature distribution throughout the WDN (Kirstein et al., 2019). Analysis of the smart meter temperature data may reveal measurements deviating highly from the majority of samples in an area, for example, useful for highlighting wrongful service pipe characteristics (e.g. length or affiliation) and wrongly affiliated consumers.

Water loss management. Unauthorized consumption and leakages increase the demand in areas and affect the time the water has spent in the WDN before being consumed. As time is a major factor of the resulting heat transfer, discrepancies between simulated and observed temperatures could indicate leakage locations as well as unaccounted-for consumers.

Consumer benefits. Consumers should be alerted in periods of too high temperatures and be advised to flush the water at the tap at the next time of usage. Likewise, consumers can be alerted when too low temperatures are measured, increasing the consumers' chance to prevent frost bursts (e.g. Rokkjær, 2018)).

Energy management. Heat pumps deployed in WDNs can be used to extract energy and deliver it to heating systems (Blokker et al., 2013; De Pasquale et al., 2017; Hubeck-Graudal et al., 2019). Here, smart meter temperature data could, for example, be used to enhance estimates of the heat transfer occurring in service lines, potentially advocating for the deployment of heat pumps.

Data improvement. Accumulated volumes from smart meters, e.g. sampled each full hour, do (most of the time) not represent the actual time of consumption, such as samples taken at lower time intervals. However, the temperature of the water inside the meter at time of transmission may be used as a proxy for indicating how long the water has been standing still since the last time of consumption, thus potentially capable of improving knowledge on the actual time of consumption.

Here, we present a way to use smart meter consumption and temperature data for WDN analysis based on data from two DMAs in Utility Brønderslev, Denmark. In this process we 1) categorize temperature data into samples representing soil temperatures and WDN temperatures, respectively.

2) Simulate the water temperature throughout the DMA and evaluate the simulated values against measured smart meter temperatures.

2 Case study

The utility of Brønderslev, Denmark, has completed a full smart meter rollout. Two DMAs from this utility were selected for detailed analysis (Table 1). DMA1 is located in an urban area with mainly single-family houses, whereas DMA2 is located in a rural area with longer distances between consumers and agricultural users. Smart meter data is sampled ‘quasi-randomly’ in the utility, as the smart meters’ data is transmitted through a custom data network consisting of multiple local cell-towers, that pick up the meters’ signal and transmit once an hour a single data point for each meter to the utility’s database. As samples may reach multiple cell-towers, it is possible for certain smart meters to reach sampling resolutions of around 8 samples every hour (Table 1).

The temperature model is only run in DMA1. The temperature at the DMA1 inlet (T_{inlet}), pressure and flow are measured in a 5 min interval. In the hydraulic model, pressure and T_{inlet} were used as boundary conditions. Moreover, the gap between each smart meter reading was linearly interpolated to have uniform sampling resolutions at a 5 min resolution available. Each smart meter consumption time series was then bundled in the nearest nodes of the hydraulic model. The hydraulic model was run in EPANET (Rossman, 2000).

Table 1. Consumer and network information of two district metered areas (DMAs) in Brønderslev for two weeks in August 2018. *Smart meters with single readings were omitted.

Name	DMA1	DMA2
Smart meters	526	579
Number of samples	393,979	318,956
Median annual consumption [L/day]* (Min/Mean/Max)	193 (1/254/6,470)	228 (0/487/22,126)
Median sampling resolution [min/sample/device]* (Min/Mean/Max)	31 (23/49/4,853)	40 (8/258/6,934)
Median sampled temperature [°C] (Mean)	24 (25)	20 (20)
Σ Mains [km] (Mean)	10 ($73 \cdot 10^{-3}$)	85 ($223 \cdot 10^{-3}$)
Mains diameter [outer diameter in mm] [%]: $\leq 80/81-120/>120$ /Unknown	16/76/8/0	34/59/4/3
Mains material [%]: PVC/PE/Other	80/14/6	57/37/6
Service lines [outer diameter in mm] [%] 32/40/50	<5/95/<1	-
Σ Service lines [km] (Mean per smart meter)	9 ($17 \cdot 10^{-3}$)	-

3 Method

3.1 Heat transfer model

The heat transfer model described in (Hubeck-Graudal et al., 2019) was used for simulating the change in water temperature, $T_{water,sim}$ [°C], over time, t [s], throughout the WDN:

$$\frac{dT_{water,sim}}{dt} = \frac{K_p}{c_{p,water} \cdot \rho_{water} \cdot \pi \cdot r_i^2} \cdot (T_{water,sim} - T_{soil}) \quad (1)$$

T_{soil} describes the undisturbed soil temperature [°C] at pipe burying depth, $c_{p,water}$ the specific heat of the water [J/kg/K], ρ_{water} the density of the water [kg/m³] and r_i the inner radius of the analysed pipe [m]. This non-transient heat model is implemented in the EPANET Multi-Species eXtension (MSX)

(Shang et al., 2007). The heat transmission coefficient K_p [W/m/K] is subject to some variations in the literature (Blokker and Pieterse-Quirijns, 2013; De Pasquale et al., 2017). We have chosen to use the term as described in Hubeck-Graudal et al. (2019):

$$K_p = (R_{soil} + R_{pipe} + R_{water})^{-1} = \left(\frac{\ln\left(\frac{4 \cdot H}{2r_o}\right)}{2 \cdot \pi \cdot \lambda_{soil}} + \frac{\ln\left(\frac{r_o}{r_i}\right)}{2 \cdot \pi \cdot \lambda_{pipe}} + \frac{1}{\alpha_{water} \cdot \pi \cdot 2 \cdot r_i} \right)^{-1} \quad (2)$$

With r_o being the outer radius of the pipe [m], λ_{soil} and λ_{pipe} the thermal conductivity of soil and pipe material [W/m/K], respectively and H [m] representing a correction factor for the convective resistance at the soil surface:

$$H = h + \varphi \cdot \lambda_{soil} \quad (3)$$

With h being the height of soil [m] above the pipe centre and φ describing a correction factor equal 0.07 [K·m²/W]. Furthermore, α_{water} describes the convective boundary coefficient of the water against the pipe wall [W/m²/K]:

$$\alpha_{water} = \frac{\lambda_{water} \cdot Nu}{2 \cdot r_i} \quad (4)$$

When computing α_{water} , the Nusselt number Nu [-] needs to be incorporated. Based on Janssen and Warmoeskerken (1987) cited in Zlatanovic et al. (2017), a stepwise function of Nu for stagnant, laminar and turbulent flows was implemented:

$$Nu = \begin{cases} Re < 10 & 5.8 \\ 10 < Re \leq 2300 & 3.66 \\ Re > 2300 & 0.023 \cdot Re^{0.8} \cdot Pr^{1/3} \end{cases} \quad (5)$$

With the latter term being the Colburn equation (e.g. Cengel, 2003). This stepwise approach for computing Nu was incorporated into EPANET-MSX, with the Reynolds number (Re) computed internally by the program. Moreover, the Prandtl number Pr [-] is a temperature dependent term (Cengel, 2003):

$$Pr = \frac{\mu_{water} \cdot C_{p,water}}{\lambda_{water}} \quad (6)$$

With μ_{water} describing the dynamic viscosity of water [kg/m/s] and λ_{water} the thermal conductivity of water [W/m/K]. Additional information about the parameters applied in the temperature model of the case study network are listed in Table 2. For some parameters in the table, a standard temperature of 15 °C was selected as representative water temperature based on the source and smart meter temperatures in the case study. This resulted, for example, in a constant Pr value equalling 8.25 (at 15 °C) in the heat transfer model. Moreover, to account for EPANET's kinematic viscosity being based on 20 °C the relative viscosity was changed to 1.13 in EPANET.

3.2 Categorisation and filtration of temperature data

3.2.1 Parameter sensitivity

The required stagnation time for the water temperature to reach the soil temperature in service lines was identified by varying heat transfer model parameters (Table 3).

Table 2. Additional information about parameter selection in applied temperature simulations.

Parameter	Symbol	Value	Unit	Reference
Dynamic viscosity of water (at 15 °C)	μ_{water}	$1.14 \cdot 10^{-3}$	[kg/m/s]	Dinçer and Zamfirescu (2016)
Specific heat of water (at 15 °C)	$C_{p,water}$	4,184	[J/kg/K]	Dinçer and Zamfirescu (2016)
Thermal conductivity of water (at 15 °C)	λ_{water}	$5.77 \cdot 10^{-1}$	[W/m/K]	Dinçer and Zamfirescu (2016)
Thermal conductivity of polyethylene pipes	λ_{PE}	$4.2 \cdot 10^{-1}$	[W/m/K]	Czichos et al. (2014)
Thermal conductivity of polyvinylchloride pipes	λ_{PVC}	$2.2 \cdot 10^{-1}$	[W/m/K]	Czichos et al. (2014)
Burying depth	h	$1.4 + r_o$	[m]	Hubeck-Graudal et al. (2019)
Pipe wall thickness	r_{wall}	$2/9 \cdot r_i$	[m]	Hubeck-Graudal et al. (2019)
Outer pipe radius	r_o	$r_i + r_{wall}$	[m]	Hubeck-Graudal et al. (2019)

Table 3. Parameter variation used to identify the required stagnation time for the water temperature to reach the soil temperature in service lines.

Parameter	Unit	Variation	Information
T_{soil}	°C	$18 \pm 30\%$	Estimate of T_{soil} variation over summer (e.g. Hubeck-Graudal et al., 2019).
T_{inlet}	°C	$13 \pm 30\%$	Estimate of T_{inlet} variation during summer.
λ_{pipe}	W/m/K	$0.42 \pm 30\%$	Pipe material conductivity based on polyethylene (Czichos et al., 2014).
λ_{soil}	W/m/K	$2.2 \pm 50\%$	Possible soil conductivity variation in area (Ditlefsen et al., 2014; Ditlefsen and Sørensen, 2014).
d_o	m	$0.04 \pm 30\%$	Typical outer service line diameter.
d_i	m	$d_o \cdot 0.8$	Typical relationship between inner and outer service line diameter (e.g. Wavin, 2019).

3.2.2 Estimating soil temperatures based on smart meter temperatures

Smart meter temperature data were categorized into soil temperature estimates and WDN temperatures. The following filtering rules were established for all smart meter samples, categorizing temperature samples as soil temperature estimates, being an important input parameter of the heat transfer model (Eq. (1)):

1. Retention time within service line $\geq x$ [hrs]
2. Demand \geq between y % and z % of estimated service line volume
3. Sample age $\leq \beta$ [hrs]

It is necessary to estimate the service line volume because the service lines' exact locations and diameters are usually not stored in the utilities' asset databases. However, in the case study, the distance between WDN mains and stop valves was stored in the utility's asset database (Figure 1). Moreover, the GPS coordinates of each home were known, but not the exact smart meter location. Thus, service lines were prolonged linearly between stop valves and each home's GPS coordinates.

3.2.3 Estimating water distribution network temperatures based on smart meter temperatures

The following filtering rules were established for categorizing smart meter temperature samples as WDN temperature estimates:

1. Demand $\geq w$ % of estimated service line volume
2. Sample age $\leq \omega$ [hrs]

3.3 Model evaluation

To evaluate the performance and resemblance between simulated WDN temperatures and estimated WDN temperatures (based on smart meter samples), the root mean squared error (RMSE) was computed for each node in the hydraulic model with temperature estimates available. A time window was introduced to allow for some variation owing to potential transient heat transfers occurring between the water, pipe and soil that are not included in the temperature model. Moreover, the time window reduces the impact of sudden drops or rises in temperature potentially leading to misleading RMSE estimates. Such deviations in the timing between simulated and sampled temperatures values may occur owing to the method behind smart meter data sampling and linking of the data to a hydraulic model.

Considering a smart meter with hourly measurements, the water consumption between measurements can be estimated by linear interpolation (dotted line, Figure 2). The actual time of consumption is, however, unknown (solid line, Figure 2). In the example, most of the consumed water between hourly readings was actually consumed right after or before sampling and not linearly distributed over the entire hour. The simulation does not include this information; thus, at times, the water is consumed too early or too late in the hydraulic model. In particular, an increasing number of smart meters bundled in the hydraulic model's nodes may affect this error. This uncertainty decreases and increases with finer and coarser sampling resolutions, respectively. Thus, the timestamps of temperature samples were allowed to vary ± 1 hr for the computation of the RMSE. The minimum distance between measurement and model within this window was used.

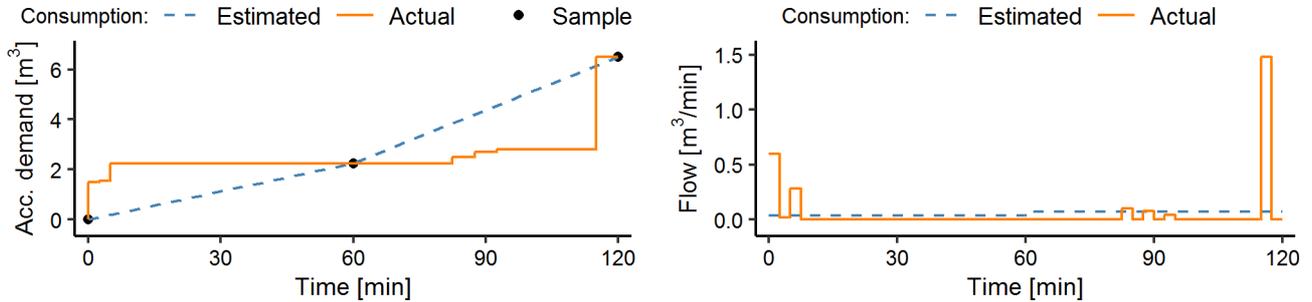


Figure 2. Linearly ‘estimated’ consumption over time based on two hourly measurements vs. actual consumption.

3.3.1 Soil conductivity estimation

The exact λ_{soil} of DMA1 was unknown, even though a tool by the Geological Survey of Denmark and Greenland (Ditlefsen et al., 2014) estimated λ_{soil} to be of 1.8 [W/m/K]. This estimate was based on three drills in the nearby area showing that 70% of the soil materials was composed of sand and gravel (λ_{soil} of 2.24 [W/m/K]) and 23% silty clay (λ_{soil} of 1.15 [W/m/K]) up to a depth of 25 m. However, as these values are based on reference conductivities and the pipe bedding material was unknown we allowed λ_{soil} to vary between 1 and 4 by steps of 0.1 [W/m/K]. The hereby identified λ_{soil} with the lowest median RMSE was then selected for further analysis.

4 Results and discussion

4.1 Raw temperature samples

For each hour, the median temperature was computed in both DMAs over the course of two weeks in August 2018 (Figure 3). In the case of DMA1, Figure 3 highlights a need for filtering the data when it should be applied within WDN analyses, as most depicted temperatures represent too high temperatures, unlikely to resemble the WDN temperature. As the water is stagnant most of the time in service lines or in in-house connections, the median was expected to represent more or less indoor temperatures. The results, however, show a clear difference between the two DMAs (Figure 3). This might be caused by the installation location of the smart meters in the DMAs. Also, an effect of the air temperature on smart meter temperature data is visible. For example, the drop in the daily air temperatures after the 8th of August led to a decrease in water temperatures approximately delayed by two days. The lag between air temperatures and smart meter temperature may indicate a potential heat buffering capacity of the soil (and homes) affecting the smart meter temperatures.

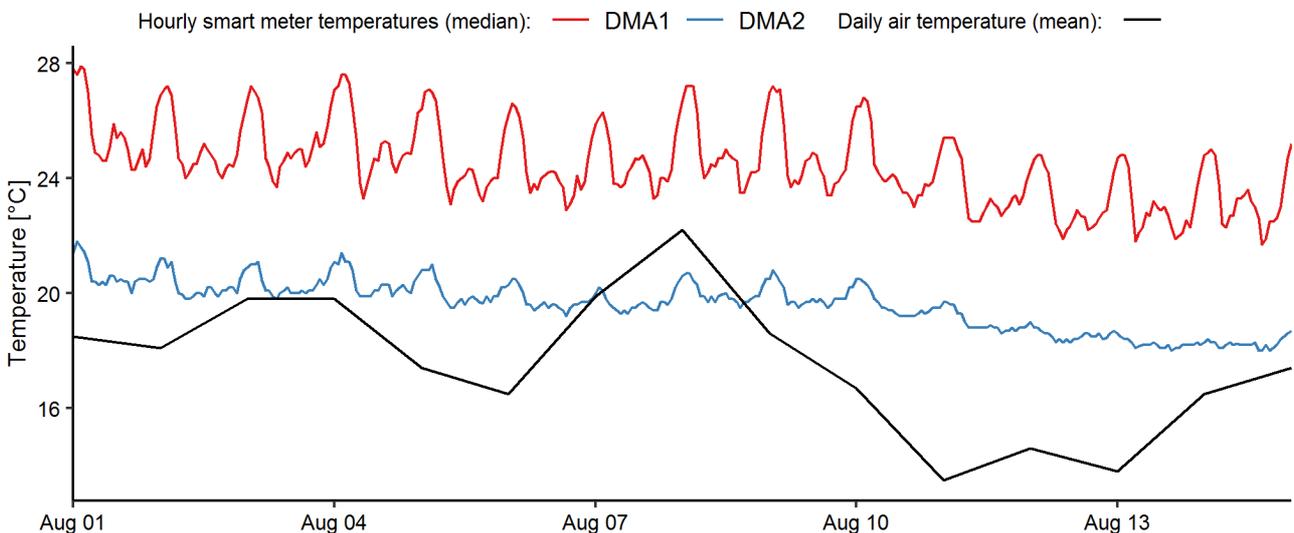


Figure 3. Temperature variation from 1,105 smart meters in two DMAs covering two weeks in 2018 and the daily mean air temperature from (DMI, 2019).

4.2 Categorisation and filtration of temperature data

4.2.1 Parameter sensitivity

Figure 4a displays the scenarios (Table 3), where the initial water temperature started at 9.1 °C and approximated a soil temperature of 12.6 °C. The figure shows that most scenarios reached the soil temperature after around $t = 2$ hr. Figure 4b displays the ratio between water and soil temperature at $t = 3$ hr for eight scenarios (denoted A–H). The two scenarios furthest away from having reached the soil temperature (B and D), included the lowest soil conductivity ($\lambda_{soil} = 1.1$ W/m/K) and largest service pipe diameters ($d_o = 52$ mm). This low soil conductivity corresponds to literature findings (Ditlefsen et al., 2014), but the true value is still unknown, among other things because at times sand and gravel are used as pipe bedding material. However, the largest pipe diameter does not represent a typical service line in the area (Table 1), and is only possible for large scale consumers not common in DMA1 and DMA2. Thus, the retention time x within service lines needed for soil temperature estimates (Rule 1, Sec. 0) was determined to be sufficient at being set to 3 hr. The remaining scenarios

with different initial water temperatures as well as soil temperatures resembled the one presented here and can be found in the supporting information (SI) A.

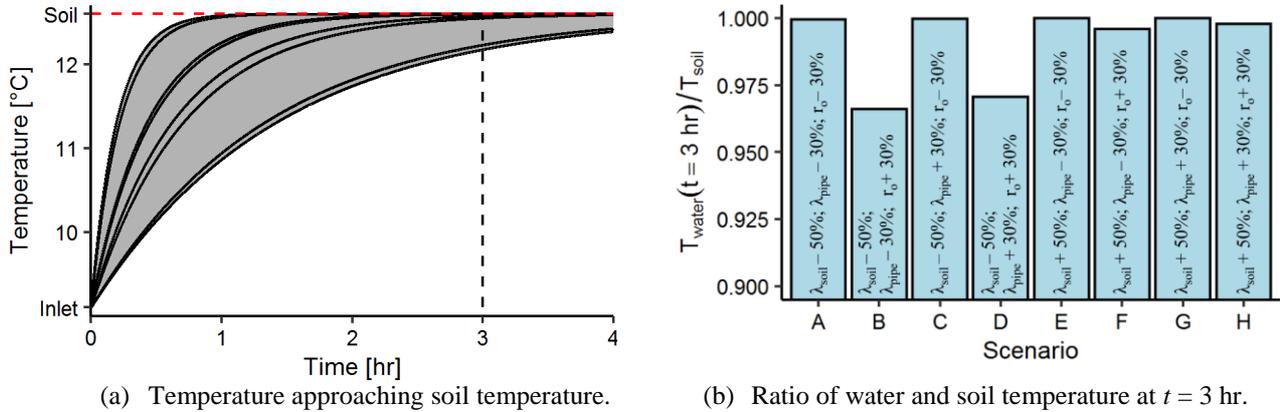


Figure 4. (a) Each line illustrates the change in water temperature of stagnant water in service lines over time based on the parameter scenarios listed in Table 3. Here, the initial water temperature was set to 9.1 °C and the soil temperature was set equal 12.6 °C. (b) Illustrates the water to soil temperature ratio after 3 hours. Remaining scenarios from Table 3 are listed in SI A.

4.2.2 Estimating soil temperature

The soil temperature was estimated based on smart meter temperature and demand samples in DMA1. These soil temperature estimates were then used as model boundary when simulating the water temperature throughout the DMA. Soil temperature estimates were obtained according to the rules defined in Sec. 0, where an estimate of the service line volume was needed. Based on this uncertain information, y and z were set to 20 % and 80 % respectively (Rule 2, Sec. 0) and β to 0.25 hr (Rule 3, Sec. 0). In total, only 393 samples (< 0.1% of all samples) passed this filtration (Figure 5) of which 13 samples were clearly lower or higher than the bulk of soil temperature estimates (Figure 5). These samples were deemed invalid as a manual inspection of all service lines in the geographic information system (GIS) revealed service lines with misleading information in the utility’s asset database (marked orange, Figure 5). Moreover, the blue estimations were located in the only 2-storey houses in the DMA with a smart meter in each storey having identical GPS coordinates. Thus, the upper storey has an additional in-house connection pipe length not included during filtration. In the temperature model, the fitted linear regression line was used as soil temperature. Disregarding the before mentioned anomalies, a mean soil temperature of 21.4 °C (median of 21.2 °C) was identified with the linear model output decreasing from 22.2 °C to 20.5 °C. The soil temperature in Figure 5 shows variations not included in the linear fit. The local environment can influence the soil temperature differently between the consumers’ homes, for example, owing to pavement, vegetation or other pipes in the area (Figure 1). Moreover, sunshine duration and intensity, the exact pipe burying depth and the time of sampling may play a substantial role. Soil temperatures underneath a hot asphalted road must be expected to be higher than temperatures under a green irrigated lawn. Thus, Figure 5 indicates that the effect of local soil temperatures throughout the DMA should be assessed to see whether they enhance modelling results.

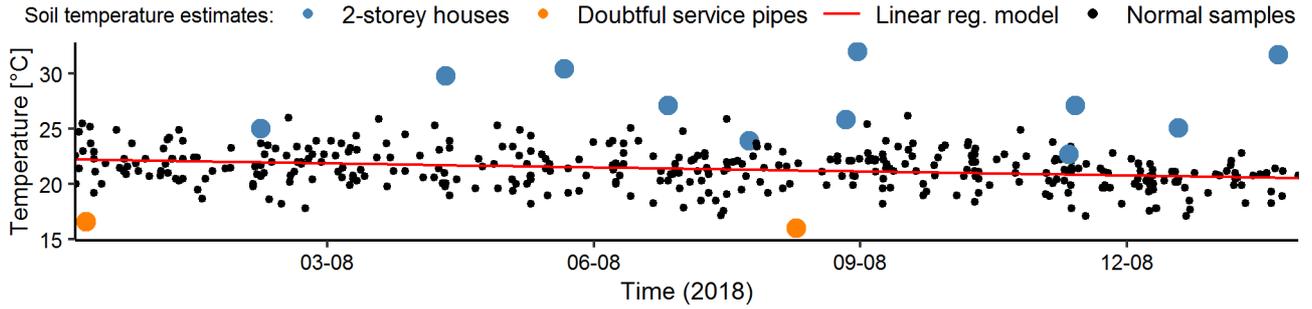


Figure 5. Soil temperature estimates and fitted linear regression model with anomalies (2-storey houses and doubtful service pipes) removed.

4.2.3 Estimating water distribution network temperature

The WDN temperature was estimated based on smart meter temperature and demand samples in DMA1. According to the rules presented in Sec. 3.2.3, a very conservative setting of $w = 400\%$ (Rule 1, Sec. 3.2.3) and a maximum sample age of $\omega = 0.25$ hr were chosen (Rule 2, Sec. 3.2.3) to identify WDN temperature estimates. At lower values of w , a marked number of samples were above the soil temperature (SI B). This indicates that local soil temperatures might be higher than expected (for instance, between 2012 and 2014 the warmest water samples representing soil temperatures in Copenhagen were sampled at $20.3\text{ }^{\circ}\text{C}$ (Hubeck-Graudal et al., 2019)) and vary greatly. However, the uncertainty about the actual location, length and diameter of service lines may still lead to samples not accurately representing WDN temperatures. Finally, the level of detail in the hydraulic model may play an important role. The consumption from smart meters is bundled in the smart meters' nearest nodes in the WDN model and not at their actual location. The identified WDN temperature estimates are plotted in Figure 6, with less than 0.1% ($n = 378$) of samples passing the conservative filtration.

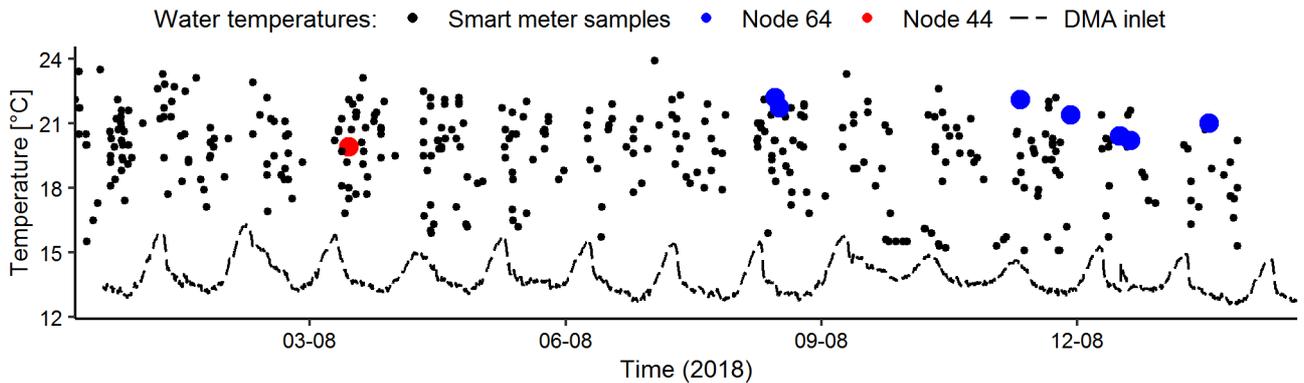


Figure 6. Water distribution network temperature estimates based on smart meter samples and the measured inlet temperature to the district metered area (DMA). Temperatures from two distinct nodes (44 and 64) resemble the bulk of temperature estimates, but differ highly when compared with simulated temperatures (Figure 7b).

4.3 Model evaluation

As the actual value of λ_{soil} was unknown, the mean and median RMSE were computed for a range of values (Figure 7a). The best median value equal $0.64\text{ }^{\circ}\text{C}$ (mean of $1.03\text{ }^{\circ}\text{C}$) was found with $\lambda_{soil} = 2$ [W/m/K], being a possible value for the case study area. Using this value of the thermal conductivity

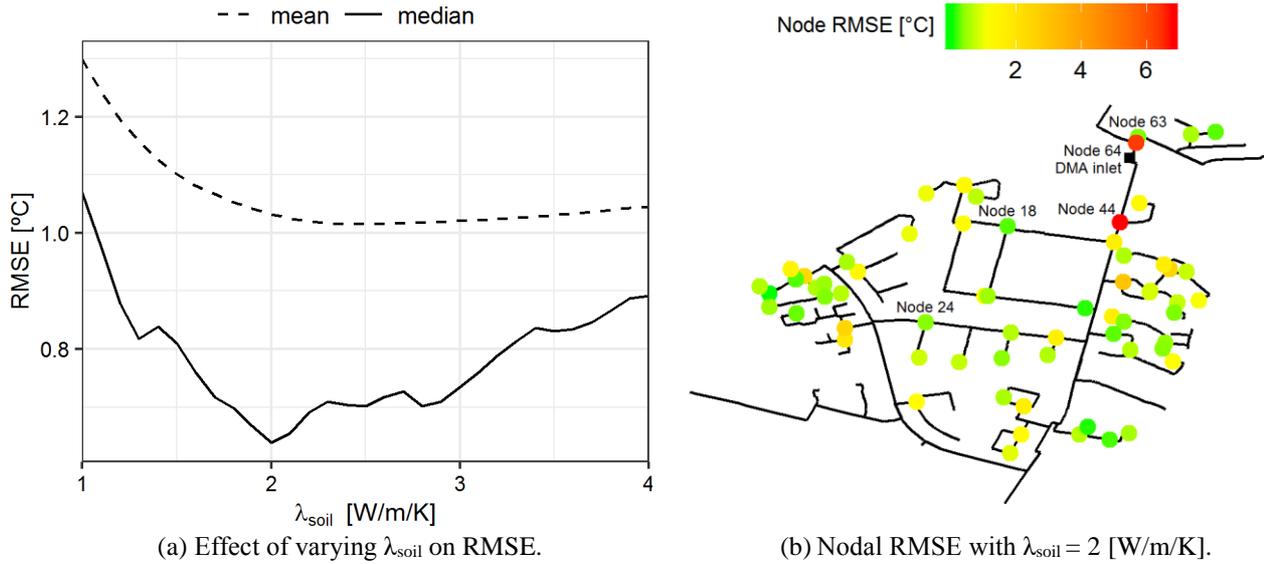


Figure 7. Varying the thermal conductivity of the soil (λ_{soil}) affects the mean and median root mean squared error (RMSE) of all nodes in the analysed network (a). The nodal RMSE for $\lambda_{soil} = 2$ [W/m/K] is illustrated in (b).

of the soil, the temperature was simulated throughout DMA1. The nodal RMSE of this setup is shown in Figure 7b.

Three examples of good model fits are shown for the three nodes with ID 18, 24 and 63 in Figure 8. In the figure, distinct colours represent unique smart meters, bundled in the same node. For example, the upper row shows three different smart meters, and highlights that only one sample per smart meter passed the filtration rules stated in Sec. 4.2.3. Moreover, Figure 7b revealed that the nodal RMSE is particular high in two nodes (> 6 °C), which is not easily seen from bulk of the temperature data (Figure 6). Discussions with the utility revealed that the smart meter linked to node 64 was incorrectly affiliated to the DMA. When only assessing smart meter demand data, this consumer would not have been identified as easily, as a single consumer's demand has little influence on the pressure loss or total demand of the DMA. Furthermore, the high RMSE of node 44 was owing to an inaccurate bundling of smart meters in a node too far away from the actual consumption. The actual position of the smart meter was on the side branch next to node 44 (Figure 7b). Removing these two nodes from RMSE computations improved the median RMSE to 0.63 °C (mean to 0.85 °C).

4.4 Limitations of temperature samples and modelling

Owing the system set-up of collecting smart meter data, only a small fraction of temperature data (less than 0.2 % during the analysed period) could be used for estimating soil and water temperatures. A more sophisticated data collection approach, for example storing the temperature values directly after a high volume of water was consumed and transmitting this data would be beneficial. Likewise, a finer sampling resolution would increase the applicability of the data. Moreover, the hydraulic model turned out to be a limiting factor in this work. As smart meters were bundled in WDN nodes, they might not represent the temperature at the actual location. Thus, a higher level of detail with

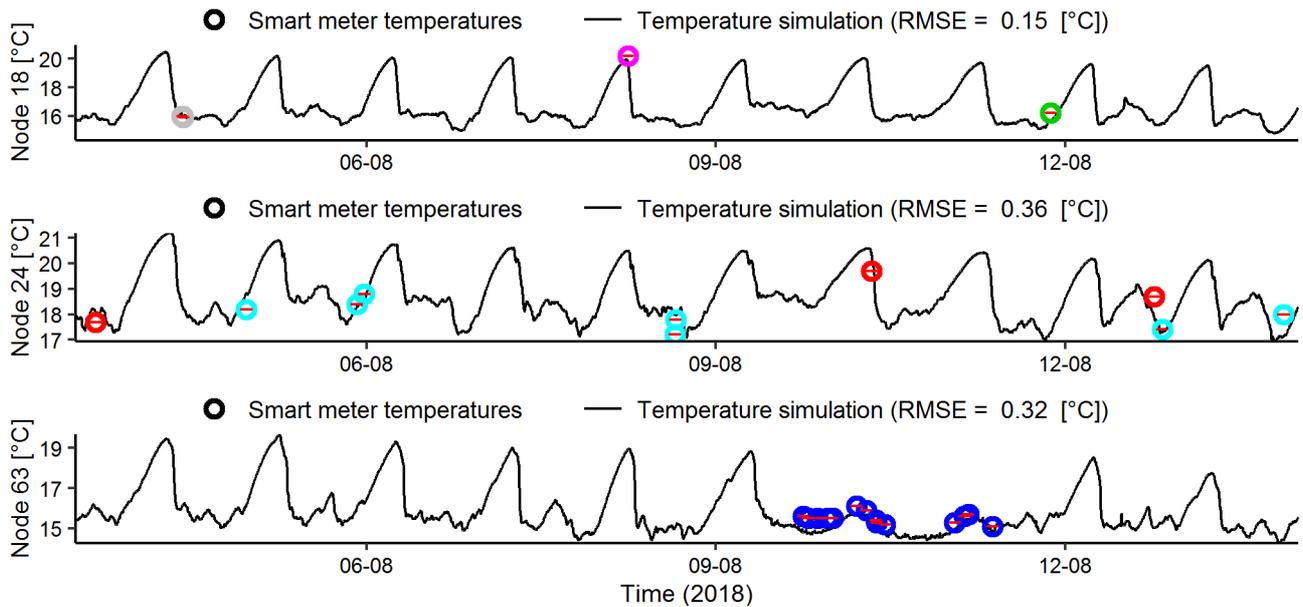


Figure 8. Smart meter temperatures (including ± 1 hr error margin) vs. simulated temperatures and the respective root mean squared error (RMSE). Distinct colours represent unique smart meters.

distinct nodes for each consumer, including service lines, could be beneficial to improve modelling results.

5 Conclusion

Application of smart meter temperature data revealed various valuable potentials within the field of water distribution system analysis. Consumption based filtration rules were run to categorize the data into samples representing soil and water distribution network temperatures. Analysis of these soil temperature estimates revealed anomalies in the utility's asset database, such as missing and incorrect service line information. Combining a hydraulic and temperature model made it possible to simulate the temperature throughout a DMA. Comparing the modelled values with filtered smart meter temperatures showed a satisfying resemblance (a mean root mean squared error, based on the hydraulic model's nodes, of around $1\text{ }^{\circ}\text{C}$), however, two nodes had a much high error ($> 6\text{ }^{\circ}\text{C}$) than the mean. The reasons for this were identified: one consumer was wrongly affiliated to the DMA, otherwise not easily detectable from the smart meter demand data only. The other nodal error revealed that the applied hydraulic model should include distinct nodes for each consumer, as temperatures otherwise may not represent the actual temperature at the node's location. Moreover, only a small fraction of the analysed data turned out to be applicable ($< 0.2\%$), as often temperature data represented a mixture between indoor and soil temperatures. Thus, more advanced sampling methods are required, increasing the applicability of the data. An easy solution could be to collect a finer sampling resolution than currently implemented (< 30 min). All in all, the application of the model and data highlighted the potential of smart meter temperatures as a valuable tool for improved asset management and water distribution system analysis.

6 Acknowledgment

We would like to thank Per Grønvald, Henrik Horsholt Christensen and Thorkil Bartholdy from Utility Brønderslev for providing data and answering questions related to their water supply.

7 References

- Blokker, E.J.M., Pieterse-Quirijns, E.J., 2013. Modeling temperature in the drinking water distribution system. *J. Am. Water Work. Assoc.* 105, E19–E28. doi:10.5942/jawwa.2013.105.0011. Erratum (2018) in *J. Am. Water Work. Assoc.* 110, 98–98. doi:10.1002/awwa.1178
- Blokker, E.J.M., van Osch, A.M., Hogeveen, R., Mudde, C., 2013. Thermal energy from drinking water and cost benefit analysis for an entire city. *J. Water Clim. Chang.* 4, 11–16. doi:10.2166/wcc.2013.010
- Blokker, M., 2019. Making water meters truly smart by adding additional features. Poster presented at the 17th International Computing & Control for the Water Industry Conference (CCWI). University of Exeter, UK.
- Boyle, T., Giurco, D., Mukheibir, P., Liu, A., Moy, C., White, S., Stewart, R., 2013. Intelligent Metering for Urban Water: A Review. *Water* 5, 1052–1081. doi:10.3390/w5031052
- Cengel, Y.A., 2003. *Heat Transfer: A Practical Approach*. Mc Graw-Hill, Boston.
- Czichos, H., Skrotzki, B., Simon, F.-G., 2014. *Das Ingenieurwissen: Werkstoffe* (in German). Springer Berlin Heidelberg, Berlin, Heidelberg. doi:10.1007/978-3-642-41126-7
- De Pasquale, A.M., Giostri, A., Romano, M.C., Chiesa, P., Demeco, T., Tani, S., 2017. District heating by drinking water heat pump: Modelling and energy analysis of a case study in the city of Milan. *Energy* 118, 246–263. doi:10.1016/j.energy.2016.12.014
- Dinçer, İ., Zamfirescu, C., 2016. *Drying Phenomena*, First Edit. ed, *Drying Phenomena: Theory and Applications*. John Wiley & Sons, Ltd, Chichester, UK. doi:10.1002/9781118534892
- Ditlefsen, C., Sørensen, I., 2014. D6 Overfladenære jordarters termiske egenskaber (in Danish). VIA-UC/GEUS. Available at http://geoenergi.org/xpdf/d6_jordarters_termiske_egenskaber.pdf (accessed November 13, 2019).
- Ditlefsen, C., Sørensen, I., Slott, M., Hansen, M., 2014. Estimating thermal conductivity from lithological descriptions – a new web-based tool for planning of ground-source heating and cooling, *Geological Survey of Denmark and Greenland Bulletin*. Copenhagen. Available at http://www.geus.dk/media/10907/nr31_p55-58.pdf (accessed November 13, 2019).
- DMI, 2019. *Vejrarkiv* (in Danish). Danish Meteorological Institute. Available at <https://www.dmi.dk/vejrarkiv/> (accessed November 13, 2019).
- Eggimann, S., Mutzner, L., Wani, O., Schneider, M.Y., Spuhler, D., Moy de Vitry, M., Beutler, P., Maurer, M., 2017. The Potential of Knowing More: A Review of Data-Driven Urban Water Management. *Environ. Sci. Technol.* 51, 2538–2553. doi:10.1021/acs.est.6b04267
- Hubeck-Graudal, H., Kirstein, J.K., Ommen, T., Rygaard, M., Elmegaard, B., 2019. Drinking water supply as low-temperature source in the district heating system: a case study for the city of Copenhagen (submitted).
- Janssen, L.P.B.M., Warmoeskerken, M.M.C.G., 1987. *Transport phenomena data companion*. Hodder Arnold.
- Kirstein, J.K., Liu, S., Høgh, K., Borup, M., Rygaard, M., 2019. Valve status identification by temperature modelling in water distribution networks. Technical University of Denmark, DTU Environment, Kgs. Lyngby. Unpublished work.
- Kyriakou, M., Eliades, D.G., Polycarpou, M.M., 2016. dbpRisk: Disinfection By-Product Risk Estimation, in:

- Panayiotou, C.G., Ellinas, G., Kyriakides, E., Polycarpou, M.M. (Eds.), *Critical Information Infrastructures Security*. Springer International Publishing, Cham, pp. 57–68. doi:10.1007/978-3-319-31664-2_7
- Larsen, S.L., Christensen, S.C.B., Albrechtsen, H.-J., Rygaard, M., 2017. GISMOWA: Geospatial Risk-Based Analysis Identifying Water Quality Monitoring Sites in Distribution Systems. *J. Water Resour. Plan. Manag.* 143, 04017018. doi:10.1061/(ASCE)WR.1943-5452.0000754
- Liu, S., Gunawan, C., Barraud, N., Rice, S.A., Harry, E.J., Amal, R., 2016. Understanding, Monitoring, and Controlling Biofilm Growth in Drinking Water Distribution Systems. *Environ. Sci. Technol.* 50, 8954–8976. doi:10.1021/acs.est.6b00835
- Makropoulos, Savić, 2019. Urban Hydroinformatics: Past, Present and Future. *Water* 11, 1959. doi:10.3390/w11101959
- Monks, Stewart, Sahin, Keller, 2019. Revealing Unreported Benefits of Digital Water Metering: Literature Review and Expert Opinions. *Water* 11, 838. doi:10.3390/w11040838
- Rokkjær, K., 2018. 6 things you didn't know you could do with a water meter. Available at <https://www.kamstrup.com/en-en/blog/6-things-you-did-not-know-you-could-do-with-a-water-meter> (accessed November 13, 2019).
- Rossman, L.A., 2000. *EPANET 2 Users Manual*. Cincinnati, OH.
- Shang, F., Uber, J.G., Rossman, L.A., 2007. *EPANET Multi-Species Extension User's Manual* 115.
- Vejen Forsyning, 2014. Typisk vandforbrug. Available at <https://www.vejenforsyning.dk/Vand/Typisk-vandforbrug> (accessed November 08, 2019).
- Wavin, 2019. Catalogue of Products: Wavin Ekoplastik System. Available at <https://www.wavinekoplastik.com/en/product-catalogue> (accessed November 10, 2019).
- Zlatanovic, L., Moerman, A., van der Hoek, J.P., Vreeburg, J., Blokker, M., 2017. Development and validation of a drinking water temperature model in domestic drinking water supply systems. *Urban Water J.* 14, 1031–1037. doi:10.1080/1573062X.2017.1325501

Supporting Information

A Parameter sensitivity

The eight scenarios, denoted ‘A-H’ in Table A-1 describe various parameter combinations based on the values listed in Table 3.

Table A-1. Eight scenarios based on parameter changes from Table 3 displayed in Figure A-1.

Scenario	$\lambda_{\text{soil}} - 50\%$	$\lambda_{\text{soil}} + 50\%$	$\lambda_{\text{pipe}} - 30\%$	$\lambda_{\text{pipe}} + 30\%$	$r_o - 30\%$	$r_o + 30\%$
A	X		X		X	
B	X		X			X
C	X			X	X	
D	X			X		X
E		X	X		X	
F		X	X			X
G		X		X	X	
H		X		X		X

Figure A-1 illustrates how the simulated stagnant water temperature slowly approximates the soil temperature based on varying parameter scenarios. The ratio between the simulated water temperature and soil temperature at $t = 3$ hr for these scenarios is shown in Table 3.

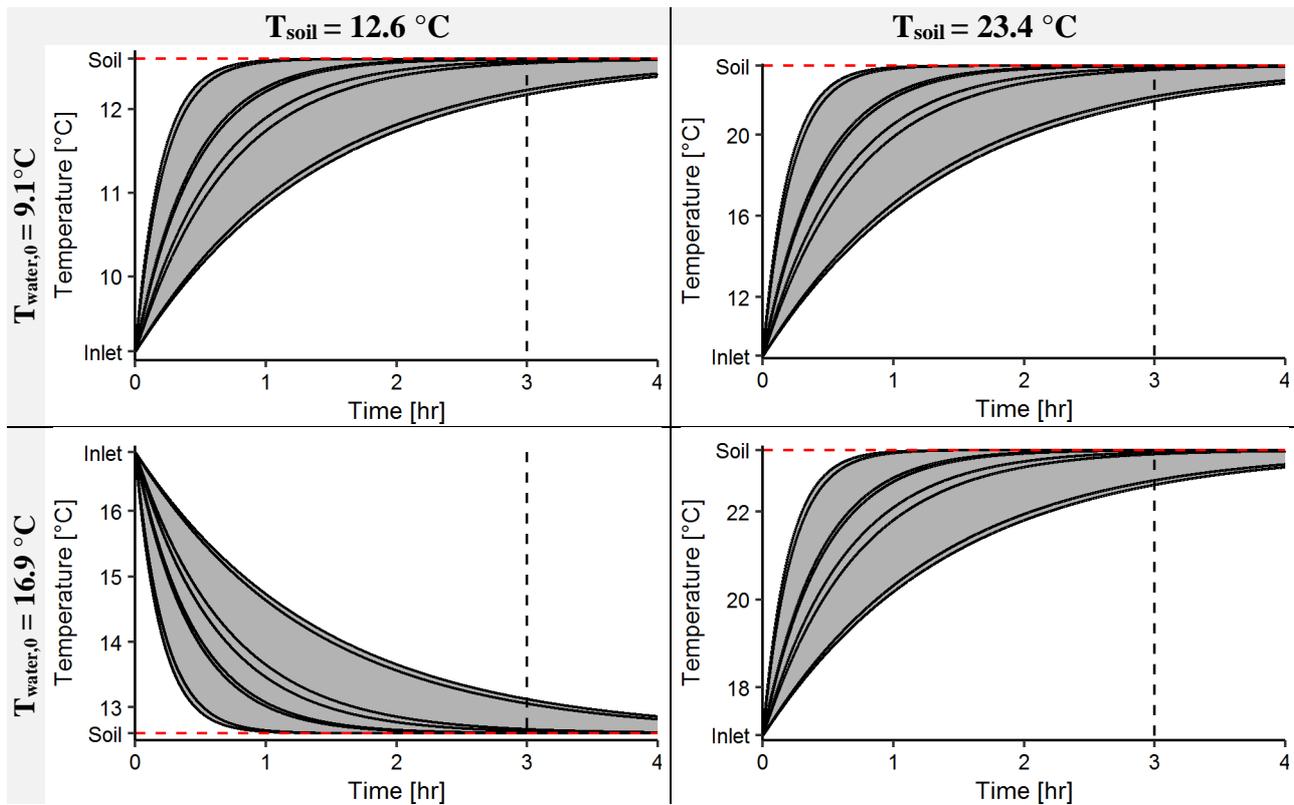


Figure A-1. Each line illustrates a parameter scenario (Table 3) leading to a change in water temperature of stagnant water in service lines over time.

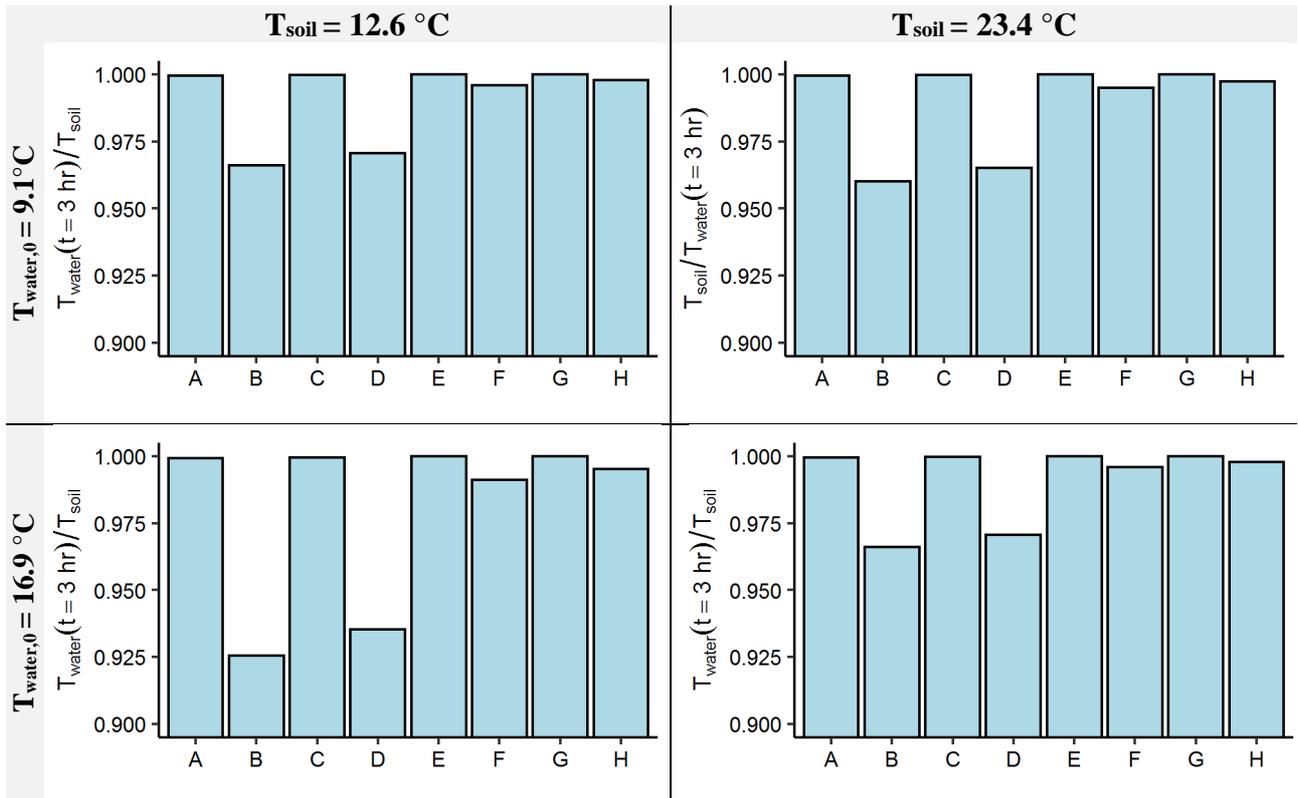


Figure A-2. Sensitivity analysis of parameters affecting the temperature in service pipes at $t = 3$ hr.

B Water temperatures

Table B-1 lists the number of samples passing the filtration rules (Sec. 3.2.3) when varying y .

Table B-1. Water temperature model results with different filtration rules (Sec. 3.2.3) and soil temperatures estimated from linear regression model shown in Figure 5.

Demand $\geq y$ %	Sample age $\leq \omega$	No. of samples	Samples above soil temperature
100	15	3091	718/3091 = 23%
200	15	1168	222/1168 = 19%
300	15	597	100/597 = 17%
400	15	387	64/387 = 17%

Figure B-1 shows the distribution of the samples at different filtration rates listed in Table B-1.

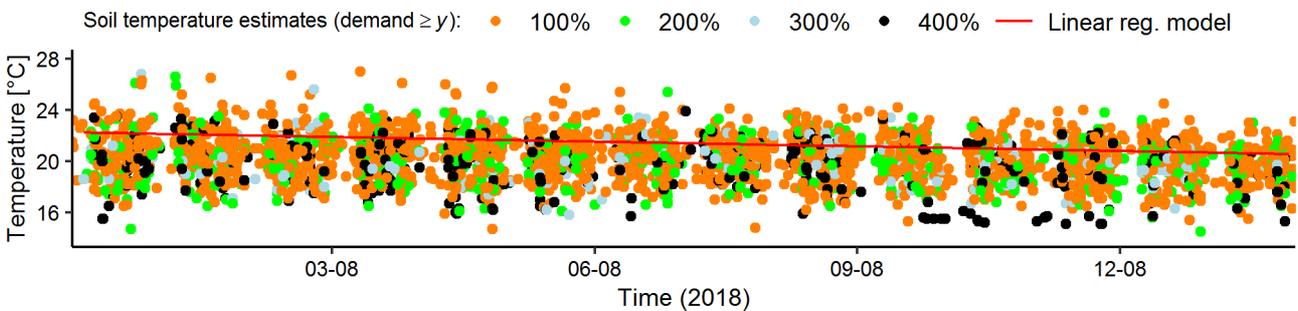


Figure B-1. Smart meter temperatures representing the water temperature in the water distribution network at different demand rates of y based on the filtering rules set in Table B-1.



Drinking water supply as low-temperature source
in the district heating system: a case study for the
city of Copenhagen

H. Hubeck-Graudal, J.K. Kirstein, T. Ommen, M. Rygaard
and B. Elmegaard

Drinking water supply as low-temperature source in the district heating system: a case study for the city of Copenhagen

Helga Hubeck-Graudal^{1,*}, Jonas Kjeld Kirstein², Torben Ommen³, Martin Rygaard² and Brian Elmegaard³

*corresponding author

¹HOFOR, Copenhagen, Denmark. helhub@hofor.dk.

²DTU Environment, Technical University of Denmark, Kgs. Lyngby, Denmark. jkir@env.dtu.dk; mryg@env.dtu.dk.

³DTU Mechanical Engineering, Technical University of Denmark, Kgs. Lyngby, Denmark. tsom@mek.dtu.dk; be@mek.dtu.dk.

Abstract This paper explores the potential for using large-scale heat pumps (HPs) to extract energy from Copenhagen's drinking water network and deliver it to its district heating system. The system involves certain losses in terms of additional heat and power consumption for end-use water heating. The net potential for energy extraction was analysed by means of an EPANET model to simulate system-wide temperatures in a piped distribution network. The model was validated against measured data from the network. Heat transfer in service lines was computed analytically and included in the net potential for energy extraction, which was determined to be 21 MW in Copenhagen. Around 38% of the HP source demand was harnessed from the ground. With HP COPs between 2.8 and 3.2, the System COP was only 1.7, thus suggesting that the choice of drinking water as a low-temperature heat source should depend on the available alternatives. Drinking water HPs have the side-benefit of preventing high drinking water temperatures; if operated in the summer they increased the share of supplied water complying with a recommended upper temperature limit of 12 °C from 42% to 81%.

Keywords: District heating; Drinking water supply; Energy efficiency; Heat pumps; Heat transfer modelling

1. Introduction

Electrically driven heat pumps (HPs) represent an excellent means of integrating the power and heating systems in areas with a significant heating demand and large amounts of intermittent renewable energy in the power system [1]. HP operation will be particularly relevant during hours when a substantial wind power production lowers the price and carbon emission of electricity. A low electricity price reduces the revenues earned by CHP plants, and thus CHP and HP technology can complement each other well in district heating (DH) systems.

To complement, or perhaps in some hours even replace individual CHP units, a significant total capacity of HPs must be installed. In dense urban environments, this poses the challenge to find low-temperature sources of sufficient capacity. An in-depth assessment of all available heat sources is needed, and based on this, the full range of heat source options should be prioritized based on socio-economic and/or private-economic analysis. An example of how to assess individual heat sources based on availability and performance is given in [2].

Since suitable large-scale heat sources are not necessarily abundant in urban areas [3], it may be worthwhile to take into account a variety of smaller sources, including wastewater networks [4,5] and groundwater reservoirs [6], in possible combination with surface water [7]. Another potential source

Nomenclature		<i>Subscripts</i>	
A	amplitude of annual surface temperature [K]	0	initialization
C_p	specific heat [J/kg/K]	A	case A
d	diameter [m]	air	air
D	damping depth of annual fluctuation [m]	ave	average
\dot{E}	electrical power demand [W]	B	case B
f	fraction of total consumption	com	compressor
h	height of soil layer above pipe [m]	demand	demand
H	corrected height of soil layer above pipe [m]	DH	district heating
K_p	heat transmission coefficient per m [W/m/K]	e	evaporator
l	height of fictitious soil layer above pipe [m]	ext	extraction
L	length [m]	i	inner
m	mass [kg]	model	modelled
\dot{m}	mass flow rate [kg/s]	net	net
\dot{Q}	heat flow rate [W]	o	outer
r	radius [m]	pipe	pipe
R	thermal resistance [m·K/W]	return	return
t	time [s]	soil	soil
T	temperature [K]	source	heat source
\dot{V}	volumetric flow rate [m ³ /s]	supply	supply
W	electrical power [W]	sys	system
		year	year
		water	water
		wall	pipe wall
<i>Greek symbols</i>		<i>Abbreviations</i>	
a	heat transfer coefficient [W/m ² /K]	CHP	combined heat and power
δ	thickness [m]	COP	coefficient of performance
Δ	change	DH	district heating
η	efficiency [-]	HP	heat pump
κ	thermal diffusivity [m ² /s]	MSX	Multi-Species Extension
λ	thermal conductivity [W/m/K]	PE	polyethylene
μ	dynamic viscosity [kg/m/s]		
ρ	density [kg/m ³]		
φ	correction factor [K·m ² /W]		
ω	frequency [s ⁻¹]		

is the reduction of flow temperature from drinking water distribution mains, which are extensively available in the immediate underground of any city.

Advantages of using drinking water as the heat source for a HP include its stable temperature around the year; its purity, which reduces fouling and maintenance needs on the HP components; and the fact that cooling is beneficial to drinking water as it limits bacterial growth [8], thereby preserving its quality for human consumption.

A comprehensive review of energy use and recovery potential in urban water systems was provided by Elías-Maxil et al. [9]. They presented different ways of using water as a source for heat pumping, but stated that only few studies relate to the application of water before end-use. However, recent studies include the work of Blokker et al. [10] who analysed a drinking water HP supplying heat for 900 homes in Almere City, Netherlands. De Pasquale et al. [11] analysed the potential for using a part of Milan's drinking water network as a source for a 4.65 MW HP, and Jadwiszczak and Niemierka studied a system for integration into the DH system of Głogów, Poland [12]. Van der Hoek et al., in turn, investigated cold recovery from drinking water and the related microbiological effects [13].

A drawback of extracting heat from drinking water mains is that reduced water temperatures downstream give rise to increased end-use heat demand in households. The extra domestic heating

demand is partly offset by additional heat transfer from the soil surrounding the water pipes, which is triggered by the increased temperature difference between soil and water upstream.

Since the heat transfer from soil to water is to a great extent what determines the energy balance of drinking water HPs, the use of this technology is comparable to ground-source heat extraction at a very large scale. In other words, the amount of additional heat transfer obtained from the soil determines the rationality of using drinking water as a heat source.

The purpose of the present study is to determine the technical potential for large drinking water HPs in the DH system of the City of Copenhagen, while taking into account a number of technical constraints related to their location. The potential is determined by modelling of the ground-source heat that is gained when HPs are installed at all water distribution mains with an average flow rate exceeding 50 L/s. The technical constraints include the operational profile of the HPs, their distance to a DH network, and differences in supply temperatures of various DH distribution networks. In addition to the HP potential, this study examines the derived water cooling effects experienced by consumers.

The study employs a hydraulic model provided by HOFOR, the Greater Copenhagen Utility [14], which was modified to include a heat transfer model to be solved in the EPANET Multi-Species Extension (MSX) [15,16], which enables computation of the water temperature change throughout the distribution system. Service lines were not included in the distribution network model, but to account for a potentially significant share of the total heat transfer in the pipelines, their contribution was computed outside the model. The EPANET-MSX approach for modelling heat transfer in drinking water networks was first used by Blokker et al [17], and later improved by De Pasquale [11]. The present work contributes with a heat transfer model validated against measurement data from the drinking water network, enabled in particular by the use of an accurate measure of the undisturbed soil temperature.

2. Methods

2.1. System description of drinking water heat pumps

A visualization of the analysed case study is presented in Fig. 1, where the energy flows in the traditional energy system (case A, the baseline) are compared to a system with drinking water HPs installed (case B). In case A, water from the waterworks is distributed in the city and exchanges heat, $\dot{Q}_{\text{soil,A}}$, with the surrounding soil along its way through the network. The water is heated by warmer soil during summer, and cooled by colder soil during winter. The water consumers of Copenhagen have a demand for DH, $\dot{Q}_{\text{demand,A}}$, to cover their hot water needs, e.g. for showering. The DH system in Copenhagen covers more than 98% of the heat demand, so there is practically full overlap between water consumers and DH consumers in the system. The DH is produced mainly by CHP plants, waste incineration and peak-load boilers. The water consumers also demand electricity, $\dot{E}_{\text{demand,A}}$, for water heating, e.g. in washing machines and for cooking. The electricity production may be based on, inter alia, wind turbines and CHP.

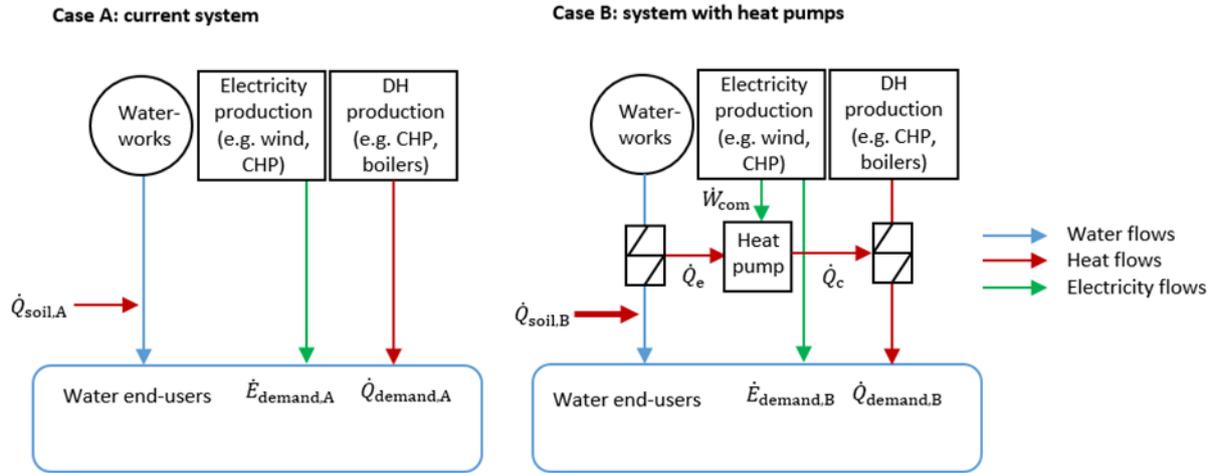


Figure 1. Energy system model without (A) and with (B) drinking water heat pumps (inspired by [11]).

In case B a number of HPs driven by electricity, \dot{W}_{com} , utilize low-temperature heat from the water supply network, \dot{Q}_e , and deliver it to the DH network, \dot{Q}_c . Due to the operation of HPs, water temperatures are reduced downstream. Compared to case A, this prompts a higher end-use DH and electricity demand, $\dot{Q}_{demand,B}$, and $\dot{E}_{demand,B}$, for water heating for hot water consumption. However, the increase in demand is partly offset by increased heat transfer from soil to water pipes, $\dot{Q}_{soil,B}$, which is triggered by the increased temperature difference between soil and water when the water is cooled upstream. The system's net gain in energy, compared to case A, is the outcome of one of two situations: 1) If the soil is warmer than the water, the heat flow from soil to water increases. 2) If the soil is colder than the water, the heat loss from water to soil is reduced.

The additional DH demand in case B compared to case A depends on the amount of additional heat transfer from the soil to the water distribution network, and on the fraction of total water consumption that is heated by means of DH:

$$\Delta\dot{Q}_{demand} = \dot{Q}_{demand,B} - \dot{Q}_{demand,A} = ((\dot{Q}_e - \dot{Q}_{soil,B}) + \dot{Q}_{soil,A}) \cdot f_{DH} = (\dot{Q}_e - \Delta\dot{Q}_{soil}) \cdot f_{DH} \quad (1)$$

The potential for heat utilization from the drinking water supply was defined as a function of the acceptable temperature reduction in the drinking water main flows, and of the additional end-use DH demand. It is the thermal capacity effectively gained by the DH system from the drinking water HPs:

$$\dot{Q}_{net} = \dot{Q}_e - \Delta\dot{Q}_{demand} \quad (2)$$

The power consumed by the HP is converted to thermal energy and contributes positively to the overall DH balance. This heat does not necessarily contribute directly to the reheating of cooled drinking water, but it will ease the load on other heat-producing units in the DH system. Leaving out this term allows us to assess a potential for the heat extraction that is independent of the concrete technology and COP of the HP applied.

The gross potential for heat extraction from the water supply is then the thermal energy yield at the HP evaporator, which is given by the potential temperature decrease of the drinking water:

$$\dot{Q}_e = \dot{V}_{water} \cdot \rho_{water} \cdot c_{p,water} \cdot \Delta T_{source} \quad (3)$$

When entering buildings, the water is thermally affected by the indoor ambient temperature, and was assumed to gain 1 K or more between the end of the service line and the first tap [18]. To warrant end-use temperatures higher than 4 °C, considered as the acceptable limit, the modelled exit temperatures would thus have to be higher than 3 °C.

The degree of heat utilization achieved from the water supply was described as:

$$\eta_{\text{util}} = \dot{Q}_{\text{net}} / \dot{Q}_e \quad (4)$$

The ratio defines a measure of the thermal yield effectively gained from the drinking water HPs against the thermal yield extracted from the evaporator. A heat utilization degree that only considers the contribution made by heat transfer from the soil was defined:

$$\eta_{\text{util,soil}} = \Delta \dot{Q}_{\text{soil}} / \dot{Q}_e \quad (5)$$

The COP of the HPs is given as:

$$\text{COP} = \dot{Q}_c / \dot{W}_{\text{com}} \quad (6)$$

The additional end-use electricity demand in case B compared to case A is given by:

$$\Delta \dot{E}_{\text{demand}} = \dot{E}_{\text{demand,B}} - \dot{E}_{\text{demand,A}} = \left((\dot{Q}_e - \dot{Q}_{\text{soil,B}}) + \dot{Q}_{\text{soil,A}} \right) \cdot f_{\text{el}} = (\dot{Q}_e - \Delta \dot{Q}_{\text{soil}}) \cdot f_{\text{el}} \quad (7)$$

To account for the additional electricity end-use and the additional DH end-use, a system COP was defined:

$$\text{COP}_{\text{sys}} = \frac{\dot{Q}_e - \Delta \dot{Q}_{\text{demand}} + \dot{W}_{\text{com}}}{\dot{W}_{\text{com}} + \Delta \dot{E}_{\text{demand}}} \quad (8)$$

where the numerator states the net thermal energy output of system B compared to system A, and the denominator states the net difference in electrical energy input.

2.2. Assumptions on water, heating and electricity use

Table 1 shows the assumed distribution of water use among different water services, and the energy source used to provide these services. The resulting f_{DH} of 0.37 was applied for summer. In winter, one third of the cold water stagnant in domestic pipes and cisterns was assumed heated to room temperature, which results in a larger f_{DH} of 0.47. In both summer and winter, f_{el} was assumed to be 0.33. The remaining water fractions of 0.2 (in winter) or 0.3 (in summer) were assumed not to be heated.

2.3. Case study network

The case study investigated the heat transfer in the drinking water network of the City of Copenhagen. The water utility of Copenhagen, HOFOR, extracts groundwater from 7 waterworks and supplies drinking water to the entire city and some neighbouring municipalities. The system includes a reservoir that balances supply and demand by directing water to the reservoir at night when consumption is low. The network consists of a mix of cast iron, steel-reinforced concrete and polyethylene (PE) pipes. It is made up of 126 km transport lines ($d_o = 650$ mm to 1,250 mm), 157 km main lines ($d_o = 300$ mm to 1,200 mm), and 787 km distribution lines ($d_o = 50$ mm to 290 mm) [20]. A hydraulic network model of this system [14] was implemented in EPANET. The total average discharge rate

Table 1. Shares of water use for different purposes, and the energy source needed to heat it (based on [19]). Numbers in parenthesis apply to the summer situation.

Water services	Share of use [%]			Heat source	Share of water service that is heated [%]
	Households	Workplaces	Total		
Shower and personal hygiene	46	10	37	district heating	100
Toilet flushing	17	50	25	district heating	33 (0)
Clothes washing	13	5	11	electricity	100
Dish washing	10	20	13	electricity	100
Cooking	7	15	9	electricity	100
Miscellaneous	7	0	5	district heating	33 (0)
Total share heated by district heating			47 (37)	district heating	
Total share heated by electricity			33	electricity	
Total share unheated			20 (30)	-	

from the 7 waterworks in the model is 2,165 L/s. The demand pattern applied in most nodes represents typical city consumption in an average 24-hour period.

2.4. Model development

EPANET-MSX is a water quality modelling tool designed to track the flow and reaction of chemical species in discrete water volumes, which are transported through the pipes until they reach the junctions and undergo complete mixing [16]. Mathematically this is done by numerical integration of differential-algebraic equations. Temperature was modelled as a chemical species in EPANET-MSX, handled in concentration units.

EPANET and MSX use a Lagrangian model to describe the system hydraulics and changes in water quality. The energy balance for water at any point in the pipes is therefore the energy balance of a closed system, i.e. of a control mass that travels through the pipes. For the energy balance for the closed system, it is assumed that changes in kinetic and potential energy are negligible. Boundary work may be neglected since EPANET assumes that the density of water is constant [15].

For a non-insulated cylindrical pipe buried in a semi-infinite solid, the rate of heat transfer into the control mass from the ambience was computed as proposed by [21]:

$$\dot{Q}_{\text{soil}} = K_p \cdot L \cdot (T_{\text{water}} - T_{\text{soil}}) \quad (9)$$

The heat transfer processes taking place in the system comprise: convection of water in the pipe, conduction through the pipe wall, and conduction through the surrounding soil (Fig. 2). Drawing on heat transfer coefficients defined by [21] for DH and flue gas pipes, K_p , was defined as:

$$K_p = (R_{\text{soil}} + R_{\text{pipe}} + R_{\text{water}})^{-1} = \left(\frac{\ln\left(\frac{4 \cdot H}{d_o}\right)}{2 \cdot \pi \cdot \lambda_{\text{soil}}} + \frac{\ln\left(\frac{d_o}{d_i}\right)}{2 \cdot \pi \cdot \lambda_{\text{pipe}}} + \frac{1}{\alpha_{\text{water}} \cdot \pi \cdot d_i} \right)^{-1} \quad (10)$$

R_{soil} includes a correction for the convective resistance at the soil surface, which translates this resistance into an additional layer of soil [22] (Fig. 2):

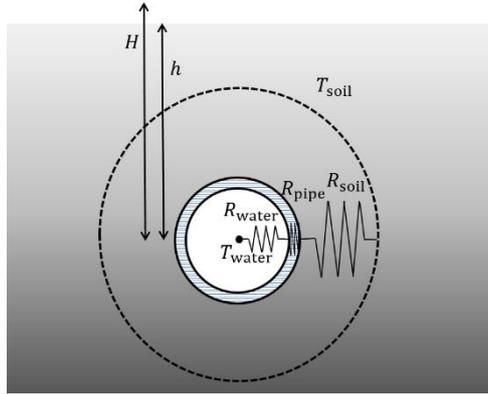


Figure 2. Thermal resistances, temperatures and correction terms included in the model [21, 22].

$$H = h + \varphi \cdot \lambda_{\text{soil}} \quad (11)$$

where φ is defined as $0.0685 \text{ K} \cdot \text{m}^2/\text{W}$.

Based on the energy balance for a closed system and the heat capacity of the control mass, the following expression was used to describe the time rate of change of the temperature of water in all pipes:

$$dT_{\text{water}}/dt = \frac{K_p}{C_{p,\text{water}} \cdot \rho_{\text{water}} \cdot \pi r_1^2} \cdot (T_{\text{water}} - T_{\text{soil}}) \quad (12)$$

At the junctions, MSX computes the final water temperature as the mass-based weighted average of the temperatures of all incoming streams.

The output $T_{\text{water}}(t)$ from EPANET-MSX was used to compute the additional heat transfer from the soil to the water distribution network:

$$\Delta \dot{Q}_{\text{soil}} = \dot{m}_{\text{water}} \cdot C_{p,\text{water}} \cdot (T_{\text{water,A}}(t) - T_{\text{water,B}}(t)) \quad (13)$$

where A and B refer to the considered cases without and with HPs (Fig. 1).

2.5. Input parameters

2.5.1 Initial water temperatures

The initial conditions for Eq. 12 – the water temperature, $T_{\text{water}}(0)$, at the inlets of the system – are the temperatures of the water exiting the 7 waterworks and the reservoir. HOFOR measures these temperatures on a daily basis, and the data were used to calculate weekly averages for each source, covering the period July 2012 – September 2014. $T_{\text{water}}(0)$ varied from $8.2 \text{ }^\circ\text{C}$ (the coolest source in its coolest state) to $11.3 \text{ }^\circ\text{C}$ (the warmest source in its warmest state). The yearly mean water temperature across all the sources was $9.3 \text{ }^\circ\text{C}$.

2.5.2 Undisturbed soil temperature

The water distribution network of Copenhagen is buried at a depth of approximately 1.5 m. In upper soil layers, the soil temperature changes with depth. The soil temperature depends on ambient air temperature, wind velocity, solar radiation, sky temperature, thermal diffusivity of the soil, and phase

changes of the soil water [17, 24, 25]. The temperature of city soils, moreover, may be affected by non-natural heat sources such as nearby basements or DH pipes [26]. The soil surrounding water pipes is also thermally affected by the convection of the water flowing through the pipes.

The soil temperature applicable to the model is the *undisturbed* ground temperature, i.e. unaffected by heat transfer from the pipe. The soil temperature is sometimes taken as the temperature at the soil surface (or the air temperature) [11,23]; sometimes at the level of the pipe centre [22,14], but due to the thermal impact of the pipe on the surrounding soil, its accurate location lies somewhere between these two points (Fig. 2) and depends on the specific system [24].

The undisturbed soil temperature was estimated to be equal to the water temperature measured at the periphery of the network at a location with low demand and where the water is in thermal equilibrium with the soil. The measured soil temperature ranged from 6.2 °C in the winter to 20.3 °C in the summer during 2012-2014.

2.5.3 Soil, pipe and water properties

Information about pipe depths, pipe wall thicknesses, and the material of each segment were based on [27]. The physical properties of water at 10 °C were used. All input parameters used in EPANET-MSX are listed in Table 2.

The thermal conductivity of soil depends on its water content, mineral content and texture (pore and particle size). For moraine clay, the predominant superficial soil type in Eastern Denmark and Copenhagen, literature values range from 1.1 W/m/K to 2.9 W/m/K [28]. The few studies conducted on Danish soils (although none in Copenhagen) indicate a thermal conductivity for moraine clay in the range 1.6 W/m/K to 2.2 W/m/K [28]. A high water content of the clay increases the conductivity. Since Copenhagen is located by the sea, the upper-end value of 2.2 W/m/K was used.

For the validation of the model, the thermal conductivities of the three existing pipe materials were used. For the case study with HP simulation, all pipes in the network were assumed to be made of PE since HOFOR is in a process of gradually converting the entire piping network into this material.

Table 2. Input parameters used in EPANET-MSX and for the soil temperature evaluation.

Parameter	Symbol	Unit	Value	Description
Initial temperatures of water leaving waterworks and reservoir	$T_{\text{water}}(t=0)$	[°C]	8.2 to 11.3	range of data measured at each waterworks/ reservoir in July 2012 – Sept. 2014
Undisturbed soil temperature	T_{soil}	[°C]	6.2 to 20.3	range of water temperatures measured at HOFOR's laboratory in July 2012 – Sept. 2014
Pipe wall thickness	δ_{wall}	[m]	$1/9 \cdot d_i$	
Pipe depth	h	[m]	$1.4 + d_o/2$	
Thermal conductivity of soil	λ_{soil}	[W/m/K]	2.2	measured value for Danish moraine clay
Thermal conductivity of pipes	λ_{pipe}	[W/m/K]	55; 0.42; 1	for pipes made of cast iron, polyethylene and steel-reinforced concrete, respectively
Thermal conductivity of water	λ_{water}	[W/m/K]	0.56	
Specific heat of water at 10 °C	$C_{p,\text{water}}$	[J/kg/K]	4188	
Specific heat of soil	$C_{p,\text{soil}}$	[J/kg/K]	1100	
Dynamic viscosity of water at 10 °C	μ_{water}	[Pa·s]	0.001308	
Density of water at 10 °C	ρ_{water}	[kg/m ³]	999.7	
Density of soil	ρ_{soil}	[kg/m ³]	2000	

2.6. Soil temperature evaluation

In order to validate the heat transfer model and determine the magnitude of heat transfer correctly, the empirically determined T_{soil} was compared with the measured air temperature, T_{air} , and a modelled soil temperature at 1.5 m depth, $T_{\text{soil,model}}$ – both commonly used approximations of the soil temperature. The measurement data for both T_{air} and T_{soil} were fitted by the least squares method to a sinusoidal function [31], which describes soil temperature at a given depth and time during the year as a function of air temperature:

$$T_{\text{soil}}(h, t) = T_{\text{air,ave}} + A_{\text{air}} \cdot e^{-h/D} \cdot \sin(2 \cdot \pi(t - t_0)/t_{\text{year}} - h/D - \pi/2) \quad (14)$$

where D is given by $(2 \cdot \kappa/\omega)^{0.5}$, κ is defined as $\lambda_{\text{soil}}/(\rho_{\text{soil}} \cdot C_{p,\text{soil}})$, and ω is $2 \cdot \pi/t_{\text{year}}$. $T_{\text{soil,model}}$ was modelled as T_{air} fitted to Eq. 14 and computed at $h = 1.5$ m.

2.7. Heat pump modelling

For the integration with the DH distribution network, each HP was assumed to re-heat a partial stream of the return flow to reach the DH supply temperature (Fig. 3). With this configuration, the HP increases the DH temperature from between 45 °C and 50 °C up to between 70 °C and 90 °C, depending on whether the distribution area is a low-temperature or high-temperature DH zone. In such configuration, the HP might be placed along the length of a DH main line. More upstream locations on the DH line, i.e. as close as possible to the local DH station, were preferred to secure larger heat demands.

HPs for DH in Copenhagen tend to use ammonia as the working fluid, while other options include CO₂, hydrocarbons, or one of the recently developed synthetic refrigerants with low global warming potential. Implementing HPs in the water supply requires special attention with respect to risk of contamination; avoiding leakage of working fluid into the water is crucial. The heat transfer from drinking water was therefore assumed to require either a double-walled heat exchanger, or a secondary heat transfer loop with a brine or heat transfer fluid such as CO₂. Both solutions require a larger temperature difference between the heat source and the working fluid. The estimation of COP was based on the assumption of a brine system with 2 K minimum temperature differences in both heat exchangers. With reasonable design of the brine system, i.e. low temperature differences for heat

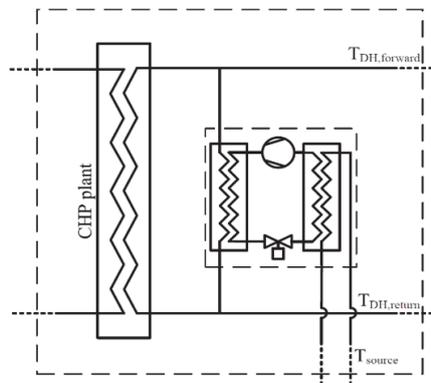


Figure 3. Configuration of heat pump connection to a DH line (from [29]). The heat pump increases the temperature of a part of the return flow ($T_{DH,return}$) to the supply temperature ($T_{DH,supply}$), via a reduction of the drinking water temperature (T_{source}).

exchange and insignificant maldistribution, this should not cause any issues related to the freezing point of clean water.

For the purpose of assessing \dot{W}_{com} and the system COP, the COP of the HPs was estimated based on characteristics for similar HP equipment recently installed in the same DH network [30]. The assumed temperatures and the resulting COP for the HPs in high-temperature and low-temperature DH zones are listed in Table 3.

A drinking water flow of 50 L/s subjected to a temperature reduction of 5 K would deliver 1 MW source to a HP, which was considered the minimum capacity of interest in this study.

The water flow would be subjected to additional pressure losses from the heat exchange process, manifold and additional piping, which requires increased pump power for the distribution of water. It is considered likely that the heat exchangers can be designed with a lower pressure loss than 50 kPa. In case of a 50 L/s water flow rate, and an increased pressure loss of 50 kPa, the isentropic pump power is approximately a factor of 100 less than the power required for the HP, and thus considered negligible in the study.

All major flows entering the city had volume flow rates exceeding the minimum value of 50 L/s. With a view to utilizing thermal energy from all of them, HPs were simulated in EPANET-MSX by application of a temperature reduction, ΔT_{source} , in 8 selected nodes on the drinking water network (Fig. 4).

The HPs were placed in spaces within a distance of 100 m from DH main lines, and priority was given to locations in low-temperature DH zones (HPs f) and g) met this criterion). Table 4 lists the average water flow rates of each HP location and the modelled source temperatures in February.

For Copenhagen, the economic benefit of integrating large HPs can mainly be reaped in winter when HPs may replace peak-load production. To assess the net potential, \dot{Q}_{net} , in the coldest month, the EPANET-MSX model was run with temperature data from February ($T_{\text{soil}} = 6.7 \text{ }^\circ\text{C}$) for a range of ΔT_{source} applied identically across the 8 nodes.

In most of the drinking water mains, the direction of the water flow is changing at night when water fills the reservoir during the low-demand period. The flow direction reverses up to 10 times over the course of 5 to 7 nightly hours. Running HPs in such situations would result in excessive cooling of the water. Consequently, HPs were simulated to operate only during periods with sustained unidirectional flow. Fig. 5 shows the daily water flow rates for the 8 locations when the ‘prohibited’ HP operation hours are left out (their water flow rates set equal to 0).

Table 3. Assumptions on temperatures, and estimation of COP for HPs in high-temperature and low-temperature zones.

Heat pump location	High-temperature	Low-temperature
Source inlet temperature [$^\circ\text{C}$]	6	6
Source outlet temperature [$^\circ\text{C}$]	1	1
Sink outlet temperature [$^\circ\text{C}$]	90	70
Sink inlet temperature [$^\circ\text{C}$]	50	45
Estimated COP [-]	2.8	3.2

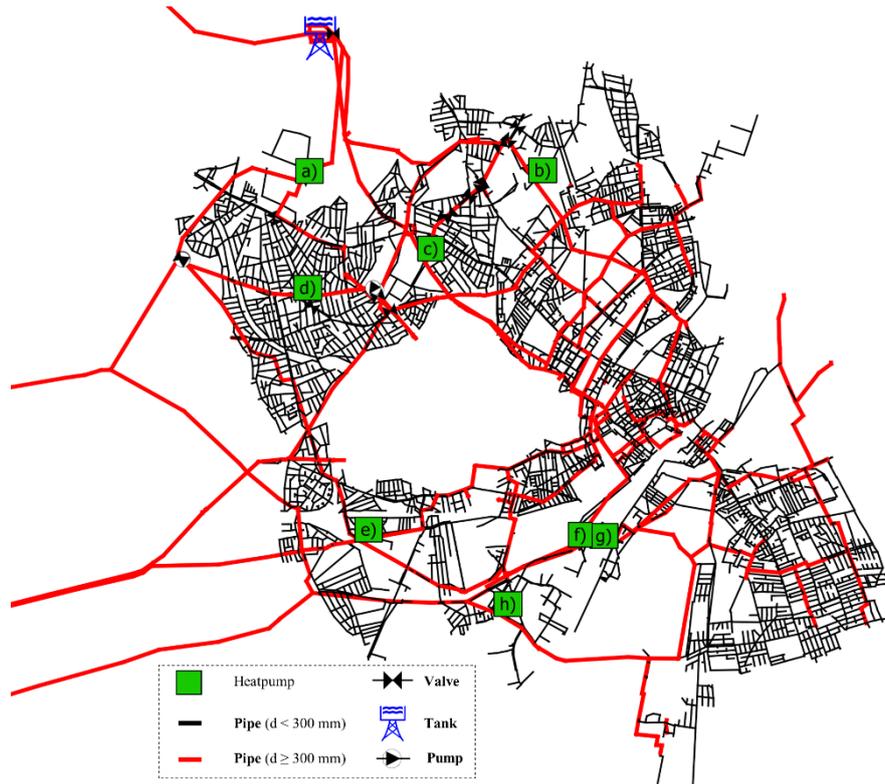


Figure 4. Locations of modelled heat pumps on the water distribution mains (red). Heat pumps f) and g) are located in low-temperature DH zone, while the other heat pumps are located in high-temperature DH zones.

Table 4. Average daily water flow rate and modelled average source temperature in February at eight HP locations a–h (see Fig. 4).

Heat pump evaporator location	a)	b)	c)	d)	e)	f)	g)	h)
February mean source temperature [°C]	8.7	8.7	8.6	8.6	8.5	8.4	8.4	8.4
Daily mean flow rate [L/s]	305	225	158	190	96	443	68	191

2.8. Heat transfer in service lines

Service lines or private ground lines contribute to the total heat transfer between soil and water, as they increase the residence time through a) additional pipe length and b) velocities being particularly low in these pipes. Service lines alone add 17% to the total pipe length of the system [20]. Moreover, heat transfer in small-diameter pipes may be relatively large due to the high ratio of surface area to mass of water. The magnitude of the extra heat transfer in service lines was estimated outside the model.

The number of service lines in the City of Copenhagen is about 34,500 with a total length of 185 km [20], i.e. 5 m on average. The 34,500 lines serve a modelled demand of 1366 L/s, so the flow rate in a service line would be 0.04 L/s as a daily average. The typical service line has a diameter of 50 mm and was assumed made of PE. 34,500 such service lines were assigned to model nodes with a demand (a total of 8,674 nodes). Given that any node may contain several end users, the service lines were assigned by division of the demand in each node by the average demand of 0.04 L/s.

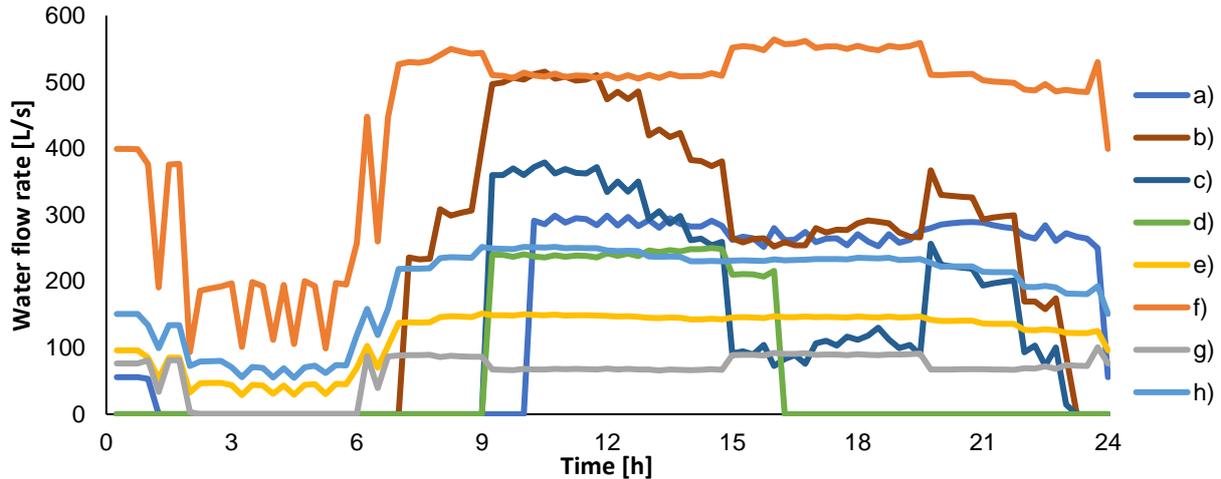


Figure 5. Daily water flow rates at the 8 selected heat pump locations. The flow rate during hours with changing direction have been set equal to zero since the heat pumps do not operate during this period. Three locations (e, f and h) have unidirectional flow for 24 hours. In order to avoid excessively cold temperatures, the operation of node d is particularly constricted since a part of its stream has already been cooled by another heat pump during several hours.

For case A and B, the additional heat transfer across the average service line was computed with the analytical solution to Eq. 12, while using the output temperatures from the EPANET-MSX simulation in February as input. The resulting new output temperatures were used to compute a new $\Delta\dot{Q}_{\text{soil}}$ that included heat transfer in service lines.

2.9. Neighbouring water distribution networks

The EPANET model contained four nodes that represented a special case since the water from these nodes is discharged to neighbouring municipalities. In reality, the water is distributed onwards in smaller, local distribution networks that were not included in the EPANET model. The four nodes had a total discharge of 184.5 L/s and were located downstream of the HPs. As all of the discharged water was affected by the HPs, these sub-networks had to be considered in the energy balance. To account for the heat transfer taking place in these local supply networks (including service lines), the results for the supply network of Copenhagen were extrapolated by assuming that the ratio of $\Delta\dot{Q}_{\text{demand}}$ to \dot{Q}_e was identical for water discharged in the city and water discharged in the local distributions networks.

3. Results

3.1. Model validation

Fig. 6 shows a comparison between the measured T_{soil} (see section 2.5.2), the measured air temperature, and a modelled soil temperature at 1.5 m depth (see section 2.6). The figure also shows the best fits of the measured air and soil temperatures to Eq. 14, as well as the yearly mean values for both.

The measured T_{soil} has the flatter curvature expected for a temperature measured at some depth. Theoretically, T_{soil} would lie somewhere between T_{air} and $T_{\text{soil,model}}$ if atmospheric air were the only heat source affecting the soil temperature (as assumed in Eq. 14). As it appears, however, the yearly mean value of T_{soil} (12.5 °C) is 3 K higher than the yearly mean air temperature in Copenhagen (9.3 °C). This difference probably reflects that the impact from direct solar radiation and/or non-natural heat

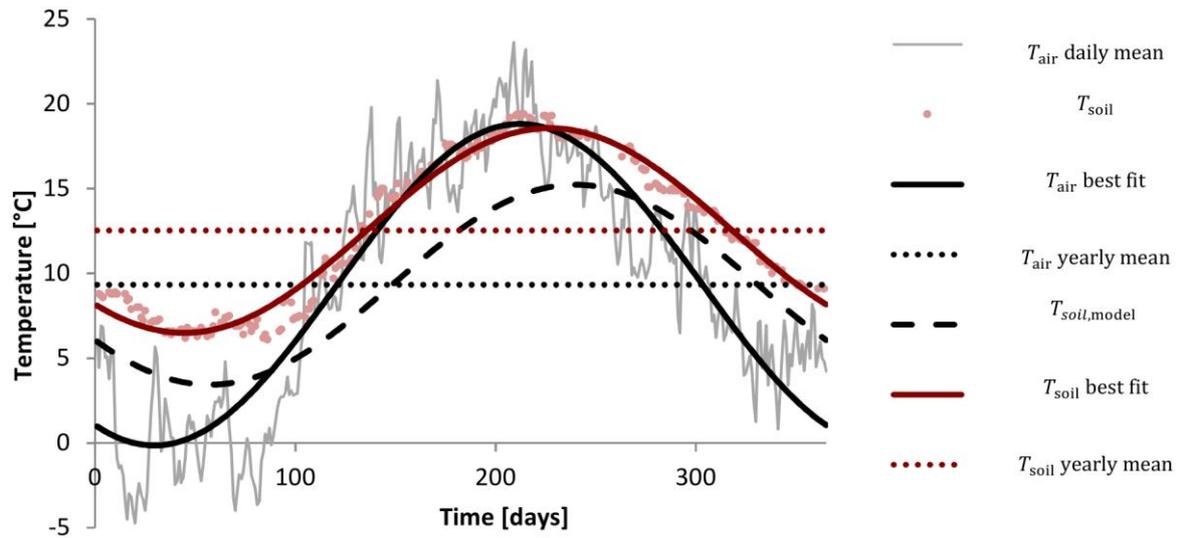


Figure 6. Daily values of measured soil temperature, T_{soil} , from 2013 (red dots), its best fit to a sinusoidal function (red, fully drawn), and its yearly mean (red, dotted); daily mean measured air temperature, T_{air} , from a private weather station [34] in Copenhagen in 2013 (grey), its best fit to a sinusoidal function (black, fully drawn), and its yearly mean (black, finer dotted); and a sinusoidal soil temperature modelled for a depth of 1.5 m, $T_{soil,model}$ (black, longer dotted).

sources is too significant to ignore. Other empirical soil temperature studies have showed similar results, i.e. of air temperature-based models underestimating the soil temperature by 2 K to 4 K [32, 33]. Based on these considerations, the measured soil temperature was deemed a plausible estimate for the true soil temperature.

The water temperature is measured every week at 20 end-use points on the distribution network. The model results were compared with measurements from 15 of these sampling points, the locations of which were all represented in the model. The full network model was thus compared with weekly measurements in 15 points from week 28 of 2012 to week 36 of 2014 (see Appendix).

The Pearson's product-moment correlation was computed to assess the relationship between the measured and modelled water temperature. The results indicate a strong positive correlation between the two variables that is statistically significant ($r = 0.95$, $n = 1497$, $p < 0.001$). The best fit of a linear regression model of the modelled vs. measured water temperatures (Fig. 7) yields an $R^2 = 0.91$. Results from weeks 30-32 of 2013 have been left out from the regression since a major interruption in operation happened during that period [35], and the corresponding measurements (the outliers clearly visible in the Appendix) appear unusual.

The regression line reveals a slight skew; the model slightly overestimates water temperatures in the winter and underestimates them in the summer. This would be the case if the model somewhat underestimates heat transfer. This may be due to the lack of heat transfer in service lines in the model, which is, conversely, reflected in the measurements. It is also possible that the thermal conductivity of the soil is larger than 2.2 W/m/K. In turn, the skew is not consistent with a different soil temperature input such as T_{air} or $T_{soil,model}$ (Fig. 6).

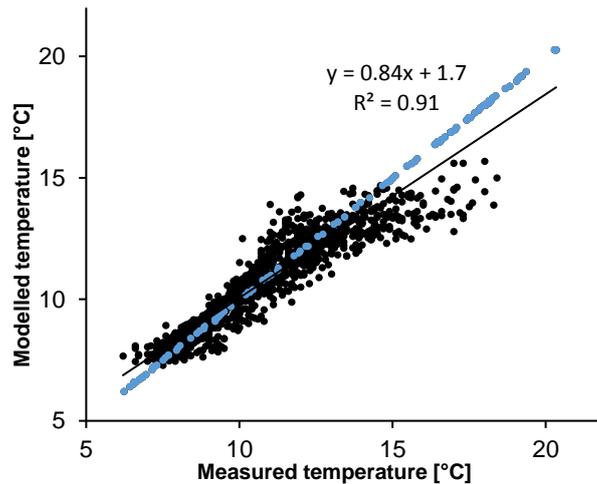


Figure 7. Modelled water temperatures vs. weekly measured temperatures (1497 samples from 15 sampling points). The blue dots represent the peripheral location used to determine the soil temperature. The main deviations between simulated and measured results as well as the largest weekly fluctuations in measured temperatures appear during the summer holiday season (see Appendix). Furthermore, the bias is particularly dominated by results for three locations that are possibly influenced by specific local conditions (see Appendix).

Since the correlation is strong, and the skew is not systematically present across sampling points, the heat transfer model was considered credible and suitable for the case study.

3.2. Characteristics of the system and degree of exploitation of the heat source

Figs. 8 and 9 show the temperature distribution from EPANET-MSX of the total amount of supplied water in February and August, respectively, in both case A and B with a ΔT_{source} of 5 K. In August, less than 1% of the distributed water reaches the soil temperature of 18.9 °C in the distribution network, whether or not HPs are used (Fig. 9). For 99% of the network the residence time of the water is shorter than the time required to heat the water to soil temperature. If the residence time of all water in the system were high enough, or if the heat flux in the network were higher, the water would approach the soil temperature, thereby decreasing the additional end-use heating demand. This substantiates the need to address the additional heat transfer occurring in service lines, where the residence time is prolonged due to additional pipe lengths and low velocities.

Case B in Fig. 8 provides a few insights on the degree to which the drinking water heat source is exploited. 69% of the supplied water (the part with temperatures below 6.7 °C) has been subjected to cooling by the HPs and has to some extent regained thermal energy from the soil. 31% of the total water supply (the part with temperatures above 6.7 °C) has not been in contact with HPs, or has been significantly mixed with unexploited flows. This is water that passes the HPs while they are not operating (at night), and water that leaves the distribution mains before reaching the HPs. 2.7% of the water supply reaches the soil temperature.

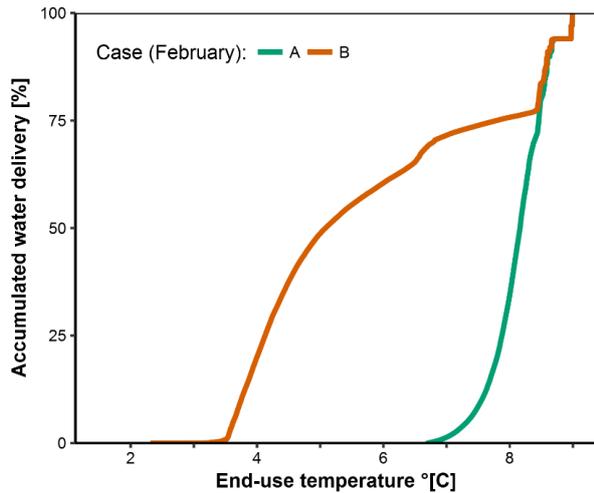


Figure 8. Temperature in the distribution mains of the accumulated water delivery in February without (A) and with heat pumps and a ΔT_{source} of 5 K (B).

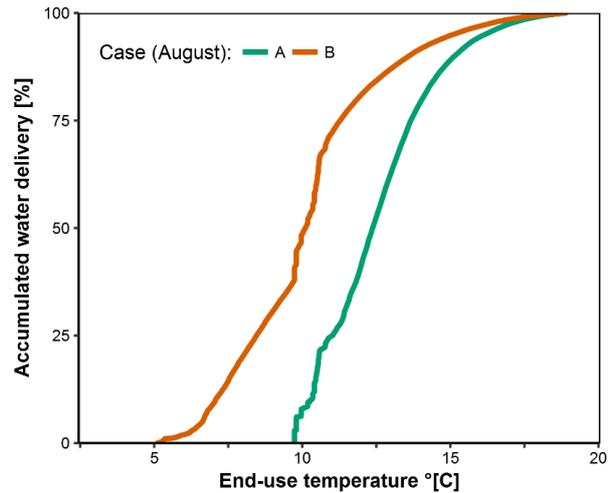


Figure 9. Temperature in the distribution mains of the accumulated water delivery in August without (A) and with heat pumps and a ΔT_{source} of 5 K (B).

3.3. Potential for heat recovery from Copenhagen's water distribution network incl. service lines

The results of simulating the heat transfer in the water distribution network and service lines are shown in Fig. 10 where \dot{Q}_{net} is plotted against \dot{Q}_e for five different values of ΔT_{source} (1 K to 5 K). The change in the system energy flows resulting from a ΔT_{source} of 5 K is shown in Fig. 11 as a daily average. \dot{Q}_{net} in this case ($\dot{Q}_e - \Delta\dot{Q}_{demand}$) is 20.7 MW. The utilization degree, η_{util} , at any ΔT_{source} is 71% of the thermal energy extracted at the evaporator (Fig. 10). The ground-source heat of Copenhagen's soils on its own returns 38% of the extracted heat to water consumers through additional heat transfer to the distribution and service lines (Fig. 12). Of the remaining 62% of the extracted energy, 20% (equal to 12%-points) is kept by the DH system because the water fraction for cold water consumption needs no reheating by end-users. Another 33% (or 20%-points) is kept by the DH system because this water fraction is not reheated by DH, but by electricity (Fig. 12).

The magnitude of the heat extraction does not affect the degree of utilization, so the technical potential for heat utilization, \dot{Q}_{net} , is the net potential resulting when the installed HP capacity is maximised without causing unacceptably low end-use temperatures. The potential of 20.7 MW is obtained with a ΔT_{source} of 5 K and an average evaporator yield of 29.2 MW. With a ΔT_{source} of 5 K, 8% of the modelled demand had temperatures between 3 °C and 4 °C, but these are assumed to increase to acceptable levels in indoor piping installations affected by the ambient air temperature.

Given a COP of 3.2 for HPs in low-temperature DH zones, and 2.8 in high-temperature DH zones (the weighted average COP being 2.9), a heating potential of 44.4 MW is obtained (Fig. 11). Subtracting the additional heating demand of 8.5 MW gives a net heating potential of 35.9 MW.

The system COP is only 1.7, which owes not least to the increased end-use electricity demand for water heating that is nearly half as large as \dot{W}_{com} (Fig. 11).

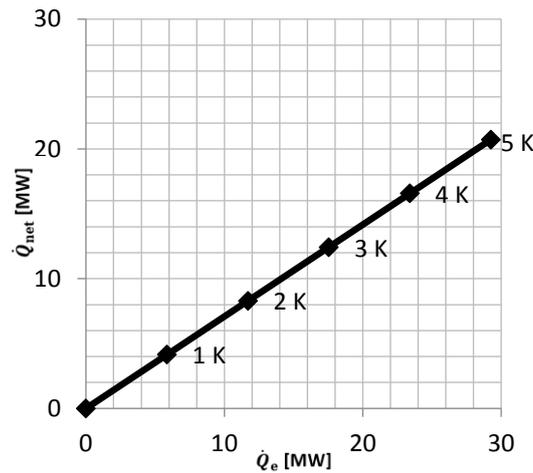


Figure 10. Potential for heat utilization (\dot{Q}_{net}) vs. heat extraction (\dot{Q}_e). The temperature reduction of the source is shown.

Case B minus A

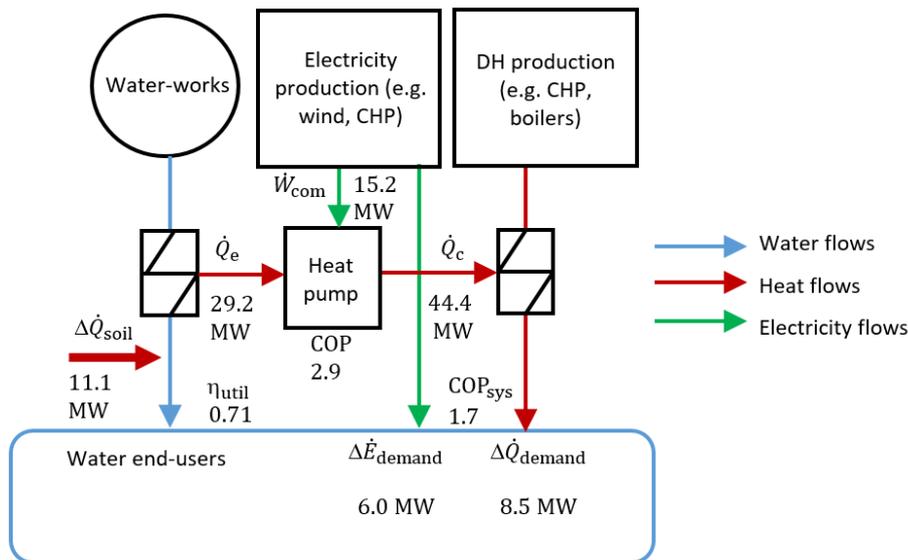


Figure 11. Change in the system energy flows resulting from a ΔT_{source} of 5 K.

The results indicate that a full-scale application of large HPs in the drinking water network would cover around 2.6% of the DH peak demand, after the additional end-use DH demand has been subtracted (Fig. 13).

3.4. Significance of heat pump COP

A HP elaborately designed for the case system could reach a COP higher than the values of 2.8 and 3.2 found in this study. However, as a number of factors affect the COP, it is also possible that a lower efficiency may be obtained.

Fig. 14 shows the system COP at various values of the HP COP. A HP COP of e.g. 2.0 reduces the system COP to 1.4, which is only a little more than that of an electric boiler. Increasing the COP to

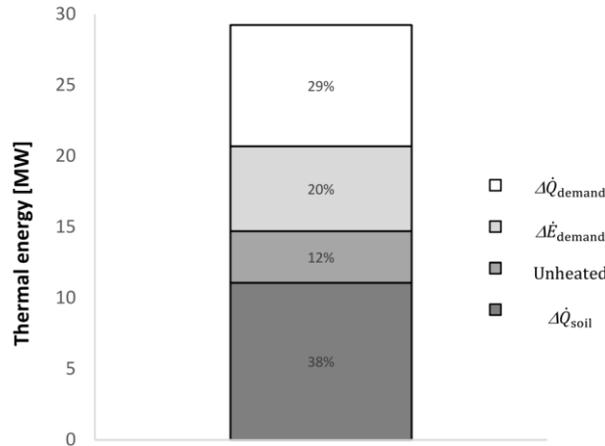


Figure 12. Heat utilization in February with heat pumps and a ΔT_{source} of 5 K. The full height of the bar is, the energy extracted from the water supply. The sum of the grey bar elements is the heat utilization from drinking water in the DH system; and their share of the full bar height is η_{util} . The white bar element is the needed water re-heating provided by the DH system. Also shown are the contributions to heat utilization from the soil, the unheated share of drinking water, and the water share re-heated by electricity.

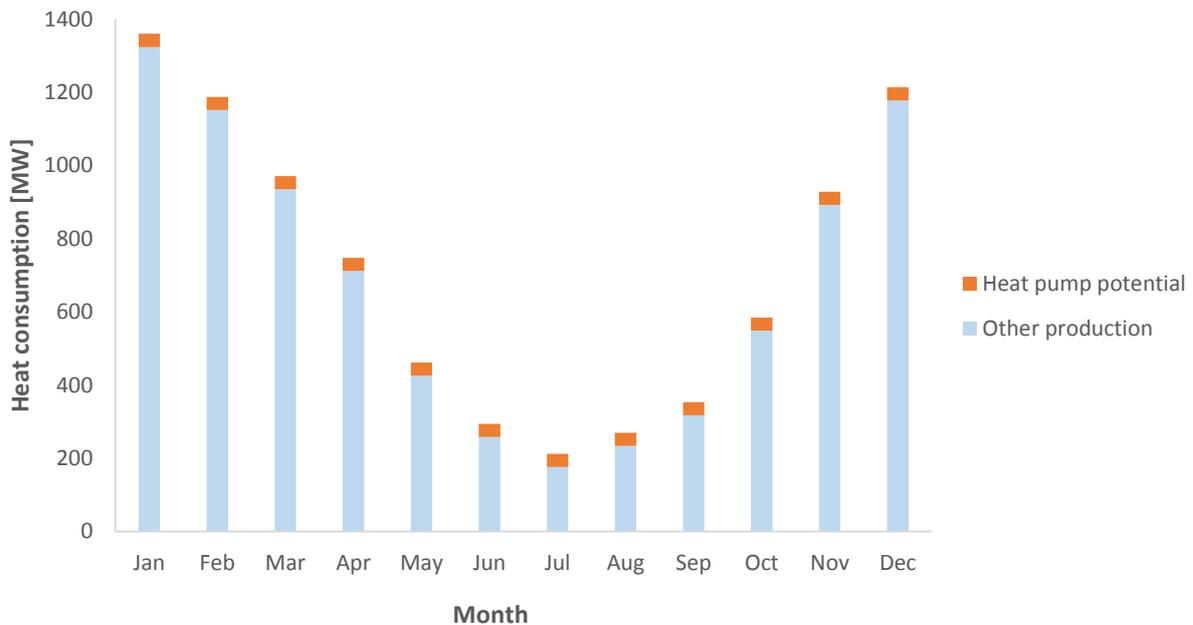


Figure 13. Drinking water heat pump potential as a share of the estimated annual DH demand [36]. The heat demand applies to Copenhagen and the other municipalities whose drinking water is utilized as a heat source in this study (the Greater Copenhagen DH system in its entirety covers a larger demand than what is shown).

4.0 increases the system COP to just 1.9. This difference is due to the fixed end-use power and heat consumption for water heating.

3.5. Significance of pipe and soil thermal resistances

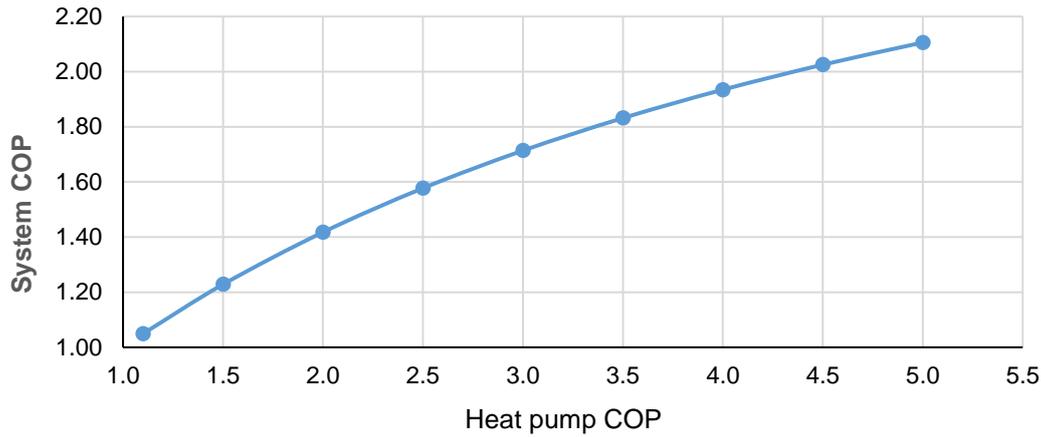


Figure 14. System COP vs. heat pump COP.

An assessment was made of the impact of converting the network from mainly cast iron to PE, which has a lower thermal conductivity. However, the resistance of the pipe material is insignificant compared to the soil resistance (Table 5). The results indicate that the thermal resistance of soil is the main impediment to heat transfer, leaving the pipe material to be of little importance.

Table 5. Sensitivity analysis for different thermal resistances included.

Thermal resistances included	\dot{Q}_{net} [MW]	η_{util}	$\eta_{util,soil}$
Soil, water, pipes of cast iron etc. (present network)	20.7	0.71	0.38
Soil, water, pipes of PE (future network; main results)	20.7	0.71	0.38
Water, pipes of PE	25.3	0.86	0.71

The exact magnitude of the soil thermal conductivity in Copenhagen is uncertain, and it appears to be of some significance. A 32% increase or decrease in the thermal conductivity of soil changes $\eta_{util,soil}$ by 11% to 15% (Fig. 15). However, η_{util} only varies between 2% and 4%, so the potential for heat utilization is not affected much.

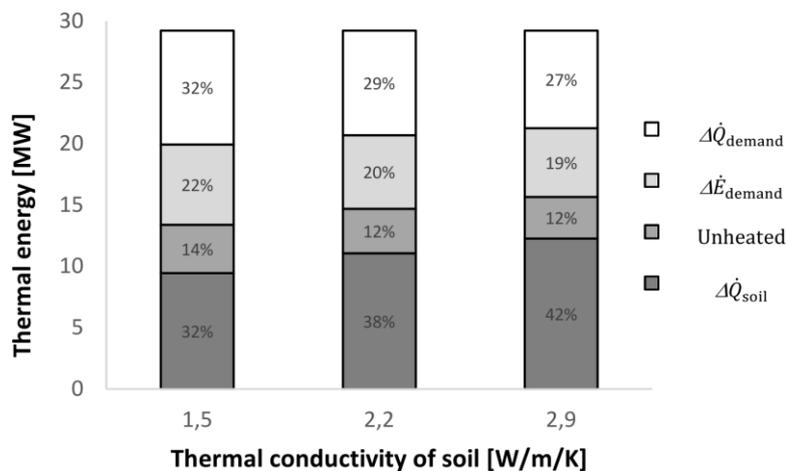


Figure 15. Heat utilization at different thermal conductivities of soil. The full height of each bar is \dot{Q}_e , the energy extracted from the water supply. The sum of the grey bar elements is \dot{Q}_{net} , the heat utilization from drinking water in the DH system; and their share of the full bar height is η_{util} . The percentage label on $\Delta\dot{Q}_{soil}$ is $\eta_{util,soil}$.

3.6. Water cooling

HP operation is currently expected to be uneconomical during the summer when Copenhagen's entire DH demand is met by waste incineration. However, with growing fractions of waste being recycled instead of incinerated, a part of the summer heat demand may be covered by other production technologies in the future. Furthermore, operating HPs during the warm season would contribute to preventing high drinking water temperatures. Danish regulation [37] recommends an upper temperature limit of 12 °C at the tap, but large areas in Copenhagen have difficulties meeting this limit in warm periods. HP operation in August with a $\Delta T_{\text{source}} = 5$ K increased the share of water delivered from distribution lines complying with the 12 °C limit from 42% to 81% (Fig. 9).

4. Discussion and potential future developments

4.1. Comparison with a previous case study

As previously mentioned, the EPANET-MSX modelling approach has been used for modelling HPs in drinking water systems in Almere [10], Głogów [12] and Milan [11]. For the case study in Almere, the heat utilization of the system was reportedly very high, but the heat transfer model used there did not take into account the thermal resistance of soil (cf. Table 5). The case study in Głogów assumed no need entirely for the reheating of water in the buildings. Moreover, the COP of 5.2 assumed in the Głogów study [12], seems optimistic, while the COPs of 2.8 and 3.2 used in our study may be slightly conservative.

In the present work, the heat utilization degree from soil (38%) is more than three times higher than found by De Pasquale et al. (10.3%) in Milan [11]. De Pasquale compared their heat transfer model with a model similar to the one used in the present study, labelled the shape factor model. They used air temperature as input to the shape factor model, as opposed to a modelled soil temperature at 1.5 m depth for the De Pasquale model. Their results indicated that the shape factor model overestimated heat transfer compared to the De Pasquale model.

Using our heat transfer model as the basis, Fig. 16 shows the predicted heat transfer per unit pipe length as a function of three different soil temperature inputs displayed in Fig. 6. A case where the term H has been replaced by a fictitious soil layer of magnitude $l = \sqrt{\kappa/\omega}$ was also assessed to mirror the main difference between the shape factor model and the De Pasquale model.

Air temperature is shown to predict larger heat transfer than the modelled soil temperature at 1.5 m depth (Fig. 16). The fictitious soil layer appears to be an insignificant alteration of the shape factor model when identical soil temperatures (at 1.5 m depth) were applied for both cases. We therefore conclude that the difference between the De Pasquale model and the model used in our study is in fact insignificant.

The difference in results is also not a consequence of the different choice of soil temperature. The empirical soil temperature in our case study provided a more accurate basis for validating the model's performance against measurement data. But whereas the soil temperature influences heat transfer, it does not influence the heat utilization degree.

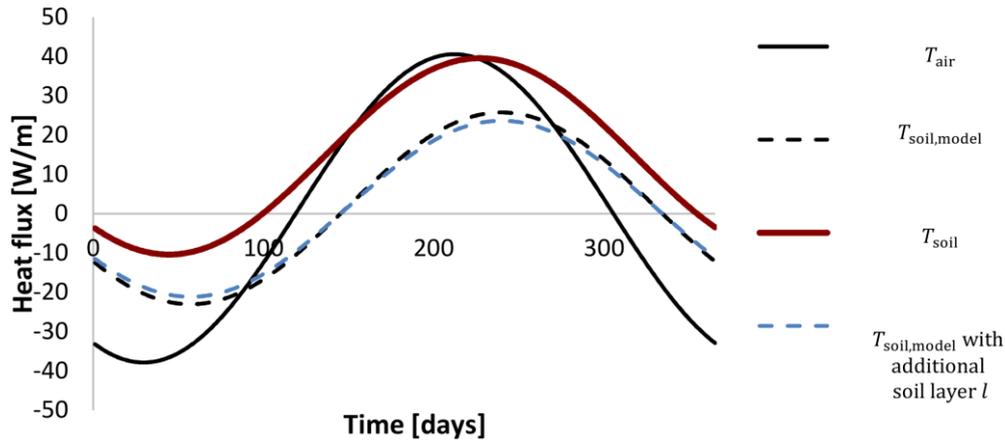


Figure 16. Heat flux per unit pipe length when the three temperature curves in Fig. 6 (T_{air} , T_{soil} and $T_{soil,model}$) are applied as input, and when (in a fourth case) an additional fictitious soil layer is added to the heat transfer model (blue, dotted). The inner pipe diameter is 0.2 m, and the pipe material is cast iron. The water temperature is 9 °C, and the water flow rate is 10 L/s.

Instead, the different results must pertain to the characteristics of the two systems. Firstly, the soil conditions in Copenhagen and Milan are different; Copenhagen has waterlogged clay soils that conduct heat relatively well, resulting in more favourable conditions for heat utilization. Sensitivity analysis showed that reducing the thermal conductivity of soil from 2.2 to 1.5 W/m/K reduced $\eta_{util,soil}$ from 38% to 32%. Disregarding the heat transfer in service lines would further reduce it to 28%.

Secondly, the residence time of water in the pipes is likely to have an impact. Our modelled system has a demand nearly 7 times greater than the system studied by De Pasquale et al., and a number of pipes almost 11 times greater. The difference between these two ratios may support the assumption that the general residence time of water is higher in the Copenhagen network, which might explain the remaining difference to the Milan results. The layout of the pipe network (dimensions, length) and water velocities are key factors determining how much heat can be reaped. Individual analysis of any given water supply system therefore seems necessary to accurately assess its potential as a ground-source heat system.

4.2. Potential further model improvements

The empirical undisturbed soil temperature used in this study is considerably higher than the modelled temperatures commonly used (e.g. by [38]) to represent soil temperature at shallow depths in Denmark, but it proved more valid for Copenhagen. One limitation of this study is that when HPs operate, the radius of the soil layer that is thermally disturbed by the water pipes may increase, and so the undisturbed temperature may change as well. Developing a transient model that accounts for this situation could be a meaningful improvement. Service lines add significantly to the total pipe length and residence time of water in a system, and in our case increased $\eta_{util,soil}$ from 34% to 38%. Due to a lack of data, our study did not include private ground lines extending from the end of service lines to the buildings. Moreover, our use of the hourly average water flow rate to compute the heat transfer in service lines is a simplification. A typical domestic demand pattern may involve great fluctuations and sustained periods with no use at all. How demand patterns and private ground lines influence the residence time and heat transfer deserves further investigation.

4.3. Competing heat sources and economic considerations

So far, the evidence points to a heat utilization degree of around 71% in Copenhagen, meaning that 29% of the evaporator capacity of a drinking water HP will be dedicated to covering the additional DH demand that arises due to its operation. The system COP is also affected, not least by the increased electricity demand. We therefore suggest that other potential heat sources should be considered before drinking water, in order to avoid the social costs involved.

Competing low-temperature sources for large HPs include air, seawater, wastewater and groundwater, but not ordinary ground-source heat. Ground-source HPs delivering 10 MW would require an area of 1 million m² for horizontal piping [39], which is hardly attainable in a dense urban environment. Ground-source HPs would also require significant investment and installation costs for new piping, which for drinking water HPs would be sunk costs.

Competing heat sources for HPs are not abundant everywhere [3], and in specific locations of interest drinking water HPs may be the only feasible option. In such cases, only an economic feasibility study will reveal if the benefits of their operation offset the cost of the additional demand. Such economic analysis depends on details related to the alternative heat supply technology, fuel and electricity costs, as well as taxation or subsidy schemes. A complete assessment considering both private-economic, such as presented in [2], and socio-economic analysis should be carried out to reflect the real value of new capacity.

4.4. Heat pump potential and system boundaries

In [3] the estimated potential heating capacity of HPs using drinking water as heat source in Greater Copenhagen was 13.5 MW. The present work shows a potential more than 50% higher. The principal reason for this difference is that [3] only included reservoirs and waterworks, and not the water distribution mains.

This study has focused on the potential for applying large drinking water HPs in the City of Copenhagen, and this implied exploitation of roughly 69% of the water supply. The remaining 31% of the water is either distributed upstream of the HPs, or passes the HPs while they are not operating at night. A real implementation of HPs might permit some operation during the night if it would involve a system control of a set minimum exit temperature of the water. This would prevent excessive cooling even as the water changes direction several times. A minimum number of full load hours for the HP would be required to cover the installation cost. Thus, for several of the HP locations, the challenge with limited potential hours of operation (Fig. 5) would have to be overcome, or they could turn out infeasible.

The water that is distributed upstream of HPs represents an unexploited potential that could be harnessed by means of HPs upstream of Copenhagen or HPs smaller than 1 MW inside Copenhagen. Potential HP locations in neighbouring municipalities might thus be worth considering, not least the reservoir, which is located outside the city but within the same DH transmission area. Not only is the reservoir more accessible for installation and maintenance works; it also circumvents the problem with water flow reversal limiting HP operation on several main lines during the night.

5. Conclusion

This study investigated the potential for large drinking water HPs in Copenhagen's DH system, based on a number of technical constraints for the HPs, including their operational profile, their distance to a DH network, and differences in supply temperatures. The potential was assessed by determining the energy balance for drinking water through simulation of the ground-source heat that is transferred to the water distribution lines after HPs have cooled the water. A heat transfer model was developed for simulation of system-wide temperatures in a piped water network using the EPANET-MSX model. Our simulations were validated against measured temperatures in an actual drinking water network, and the use of empirical urban soil temperatures proved key to achieving a valid model.

Model simulations showed that the ground-source heat of Copenhagen's soils returned 38% of the extracted heat to water consumers through additional heat transfer to the distribution lines. This is more than found in a previous study, probably due to better thermal properties of the soil and higher residence times of water in the supply system.

A share of the drinking water is not re-heated by end-users. Taking that share into account, the results showed a potential 71% utilization degree for water distribution mains used as a heat source for HPs in Copenhagen's DH system.

The simulations showed a total net potential of 20.7 MW (evaporator capacity) for eight large drinking water HPs. With HP COPs ranging from 2.8 to 3.2, the resulting net heating potential was 35.9 MW corresponding to 2.6% of the DH peak demand. The corresponding system COP of the drinking water HPs was 1.7. The results indicate that drinking water HPs should only be considered where other heat sources are unavailable.

The study also investigated the derived cooling effect experienced by end-users. This is an added value to be considered from drinking water HPs in warmer periods. Simulation of HP operation in August increased the share of distributed water complying with the 12 °C recommended upper limit from 42% to 81%.

Acknowledgements

We thank HOFOR A/S for providing the hydraulic network model and water temperature measurement data used in this study. We thank CTR I/S for providing air temperature measurement data.

Appendix

The heat transfer model results were compared with water temperature measurements from 15 sampling points at end-use nodes across the network. The full network model was thus compared with weekly measurements in 15 points from week 28 of 2012 to week 36 of 2014. Since sampling is undertaken in the morning, averaged model results for the five hours between 7 a.m. and 12 p.m. were used.

The model predicts the water temperature well for the 15 sampling locations, with a few exceptions. Locations m), n) and o) have measured summer temperatures somewhat higher than the modelled ones. It is possible that these test locations are influenced by specific local conditions such as very long service lines. A map of the three locations indicates service line lengths between 23 m and 50 m

[40]. For these locations, the exact timing of the water sampling also seems key; the model predicts significant temperature differences between early morning hours and noon, which may explain the disparity.

Location f) has modelled temperatures generally higher than the measured ones. The locality of this sampling point has undergone significant development in recent years, so it is quite possible that the water demand has increased and the residence time has decreased since the hydraulic model was developed. This would explain the difference. In general, small alterations and simplifications in the network model compared to the real set-up can explain deviations between simulated and observed results.

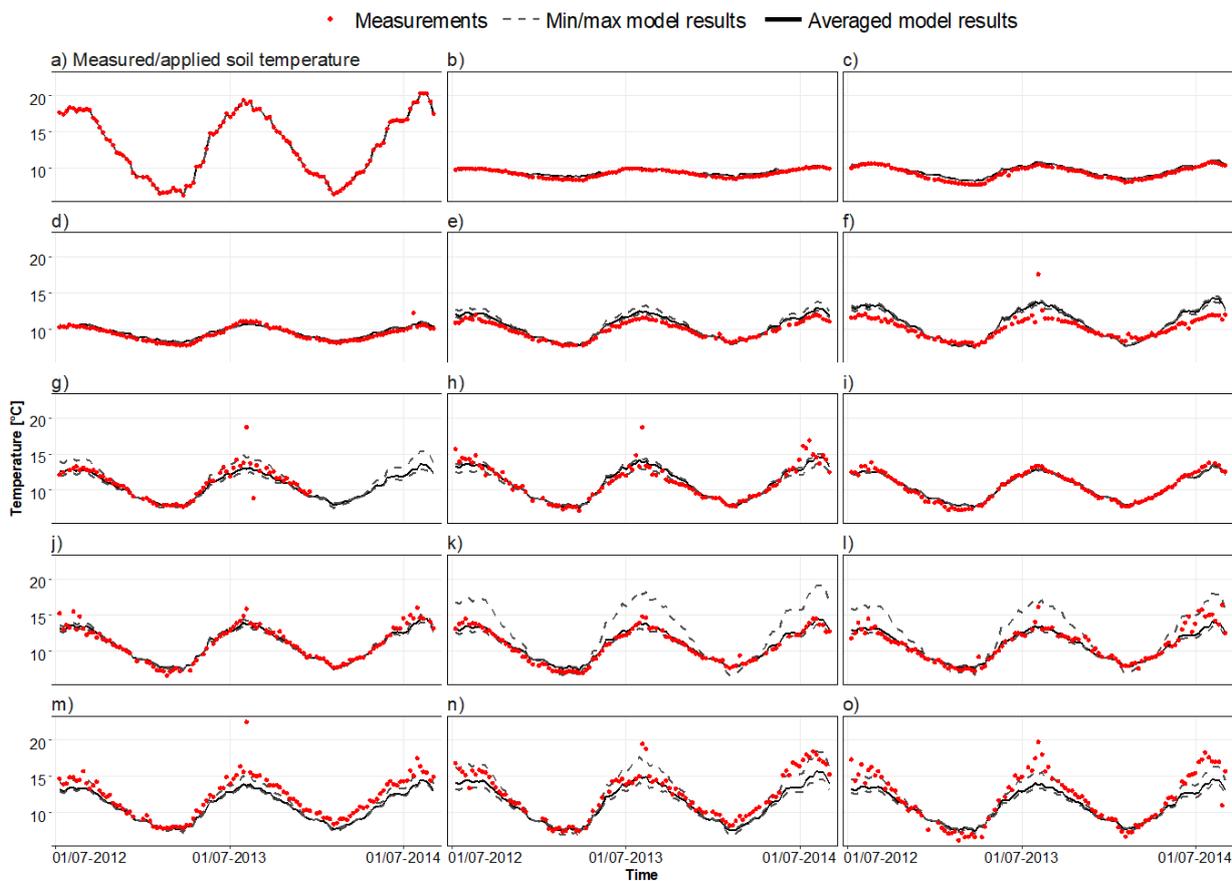


Figure 17. Measured and modelled water temperature data from 15 locations in Copenhagen from July 2012 to August 2014. Model averages are shown together with the minimum and maximum values obtained between 7 a.m. and 12 p.m.

References

- [1] H. Lund, S. Werner, R. Wiltshire, S. Svendsen, J. E. Thorsen, F. Hvelplund, B. V. Mathiesen, 4th Generation District Heating (4GDH) Integrating smart thermal grids into future sustainable energy systems, *Energy* 68 (2014) 1-11.
- [2] H. Pieper, V. Mašatin, A. Volkova, T. Ommen, B. Elmegaard, W.B. Markussen, Modelling framework for integration of large-scale heat pumps in district heating using low-temperature heat sources: A case study of Tallinn, Estonia, *International Journal of Sustainable Energy Planning and Management* 20 (2019) 67-86.
- [3] B. Bach, J. Werling, T. Ommen, M. Münster, J.M. Morales, B. Elmegaard, Integration of large-scale heat pumps in the district heating systems of Greater Copenhagen, *Energy* 107 (2016) 321-334.

- [4] S.S. Cipolla, M. Maglionico, Heat recovery from urban wastewater: Analysis of the variability of flow rate and temperature, *Energy and Buildings* 69 (2014), 122-130.
- [5] M. Abdel-Aal, A. Schellart, S. Kroll, M. Mohamed, S. Tait, Modelling the potential for multi-location in-sewer heat recovery at a city scale under different seasonal scenarios, *Water Research* 145 (2018), 618-630.
- [6] P. Fleuchaus, B. Godschalk, I. Stober, P. Blum, Worldwide application of aquifer thermal energy storage – A review, *Renewable and Sustainable Energy Reviews* 94 (2018), 861-876.
- [7] R. de Graaf, F.H.M. van de Ven, I.J. Miltenburg, B. van Ee, L.C.E. van de Winckel, G. van Wijk, Exploring the technical and economic feasibility of using urban water system as a sustainable energy source, *Thermal Science* 12(4) (2008), 35-50.
- [8] S. Liu, C. Gunawan, N. Barraud, S. A. Rice, E. J. Harry, R. Amal, Understanding, monitoring and controlling biofilm growth in drinking water distribution systems, *Environmental Science & Technology* 50 (2016) 8954-8976.
- [9] J.A. Elías-Maxil, J.P. van der Hoek, J. Hofman, L. Rietveld, Energy in the urban water cycle: Actions to reduce the total expenditure of fossil fuels with emphasis on heat reclamation from urban water, *Renewable and Sustainable Energy Reviews* 30 (2014), 808-820.
- [10] E. J. M. Blokker, A.M. van Osch, R. Hogeveen, C. Mudde, Thermal energy from drinking water and cost benefit analysis for an entire city, *Journal of Water and Climate Change* 4:1 (2013) 11-16.
- [11] A.M. De Pasquale, A. Giotri, M.C. Romano, P. Chiesa, T. Demeco, S. Tani, District heating by drinking water heat pumps: Modelling and energy analysis of a case study in the city of Milan, *Energy* 118 (2017) 246-263.
- [12] P. Jadwiszczak, E. Niemierka, Water pipe network as a heat source for heat pump integrated into a district heating, *E3S Web of Conferences* 22 (2017), art. no. 00210
- [13] J.P. van der Hoek, S. Mol, S. Giorgi, J.I. Ahmad, G. Liu, G. Medema, Energy recovery from the water cycle: Thermal energy from drinking water, *Energy* 162 (2018), 977-987.
- [14] Greater Copenhagen Utility, 2012. HOFOR, EPANET model of Copenhagen, version 2009, modified 2012. Received from C. Hansen on April 11 2014.
- [15] L. A. Rossmann, EPANET 2 Users Manual, National Risk Management Research Laboratory, Cincinnati, 2000.
- [16] F. Shang, J. G. Uber, L. A. Rossmann, EPANET Multi-Species Extension User's Manual, University of Cincinnati, Cincinnati, 2008.
- [17] E. J. M. Blokker, E. J. Pieterse-Quirijns, Modeling temperature in the drinking water distribution system, *Journal American Water Works Association* 105 (2013) E19-E28. Erratum (2018) in *Journal American Water Works Association* 110 (2018), 98-98. doi:10.1002/awwa.1178
- [18] A. Moerman, M. Blokker, J. Vreeburg, J.P. van der Hoek, Drinking water temperature modelling in domestic systems, *Procedia Engineering* 89 (2014) 143-150.
- [19] M. Rygaard, L. Alsbjörn, M. Ejsing, B. Godskesen, R. Hansen, B. Hoffmann, C. Jørgensen, K. Ledgaard, H.-M. F. Møller, M.-B. B. Poulsen, M. Schmidt, L. Tarp-Johansen, S. Vigsø, K. Zambrano, Sekundavand i Nordhavn - En forundersøgelse til strategi for alternativ vandleverance, Technical report, Technical University of Denmark, Lyngby, 2013.
- [20] Københavns Kommune, Vandforsyningsplan 2012, Teknik og Miljøforvaltningen, København, 2012.
- [21] A. B. Lauritsen, Varme Ståbi, Nyt Teknisk Forlag, Copenhagen, 2009.
- [22] B. Kvisgaard, S. Hadvig, Varmetab fra fjernvarmeledninger. Heat loss from pipelines in district heating systems, Technical report, Technical University of Denmark, Lyngby, 1980.
- [23] Y. A. Çengel, Heat and mass transfer. A practical approach, McGraw-Hill, New York, 2007.
- [24] B. Bøhm, On transient heat losses from buried district heating pipes, *International Journal of Energy Research* 24 (2000), 1311-1334.
- [25] M. Ouzzane, P. Eslami-Nejad, M. Badache, Z. Aidoun, New correlations for the prediction of the undisturbed ground temperature, *Geothermics* 53 (2015) 379-384.
- [26] A. D. Rosa, H. Li, S. Svendsen, Method for optimal design of pipes for low-energy district heating with focus on heat losses, *Energy* 36 (2011) 2407-2418.
- [27] C. Hansen, personal communication. Emails to author from May 6, September 24 and October 29 2014.
- [28] C. Ditlefsen, I. Sørensen, D6 Overfladenære jordarters termiske egenskaber, 2014. (Technical report, Energistyrelsen, Energiteknologisk Udviklings- og Demonstrations-Program).

- [29] T. Ommen, W.B. Markussen, B. Elmegaard, Heat pumps in combined heat and power systems, *Energy* 76 (2014) 989-1000.
- [30] T. Ommen, J.K. Jensen, W. Meeseburg, P.H. Jørgensen, H. Pieper, W.B. Markussen, B. Elmegaard, Generalized COP estimation of heat pump processes for operation off the design point of equipment, *Proceedings of the 25th IIR International Congress of Refrigeration*, 2019.
- [31] D. Hillel, *Introduction to soil physics*, Academic Press, San Diego, CA, 1982.
- [32] J. Wu, D. L. Nofziger, Incorporating temperature effects on pesticide degradation into a management model, *Journal of Environmental Quality* 28 (1999) 92-100.
- [33] T. Kurevija, D. Vulin, Determining undisturbed ground temperature as part of shallow geothermal resources assessment, *Rudarsko-geološko-naftni zbornik* 22 (2010) 27-36.
- [dataset] [34] CTR I/S, Data from weather station, 2013.
- [35] Greater Copenhagen Utility, HOFOR, Rapport HOFOR Vand og Spildevand 2013, Technical Report, HOFOR A/S, København, 2014.
- [dataset] [36] Greater Copenhagen Utility, HOFOR, Estimated annual DH demand distribution, 2019.
- [37] BEK nr 1068 af 23/08/2018. Retrieved on <https://www.retsinformation.dk/forms/R0710.aspx?id=202768>, on February 1, 2019
- [38] Danmarks Meteorologiske Institut, Teknisk Rapport 2001-2010 Dansk Design Reference Year Supplerende Datasæt, Technical Report, Danmarks Meteorologiske Institut, København, 2013.
- [39] B. Bach. Integration of heat pumps in Greater Copenhagen. MSc thesis at the Technical University of Denmark, 2014.
- [40] Greater Copenhagen Utility, 2019. Service lines maps retrieved from <https://www.hofor.dk/privat/vand/kort-overstoppaner/>

VI

Using smart meter water consumption data and in-sewer flow observations for model based sewer system analysis

N.S.V. Lund, J.K. Kirstein, H. Madsen, O. Mark, P.S. Mikkelsen and M. Borup

Using smart meter water consumption data and in-sewer flow observations for model based sewer system analysis

Nadia Schou Vorndran Lund^{1,*}, Jonas Kjeld Kirstein¹, Henrik Madsen², Ole Mark², Peter Steen Mikkelsen¹ and Morten Borup¹

*corresponding author

¹DTU Environment, Technical University of Denmark, Kgs. Lyngby, Denmark. nalu@env.dtu.dk; jkir@env.dtu.dk; psmi@env.dtu.dk; morb@env.dtu.dk.

²DHI, Hørsholm, Denmark. hem@dhigroup.com; omj@dhigroup.com.

Abstract ‘Smart meters’ measuring the water consumption with a high temporal resolution at the consumers’ households are globally deployed at an increasing rate. This may provide detailed knowledge of the wastewater inflow to the sewer systems in space and time, which can open up for new types of system analyses. In this study, the quality of the smart meter data is first validated by comparing with data from various sources. Subsequently, we use a detailed hydrodynamic sewer system model to link the smart meter data from almost 2,000 consumers with in-sewer flow observations to simulate the wastewater component of the dry weather flow and to identify anomalies. The results show that smart meter data is suitable as input to a distributed sewer model since the temporal dynamics of the model results and in-sewer observations match well. Furthermore, the study indicates that the biggest uncertainty is in fact related to the in-sewer flow observations, which prevents more advanced investigations of the dry weather flow composition. The study also shows that digital system integration may be complicated by data and models being hidden in different silos within and across organisations. Overcoming the obstacles may, however, improve both offline and real-time urban drainage management.

Keywords: Anomaly detection; Distributed model; Dry weather flow; Smart meters; Urban drainage; Wastewater flow

1. Introduction

The dry weather flow (DWF) in sewer systems may consist of wastewater from households, industry and institutions; groundwater infiltration; and rain-induced infiltration. There may also be exfiltration out from the system. Knowledge of the wastewater flow component can, by comparison with in-sewer flow data, be used to estimate the addition or loss of water through infiltration and exfiltration (Cole and Stewart, 2013) and thus contribute to closing the water balance of the urban drainage system. Accurate knowledge of the sources of DWF throughout the system can be used for a wide range of management objectives:

- Optimised dimensioning (Cole and Stewart, 2013) and operation (Brito et al., 2017) of sewer systems, hereunder control during dry weather with focus on energy minimisation, or CSO minimisation if wastewater makes up a substantial part of the flow. Such control could be carried out using model predictive control (Lund et al., 2018).

- Optimised dimensioning and operation (for example, chemical dosing) of wastewater treatment plants (WWTPs) based on information on water consumer types and thus the likely constituents in the wastewater (Nguyen et al., 2018).
- Individual and more equitable household wastewater billing based on water consumption type and constituents (Nguyen et al., 2018).
- Asset management (Brito et al., 2017; Djebbar and Kadota, 1998) based on, for example, knowledge of leaky sewer pipes.
- Regulatory compliance (Brito et al., 2017).
- Estimation of total suspended solids (TSS), chemical oxygen demand (COD), organic matter, drugs, and nitrogenous pollution loads, if the DWF estimation is supported by compound measurements/estimations (Métadier and Bertrand-Krajewski, 2011; Plósz et al., 2013; Schilperoort et al., 2012).
- Improved heat recovery from sewage systems (for example, Abdel-Aal et al., 2018).
- Planning of decentralised wastewater reuse (Elías-Maxil et al., 2014).

In residential areas, the wastewater flow is dominated by the citizens' behaviour. In upstream parts of the system or in small systems, the flow is intermittent (i.e. occurring at irregular intervals) and the detailed flow dynamics may not be captured by the coarse sampling resolutions of the measurement equipment (Butler and Graham, 1995; Elías-Maxil et al., 2014). The DWF of smaller areas is thus in general more difficult to estimate (Djebbar and Kadota, 1998). Further downstream, the aggregation of different inputs from many small upstream sources results in a change in the nature of the wastewater flow and may appear less random (Butler and Graham, 1995). The DWF out from a catchment has previously been studied in the open literature. Métadier and Bertrand-Krajewski (2011) analysed DWF data and recognised different flow patterns depending on the weekday and date, but also found a relative large variation within each pattern. Djebbar and Kadota (1998) estimated DWF peaks and average DWF using a neural network model based on the land use and population, and Brito et al. (2017) fitted a partial least squares model to DWF data to be able to estimate the DWF in situations with missing data. All of these methods may be used to establish generalised DWF patterns, but will not give a real-time picture of the DWF. Such real-time information could be obtained from in-sewer sensing, but flow sensors are often scarcely distributed since it is both expensive and impractical to cover a large urban drainage system (Djebbar and Kadota, 1998). Highly spatially distributed real-time DWF information is therefore not realistic to obtain using only in-sewer observations.

Water supply and urban drainage systems are intrinsically linked since most of the consumed water ends up in the sewer system. The wastewater flow can thus be approximated by estimating the water consumption. Butler and Graham (1995) and Elías-Maxil et al. (2014) used a questionnaire and a probabilistic model to estimate the water consumption. Both studies subsequently modelled the resulting flow in the sewer system to obtain a spatial and temporal distribution of the wastewater flow. However, these methods also only provide generalised flow patterns. Contrarily, smart meters measure the real-time water consumption in each household every hour or even more frequently, and are increasingly being implemented as part of the digitalisation of the water sector, mainly for billing and leakage detection purposes (Boyle et al., 2013; Monks et al., 2019). In general, data col-

lected from the water supply system should be pre-processed, including for example validation and re-estimation of missing and invalid data, before use (Kirstein et al., 2019a). After this, the smart meter data can be used to estimate the wastewater flow with a high spatial and temporal resolution, and may potentially also be used to estimate the wastewater constituents. Even though the concept of applying smart meter data to estimate wastewater flows has been mentioned in recent literature (Cole and Stewart, 2013; Monks et al., 2019; Nguyen et al., 2018), examples of this have to our knowledge not been investigated and documented yet.

This study investigates the hypothesis that smart meter consumption data can be used for estimating the magnitude, timing and spatial distribution of the wastewater flow. This hypothesis is tested using data from the city of Elsinore, Denmark, by validating the smart meter data against observations from the waterworks outlet, observations from the WWTP inlet, and annual water consumption data from a database. The smart meter data is also routed through a 1D hydrodynamic (high-fidelity (HiFi)) urban drainage model for simulating the flow dynamics. Both the smart meter data and the simulated flow are compared to in-sewer flow observations in five locations in Elsinore city centre to identify system anomalies, hereunder the possibility of assessing other DWF components and the reliability of the data.

2. Case study area and data

In the city of Elsinore, Denmark, the utility company has installed smart meters at the consumers' homes and in-sewer sensors in the urban drainage system in an overlapping time period, which enables the type of investigations performed in this study.

Data from the water supply system

Around 19,000 MULTICAL21 smart meters with a temporal resolution of one hour and an uncertainty of up to $\pm 5\%$ (Kamstrup, 2019) are permanently installed in households and industries for billing purposes in Utility Elsinore (Figure 1). These smart meters cover the entire city, with only few consumers left without smart meters. 1,970 of these smart meters are situated in the area upstream of the in-sewer sensors (Figure 1), which is the main area of interest in this study. The installation and operation of the smart meters is undertaken by Contractor 1 who provided data from three separate weeks in 2018 and 2019 (Table 1). This smart meter data originates from hourly accumulated volume readings. In case of non-uniformly distributed time steps or missing meter readings, Contractor 1 fill the data gaps primarily using linear interpolation to obtain a uniform interval between flow values of one hour. Kirstein et al. (2019b) showed that hourly readings are sufficient to estimate the water consumption if a great proportion of the smart meter data is not missing and if the original ("non-cleaned") data resolution is less than two hours. The time periods in Table 1 were selected based on the availability of data as well as the absence of rain. Hourly outflow data from the waterworks were present for specific days in each time period (Table 1). Furthermore, a database with the annual water consumption for each individual consumer in Elsinore from 2012 was provided by Contractor 2, who maintains the utility's hydraulic water distribution network model. The utility also provided a list of known pipe bursts in the three time periods.

Table 1. Time periods with data. WWTP = wastewater treatment plant, SVK = Spildevandskomiteen.

Data	Provider	Period 1	Period 2	Period 3
Smart meter data [m ³ /hr]	Contractor 1	Oct 8-14, 2018	Nov 19-25, 2018	Feb 23-Mar 1, 2019
Waterworks outflow data [m ³ /hr]	Utility Elsinore	Oct 11, 2018	Nov 18, 2018	Feb 24, 2019
WWTP inflow data [m ³ /day]	Utility Elsinore	Oct 8-14, 2018	Nov 19-25, 2018	Feb 23-Mar 1, 2019
Water consumption database [m ³ /yr]	Contractor 2	2012	-	-
In-sewer flow data [m ³ /2 min]	Contractor 4	Oct 8-14, 2018	Nov 19-25, 2018	Feb 23-Mar 1, 2019
Last day with rain before observations	SVK data	Oct 7, 2018	Nov 14, 2018	Feb 21, 2019

Data from the urban drainage system

Utility Elsinore has divided Elsinore into wastewater catchments depending on the flows in the sewer system. These catchments are described in the ‘wastewater plans’ together with the asset data such as the size and location of manholes, pipes, etc. (Figure 1 only shows the location of pipes). The asset data has been used as basis for a HiFi model of the sewer system, constructed in MIKE URBAN. The model has a drainage area of 386 hectares and contains 1,959 nodes of which seven are basins, 1,986 links, 30 pumps, 33 weirs and one orifice. Contractor 3 is responsible for maintaining the model.

Level and velocity sensors were installed in Elsinore city centre as part of a measurement campaign running from May 2018 to June 2019 initiated to calibrate the HiFi model. The sensors were present at five locations: east upstream (‘EU’), east downstream (‘ED’), central downstream (‘CD’), west upstream (‘WU’) and west downstream (‘WD’) (Figure 1). These abbreviations will be used both to refer to the sensor locations and to their upstream sub-catchments. The level and velocity observations stem from pressure and Doppler sensors, respectively. The uncertainty of the resulting flow data for the given sensors is estimated to around 12% (Franck, 2019). Two of each sensor type were installed at each of the five locations to ensure that data was always present as well as for validation purposes. The installation and maintenance of these sensors were undertaken by Contractor 4. This contractor combined the level and velocity data with the geometry of the pipes (multiplying the velocity with the wetted area based on the geometry) to obtain flow observations. The pipes are circular in four of the locations (‘WU’, ‘ED’, ‘CD’, and ‘WD’) while one location has an egg-shaped geometry (‘EU’). Daily inflow data was also available for the WWTP (Table 1). These were aggregated from an ultrasound inlet flow sensor with an uncertainty of ± 10 m³/hr (Laursen, 2019). Furthermore, information on last calibration dates for the WWTP inlet sensor as well as the in-sewer sensors were obtained.

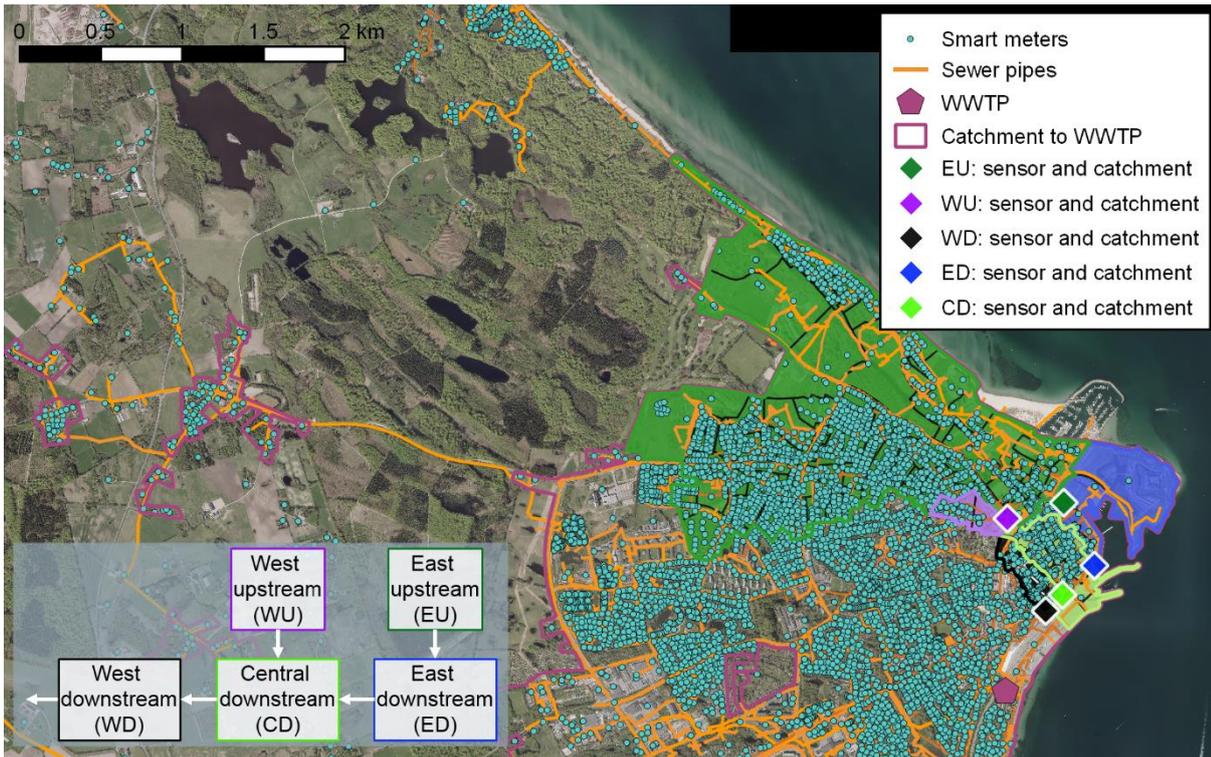


Figure 1. Map of Elsinore, Denmark, showing smart meters, the sewer system, five in-sewer flow observation locations with related sub-catchments (sub-divided into 111 smaller groups for smart meter aggregation), the wastewater treatment plant (WWTP), and the catchment discharging water to the WWTP. The flow between the sub-catchments in the sewer system is displayed in the lower left corner.

3. Methods

Extraction of smart meter consumption data

Utility Elsinore supplied entries from 18,804 meters, of which 18,449 were georeferenced (i.e. their location was mapped). 338 of the remaining 355 meters were successfully referenced using the QGIS tool ‘MMQGIS’. The 17 non-referenced meters were manually verified to be located outside the area upstream of the in-sewer sensors, and thus outside the main area of interest. Due to the general data protection regulation, we were not allowed to obtain water consumption data on household level for the 18,787 georeferenced meters, and they were therefore aggregated into groups. It is important that all smart meters within a group discharge water to the same part of the sewer system. Thus, the groups were initially based on the catchments from the Elsinore wastewater plans. 6% of the 18,787 meters were located outside a wastewater catchment and were thus discarded. Likewise was 14% of the meters were receiving water from other waterworks than Elsinore Waterworks (these neither discharged water to the WWTP of interest). The result was 15,011 meters distributed in groups containing each up to 869 meters. We were interested in the flow dynamics in especially the area upstream of the in-sewer sensors, and these wastewater catchments were therefore further manually sub-divided into in total 111 smaller groups based on the general flow paths in the sewer system, as described by the asset database (see the group divisions in Figure 1) with between 2 and 38 meters in each group.

A list with the meters in each group was provided to Contractor 1 who returned the aggregated smart meter water consumption for each group in hourly time steps. 2.4% of the 15,011 meters were missing in the files provided by the contractor, of which 45 were located in the area upstream of the in-sewer sensors. These meters were missing because they had either not yet been replaced with smart meters, or because they were older meters no longer in use. The total water consumption for these 45 meters in 2018 was estimated to around 9,500 m³ (corresponding to a yearly average of 0.3 L/s), which is less than 2% of the water consumption in the area. The final result included 14,694 smart meters in 245 groups, of which 1,970 smart meters belonged to 111 groups in the area upstream of the in-sewer sensors.

Summed smart meter data

The ‘summed smart meter data’ was calculated for each of the five in-sewer sensors sub-catchments, $Q_{SM,loc}$, by summing the smart meter consumption, $Q_{SM,loc,i}$, of the N,loc groups belonging to these sub-catchments at all time steps, t :

$$Q_{SM,loc}(t) = \sum_i^{N,loc} Q_{SM,loc,i}(t) \quad (1)$$

In this process it was found that a single catchment was missing in the wastewater plans describing the ‘EU’ sub-catchment, which would otherwise have led to a mistake in the summed smart meter data. The missing catchment is, however, of a limited size.

Simulated flow

The sewer system model entry node for the flow from each of the 245 wastewater groups should be determined to use the smart meter data for simulating the wastewater flow with the MIKE URBAN model, Q_{sim} . Manual identification would be a tedious task; thus, we developed an automated procedure for affiliating nodes and consumption data. The aggregation of a set of smart meters into one group would cause some water consumption to be delayed and some to arrive too fast to the sewer system depending on the consumers’ distance to the entry node in the sewer system model. Thus, it was decided to add the aggregated smart meter data to the node in the MIKE URBAN that was closest to the geometrical centre of each group. The error introduced here should be seen in the light of the error present as a result of the relative coarse time step of the smart meter data which is also uncertain due to the interpolation. We first imported the nodes from the MIKE URBAN model into QGIS and extracted only the nodes (i.e. manholes) labelled as main nodes carrying either wastewater or combined waste- and stormwater (excluding parts carrying only stormwater). Subsequently, QGIS was used to find the node nearest to the centre of the shape of each group. This information was extracted using the Distance Matrix tool in QGIS and used to automatically connect the 245 water consumption files as boundary conditions to the relevant nodes in the MIKE URBAN model using the MID API (DHI, 2019) in C#. Subsequently, the model was also run using C#, which allowed us to calculate the simulated wastewater flow in the entire sewer system model, hereunder at the five sensor locations ($Q_{sim,loc}$), with a one-minute resolution.

Phenomena affecting the observed flow

Q_{SM} and Q_{sim} may differ from Q_{obs} if the system is affected by other water inflows and outflows than the smart meter-measured water consumption or due to erroneous data.

To assess these phenomena, the average flows over the one-week time periods (168 hours or 10,080 minutes) was calculated, where $\Delta t_{obs}=2$ min, $\Delta t_{SM}=1$ hr and $\Delta t_{sim}=1$ min (Table 1):

$$\begin{aligned}\bar{Q}_{obs,loc} &= \frac{\sum_{t=1}^{10,080} Q_{obs,loc}(t)}{10,080 \text{ min}/\Delta t_{obs}} \\ \bar{Q}_{SM,loc} &= \frac{\sum_{t=1}^{168} Q_{SM,loc}(t)}{168 \text{ hrs}/\Delta t_{SM}} \\ \bar{Q}_{sim,loc} &= \frac{\sum_{t=1}^{10,080} Q_{sim,loc}(t)}{10,080 \text{ min}/\Delta t_{sim}}\end{aligned}\quad (2)$$

Furthermore, the mass balances (change in volume, ΔV) were calculated for each of the five sub-catchments based on the water consumption in the given sub-catchments as well as inflows from upstream sub-catchments ($Q_{obs,in,loc}$) and the outflow ($Q_{obs,out,loc}$) in the three investigated one-week time periods:

$$\Delta V_{loc} = \sum_{t=1}^{168} (Q_{SM,loc}(t)) \cdot \Delta t_{SM} + \sum_{t=1}^{10,080} (Q_{obs,in,loc}(t) - Q_{obs,out,loc}(t)) \cdot \Delta t_{obs} \quad (3)$$

A positive mass balance means that more water enters than leaves the catchment (thus, there may be a loss of water in the sewer system); a negative mass balance contrarily means that there may be an additional source of incoming water. The relative importance of the difference in volume as a function of the outflow from each sub-catchment was calculated as (Eq. 4):

$$\text{Rel}\% = \frac{\Delta V_{loc}}{\sum_{t=1}^{10,080} Q_{obs,out,loc}(t)} \cdot 100 \quad (4)$$

The difference between observed and simulated flows in the five sub-catchments, $Q_{res,loc}(t)$, were calculated based on the observed and simulated flows:

$$Q_{res,loc}(t) = Q_{obs,loc}(t) - Q_{sim,loc}(t) \quad (5)$$

These residuals may both be positive and negative, and exhibit constant, diurnal or seasonal variations depending on whether the error is constant, proportional to daily patterns or proportional to seasonal patterns. They may also vary according to the outside temperature, previous rainfall and pipe geometry. Table 2 lists three overall possible reasons for deviations as well as expected residual patterns while the remaining sections elaborate on the three deviation types.

Table 2. Potential reasons for the deviation between observed and simulated sewage flow.

Deviation type	Potential reason	Expected residual pattern
Observed flow smaller than simulated flow (negative residuals, Eq. 5)	Consumed water not discharged to the sewer	Diurnal
	Exfiltration	Seasonal
Observed flow larger than simulated flow (positive residuals, Eq. 5)	Unaccounted for consumers	Diurnal
	Pumping of groundwater to the sewer system	Constant or seasonal
	Re-use of rainwater	Diurnal
	Melting of snow (only in winter)	Temperature dependent
	Infiltration	Constant, seasonal or rainfall dependent
General reasons	Sedimentation	Pipe geometry dependent
	Erroneous smart meters, data transmission or data handling	Constant or diurnal
	Wrong conceptualization of the sewer system	Diurnal
	Erroneous in-sewer sensors	Constant or diurnal

Observed flow smaller than simulated flow (negative residuals, Eq. 5)

Consumed water not discharged to the sewer. Parts of the consumed water will never enter the urban drainage system due to the water being used for drinking, cooking, gardening, etc. Some of these, as for instance gardening, is expected to be season dependent. In Denmark, it is expected that more than 83% of the consumed water is discharged to the sewer system (Rygaard et al., 2013). Since more water is consumed during the day than at night, the residuals will be larger during the day and thus exhibit a diurnal pattern. In Elsinore, discrepancies may also arise if ferries, boats, trains, etc. take on water in Elsinore but discharge it somewhere else. This may lead to irregular residual patterns.

Exfiltration. Exfiltration from the sewer system may occur when the sewer system is leaky and the groundwater level is below the sewer system. The groundwater level is generally lower in the summer time than during winter, meaning that the exfiltration rate is expected to follow a seasonal pattern with peaks in the summer.

Observed flow larger than simulated flow (positive residuals, Eq. 5)

Unaccounted for consumers. Q_{sim} only contains data from consumers that have smart meters installed. Some consumers may, however, have older meters requiring manual readings. In fact, large consumers will often have such meters installed. Such a deviation would exhibit a diurnal pattern unless missing industrial users consume water around the clock. In Elsinore, some of the water discharged to the sewer system may furthermore originate from somewhere else, for example, from ferries, boats and trains. This may result in irregular residual patterns.

Pumping of groundwater to sewer system. Elsinore is located by the ocean, and buildings and construction sites that are located partly underground may therefore need to pump away groundwater (and possibly intruding seawater). This water may be discharged into the ocean, infiltrated somewhere else, or discharged to the sewer system. Such pumping would most likely occur both

day and night leading to a constant deviation. If the groundwater level is only too high in the winter, the residual pattern will be seasonal. Sea level variations may, however, also affect the groundwater level and thus impact the amount of pumping.

Re-use of rainwater. Rainwater re-use for toilet flushing, laundry, etc. will increase the flow in the urban drainage system compared to what is measured by the smart meters. Rainwater can only be used as a substitute water supply in the time following a rain event. Since more water is consumed during the day, the residual pattern will follow a diurnal pattern. An increase in consumption from the distribution system will be seen again when the rainwater tanks are empty.

Snow melting. In winter, precipitation falling as snow can cause a delayed runoff into the sewer system. The temperature determines when the snow melts, and the residuals will thus neither follow a constant, diurnal or seasonal pattern.

Infiltration. Infiltration into leaky sewers can arise from either groundwater, rain, pumped groundwater, or leaky water supply pipes. Groundwater levels change slowly over the course of the year, and groundwater infiltration will thus exhibit a seasonal change in the infiltration rate, which may also be affected by the sea water level. Rain-induced infiltration will contrarily be visible as an increase in the sewage flow only after rain events, which will slowly decline with time. Pumping of groundwater to places from where it can infiltrate into the sewer is expected to exhibit a constant or seasonal pattern. Leakage from the water distribution system into the drainage system is expected to be a function of the pressure in the water distribution system. If the system is pressure controlled, the leakages could to some extent display a diurnal pattern, but it will predominantly be constant over the day.

Sedimentation. The flow observations in Elsinore were obtained by multiplying the velocity with the wetted area of the cross-section of the pipe. If sedimentation is present at the location of the sensor, the water level will raise as a result; thus, leading to a larger calculated than actual flow. The residuals will vary in size depending on the geometry of the pipe.

General reasons for deviations between observed and simulated flows (negative and positive residuals, Eq. 5)

Erroneous smart meters, data transmission or data handling. Deviations may occur due to erroneous smart meter observations. In general, individual meter errors are not likely to have noticeable impact on the simulated results since each consumer only uses a minor fraction of the water in each sub-catchment. A general bias of the smart meters could, however, impact the results. Depending on whether this bias is proportional or additive, it may both show as diurnal pattern or a constant deviation. The smart meters in Elsinore are reported to have a maximum uncertainty of $\pm 5\%$ (and thus a proportional, diurnal residual pattern). Furthermore, delayed data transmission and re-estimation of data due to either missing transmission or time wise unequally distributed data can affect the smart meter based results, as well as wrong mapping of the geographical location of the smart meters.

Wrong conceptualization of the sewer system. The conceptualisation of the urban drainage system layout is essential to make the correct coupling between the smart meter data and in-sewer sen-

sors. In this study, the layout is contained in the asset database, which has been used to construct both the utility’s wastewater plans and the MIKE URBAN model. If one or more catchments are falsely connected to the sewer system upstream of the in-sewer sensors, the simulated flow would be larger than the observed. Contrarily, missing connections in the asset database would lead to a smaller simulated than observed flow. This deviation would follow a diurnal pattern.

Erroneous in-sewer observations. In-sewer observations of wastewater velocity, level, and flows may be erroneous due to turbulence, presence of solids, aggressive environment (Hager 1994), low water depth (Larrarte et al. 2008), or poor calibration. The residuals may both follow a diurnal pattern or be constant. Normally, flow observations are assumed to have uncertainties up to 20% (Bertrand-Krajewski et al., 2003), leading to a diurnal residual pattern.

4. Results and discussion

Here, we 1) assess if smart meter data can be used to estimate the wastewater component of DWF, 2) identify anomalies, hereunder other sources to the DWF and erroneous data, by comparing the wastewater component with the observed in-sewer flow, and 3) discuss the added value from using a HiFi model.

Assessment of smart meter data for estimation of wastewater flow

The smart meter data can be used to estimate the wastewater component of the DWF if the smart meter data set is trustworthy.

Figure 2 compares the total outflow from the waterworks, $Q_{WW,obs}$, with the aggregated smart meter data from entire Elsinore, $Q_{SM,Elsinore}$. This shows that the smart meters registered on average 14-15 L/s (15-16%) less water than $Q_{WW,obs}$. This deviation could be due to leakage in the water distribution network, errors in the waterworks’ flow sensor or in the smart meter data set. The smart meter data set may be erroneous either if the data set is incomplete (i.e. there are unaccounted for consumers), if the meters are erroneous, or if the data transmission or subsequent data handling is erroneous.

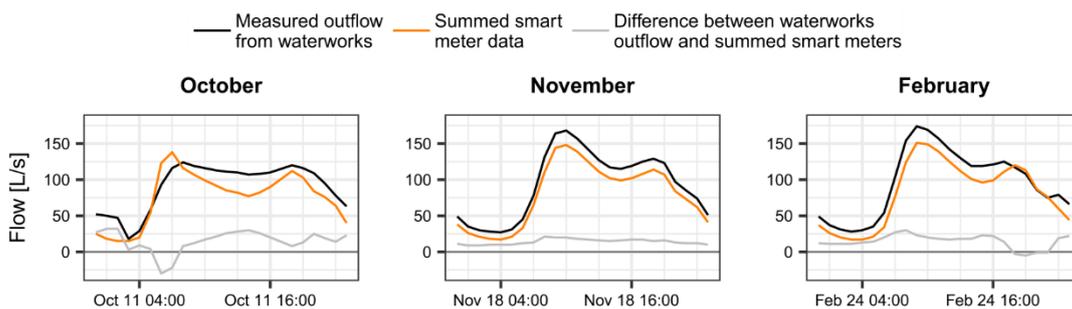


Figure 2. Hourly outflow data from the waterworks ($Q_{WW,obs}$) compared to summed smart meter data ($Q_{SM,Elsinore}$) for entire Elsinore (beyond what is shown in Figure 1) for single days in the three time periods.

The completeness and quality of the smart meter data set was further evaluated by comparing the daily inflow to the WWTP, $Q_{WWTP,obs}$, with the simulated inflow to the WWTP, $Q_{WWTP,sim}$, using all smart meters in the WWTP catchment (Figure 3). The flow out of the area with in-sewer sensors accounted for 20-25% of this inflow. $Q_{WWTP,sim}$ is around 1.5 times the one of $Q_{WWTP,obs}$ in October

and November. This corresponds to on average more than 1,500 m³/day, which seems as an extremely unlikely amount of exfiltration or an unrealistic volume of consumed water not ending up in the urban drainage system. In February, however, $Q_{WWTP, sim}$ and $Q_{WWTP, obs}$ are very similar. $Q_{WWTP, sim}$ is in the same range as in October and November, but the $Q_{WWTP, obs}$ has increased notably. This may be due to infiltration of groundwater in February (winter), but it may also be due to erroneous WWTP inlet observations (the WWTP inlet sensor was last calibrated in March 2018) or biased smart meter data.

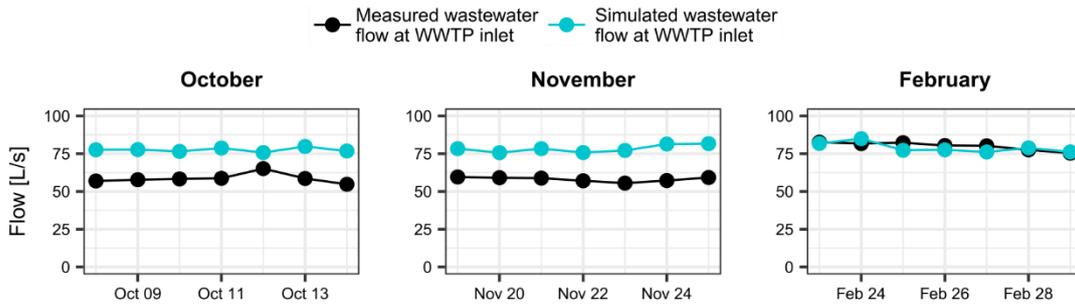


Figure 3. Observed ($Q_{WWTP, obs}$) and simulated ($Q_{WWTP, sim}$) wastewater treatment plant (WWTP) inflow in the three time periods.

There are no consumers in Elsinore city centre that has a dominating impact on the wastewater flow. Furthermore, there were no major changes in the magnitude or composition of the city's population between 2012 and 2019. It is therefore assumed that the completeness and quality of the smart meter data set could also be assessed by comparing the average flow data (Eq. 2) from the three time periods to those in the database containing annual water consumption rates from 2012 (Figure 4). \bar{Q}_{sim} and \bar{Q}_{SM} differ from each other due to water generated in empty pipes of the model for the sake of numerical stability and because \bar{Q}_{sim} contains the routing time in the sewer system. Overall, both match the database values from 2012 well even though the data is from different years. The small decrease from 2012 to 2018/2019 fits well with a general decrease in water consumption in Denmark (DANVA, 2017; Elsinore Municipality, 2019). Figure 4 also shows a large difference between the database values and \bar{Q}_{obs} . Consumers that have private wells will neither be in the database nor have smart meters installed. However, there are no private or commercial wells from where the water is subsequently discharged to the sewer system in the area upstream of the in-sewer flow sensors (Elsinore Municipality, 2019; Pratt, 2019). Overall, the comparison with the database indicates that the amount of unaccounted for consumers and uncertainty related to the smart meters, data transmission and data handling (Table 2) is limited; meaning that the smart meter data set is complete and sufficiently correct to represent the wastewater component of the DWF.

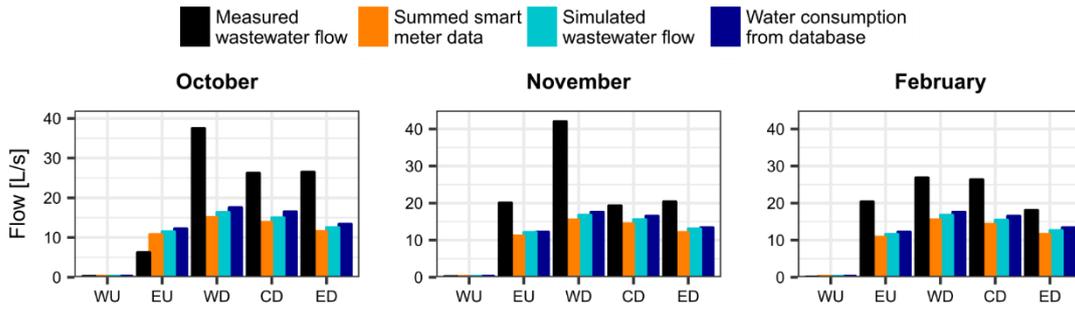


Figure 4. Average observed flow (\bar{Q}_{obs}), summed smart meter data (\bar{Q}_{SM}) and simulated flow (\bar{Q}_{sim}) (Eq. 2) in each of the three time periods compared to the average consumption from the water consumption database from 2012.

Assessment of anomalies, hereunder other DWF components

Figure 4 shows that there is a discrepancy between the smart meter based wastewater flows, Q_{sim} and Q_{SM} , and the observed in-sewer flow, Q_{obs} . To assess possible anomalies, the mass balances for the five sub-catchments were calculated for the three time periods (Eq. 3). From Figure 5 it is evident there are no clear pattern in which catchments are losing or gaining water in the three periods. Trusting the observations, 2,000 m³ more water would exit than enter ‘ED’ in October (negative mass balance), which could indicate a large external water inflow, for example, infiltration. From the percentage (Eq. 4) shown in Figure 5, it is clear that this possible “external inflow” would be the main contributor of the outflow stemming from the sub-catchment. This trend corresponds poorly with the results for November and February that show a small positive mass balance for ‘ED’. The natural variations in groundwater levels and soil moisture that could explain month to month variations in infiltration would give the opposite pattern since the summer of 2018 was very dry followed by a wet winter. ‘CD’ displays some of this expected behaviour but it would be premature to draw any conclusions from this considering the large unexplainable variation for the remaining sub-catchments.

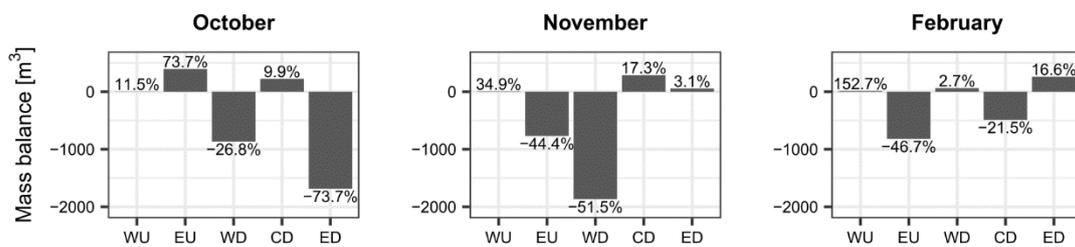


Figure 5. Mass balances for each sub-catchment (Eq. 3) and the relative size of the change in volume compared to the outflow (Eq. 4) for the three time periods.

Figure 6 shows the time dynamics of Q_{sim} , Q_{SM} and Q_{obs} , and is arranged to resemble the flow through the sewer system as conceptualised in Figure 1. The flow is naturally smallest in ‘WU’ since this is the smallest of the five sub-catchments and does not receive water from any further upstream catchments. The axis for this plot is thus scaled differently than the remaining plots. The flow increases through the system as more water aggregates. Most of the outflow from the most downstream catchment (‘WD’) is generated in the most upstream catchment, ‘EU’, which also by

far has the most consumers. The residuals between Q_{obs} and Q_{sim} (Eq. 5) are shown in the upper left corner of Figure 6 for all five sub-catchments.

Figures 4 and 6 show that, for example, the flows in 'WU' match well in October and November, but that there are negative discrepancies between Q_{obs} and the smart meter based flows (Q_{SM} and Q_{sim}) in 'EU' in October and in 'WU' in February, and many time periods and sub-catchments with positive residuals. In the following we seek to deduct what could be viable courses for the observed discrepancies between the various data sources and model results based on Table 2 and Figure 6.

Exfiltration. $Q_{res,EU}$ exhibits a clear diurnal pattern in October. Exfiltration is, however, not expected to vary diurnally, but according to season. Furthermore, net exfiltration is not expected to occur in February (where the groundwater level generally is higher without also occurring in October and November. Exfiltration can therefore not explain the negative residuals.

Consumed water not discharged to the sewer system. The diurnal variation in residuals could indicate that part of the consumed water is not discharged to the urban drainage system. $Q_{obs,EU}$ on average corresponds to less than 60% of $Q_{sim,EU}$ in October, and $Q_{obs,WU}$ to less than 40% of $Q_{sim,WU}$ in February. 'EU' and 'WU' are, however, mostly residential and a greater proportion of the consumed water is therefore expected to end up in the urban drainage system. It is therefore unlikely that this is the (main) cause of the negative residuals.

Pumping of groundwater to the system. Due to the diurnal residual pattern, pumping of groundwater is not expected to be the cause of the deviation between Q_{obs} and Q_{sim} . Furthermore, the Capital Region of Denmark (Vormbak, 2019), who stores records of such pumping sites, confirmed that no such pumping of groundwater to the sewer is taking place in the area of Elsinore where the in-sewer sensors are placed.

Re-use of rainwater. The diurnal residual pattern could stem from rainwater re-use since it was raining one, five and two days before the three analysed time periods, respectively (Table 1). Rainwater re-use is, however, not expected to be the cause of the deviations since there are very large variations in the in-sewer flows, and thus in the potential amount of re-used rainwater, between the three time periods within the same sub-catchments.

Snow melting. There was no snow in Elsinore in neither of the three considered time periods. Snow melting can thus not explain the deviations between Q_{obs} and Q_{sim} .

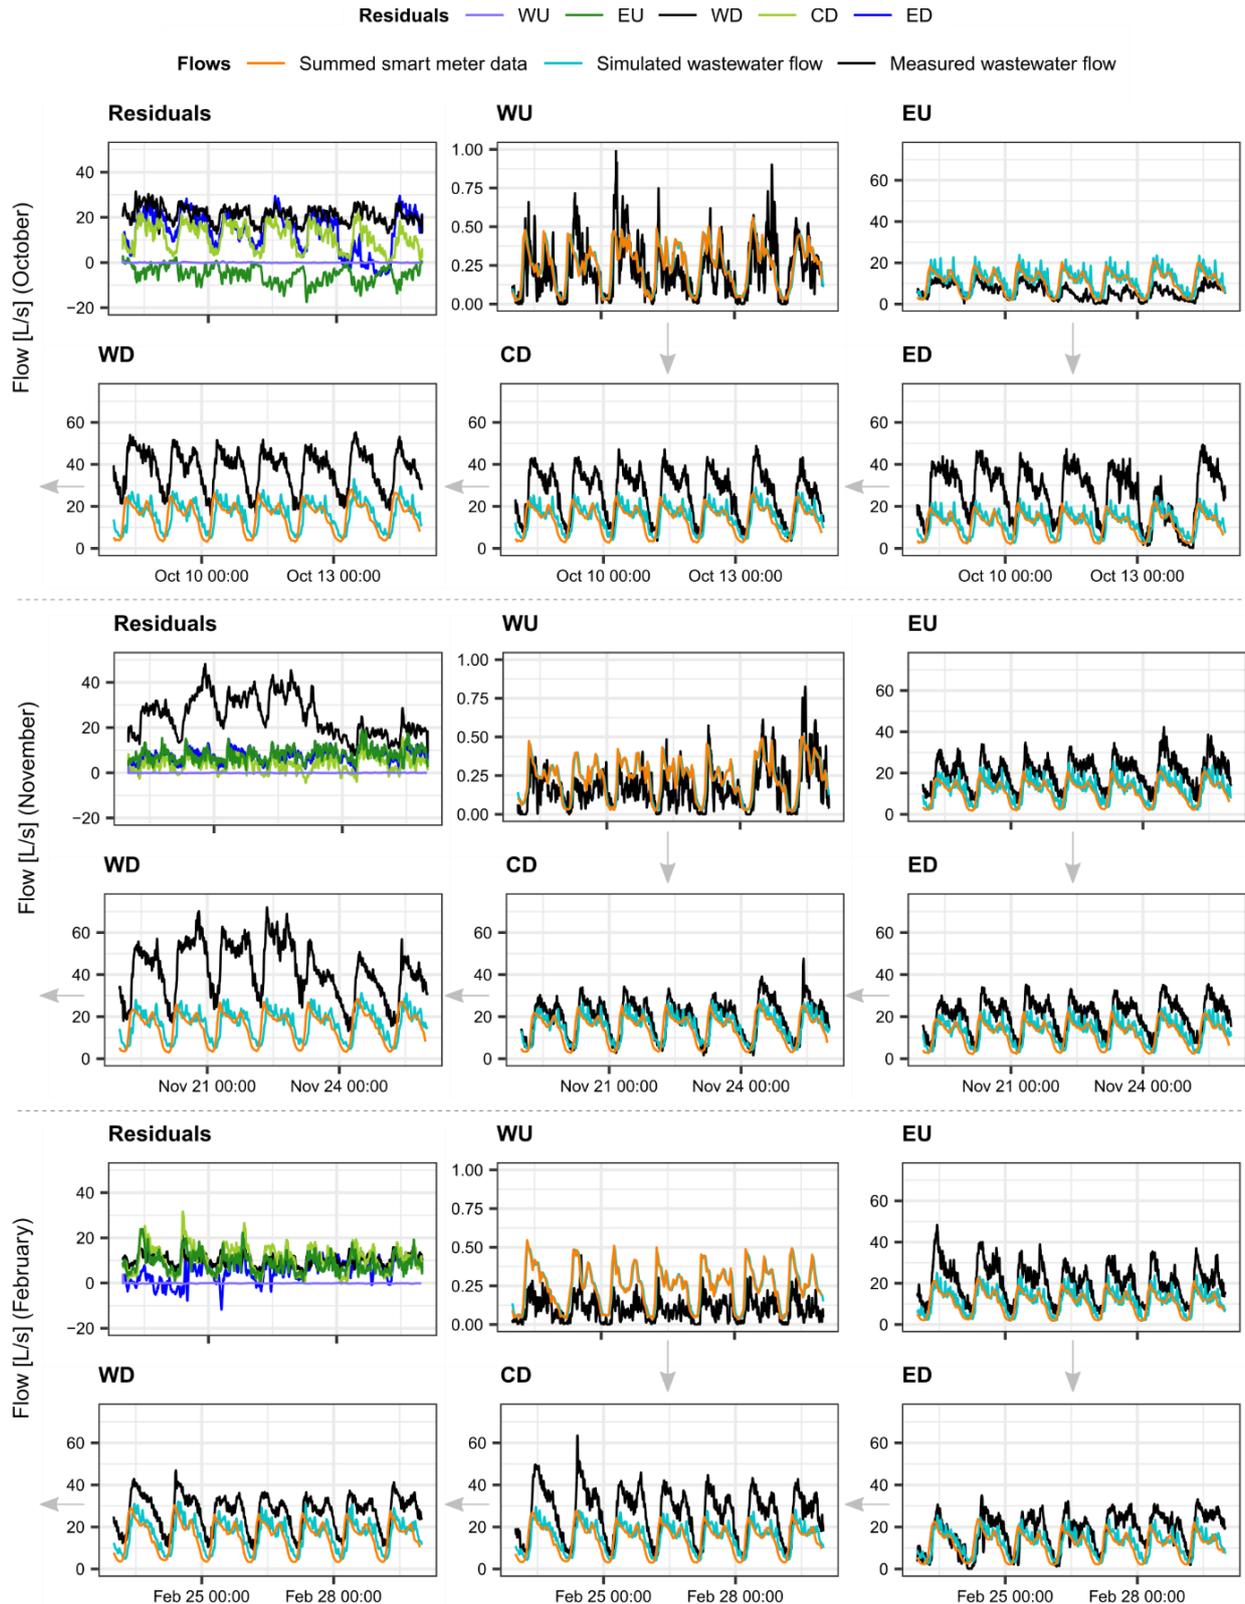


Figure 6. Summed smart meter data (Q_{SM}) and simulated (Q_{sim}) and observed (Q_{obs}) sewage flows in October (top panel, Mon-Sat), November (middle panel, Mon-Sun) and February (lower panel, Sat-Fri) (Eq. 1). All data is in local time and Q_{sim} and Q_{obs} are displayed with a 30 minutes moving average filter to be able to distinguish graphs. Notice the different scales used for ‘Residuals’, ‘WU’ and the four remaining plots.

Infiltration. It was raining in the days up to the three time periods (Table 1), but neither rain-induced infiltration or infiltration from pumped groundwater exhibit diurnal residual patterns; however, infiltration from a leaking water distribution system might. The gap between $Q_{obs,ED}$ and $Q_{sim,ED}$ is on average 14 L/s, 7 L/s and 1.5 L/s in October, November and February, respectively, whereas the deviation between the waterworks outlet and consumed water is more or less constant over the three time periods (Figure 2). Leakages in ‘EU’ has been recorded in October and February, but Figure 6 shows signs of outflowing, not inflowing, water in the former of these periods, and there is only minor deviations between the flows in the latter period. It is therefore not likely that leakage from the water distribution system is the main contributor to the deviations between Q_{obs} and Q_{sim} . Besides as diurnal pattern, the residuals for ‘WD’ has a more or less constant offset from zero which could indicate a constant infiltration from groundwater. The offset is around 20 L/s in October and November and 10 L/s in February. Q_{sim} shows that around 5 L/s of this offset stems from night consumption. Infiltration is contrarily expected to be larger in February than in October. No repair work was carried out for the sewer pipes in ‘WD’ between November and February to explain this drop in infiltration. Infiltration does thus not seem to be able to explain the elevated observed flows.

Sedimentation. The pipe in which the ‘WD’ flow sensor was located is heavily influenced by sedimentation, which could result in the constant offset from zero. The sediment was, however, not removed between November and February to explain the drop in the offset level, but the amount of sediment may have been affected by the stormwater flow in the sewer system.

Wrong conceptualization of the sewer system. A missing pipe or pump in the wastewater plans, which in reality connects another catchment to one of the five sub-catchments, would expectedly lead to similar residuals in all three time periods. This is not the case, and the conceptualization is thus not considered to cause the deviations.

In-sewer observations. Erroneous in-sewer observations remain the only likely explanation for the deviations between Q_{obs} and Q_{sim} . The observed flow changes greatly from time period to time period; for example, $Q_{obs,CD}$ is much larger than $Q_{sim,CD}$ in October and February but fits well in the intervening period in November. Looking at October, there is only one day where $Q_{obs,ED}$ and $Q_{sim,ED}$ match; however, this is due to faulty velocity observations. Contrarily, the water consumption, as expected, more or less follows the same diurnal pattern throughout October, November and February. Additionally, the erroneous nature of in-sewer observations is supported by the fact that Q_{obs} out of ‘WD’ on average is 21 L/s and 25 L/s larger than Q_{sim} in October and November, respectively. Adding this amount on top of $Q_{WWTP,sim}$ would only exacerbate the difference to $Q_{WWTP,obs}$ (Figure 3). It is therefore expected that the in-sewer sensors are the main contribution to the deviation between simulated and observed sewage flows. These deviations are at times even larger than what is considered ‘normal’ deviations of 10–20%.

Figure 7 shows the relationship between the deviation between Q_{obs} and Q_{sim} in the three time periods and the time since the last calibration of the in-sewer sensors. It is clear that no such correlation exists, neither within each time period or when comparing each location across the three time periods.

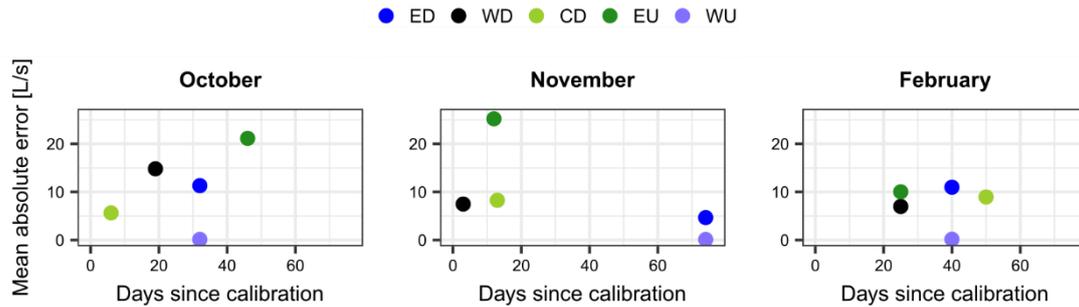


Figure 7. Days since last calibration of in-sewer sensors compared to the difference between observed flow (Q_{obs}) and simulated flow (Q_{sim}) calculated as the mean absolute error for the three time periods.

Added value from using a HiFi model

A HiFi model pose an easy way of coupling smart meter data and in-sewer observations and can eliminate errors stemming from incorrect wastewater plans and s. Furthermore, it is possible to include infiltration/exfiltration dynamics and thus make a better comparison between the simulated and observed data, which will not be possible when simply summing the smart meter data.

Figure 8 shows $Q_{obs,CD}$, $Q_{SM,CD}$ and $Q_{sim,CD}$ for three days in each of the considered time periods. The timing of the observed flow is considered reliable, despite the magnitude of the wastewater flow being off as a result of the erroneous in-sewer observations. It is clear that the routing of the smart meter data through the MIKE URBAN model catches the timing of the observed peaks and low flows much better than simply aggregating the smart meter data. The dynamics shown in Figure 8 are representative for the entire dataset for all five flow observation locations. The difference in time between the summed smart meter data and the simulated flow will naturally be more pronounced the longer the water has travelled in the system and thus depends on the size of the catchment upstream of the flow observations.

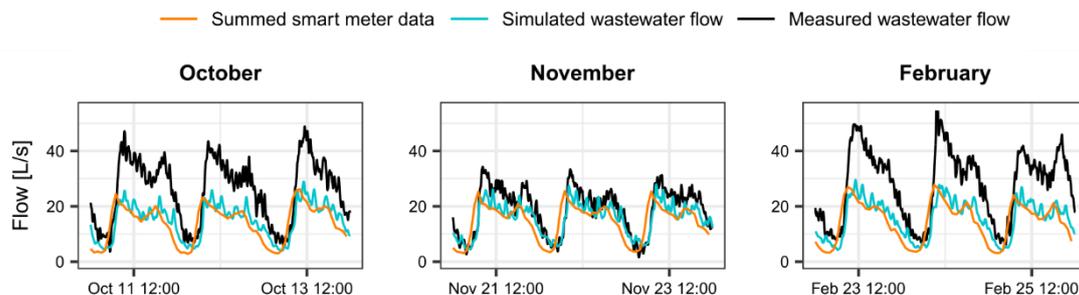


Figure 8. Observed sewage flow (Q_{obs}) compared to simulated flow (Q_{sim}) and summed smart meter data (Q_{SM}) for three days in each of the three time periods for the ‘CD’ catchment.

Limitations and outlook

This study shows that smart meter data may be used to represent the wastewater flow. This approach may be more robust than using in-sewer observations since the failure of one sensor will likely not disturb the bigger picture (unless it is a very large consumer), whereas the failure of one in-sewer sensor leaves one left with no real-time data in that specific location. The subsequent search for system anomalies by linking water supply smart meter data with in-sewer measurements is a tedious process due to the many aspects affecting both the water distribution and urban drainage

systems, such as missing consumers and sedimentation (see Table 2), and it is laborious to systematically gain access to and compare all relevant data sources. This was further complicated by the fact that not one person could access all the models and data as they were managed and stored in different silos both within the utility and by different contractors.

It would be necessary to perform independent and reliable flow observations at the five locations, possibly manually, to validate that the in-sewer flow observations caused the deviation to the smart meter simulated wastewater flow. This is, however, not possible since the measurement campaign in the urban drainage system was terminated in June 2019. It would also require independent observations to investigate if the offset in 'WD' is due to sedimentation or infiltration. These could be measurements of the sedimentation height (which could be incorporated into the urban drainage model and thus included in the simulated flow results), measurements of the groundwater level could (which could show the potential for infiltration and exfiltration), or CCTV footage.

Many utilities are currently taking first steps towards using data in new, and more integrated, ways. This study shows that digitalisation is not easy, and sometimes the actual data quality remains unknown until the data is actually used and compared with other data sources. This process will provide important learnings regarding good practices within sensing and data accessibility, and enable the utility and their contractors to further refine their work processes.

5. Conclusions

The current study aimed at using smart meter water consumption data to simulate the wastewater flow, and to combine this information with in-sewer observations to detect system and data anomalies, such as infiltration, exfiltration, sensor errors or a wrong conceptualization of the system.

The smart meter data was validated with data from other independent sources, hereunder data from the waterworks outflow, WWTP inflows, and households' annual water consumption audits. Subsequently, it was used as input to a 1D hydrodynamic urban drainage model to simulate the wastewater component of the DWF flow. The wastewater flow was thus obtained with a large spatial and temporal resolution. A 1D hydrodynamic model presented an easy way of coupling smart meter data and in-sewer flow observations compared to manual summation of smart meter data based on wastewater plans; a procedure prone to human errors. The main difference in the results between simply summing up the smart meter data and using an urban drainage model was the inclusion of the routing time in the sewer system. If time dynamics are insignificant, a simple summation may be sufficient.

Time dynamics are important when comparing the smart meter based wastewater flow with observed in-sewer DWF to estimate the contribution from other input sources (infiltration, pumping of groundwater to the system, rainwater re-use or snow melts) or outputs (exfiltration). This requires at least one reliable downstream in-sewer observation, and preferably more. In this study, the validated smart meter data was used to identify erroneous in-sewer flow observations. This was supported by the comparison with WWTP inlet data. The degree of uncertainty of these in-sewer observations at times exceeded the expected uncertainties of 10-20%.

The erroneous in-sewer observations prohibited the assessment of other input and output water sources. Still, validated smart meter data coupled with a 1D urban drainage model provides a first ‘best guess’ of the flow in the system during dry weather, and the study presents an important step towards closing the water balance in urban drainage systems.

The coupling between smart meter data and urban drainage models can be done offline for post-analysis of the data and system performance (as done in this study), but the coupling could also be done, and the urban drainage model run, in real time to get an up-to-date picture of the system state for enhanced real-time decision-making.

Acknowledgements

We would like to thank Andreas Mikkelsen, Lotte Krogh, Brian Laursen and their colleagues at Utility Elsinore for contributing with data, models and expertise. Furthermore, thank you to Lina Jensen (LNH water), Kristoffer Andersen (Kamstrup), Oliver Jodehl and Tommy Petersen (Blue-North), and Rasmus Halvgaard (DHI) for professional discussions.

References

- Abdel-Aal, M., Schellart, A., Kroll, S., Mohamed, M., and Tait, S. (2018). Modelling the potential for multi-location in-sewer heat recovery at a city scale under different seasonal scenarios. *Water Res.* **145**, 618–630.
- Bertrand-Krajewski, J.L., Bardin, J.P., Mourad, M., and Béranger, Y. (2003). Accounting for sensor calibration, data validation, measurement and sampling uncertainties in monitoring urban drainage systems. *Water Sci. Technol.*, **47**, 95–102.
- Boyle, T., Giurco, D., Mukheibir, P., Liu, A., Moy, C., White, S., and Stewart, R. (2013). Intelligent metering for urban water: A review, *Water* **5**, 1052–1081.
- Brito, R.S., Almeida, M.C., and Matos, J.S. (2017). Estimating flow data in urban drainage using partial least squares regression. *Urban Water J.* **14**, 467–474.
- Butler, D., and Graham, N.J.D. (1995). Modeling dry weather wastewater flow in sewer networks. *J. Environ. Eng.* **121**, 161–173.
- Cole, G., and Stewart, R.A. (2013). Smart meter enabled disaggregation of urban peak water demand: precursor to effective urban water planning. *Urban Water J.* **10**(3), 174–194.
- DANVA (2017). Water in Figures – DANVA statistics & Benchmarking. Report by DANVA (the Danish Water and Wastewater Association). Available at https://www.danva.dk/media/4662/water-in-figures_2017.pdf (accessed September 26, 2019).
- DHI (2019). MIKE 1D API. Available at docs.mikepoweredbydhi.com/engine_libraries/mike1d/mike1d_api/ (accessed September 25, 2019).
- Djebbar, Y., and Kadota, P.T. (1998). Estimating sanitary flow using neural networks. *Water Sci. Technol.* **38**, 215–222.
- Elías-Maxil, J.A., van der Hoek, J.P., Hofman, J., and Rietveld, L. (2014). A bottom-up approach to estimate dry weather flow in minor sewer networks. *Water Sci. Technol.* **69**, 1059–1066.

- Elsinore Municipality (2019). *Vandforsyningsplan 2019 - 2030*. Helsingør Kommune. Center for by, land og vand, natur og miljø. June, 2019. Available at https://www.fh.dk/files/media/dokumenter/Om%20os/Lovgivning%20og%20planer/vandforsyningsplan_19-30_helsingoer_kommune_juni_2019_132374-19_v1.pdf (accessed October 15, 2019)
- Franck, T. (2019). Personal communication, BlueNorth.
- Kamstrup (2019). Data Sheet. Multical® 21 & flowIQ® 2101. Available at <https://products.kamstrup.com/ajax/downloadFile.php?uid=515d4ab700278&display=1> (accessed August 29, 2019).
- Kirstein, J.K., Høgh, K., Rygaard, M., and Borup, M. (2019a). A semi-automated approach to validation and error diagnostics of water network data. *Urban Water J.* **16**, 1–10.
- Kirstein, J.K., Høgh, K., Rygaard, M., and Borup, M. (2019b). Effect of data sampling resolution of water meter readings in water distribution network simulations. Under Review.
- Laursen, B. (2019). Personal communication, Utility Elsinore.
- Lund, N.S.V., Falk, A.K.V., Borup, M., Madsen, H., and Mikkelsen, P.S. (2018). Model predictive control of urban drainage systems: A review and perspective towards smart real-time water management. *Crit. Rev. Environ. Sci. Technol.*, **48**(3), 279–339.
- Métadier, M., and Bertrand-Krajewski, J.-L. (2011). Assessing dry weather flow contribution in TSS and COD storm events loads in combined sewer systems. *Water Sci. Technol.* **63**, 2983–2991.
- Monks, I., Stewart, R.A., Sahin, O., and Keller, R. (2019). Revealing Unreported Benefits of Digital Water Metering: Literature Review and Expert Opinions. *Water* **11**(838), 1–32.
- Nguyen, K.A., Stewart, R.A., Zhang, H., Sahin, O., and Siriwardene, N. (2018). Re-engineering traditional urban water management practices with smart metering and informatics. *Environ. Modell. Softw.* **101**, 256–267.
- Plósz, B.G., Reid, M.J., Borup, M., Langford, K.H., and Thomas, K.V. (2013). Biotransformation kinetics and sorption of cocaine and its metabolites and the factors influencing their estimation in wastewater. *Water res.* **47**(7), 2129–2140.
- Pratt, A. (2019). Personal communication, Elsinore Municipality.
- Rygaard, M., Alsbjörn, L., Ejsing, M., Godskesen, B., Hansen, R., Hoffmann, B., Jørgensen, C., Ledgaard, K., Møller, H.-M.F., Poulsen, M.-B.B., Schmidt, M., Tarp-Johansen, L., Vigsø, S., and Zambrano, K. (2013). *Sekundavand i Nordhavn. En forundersøgelse til strategi for alternativ vandleverance*. Kgs. Lyngby, Denmark: Department of Environmental Engineering, the Technical University of Denmark.
- Schilperoort, R.P.S., Dirksen, J., Langeveld, J.G., and Clemens, F.H.L.R. (2012). Assessing characteristic time and space scales of in-sewer processes by analysis of one year of continuous in-sewer monitoring data. *Water Sci. Technol.* **66**, 1614–1620.
- Vormbak, F. (2019). Personal communication, the Capital Region of Denmark.

The Department of Environmental Engineering (DTU Environment) conducts science based engineering research within five sections: Air, Land & Water Resources, Urban Water Systems, Water Technology, Residual Resource Engineering, Environmental Fate and Effect of Chemicals. The department dates back to 1865, when Ludvig August Colding gave the first lecture on sanitary engineering.

Department of Environmental Engineering
Technical University of Denmark

DTU Environment
Bygningstorvet, Building 115
2800 Kgs. Lyngby
Tlf. +45 4525 1600
Fax +45 4593 2850

www.env.dtu.dk