WATERWISE@SG: A TESTBED FOR CONTINUOUS MONITORING OF THE WATER DISTRIBUTION SYSTEM IN SINGAPORE

Andrew J. Whittle[1], Lewis Girod[1], Ami Preis[2], Michael Allen[2], Hock Beng Lim[3], Mudasser Iqbal[3], Seshan Srirangarajan[3], Cheng Fu[3], Kai Juan Wong[4], Daniel Goldsmith[5]

[1] Massachusetts Institute of Technology (MIT), Cambridge, MA, USA
[2] Center for Environmental Sensing and Modeling, MIT-SMART Center, Singapore
[3] Intelligent Systems Center, Nanyang Technological University, Singapore
[4] School of Computer Engineering, Nanyang Technological University, Singapore
[5] Faculty of Engineering and Computing, Coventry University, Coventry, UK

Abstract

This paper describes the development of WaterWiSe@SG, a wireless sensor network to enable real-time monitoring of a water distribution network in Singapore. The overall project is directed towards three main goals: 1) the application of a low cost wireless sensor network for high data rate, on-line monitoring of hydraulic parameters within a large urban water distribution system; 2) the development of systems to enable remote detection of leaks and prediction of pipe burst events; 3) the integrated monitoring of hydraulic and water quality parameters. In this paper we will describe the current state of the WaterWiSe@SG testbed, and report on experimentation we have performed with respect to leak detection and localization. Furthermore, we describe how we have assimilated real time pressure and flow measurements from the sensor network into hydraulic models that are used to improve state estimation for the network. Finally, we discuss the future plans for the project.

Keywords
Hydraulic modeling, Continuous Monitoring, Sensor Networks, Water Distribution Systems

1. INTRODUCTION

Continuous online monitoring using ad hoc wireless networks of low cost autonomous, intelligent sensor nodes offers a new paradigm for the operation and control of large-scale urban infrastructure such as water distribution systems. The integration of near real-time data with accurate analytical models can be used in a variety of applications ranging from optimization of pump scheduling (efficient power management and water conservation), to the detection and quantification of leaks, and the implementation of an early warning system for contaminant intrusion. The principal challenges in advancing these concepts relate to the design of low cost, robust sensor technologies (especially for water quality measurands), the development of a generic cyber-infrastructure to enable efficient scaling for large networks of sensor nodes, and integration with existing simulation models (and ultimately with decision tools). These problems are compounded by the practical difficulties of limited access to underground pipeline networks.

Much recent research in the water industry relies on sophisticated modeling techniques (with intelligent reasoning such as heuristic search, genetic algorithms, fuzzy logic and neural networks). While these techniques have brought advantages in model calibration and design optimization, there is a serious lack on monitoring data to evaluate the actual performance of complex network systems. Recent work at MIT (Stoianov et al., 2006) led to the development and installation of a prototype wireless sensor network in Boston and demonstrated a proof-of-concept system for near-real-time monitoring of i) hydraulic and water quality parameters (pressure and pH); and ii) monitoring of water levels in sewer collectors and combined sewer outflows.

The Wireless Water Sentinel project in Singapore (WaterWiSe@SG) aims to demonstrate the concept of pervasive sensing to enable data driven simulation of network performance, operations and control. The goal of this research is to develop generic wireless sensor network capabilities to enable real time monitoring of a water distribution network. The project involves a collaboration between the Center for Environmental Sensing and Modeling (CENSAM), part of the Singapore-MIT Alliance for Research and Technology; the Singapore Public Utilities Board (PUB); and the Intelligent Systems Centre (IntelliSys) at the Nanyang Technological University (NTU). Research is directed towards three main applications:

1) To demonstrate the application and control of a low cost wireless sensor network for high data rate, on-line monitoring of hydraulic parameters within a large urban water distribution system. Real time pressure and flow measurements are assimilated into hydraulic models and used to improve state estimation for the network.

2) The development of systems to enable remote detection of leaks and prediction of pipe burst events. The detection of water leakage represents a critical problem in water conservation worldwide. Many older distribution networks have water losses that exceed 30% of supply. Although much smaller losses occur in Singapore (less than 5%), the development of remote leak detection capabilities can have enormous impacts on long term maintenance costs and reduce risks associated with pipe burst events. The proposed research uses high frequency pressure measurements (sampling up to 2kHz) of hydraulic transient events with a dynamic state estimation method to detect and quantify leaks, together with acoustic monitoring for accurately locating the leaks.

3) The integrated monitoring of hydraulic and water quality parameters. This task, which is still at an early stage of development, will comprise a detailed evaluation of the long term performance and robustness of non-specific water quality sensors (i.e., for measurands such as pH, chlorine residual, turbidity, conductivity and dissolved oxygen), the use/development of multi-parameter sonde technologies (combined measurements in a single chip), and the application of cross-correlation techniques to interpret water quality signatures locally within the network (i.e., through local signal processing at the node level).

WaterWiSe@SG is a long-term project spanning at least five years, and organized into three distinct phases: the first phase is an exploratory, proof of concept phase, where a small network of wireless sensor nodes are deployed to gather hydraulic data to validate the mechanical, hardware and software components of the system, to determine the processing techniques that can be applied to event detection and localization, and inform the placement of further sensors. The second phase expands the network to twenty-five nodes deployed in optimal locations, incorporating on-node and collaborative processing techniques, while introducing continuous measurement of water quality parameters. The third phase concentrates on expanding the network to one hundred nodes whilst optimizing the sensing platforms for minimal power consumption and size.

This paper describes progress from the first phase of the project: system architecture, deployment and exploratory data collection for event detection and localization analysis. In completing this first phase of WaterWiSe@SG, a WSN-based hydraulic test bed has been deployed in a section of downtown Singapore. A network of eight wireless sensor nodes continuously sample pressure (and optionally flow and acoustic data) at 2kHz, transmitting it to lab-based servers for processing and archiving. Data streams from the sensor nodes have been integrated into the on-line hydraulic modeling subsystem, responsible for on-line estimation and prediction of the water distribution system's hydraulic state. Several controlled leak-off experiments have been performed using the WaterWiSe@SG testbed, the results of which have been used to inform event detection and localization algorithms. In addition, several significant operational events have been observed within the WDS as well as real burst and leak events.

To our knowledge, WaterWiSe@SG is the first in-situ hydraulic test bed deployed to gather and process data continuously on a real water distribution system.

The rest of this paper is structured as follows: Section 2 describes the wireless sensing hardware used for in-situ monitoring, Section 3 describes the overall system architecture, Section 4 gives case studies for three uses of WaterWiSe@SG: hydraulic modeling, leak-off testing and operation event detection. Section 5 concludes the paper and outlines the next steps of the WaterWiSe@SG project.

2. SENSOR NODE DESIGN AND DEPLOYMENT

The current WaterWiSe@SG sensor node (Figure 1) is designed to continuously gather data at high (kHz) rates and transmit in real-time to the WaterWiSe@SG server. This means nodes in the field can provide full datasets that can be analyzed at a central location. However, it is also capable of performing data processing locally.

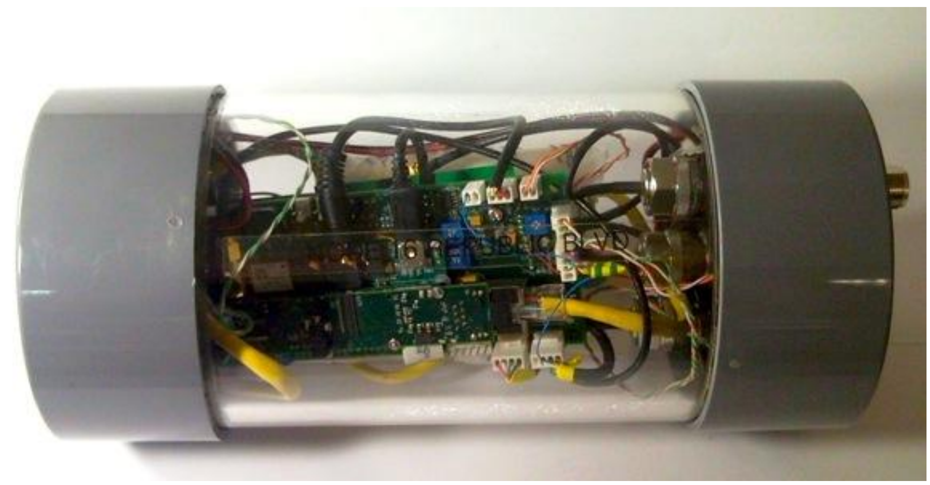

Figure 1. The WaterWiSe@SG sensor node in its water-resistant packaging

As shown in Figure 1, the sensor node is packaged in a clear plastic acrylic tube, with PVC caps at each end. These end caps house waterproof debugging ports, sensor ports, and attention buttons. The water resistant packaging facilitates the deployment of the sensors inside a manhole if needed.

The node's construction is a combination of off-the-shelf hardware and a custom-designed sensing board. The main processing board is a Gumstix Verdex Pro running the Linux operating system. The sensor node currently supports the simultaneous attachment of three types of sensor: a pressure sensor, a hydrophone and a flow meter. Both the pressure sensor and hydrophone are sampled at 2 kHz while the flow meter provides data at 1Hz. The node has 2GB of disk storage space, enough to buffer several days' worth of data in the event of any communication problems.

For wireless communication, a USB 3G modem and a USB Wi-Fi radio are connected to the main processing board; the power to each of these devices can be triggered on the main board to save energy. The USB 3G modem is used as the node's primary communication mechanism, providing an Internet connection with an average upload rate of 4-8KB/s. The Wi-Fi radio can be used for short-range communication where appropriate.

Each node is also equipped with a Pulse-Per Second capable Global Positioning System (GPS) unit. The GPS unit is used to synchronize the clock of the sensor node to UTC time, providing a global time-base for all data gathered across sensor nodes. Time synchronization accuracy of up to ±50μs accuracy is theoretically possible, and a worst case of ±1.5ms has been empirically observed under normal operation. The high accuracy of time synchronization is suitable for leak/burst location schemes based on the relative arrival times of transient pressure wave front at different points in the network.

To date, eight sensor nodes have been deployed across a 60km^2 area of Singapore on a Water Distribution System (WDS) consisting of two service reservoirs, over 19000 junctions and over 20000 pipes. The average distance between sensors is 1km. For confidentiality purposes, the exact sensor node locations with respect to the water system layout are not shown in this paper. The locations of the eight sensor nodes were determined by the water utility experts at PUB, based on their knowledge of the system and specific needs at the time of deployment. Future sensor node deployments are to be determined based on the results of simulations searching for optimal sensor placement based on system hydraulics (pressure sensor placement model, Bush and Uber 1998) and water quality (contamination event early warning system design, Ostfeld et al. 2008). The data gathered from the currently deployed nodes will aid the simulation process, as discussed in Section 4.1.

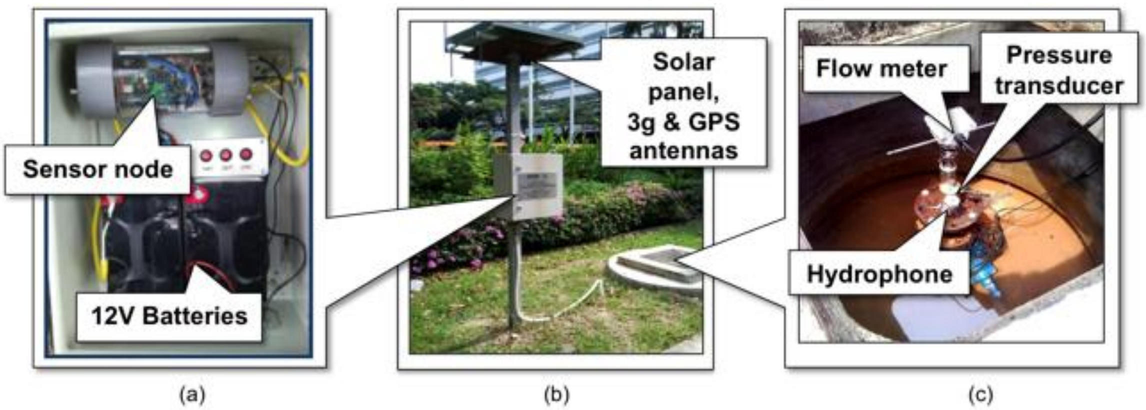

Figure 2. WaterWiSe@SG sensor node deployment: (a) shows the node and batteries, (b) shows the enclosure and solar panel mounting and (c) shows the tapping point attaching the sensors to the pipe

Each of the installed nodes is connected via a manhole to a water main through a standard tapping point that can house several types of sensor, shown in Figure 2(c). Pressure and acoustic signals (hydrophone) are attached at the periphery (side wall) of the pipe, whilst data inside the pipe (for example flow measurements) are obtained via an insertion port. The insertion port will also be used to facilitate in-situ water quality measurements.

The sensor nodes are housed above ground within an electrical enclosure, attached to a pole, shown in Figure 2 (a) and Figure 2(b). Each node is powered by a pair of 12V 33Ah batteries that are recharged during the day via a 50W solar panel attached to the top of the pole. If the site is obstructed, a wired power source such as a lamppost is used to recharge the batteries. The nominal power consumption of the current WaterWise@SG node is around 6W when acquiring data and wirelessly transmitting it over 3G. Power consumption has not been a primary concern for the current sensor node, but will become a focus in phases two and three of the project.

3. WATERWISE@SG CYBER-INFRASTRUCTURE

WaterWise@SG is an end-to-end system, dealing with the node-level acquisition and transmission of data and server-based archival, processing and visualization of data. The overall architecture of the system is shown in Figure 3. At the lowest level, sensor nodes deployed across the WDS acquire and transmit data to a lab-based group of servers via the Internet, using a 3G connection. Of these servers, the data archive is responsible for long-term storage and access to the raw data; the processing server facilitates hydraulic modeling based on the archived data, and the web server hosts the web-portal that represents the interface between the WaterWiSe@SG system and the user, facilitating visualization of data streams in near real-time. Sensor nodes can be accessed over the Internet, meaning software can be debugged, upgraded and reconfigured remotely with minimal disruption to operation. In addition sensors nodes can be accessed on-site via a direct, wired connection between a laptop and the node.

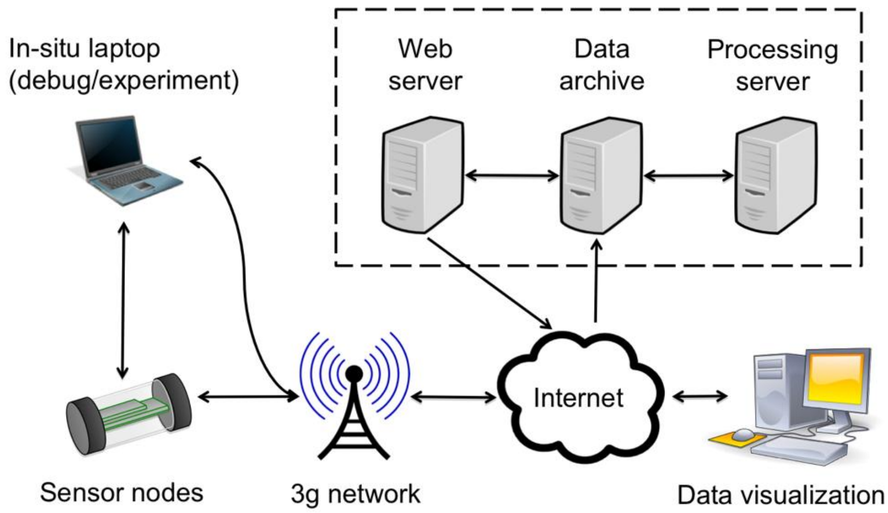

Figure 3. The WaterWiSe@SG system architecture

Figure 4 shows the flow of data and information between key components in the system. Data acquisition, reduction and transmission are performed at each sensor node, and are implemented by the middleware running on-node. In this flow, the node samples sensors continuously, and resulting streams of data are windowed into discrete 30-second files, which are then compressed before being transmitted.

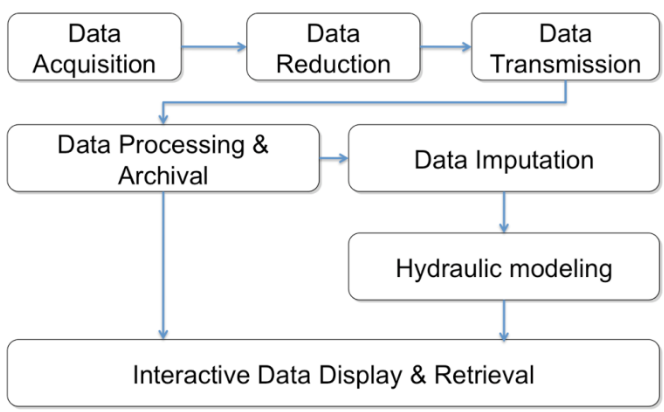

Figure 4. The system workflow

In addition to supporting this flow, the middleware periodically records data about the health of the sensor node: battery level, disk space, network transmission latencies and debugging information are all logged. Fault tolerant software components have been designed to enable automatic recovery from hardware and software failures, allowing the node to maximize its uptime.

Referring back to Figure 4, the data flow continues to the server side of the architecture where each file received is archived to disk. In addition summary statistics (mean, min, max and standard deviation) are computed for each file with the resulting values calibrated and stored into a database for visualization. From this point on, the data is available to any services that wish to make use of it. Currently in the system, there are three active services: visualization, on-line event detection and an on-line hydraulic model of the system. The on-line hydraulic model expects input data to be provided at hourly intervals, and if data is not available, then these input values are predicted using a data imputation service based on Gaussian Process Regression. The event detection component operates directly on the archived data files. Both the hydraulic modeling and event detection components are discussed in more detail in Section 4.

A vital component of the system is interactive data display and retrieval. This is provided through a web-based control panel. This control panel gives an on-line graphical interface to the data sent by the nodes, allowing the visualization of real-time sensor data for the WDS engineer/researcher and node health for the system maintainer. Near real-time graphs can be viewed on the web-portal as well as historical data from previous days or months. Figure 5 shows a screenshot from the portal, where pressure data from two sensor nodes has been overlaid to allow a relative comparison of their daily trends.

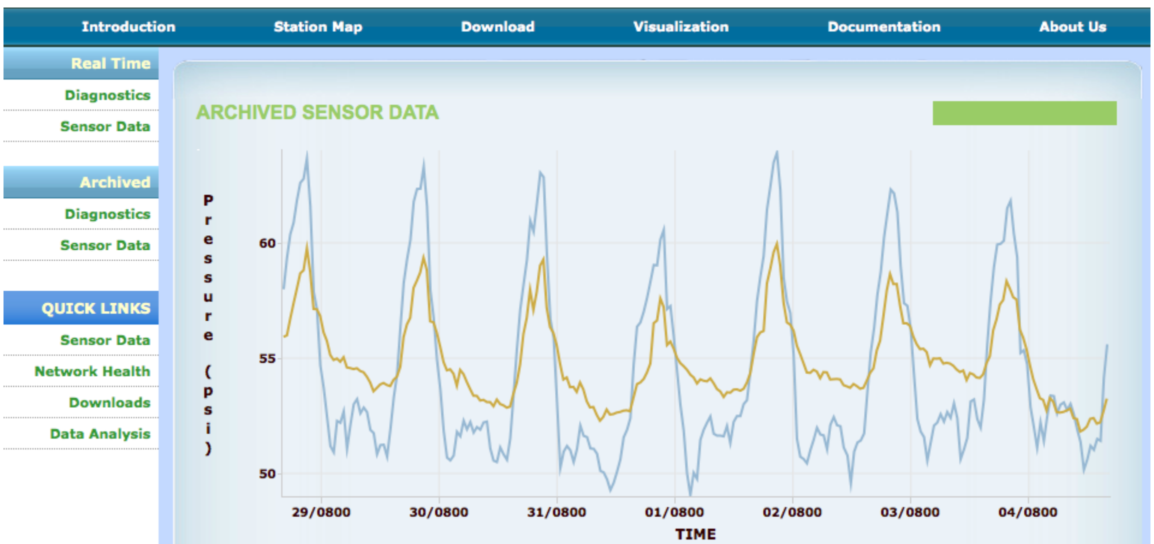

Figure 5. A screenshot from the WaterWiSe@SG portal showing two pressure traces

Data displayed on the WaterWiSe@SG portal is primarily drawn from the summary statistics stored in the database, making it suitable to understand trends on the order of minutes to months. In addition, a facility exists to show the higher-resolution 2kHz data files. To allow users to process data locally using their preferred applications, the summary data shown in graphs is also made available for download in ASCII format.

4. WATERWISE@SG SENSOR NETWORK APPLICATIONS

This section describes several applications that have been enabled by the WaterWiSe@SG system: on-line hydraulic modeling, leak and burst detection experimentation and operational event analysis. These applications represent three distinct usage models of WaterWiSe@SG: a decision support system, a hydraulics test bed and a real-time monitoring system.

4.1 On-line hydraulic modeling of a WDS

Integration of near real-time hydraulic data with hydraulic computer simulation models allows water utility engineers to operate and control their large-scale urban water distribution systems in real time. In conventional practice, hydraulic models are calibrated off-line (USEPA 2005), typically using a one-week sample of flow rate and pressure measurements within the network. Thereafter, uncertain system parameters (e.g., water demands and pipe roughness) are adjusted until an acceptable match is achieved between the model outputs and physical observations.

The main limitation of all off-line calibration procedures is that they approximate the unknown parameters using a short-term sample of hydraulic data. The calibration results may represent the system hydraulics during the short period of the sampling procedure but they are not expected to represent accurately the system conditions for the full range of operational conditions that can occur. Figure 6 illustrates the dynamic behavior of the daily-averaged pressure trend over a period of one year (May 09 - May 10) at one of the WaterWiSe@SG sensor nodes. The pressure trend is unsteady and fluctuates over a range of ~10m (~14.5psi) due to the dynamic/stochastic water consumption pattern variations and changes in the system operation. Clearly, an offline calibrated hydraulic model of this water system using a short-term data sample will not accurately represent the long-term behavior of the system. In principle, much more realistic predictions can be achieved by updating the hydraulic state-estimation using continuous on-line hydraulic measurements provided by a sensor network installed on the distribution system.

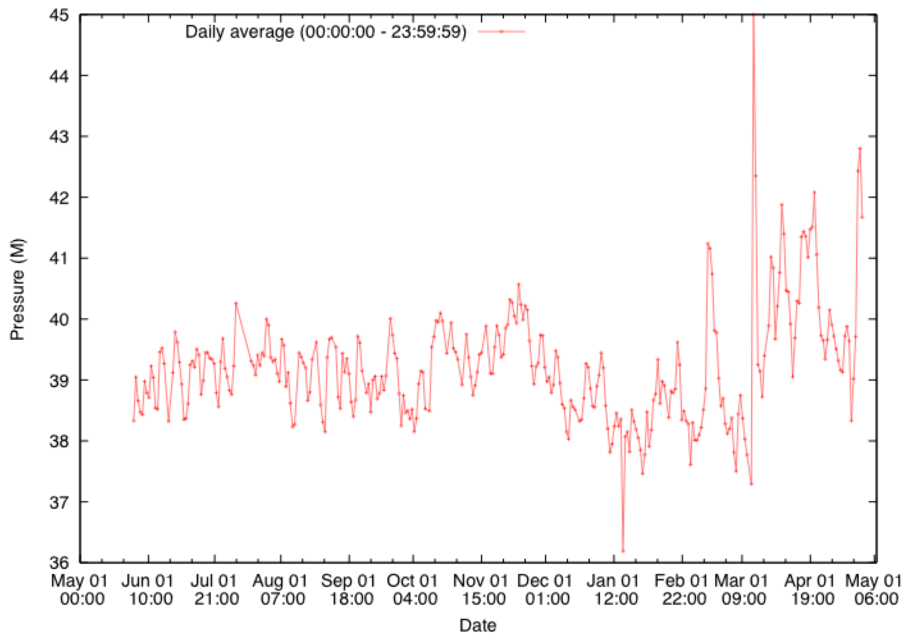

Figure 6: Fluctuation in pressure trend (i.e., daily average pressure) at one of the WaterWiSe@SG sensors over a period of one year as a reflection of the water system's dynamic behavior

An on-line hydraulic model of the urban water distribution system monitored by the WaterWiSe@SG wireless sensor network has been implemented, and a full description of the developed methodology can be found in Preis et al. (2009). The proposed method starts with identifying demand zones (i.e., clusters of water consumers) within the complex topology of the urban water supply system. The demand zone identification method implements optimization tools and graph algorithms to partition the system into homogenous clusters (i.e., water demand zones) such that (1) the within-cluster homogeneity of water consumers' characteristics is maximized - resulting with homogeneous demand zones; (2) the overall variance between total water consumption of the system's clusters is minimized - resulting with demand zones with equal base demand; and (3) the number of connecting links between neighboring clusters is minimized - resulting with demand zones that are densely linked inside and sparsely linked to the outside. Thereafter, an on-line Predictor-Corrector (PC) procedure is employed for forecasting future water demands of each zone, as shown in Figure 7.

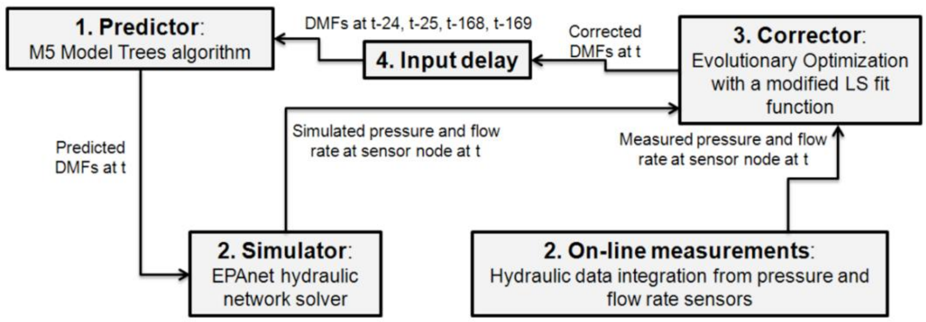

Figure 7: Predictor- Corrector loop for DMF prediction at the *t*th time step

A statistical data-driven algorithm [M5 Model-Trees algorithm (Quinlan 1992)] is applied to estimate future hydraulic states and an evolutionary optimization technique [Genetic Algorithms (Holland 1975)] is used to correct these predictions with near real-time monitoring data provided by the WaterWiSe@SG wireless sensor network. The calibration problem is solved using a modified Least Squares (LS) fit method [the Huber function (Huber 1973)] that accounts for noisy measurements in which the objective function is the minimization of the residuals between predicted and measured pressure and flow rate at several system locations, with the decision variables being the hourly variations in the zones/clusters water demands.

The calibration parameters represent variations in water demands, defined as demand multiplication factors. The Demand Multiplication Factors (DMFs) are multiplied with the baseline demands of the consumption nodes (calculated from monthly or quarterly meter readings and billing records) to obtain the actual water consumption (i.e., $D_t = D_{base} \times DMF_t$; where D_t is the actual nodal demand at time step t and DMF_t is the demand multiplication factor at the same time step). The DMFs are calibrated at each time step of the overall process. Other uncertain variables are less dynamic and their values are assumed to be constant for a certain period of time or having low impact on the hydraulic model performance. Valve and pump settings are known inputs with reasonable level of confidence, and a sensitivity analysis has shown the uncertainty of pipe roughness coefficients has a negligible impact on the hydraulic outputs (flow rates in pipes and pressures in junctions) in this particular WDS.

The current implementation of the on-line hydraulic model has been operational since January 2010. The model receives hourly averaged pressure data from the eight sensor nodes, as well as online updates from the water utility's SCADA system on the boundary conditions of the system (i.e., the service reservoirs water elevations and outflows). Running the on-line predictor-corrector hydraulic model requires continuous data from all sensor nodes. If data is temporarily unavailable, a data imputation technique is implemented where data trends in each node's data stream are tracked and data is predicted using a technique based on Gaussian Process Regression (Goldsmith et al. 2010).

The modeling considers 16 demand zones (3 residential, 4 commercial, 3 industrial, 1 mixed industrial-commercial, and 4 mixed residential-commercial) based on the demand zones identification algorithm. In order to gain confidence in the model, and also to identify possible shortcomings, several measures were used to evaluate its performances such as cross-validation with supplementary independent pressure and flow-rate measurements in different locations across the water distribution system. The results for the supplementary pressure sensors showed a mean absolute error of ~1.5 psi and for the supplementary flow meter the mean absolute error was 23m^3/hr.

The preliminary assessment of the model performance at the end of Phase 1 of the WaterWiSe@SG project has shown promising results (Preis et al. 2010), taking into consideration the limited number of sensor nodes that may not provide complete coverage of this water distribution network. This analysis of the model performance provides framework for a much more comprehensive and elaborate on-line hydraulic model implementation during phase 2 of the WaterWiSe@SG project at which the size of the sensor network will increase to 25 sensor nodes and much better coverage of the water system will be available.

4.2 Visualization of hydraulic states within the water network

A real-time three-dimensional (3D) visualization tool was developed to display the hydraulic model outputs. The tool developed with the Center for Advanced Media Technology at Nanyang Technological University of Singapore, in order to provide the on-line synchronized 3D display of water consumption data in demand zones, pressure data at the system junctions, and flow rates in the network pipes. The 3D visualization of pressure and flow rates is not shown in this paper due to confidentiality requirements. Min-max values can be set as boundaries for normal pressure conditions and in case of abnormal pressure data (i.e., pressure values exceeds these boundaries) the location of the outlier is pointed so that it can be identified and further investigated by the water utility engineers.

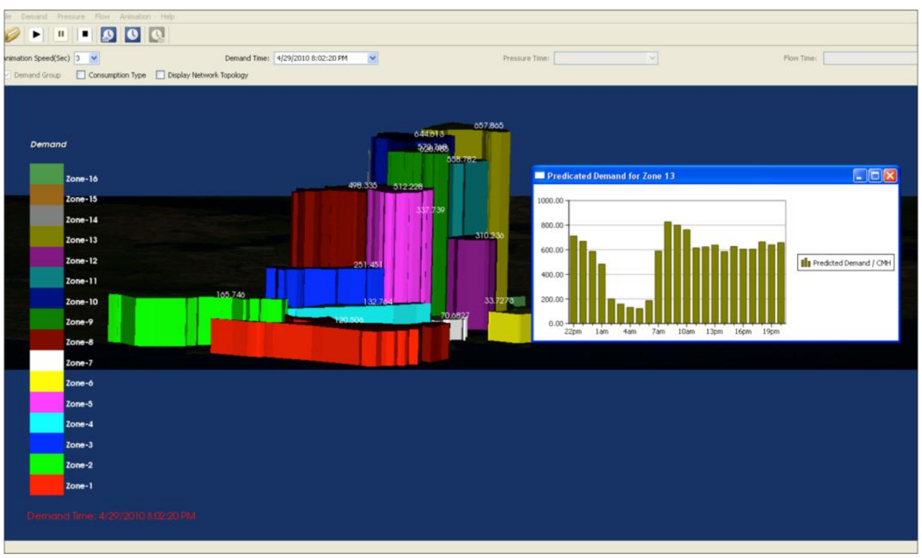

Figure 8: 3D visualization of the online hydraulic model predictions (developed with the Center for Advanced Media Technology at Nanyang Technological University of Singapore)

Figure 8 shows the main interface of the 3D visualization tool where only water demands in the demand zones are visualized. The software receives continuous predictions of demand variations from the online hydraulic model and then allows a user to run extended period hydraulic simulations with the EPAnet as the hydraulic engine. The predicted demands (24 hrs ahead) are used as inputs to the system hydraulic model and the model outputs (e.g., pressure and flow rates across the system) are displayed in 3D alongside the demands. A water utility operator can run the visualization software at any time of the day and see the 3D synchronized predictions of demand, pressure, and flow rates across the system 24 hours ahead at one hour time-step intervals.

Additional features of the software are: (1) displaying demands according to user type (e.g., residential, commercial, etc.); (2) displaying bar charts of the predicted demand (24 hrs ahead) for each demand zone; (3) displaying in 3D extended period simulation outputs of historical hydraulic data including demands, pressure and flow rates across the network; (4) allocating points of interest such as low pressure points on a popup Google map to identify the exact location and address of these points of interest; and (5) showing the 3D data on a map in different angles and point of views allowing the user to look at the data in different scales (zooming in and out/ etc.)

4.3 Online modeling of maintenance operations in the water utility

The online hydraulic model was found to be useful in modeling the system response to various maintenance operations that took place in the distribution system. In the following example, a water main was isolated for maintenance operations by the water utility. Closing several valves within the system isolated the pipe. Figure 9 provides a schematic description of this operation at which the link L1 was isolated by closing valves V1 and V2; a pressure sensor M1 is located upstream of the isolated link and pressure sensors M2, M3, and M4 are located downstream of the closed pipe.

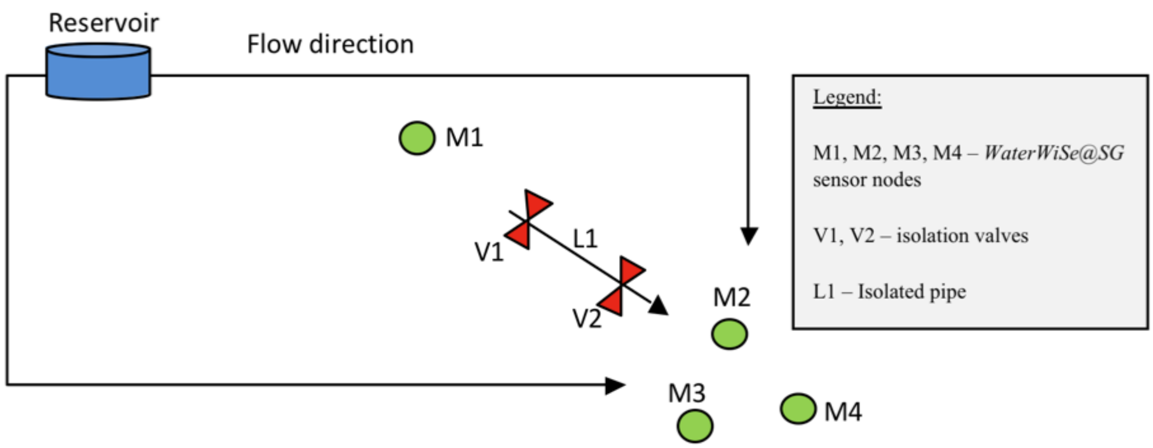

Figure 9: Schematic description of a maintenance operation in monitored water system

During the five-hour operation, the pressure at a sensor node that was located upstream of the isolated pipe increased while pressure records at sensor nodes which were located downstream of the isolated section decreased. At the end of the maintenance operation, when the isolation valves were open again, the pressure in all sensors nodes returned back to normal. The pressure data at the relevant sensor nodes is presented in Figure 10.

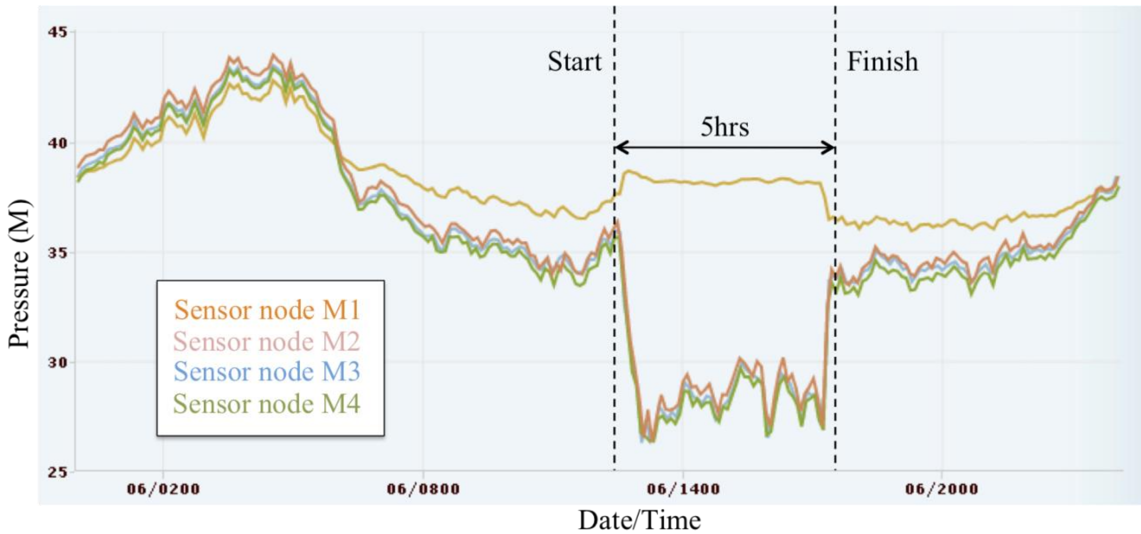

Figure 10: pressure data in relevant sensor nodes during a maintenance operation (pipe shutdown) in monitored water system

The on-line hydraulic model was used to predict the system response to this event by updating the valve settings for V1 and V2 in the system's calibrated EPAnet model from open to closed status during the expected maintenance operation period. Figure 11 shows the predicted pressure values for sensor nodes M1 and M2 alongside the observed pressure values at these same sensor nodes.

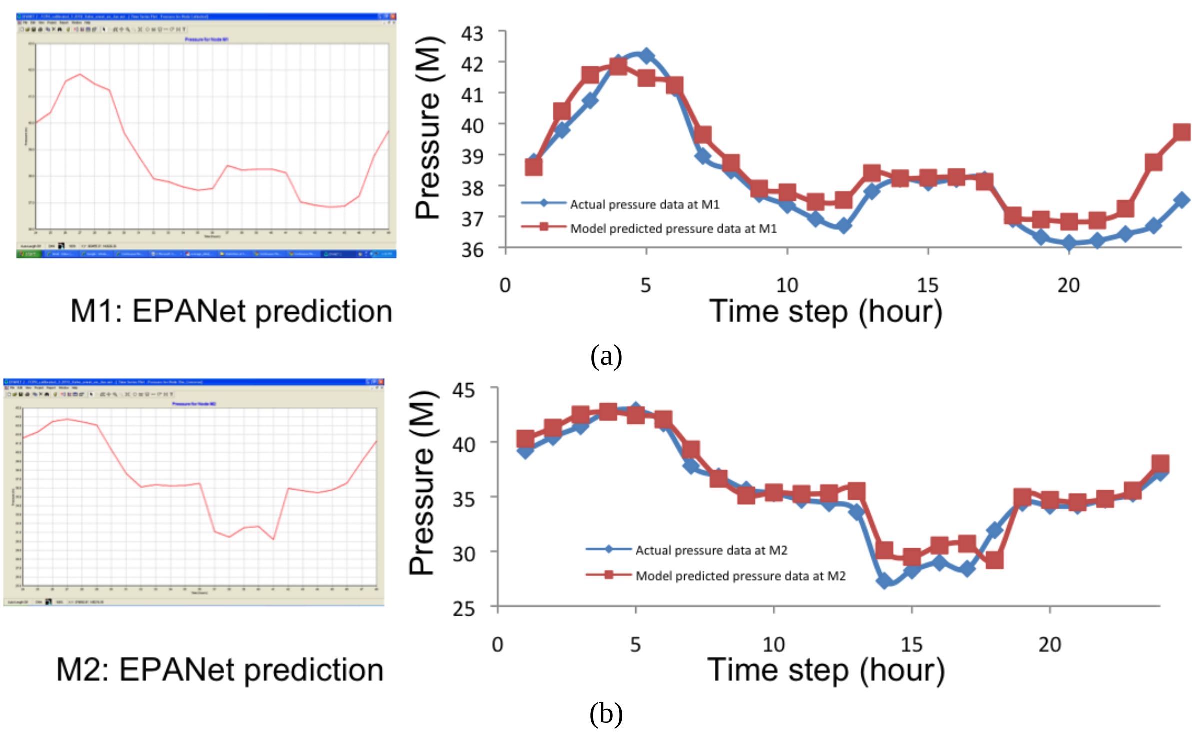

Figure 11: Comparison between actual and predicted pressure data in M1 (a) and M2 (b) sensor nodes during a maintenance operation (pipe shutdown) in the monitored area of the water system

This example shows how useful an on-line calibrated hydraulic model can be in predicting the system response to hydraulic events within the water utility. This hydraulic state information is useful to the water utility operator to facilitate the forecasting of min-max pressures that would be experienced in the system during proposed maintenance operations.

4.4 WDS field experiments and engineered tests

The use of WaterWiSe@SG as a hydraulic test bed is a key component of meeting the project's research goals. An in-situ test bed deployed on a real WDS, continuously streaming data in real-time allows experimentation that has not previously been possible. In the first phase of the WaterWiSe@SG project, we have used the test bed to conduct several controlled leak-off experiments. These experiments were carried out to gather data for developing and analyzing leak detection algorithms.

(a) (b)

Figure 12. During leak-off experiments artificial bursts and leaks are generated by fast-opening solenoid valves connected to air valves (a) or fire hydrants (b)

During leak-off experiments a solenoid-controlled valve is connected to the pipeline via an air-valve or a fire hydrant, as shown in Figure 12. The solenoid-controlled valve is triggered to create transient events emulating instantaneous pipe ruptures. The fast opening (0.1s) creates a sudden pressure drop over less than a second that can travel considerable distances around the network before being fully damped. Since sensor nodes are gathering high-rate data continuously, they capture both the pressure front and subsequent reflections. To date, four leak-off tests have been carried out at different system locations across the downtown area. Figure 13 shows a typical experimentation set-up and also shows the pressure records at sensor nodes during an artificial burst event.

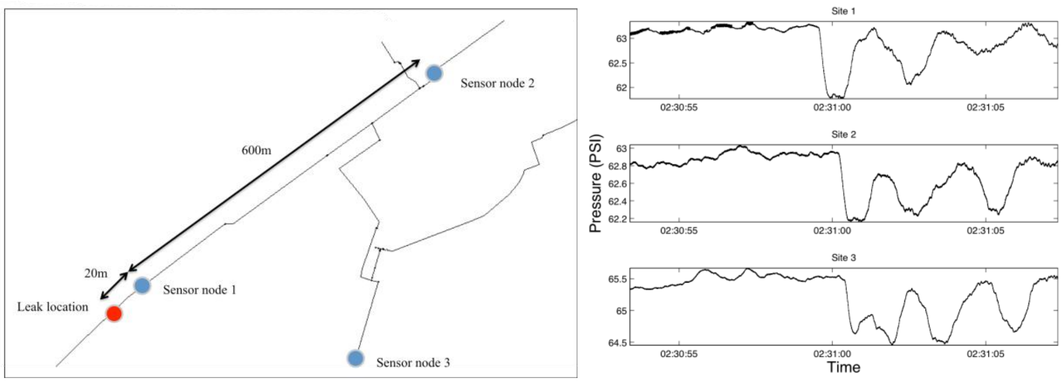

(a) (b)

Figure 13. A typical leak-off test set-up (a) and associated pressure records taken during the event (b)

This transient data has been used to test and tune event detection techniques with the eventual goal of running an automated end-to-end detection and localization system across the sensor nodes and server tiers. Thus far, we have experimented with two approaches to event detection: simple analysis of the rate of change of the pressure data in the time-domain, and wavelet decomposition. The time-domain approaches are suitable for deployment on sensor nodes, whereas the wavelet analysis requires more computational capability and thus is currently more suited to server-side operation.

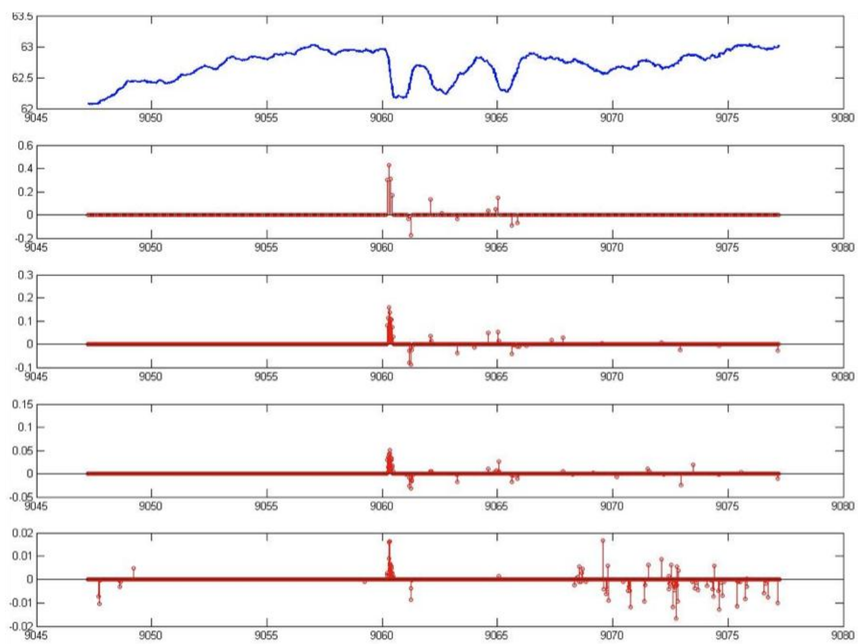

Figure 14. The graph shows the wavelet decomposition of a controlled pressure transient into four levels of detail coefficient. The transient is consistent across these four levels, thus easy to classify against noise.

The wavelet detection algorithm determines the approximate time of the pressure front arrival (due to the down-sampling operation during the wavelet decomposition) at the various sensor locations. Figure 14 shows a multi-scale wavelet decomposition of a controlled pressure transient, where pressure signal is decomposed into several levels of detail coefficients. Analyzing the coefficients over the different levels of detail and picking the most consistent signal determines the onset of the transient event. The arrival time estimates of the pressure front are used as input to a burst/leak localization algorithm that can be used to localize pipe burst events (Misiunas et al. 2005). The localization algorithm uses a graph-based search procedure to determine the physical location that best matches the relative time difference of arrivals of the pressure front between sensor nodes.

4.5 Online monitoring of maintenance operations in the water utility

Although water losses due to leaks or burst events in Singapore are very small (less than 5%), leaks and/or pipe breaks may occur occasionally. Such an event took place when a 500mm water main ruptured during a road re-pavement operation above the buried pipeline. A nearby WaterWiSe@SG sensor node measured a significant pressure drop, as shown in Figure 15. This demonstrates the benefit of gathering continuous high-rate data, providing assistance to water utility engineers in the management and control of the water network and in responding quickly to events such as leaks and pipe bursts.

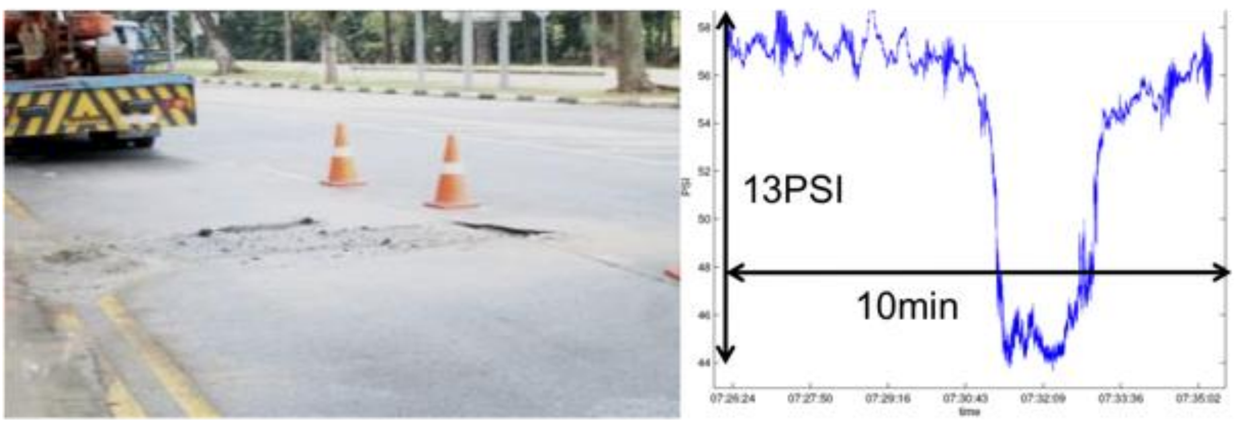

(a) (b)

Figure 15. A water-main break caused by road re-pavement operation (a) and the pressure trace (b)

5. CONCLUSIONS AND FUTURE WORK

This paper has described phase one of the Wireless Water Sentinel in Singapore (WaterWiSe@SG), a wireless sensor network to enable real-time monitoring of a water distribution network in Singapore. The overall project is directed towards three main goals: 1) the application of a low cost wireless sensor network for high data rate, on-line monitoring of hydraulic parameters within a large urban water distribution system; 2) the development of systems to enable remote detection of leaks and prediction of pipe burst events; 3) the integrated monitoring of hydraulic and water quality parameters. To our knowledge, WaterWiSe@SG is the first in-situ hydraulic test bed deployed to gather and process data continuously on a real water distribution system.

Future Research Plans are 1) to refine the sensor node design to achieve greater efficiencies in power consumption; 2) evaluate the performance of non-specific water quality sensors for in-line monitoring (current plans are to focus on pH/ORP and chlorine sensors); 3) to complete the next phase deployment of 20-25 nodes (with locations optimized using the calibrated hydraulic model of the water network); 4) to implement a software system for managing the data streams (including a data query system based on the MIT Wavescope project, Girod et al. 2007); and 5) implement and evaluate local algorithms for leak detection and localization on the sensor nodes.

6. ACKNOWLEDGMENT

This work is supported by funding from the National Research Foundation of Singapore (NRF) and the Singapore-MIT Alliance for Research and Technology (SMART), through the Center for Environmental Modeling and Sensing (CENSAM). The authors are grateful to the Singapore Public Utility Board's water network engineers and field operations staff, for their tremendous help during deployment, installation and field experimentation. The authors also wish to acknowledge the collaboration with CAMTech at Nanyang Technological University, Singapore, in particular Mr. Cheng Zhi-Gao, Mr. Gerrit Voss and Dr. Wolfgang Müller-Wittig.

References

Alvisi, S., Franchini, M., Marinelli, A., (2007), *A short-term, pattern-based model for water-demand forecasting*, Journal of Hydroinformatics, Vol. 9, No. 1

Bush, C. A., and Uber, J. G. (1998). *Sampling design methods for water distribution model calibration*, J. Water Resour. Plann. Manage.,124(6), 334–344

EPANET. (2002), *Drinking Water Research*, URL: www.epa.gov/ORD/NRMRL/wswrd/epanet.html

Girod, L., Jamieson, K., Mei, Y., Newton, R., Rost, S., Thiagarajan, A., Balakrishnan, H., Madden, S. *WaveScope: A Signal-Oriented Data Stream Management System*, CIDR 2007

Goldsmith, D., Preis, A., Allen, M. &Whittle, A.J. (2010) *Virtual sensors to improve on-line hydraulic model calibration*, Proc. 12th Water Distribution Systems Analysis Symposium (WDSA10), Tucson, Arizona, (to appear)

Holland J. H. (1975). *Adaptation in natural and artificial systems*, The University of Michigan Press, Ann Arbor

Huber, P. J., (1973), *Robust regression: Asymptotics, conjectures, and Monte Carlo*, Ann. Statist., 1, 799–821

Misiunas, D., Vitkovsky, J., Olsson, G., Simpson, A. & Lambert, M. (2005) *Pipeline break detection using pressure transient monitoring*, ASCE Journal of Water Resources Planning and Management, 131(4), 316-325

Ostfeld A. et al. (2008). *The battle of the water sensor networks: a design challenge for engineers and algorithms*, Journal of Water Resources Planning and Management Division, ASCE, Vol. 134, No. 6, pp. 556–568

Preis, A., Whittle, A.J. & Ostfield, A. (2009) *On-line hydraulic state prediction for water distribution systems*, Proc. 11th Water Distribution Systems Analysis Symposium (WDSA09), World Environmental & Water Resources Congress, Kansas City, MO, May

Preis, A., Allen, M., Whittle, A.J. (2010) *On-line hydraulic modeling of a Water Distribution System in Singapore*, Proc. 12th Water Distribution Systems Analysis Symposium (WDSA10), Tucson, Arizona, (to appear)

Quinlan J. R. (1992) *Learning with continuous classes*, Proceedings 5th Australian Joint Conference on Artificial Intelligence. World Scientific, Singapore, 343–348

Stoianov, I., Nachman, L., Whittle, A.J., Madden, S. & Kling, R. (2006) *Sensor networks for monitoring water supply and sewer systems: Lessons from Boston*, Proc. 8th Intl. Water Distribution Systems Analysis Symposium, (WDSA2006), Cincinnati, 19p

Shang, F., Uber, J., van BloemenWaanders, B., Boccelli, D., Janke, R.(2006) *Real Time Water Demand Estimation in Water Distribution System*, 8th Annual Water Distribution Systems Analysis Symposium, Cincinnati, Ohio, USA, CD-Rom

Srirangarajan, S., Allen, M., Preis, A., Iqbal M., Girod, L., Fu, C., Wong, K-J., Lim, HB. &Whittle, A.J. (2010) *Water main burst event detection and localization*, Proc. 12th Water Distribution Systems Analysis Symposium (WDSA10), Tucson, Arizona, (to appear)

USEPA (2005), *Water Distribution System Analysis: Field Studies, Modeling and Management. A Reference Guide for Utilities*, USEPA Cincinnati, Ohio, USA