

## MIT Open Access Articles

*Water Distribution System Monitoring and Decision Support Using a Wireless Sensor Network*

The MIT Faculty has made this article openly available. **Please share** how this access benefits you. Your story matters.

**Citation:** Allen, Michael, Ami Preis, Mudasser Iqbal, and Andrew J. Whittle. "Water Distribution System Monitoring and Decision Support Using a Wireless Sensor Network." 2013 14th ACIS International Conference on Software Engineering, Artificial Intelligence, Networking and Parallel/Distributed Computing (July 2013).

**As Published:** <http://dx.doi.org/10.1109/SNPD.2013.97>

**Publisher:** Institute of Electrical and Electronics Engineers (IEEE)

**Persistent URL:** <http://hdl.handle.net/1721.1/92758>

**Version:** Author's final manuscript: final author's manuscript post peer review, without publisher's formatting or copy editing

**Terms of use:** Creative Commons Attribution-Noncommercial-Share Alike



# Water Distribution System Monitoring and Decision Support Using a Wireless Sensor Network

Michael Allen\*, Ami Preis\*<sup>†</sup>, Mudasser Iqbal\*<sup>†</sup>, Andrew J. Whittle<sup>‡</sup>

\*Visenti Pte Ltd, Singapore

<sup>†</sup>Singapore MIT Alliance for Research and Technology, Singapore

<sup>‡</sup>Massachusetts Institute of Technology, Cambridge, MA, USA

michael@visenti.com, {amipreis, mudasser}@smart.mit.edu, ajwhittle@mit.edu

**Abstract**—Water distribution systems comprise labyrinthine networks of pipes, often in poor states of repair, that are buried beneath our city streets and relatively inaccessible. Engineers who manage these systems need reliable data to understand and detect water losses due to leaks or burst events, anomalies in the control of water quality and the impacts of operational activities (such as pipe isolation, maintenance or repair) on water supply to customers. WaterWiSe is a platform that manages and analyses data from a network of wireless sensor nodes, continuously monitoring hydraulic, acoustic and water quality parameters. WaterWiSe supports many applications including rolling predictions of water demand and hydraulic state, online detection of events such as pipe bursts, and data mining for identification of longer-term trends. This paper illustrates the advantage of the WaterWiSe platform in resolving operational decisions.

**Index Terms**—Wireless Sensor Network, event detection, decision support systems, distributed processing

## I. INTRODUCTION

Drinking water distribution systems carry potable water from water sources (reservoirs, water tanks) to industrial, commercial and residential consumers through complex pipe networks. An adequate water supply (in terms of availability and quality) is vital to performing many daily tasks, and disruptions to a water supply through planned maintenance or leakage repairs can cause significant problems.

A drinking water distribution system's operation is planned and managed by a department of engineers and physically maintained by field crews. These teams work together to resolve immediate system failures such as pipe leaks and bursts and longer-term upgrades. Many decisions about how best to manage the network are complex and cannot readily be automated; in these cases engineers must rely on their experience, knowledge of the current state of the system and decision support tools to help them make informed choices about the best courses of action.

The most detailed data about the current state of the water network (flow, pressure, water quality) is gathered using SCADA systems at entry points (reservoirs, water tanks), rather than within the network itself. There are generally very limited on-line monitoring and analysis capabilities deployed on the pipes and valves within water distribution networks—pressure and water quality are spot checked to comply with regulations or deal with customer complaints, and background

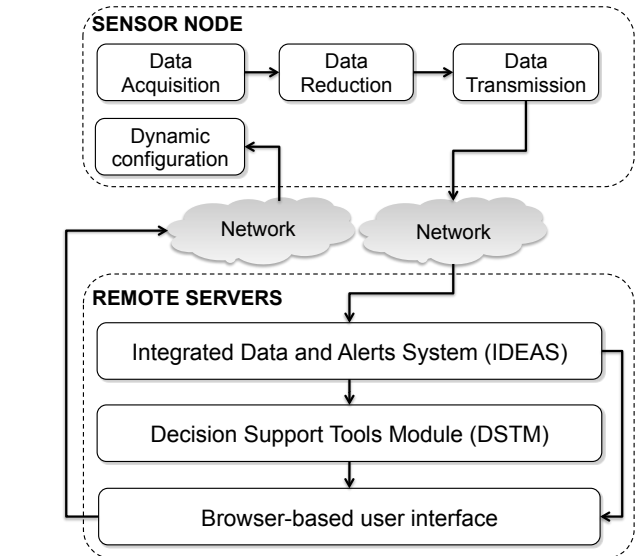


Figure 1. Data flow between WaterWiSe's WSN, IDEAS and DSTM components. Sensor nodes can be dynamically reconfigured through the browser-based interface, which also provides a means to interact with DSTM and IDEAS.

leakage surveys are performed periodically around the system. Commercial on-line monitoring solutions tend to be application and parameter specific, meaning that utilities have to find ways to integrate each solution into their overall software systems and processes. Even if data is available, there is often a lack of expertise in-house to be able to analyze the data in order to make decisions.

A lack of data and analysis capabilities is a problem when there is an increasing need to reduce leakage and operational inefficiencies (such as power consumption) while providing an adequate water supply.

WaterWiSe is a platform for water distribution systems operators that provides real-time monitoring and decision support tools [1] that can be used to help improve system management and operation. WaterWiSe uses a combination of model-based prediction and data stream analysis as a basis for its decision support tools. The primary data source is a wireless sensor network (WSN) that is deployed on the pipes within the distribution system, providing on-line updates of the hydraulics and water quality. The increase of deployment

density, number of parameters and on-line data collection enabled by deploying the WaterWiSe WSN allows a water utility to have a more detailed picture of the water distribution system across a variety of important applications, such as leak detection, water quality monitoring, on-line hydraulic calibration and others.

This paper describes WaterWiSe’s hardware and software design, and presents two indicative scenarios that demonstrate the importance of WaterWiSe at every step in the decision-making process: burst repair and pump optimization. The rest of this paper is divided up as follows: Section II describes the hardware and software of the WaterWiSe platform in detail, Section III presents two decision support scenarios, Section IV discusses related work and Section V concludes the paper.

## II. THE WATERWISE PLATFORM

The WaterWiSe platform has three components: i) an on-line wireless sensor network, ii) the Integrated Data and Electronic Alerts System (IDEAS), and iii) the Decision Support Tools Module (DSTM). The WSN provides on-line streams of data or events; IDEAS processes raw data streams to detect and localize abnormal events (such as leaks or contamination events); and DSTM provides decision support tools based around a hydraulic model of the water distribution network that is periodically calibrated using data from the WSN. These components provide key services to help both water supply network planning and operations teams in the office and in the field.

WaterWiSe can operate as a stand-alone system as well as a component in an integrated water management system. In stand-alone mode, a map-based web interface and dashboard is provided to both IDEAS and DSTM (see Figure 3). This user interface is accessible through the web-browser on regular desktop PCs as well as tablet and smart phone, allowing for in-field analysis and validation by the field crew. WaterWiSe is either hosted locally within a water utility or over the internet via the cloud.

WaterWiSe has been deployed as a system with the Public Utilities Board (PUB) of Singapore since 2009, with the sensor network scaling from eight nodes in one supply zone to tens of nodes across several supply zones (over 80km<sup>2</sup>). The WaterWiSe platform is linked in to PUB’s SCADA system to exchange sensor data, and operates remotely with a stand-alone, browser-based interface.

The rest of this Section describes the WSN, IDEAS and DSTM in more detail.

### A. Wireless Sensor Network

The sensor nodes in the WaterWiSe network need to sample multiple parameters and provide time-synchronized data for analysis. Figure 1 shows the data flow between the WSN, IDEAS and DSTM. The hardware, software and mechanical design of the WaterWiSe WSN are described below.

1) *Hardware and mechanical design:* Each sensor node in the WaterWiSe network is based around a 72MHz ARM Cortex M3 CPU with 64KB of RAM, a 2GB SD card for




Figure 2. The WaterWiSe multi-probe holding hydraulic and water quality sensors (left) and the assembly installed on a pipe through a gate valve (right)

storage, a GPS with pulse-per-second functionality for time synchronization, and a 3G modem for data transmission. Nodes do not currently operate in an ad-hoc multi-hop network because in practice, the average inter-node distance of between 500 to 1500m tends to be greater than 802.15.4 or 802.11 class radios can reliably transmit in an urban environment. The 3G connection gives typical upload speeds of around 4-5KB/s, which is more than adequate for the data generated.

Sensors are interfaced to the node using a custom I/O board that supports hydraulic parameters (pressure, hydrophone, flow), and water quality (pH, conductivity, temperature and Oxidation Redox Potential [ORP]). Each node is powered by a 33Ah 12V battery that is charged either via solar panels during the day, or via lampposts when they switch on at night. The sensor node consumes 0.036W when in sleep mode, 0.75W when sampling, and 4.5W when transmitting data over 3G. Whilst the device is not intended to be ultra-low power due to high sampling rates, energy can be conserved through software configuration when battery charging options are limited.

The sensor node is stored in an enclosure with its batteries and mounted on a pole. Hydraulic and water quality sensors are mounted on an integrated multi-probe that is inserted into the water pipe via a gate valve. Figure 2 shows an example multi-probe and sensor node deployment.

2) *Software services:* WaterWiSe nodes can be remotely and dynamically configured to sense and transmit data in a variety of ways. Individual parameters are typically sampled at intervals relevant for their expected rates of change, and transmitted at rates relevant for the applications they may serve. Hydrophone data is used to detect background leakage, and is sampled at kHz rates, but only gathered for a few seconds every hour (120KB per hour). Pressure data is continuously sampled at hundreds of Hz in order to capture transients that are indicative of pipe bursts. Data is either transmitted in 30-second windows of raw samples (around 15KB per file), or processed locally using an IDEAS event detection algorithm so that only event notifications and short data traces are sent. Water quality parameters, flow and battery data are typically sampled at 30 second intervals, aggregated, and transmitted every 5 minutes (hundreds of bytes per file).

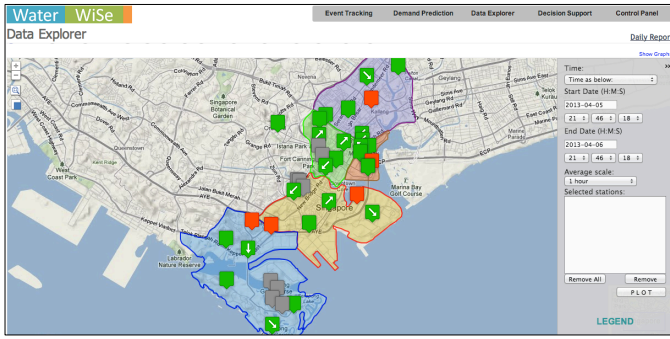


Figure 3. The IDEAS interface to WaterWiSe, showing deployment layout for per-node data-stream inspection.

Data transmission is acknowledged at application level by the remote server using checksums—unsuccessful transmissions are retried, and network or radio failures trigger reconnections and eventually radio power cycling. When commands need to be passed to nodes, they are passed at the end of a successful upload. Nodes can buffer several days of data in the event of a network outage, and will prioritize lower-rate data when network connections have low-bandwidth.

A custom time synchronization service uses NMEA messages and pulse-per-second signals to relate the on-board system clock to GPS time, providing a global frame of reference for all data collected by sensors. Sensor node data is typically synchronized to within 0.5ms of GPS.

### B. Integrated Data and Electronic Alerts System (IDEAS)

IDEAS is responsible for data stream management, processing and alert notification. On-line data from the WSN is archived as it arrives along with data from other sources such as a utility’s SCADA system, or operational information such as valve closures and pump schedules.

IDEAS has a set of analytics that are applied to data streams in order to detect abnormal events and provide location estimates. Detection of an abnormal event triggers an alert to be sent via e-mail or SMS to subscribed users of WaterWiSe. Figure 3 shows the user interface to IDEAS, where real-time data streams from all sensors can be reviewed, as well as alert notifications. The specific analytics applied depends on the parameters being processed (pressure, acoustics or water quality), but the basic data flow is common: abnormal events are detected in the data stream of each sensor node and are grouped together and used to estimate the location of the abnormal event.

1) *Transient detection and localization*: Certain pipe bursts release characteristic pressure transients lasting less than 1s that travel around a pipe network at over 1km/s. In IDEAS’ burst detection analytics, the high rate synchronized pressure data gathered from each sensor node is filtered in the time domain to search for indicative drops in pressure that deviate significantly from the background pressure levels. If a pressure transient is detected by enough nodes, the arrival times are used to estimate the location of the burst. This is accomplished

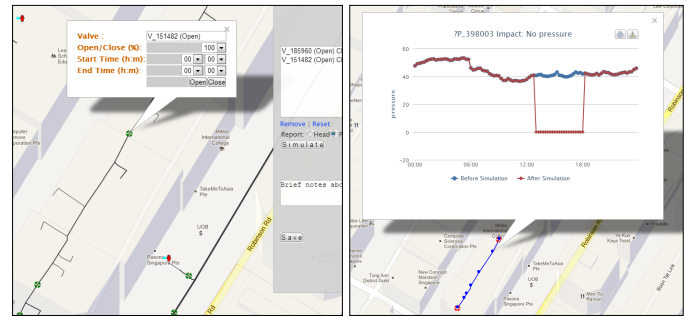


Figure 4. Example output from an operational simulation. The left hand side shows the selection interface for valve closure, and the right hand side shows the resulting impact after running the simulation. The pop-up shows a predicted data trace at a given junction with and without valve closure.

by matching the arrival times from the detecting sensors to theoretical arrival times of all junctions in a graph model of the pipe network (where pipe lengths are arcs and junctions are vertices), taking into account the speed of sound in water (around 1200m/s) [7]. The set of most likely junctions are indicated to the user on the WaterWiSe map for further analysis (see Figure 6).

IDEAS’ leak and burst detection and localization has been validated with a series of controlled leak-off experiments, where leaks are created through fire hydrants by opening a solenoid valve into open air. This approach led to an average localization accuracy of 37.5m over 9 trials on a 1km<sup>2</sup> section of pipe network [7].

### C. Decision Support Tools Module (DSTM)

DSTM uses the sensor data streams provided by IDEAS as input to decision support tools that model the water network as a set of demand-zones. At the core of DSTM is a real-time hydraulic model of the water distribution system that is calibrated at intervals as short as every 15 minutes [6]. The calibration process uses the most recent pressure and flow data from sensor nodes along with a baseline demand pattern and seasonal information to estimate the consumption for each demand-zone in the network model. To improve user confidence in the model, the user can compare model predictions with the actual data that was collected in a given time period. The on-line hydraulic model enables several decision support tools, discussed below.

1) *On-demand valve operation simulation tool*: The operation simulation tool can be used by engineers to analyze the potential impact on the network of an operational event such as valve closures and pipe isolations. The simulation interface allows the engineer to select one or more valves to isolate and determine the times for which these valves will be closed. When it has completed, the simulation results are presented, identifying pipes that will have low or reversed flow, areas that will have abnormally low or high pressure, and customers that will be isolated by the operation. Figure 4 shows a screenshot of the valve operation tool.

2) *Demand prediction tool*: Water consumption can be predicted in advance for a 24-hour rolling window. Predictions



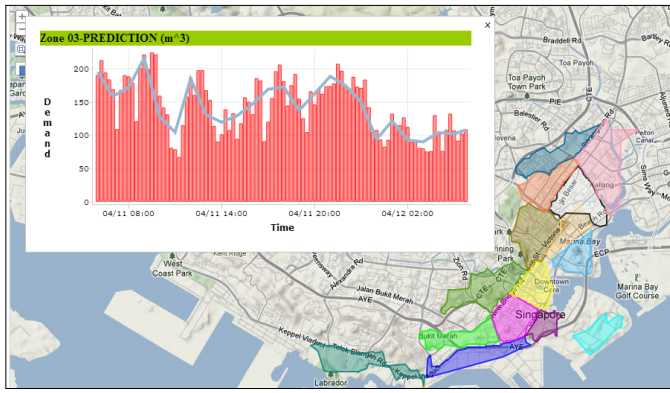


Figure 5. DSTM's demand prediction tool showing selectable sub-zones; the pop-up shows predicted consumption within a given sub-zone 24 hours ahead. The interface can also show daily consumption vs. predicted daily consumption across the whole supply zone and predicted for the next 24-hour period.

can be shown on a daily summary, with comparison to the actual consumption, across the whole zone or in specific sub-zones. Figure 5 shows a screen shot of the demand prediction interface.

3) *Water age and water source analysis tool*: Using the hydraulic model, water age in the system can be predicted and compared against the realtime water quality measurements being taken in the system. This helps to identify areas of high water age that may be of concern. The mixing of water in the system from different reservoirs can also be predicted and visualized, showing relative percentages of water sources at any given location over user-defined time periods.

4) *Pump optimization tool*: In systems where water is pumped to consumers or storage tanks, WaterWiSe can help reduce cost and energy consumption by optimizing the pump operation schedule based on electricity tariffs, pump efficiency and predicted consumption. The pump optimization interface is shown in Figure 7.

#### D. Use of WaterWiSe

As previously noted, WaterWiSe has been used by the PUB of Singapore since 2009, and has scaled from a proof-of-concept to a larger sensor network covering several supply zones and providing decision support tools that are used in day-to-day operation of the water system. WaterWiSe has also been used as a live research testbed for leak and burst detection experimentation that informs development and tuning of detection and localization algorithms. During its development and expansion over four years, WaterWiSe has accumulated a large archive of sensor data from sensor nodes. This archive has proved an invaluable resource on which to perform exploratory data mining to identify common patterns and trends, and to help identify areas of concern in the network that warrant further, in-depth investigation. It is also helpful for post-mortem analysis of response to events, where the actions that were taken can be matched to data streams and used to help inform future actions.

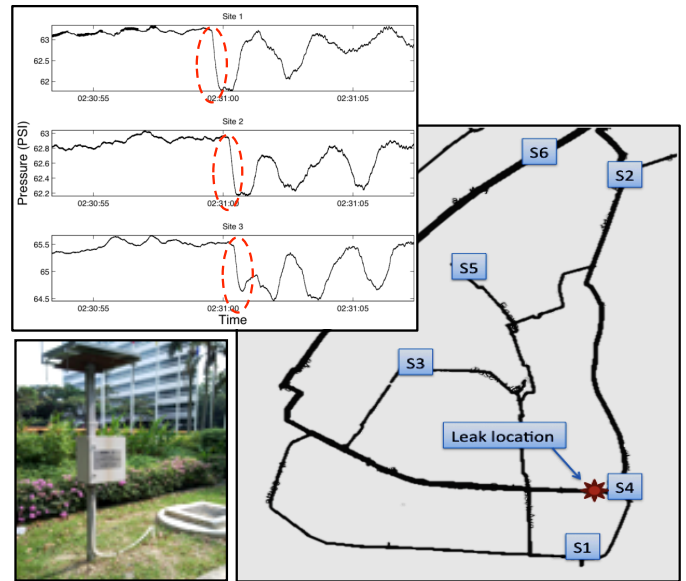


Figure 6. An overview of burst detection—an event occurs and is detected by several sensor nodes deployed around the network. The top left-hand corner shows the sub-second transients in the data trace. The detection times are used to track back to the most likely location using a graph model of the network.

### III. DECISION SUPPORT SCENARIOS

This section describes two real decision support scenarios that illustrate the usage of WaterWiSe, showing how the WSN, IDEAS and DSTM components combine together to facilitate an end-to-end approach.

#### A. Detecting and resolving a leak

Leaks and bursts are an unavoidable aspect of water distribution systems management, and can account for significant water loss within a distribution network if left undetected for long periods. Leaks often occur through a build up of corrosion that causes structural failures in aging pipes, particularly at joints. Large, obvious leaks are often easy to spot due to water appearing at surface level and are usually reported by field crews or members of the public; existing background leakage is often pro-actively surveyed by leak-detection crews who deploy acoustic listening devices in areas for short periods, or take flow measurements during periods of low consumption such as early morning. In all cases, if a leak is detected or reported, it must be confirmed and then repaired. It is important that the engineers and field crews work to reduce the amount of time taken to repair the leak, reduce the impact to customers in terms of supply loss and make sure the repair has been completed satisfactorily.

The following text describes a simplified version of detecting and repairing a leak from the perspective of the water utility. It highlights decisions that must be made and shows how WaterWiSe helps in making these decisions.

1) *Leak detection*: IDEAS automatically detects transients at several sensor node data streams minutes after the data has been sampled. Several detections are grouped together and show a cluster of likely junctions that match the relative times

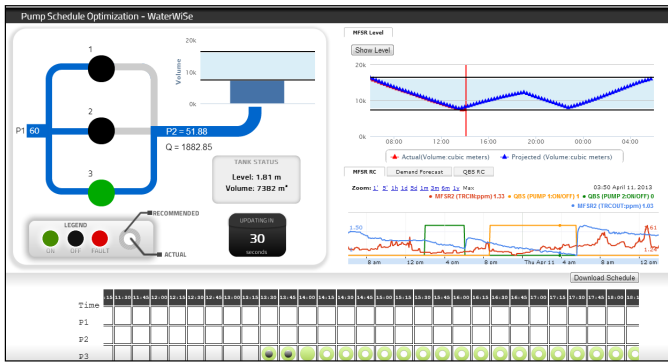


Figure 7. DSTM's pump optimization module interface showing the planned pump schedule (bottom) and real-time data feeds of pump, tank and in-network data.

of the detections. An alert is sent to all subscribed engineers about the possible burst and the estimated location.

2) *Leak validation*: The coordinating engineer sends a field crew to the estimated location to look for obvious signs of a leak. The field team remotely contact the DSTM network map on a tablet device to pinpoint the location of the junctions based on their names and locations, and find water above ground close to where IDEAS indicated. The leak-detection crew arrive to find the exact leak location.

3) *Solution testing*: The team of engineers in the supply network office determine the set of valves required to isolate the leaking pipe for repair. The engineer uses DSTM's operational simulator to find the relevant valves in the network model and run a 24 hour-ahead simulation to determine the expected impact on customers. Satisfied that the valve operations will not adversely affect customers, the engineer gives the go-ahead to the field crew to run the operation. The engineer shares the simulation results with the field crew via their tablet logged into DSTM, who begin looking for the correct valves within the area to close.

4) *Execute repairs*: Over the following hours, the ground around the leak is excavated and the pipe isolated. During this time, the engineers use the IDEAS interface to keep a close eye on the data streams of the sensor nodes that detected the burst and make sure the data match the expected behavior, paying close attention to any new alerts that arrive over email or SMS. The burst pipe is repaired and the pipe is decommissioned. The field crew verify the pressure and flow data have returned to the expected values before the leak and report back to the engineers in the supply network office.

5) *Validate repairs*: Over the following hours and days, the sensor node streams are closely monitored on IDEAS to make sure the repairs have fixed the leak as expected.

### B. Pump optimization

The cost of electricity used to transmit water across a distribution network is one of a water utility's largest expenses, and climate change regulations may cause an increase in electricity costs to water utilities of as much as 10% over market rates over a decade [4]. Reducing the amount spent

on energy, and reducing the energy consumed in transmitting water are therefore two important ways in which utilities can optimize system operation. When there are peak and off-peak tariffs for electricity usage, pump operation can be optimized by finding a schedule that keeps reservoirs and tanks at optimal levels to supply customers whilst pumping water at cheaper times of day [3]. Pumps can also be optimized to run at the most efficient part of their operating curves, such that energy consumed will be optimally used. When trying to optimize energy and electricity usage, it is important that operators maintain adequate supply to consumers (availability and quality), and validate that pump operations are not causing any problems during implementation.

The following text describes a simplified version of optimizing and implementing a pump schedule to save energy and/or electricity, highlighting the decisions the engineers must make and how WaterWiSe supports this.

1) *Optimize and validate pump schedule*: Each morning at a set time, DSTM's demand predictor forecasts consumption 24 hours ahead. The pump optimization algorithm takes this forecast along with the price tariff for energy use and constraints on reservoir water levels, system pressures and water quality and uses them to determine a set of rules for the pumping schedules over the next 24 hours to minimize energy use (see Figure 7); this is run in a simulation to verify the optimal pump schedule found does not violate any operational constraints.

2) *Implement pump schedule*: The utility manually or automatically implements the pump schedule according to the rules determined by the pump optimization algorithm. Engineers can verify through IDEAS' WSN and SCADA data that the pumps are operating according to schedule; this is presented in the pump optimization interface in Figure 7.

3) *Monitor impact of schedule on supply*: IDEAS monitors data streams in the areas supplied by pumping and makes sure that the pressure and water quality readings are within the constraints specified in the pump optimization algorithm. If the parameters go out of bounds, an alert is issued to subscribed users. Water age is monitored using DSTM's water age tool and compared to on-line water quality measurements to ensure the quality of the water being supplied is within age limits.

4) *Adjust behavior*: The coordinating engineer determines whether the IDEAS alert is sufficient cause to abort the current schedule and return to a default operation to control the unexpected behavior. Unexpected changes in water age and quality may also prompt the engineer to pro-actively modify the pumping schedule.

### C. Summary

The two scenarios presented in this section showed how WaterWiSe aids decision support in water system operation. Detecting and resolving a leak requires several different interactions between the engineers, field crews and IDEAS and DSTM. Organizing and executing a pump schedule requires close monitoring of the current system behavior in response to the new pump schedule, so that it may be compared to supply

expectations and aborted if the two deviate unacceptably. Both scenarios heavily rely on up-to-date information from the WSN to validate at each stage that the system is behaving how it should and to predict how it should behave in the future.

#### IV. RELATED WORK

WaterWiSe is most closely related to PipeNet [8], a low power wireless sensor network for monitoring large diameter bulk-water transmission pipelines. PipeNet was deployed for 22 months, with much analysis and experimentation performed off-line and in laboratory environments. The WSN used a relatively limited sensing platform (iMote2) to optimize for power consumption, and this restricted the sampling policy to 100Hz for five seconds at five minute intervals. WaterWiSe has seen long-term continuous deployment and provides an on-line system to users. WaterWiSe's WSN is more capable, and supports dynamic reconfiguration to trade-off to power consumption relative to application requirements.

The application of cloud-based IT solutions and generic data analytics to water distribution system monitoring has begun to gather favor, notably with IBM's Smarter Planet initiative [5]. This coincides with the emergence of the concept of a Smart Water Grid, where improved data collection coupled with more intelligent analytics will enable water utilities to better operate their networks to reduce costs [2]. The WaterWiSe platform is a clear example of end-to-end supply-side monitoring and decision support in the context of a Smart Water Grid.

#### V. CONCLUSION

This paper has shown how the WaterWiSe platform allows water utilities to better understand the behavior of their distribution network and thus improve its operation. To more efficiently operate and manage water distribution systems, it is important not just to intelligently increase the amount of measurements being made within the system, but also to use intelligent analytics to extract information from the data. WaterWiSe has been deployed on a live water distribution system in Singapore since 2009.

The scenarios presented in this paper show how a platform such as WaterWiSe is important for decision support management in an end-to-end context, from operation initiation to resolution, providing timely feedback through data.

The WaterWiSe platform is built around using a WSN to provide an on-line source of data. This is necessary because existing technology to perform in-network monitoring of hydraulic and water quality parameters is often application-specific; enabling in-network sensing that provides data for a number of applications is important for integrated operation. The data collected by the WSN also acts as an archive for tuning detectors, future data-mining and event response post-mortem.

In general, WSNs are seen as a good fit for in-situ, continuous monitoring applications due to being battery powered, having relatively minimal infrastructure requirements and being low-cost. In WaterWiSe, a WSN is particularly important in order to provide timely updates of multiple data

parameters. In addition, WSNs bring computational capability to the sensing device, allowing potential for in-network processing and dynamic re-configurability depending on current application requirements. At this stage, the WaterWiSe WSN is a starting point for improving network monitoring that provides a compelling trade-off between local processing capability, power consumption and increased sampling rates across a variety of important parameters for water distribution systems monitoring. Future work will address further the power consumption/processing capability trade-off in the WSN, with specific focus on how ad-hoc multi-hop networks can be implemented in urban environments without increased deployment density.

#### REFERENCES

- [1] M. Allen, A. Preis, M. Iqbal, S. Srirangarajan, H.B. Lim, L.D. Girod, and A.J. Whittle. Real-time in network distribution system monitoring to improve operational efficiency. *Journal AWWA*, 103(7), 2011.
- [2] Anonymous. Water 20/20: Bringing smart water networks into focus. Technical report, Sensus, 2012.
- [3] D. Beyer, J. Young, E. White, and S. Bunn. Implementing an energy management system for distribution pumping. In *AWWA Annual Conference and Exposition (ACE)*, 2005.
- [4] J. Gledhill, C. Seidel, and J. Oxenford. Building a climate-ready regulatory system. Technical Report 4129, Water Research Foundation, 2012.
- [5] M. Kehoe, M. Cosgrove, S. De Gennaro, C. Harrison, W. Harthoorn, J. Hogan, J. Meegan, P. Nesbitt, and C. Peters. Smarter cities series: A foundation for understanding ibm smarter cities. Technical report, IBM, 2011.
- [6] A. Preis, M. Allen, and A.J. Whittle. On-line hydraulic modeling of a water distribution system in singapore. In *Proc. 12th Water Distribution Systems Analysis Conference (WDSA)*, 2010.
- [7] S. Srirangarajan, M. Allen, A. Preis, M. Iqbal, H.B. Lim, and A.J. Whittle. Water main burst event detection and localization. In *Proc. 12th Water Distribution Systems Analysis Conference (WDSA)*, 2010.
- [8] I. Stoianov, L. Nachman, S. Madden, and T. Tokmouline. Pipenet: a wireless sensor network for pipeline monitoring. In *Proceedings of the 6th international conference on Information processing in sensor networks (IPSN '07)*, pages 264–273, 2007.