AI–BASED EVENT MANAGEMENT AT UNITED UTILITIES

BY MICHELE ROMANO, SHAUN BOATWRIGHT, STEVE MOUNCE, EIRINI NIKOLOUDI AND ZORAN KAPELAN

Nowadays, water companies in the UK and worldwide face a significant challenge as they have ageing assets, have to deal with budget and resource constraints and, yet, they need to meet increasing customer expectations. To effectively respond to network events (e.g. pipe bursts/leaks, equipment failure, etc.), water companies must proactively manage the full life-cycle of events in the right priority and in a speedy manner. This will drive a higher efficiency in water network operations and result in much higher customer satisfaction. As digital technologies are penetrating every aspect of our society, the water sector is starting to leverage them to enable the move from reactive to proactive event management. This article presents three examples of the work that United Utilities has carried out in collaboration with two leading UK Universities to improve event management practices by using Artificial Intelligence (AI), Machine Learning (ML) and other advanced analytics techniques. These examples demonstrate not only the power of these technologies, but also that water companies can benefit from their adoption as they enable them to efficiently take a holistic, fully managed life-cycle of events approach.

Within the UK and worldwide water industry, pipe bursts/leaks and other similar failure events are recognised as high priority issues. These events cause economic losses to the water companies, represent an environmental issue and have a negative impact on the water companies' operational performance, customer service and reputation. Water companies currently allocate a vast amount of resources to manage these events, but with limited success. The largest barriers to progress in the UK are the complexity of Water Distribution Systems (WDSs), ageing water supply infrastructure and unknown/unknowable condition of assets which make these events impossible to eliminate/avoid completely. In their day-to-day operations, water companies are tasked with operating their WDSs optimally to minimise costs and meet the required standards of service and, therefore, also managing contingency situations when events occur. In this scenario, an efficient event management process provides opportunities to improve the situation (e.g. by reducing the number/ duration of supply interruptions, conserving water and reducing the overall carbon footprint).

Event management in WDSs can be divided into three principal stages[1]: 1) event detection, 2) event location and 3) event response. The first two stages involve detecting and localising the event in the network and raising the relevant alarm. The third stage is associated with the decisions and actions required to reduce and, ultimately, eliminate the negative impact of the event on the water company and its customers.

In the last decade the importance of a proactive approach to event management, supported by near real-time assets monitoring, has become apparent as water companies in the UK have had to deal with tightening regulatory and budgetary constraints. Instrumentation, data gathering and communication technologies have also improved over the years and become less expensive to own and operate. As a result, a vast array of pressure and flow data originating from the many District Metered Areas (DMAs) that typically form a UK WDS is now frequently available and expected to quickly grow over time (especially data from pressure sensors, because of their lower cost and easier installation and maintenance when compared to flow sensors). The flow is nowadays typically measured at the DMA entry and exit points to allow the volume of water consumed in each DMA to be tracked over time and pressure is measured at a limited number of DMA critical monitoring points to ensure adequate pressure at the customers' taps.

The above monitored data can give insights into the operation and current/future status of WDSs (including pipe bursts/leaks and other similar events), especially when coupled with suitable data driven techniques. Advances in these techniques utilising advanced statistical tools, Machine Learning (ML) and Artificial Intelligence (AI) have led to the development of pioneering techniques that automatically manage and analyse increasing numbers of near real-time data streams aiming at enabling the detection[2-6], approximate location[7-9] and response[10,11] to pipe bursts/leaks and other similar network events. These techniques are very promising for alerting the water company personnel as soon as an event occurs, guide them to the problem area (i.e. for narrowing down the event search area within a DMA) and for supporting the control room operators in the identification of a suitable strategy to respond to those events in near real-time. This is mainly because they automate the mundane tasks involved in the data analysis process, provide more consistent analysis of the data and because they can efficiently deal with the vast amount of, and often imperfect, sensor data collected by modern supervisory control and data acquisition (SCADA) systems and extract information useful in making reliable operational decisions.

United Utilities has had a longstanding relationship with some of the, water systems engineering and hydroinformatics, leading UK Universities and in recent years has initiated a number of collaborative innovation projects with them. In some cases, these collaborations have taken advantage of programmes such as STREAM (the Industrial Doctoral Centre for the Water Sector - http://www.stream-idc.net) and WISE (Water Informatics: Science and Engineering Centre for Doctoral Training - http://wisecdt.org.uk) that are partially funded by the Engineering and Physical Sciences Research Council (EPSRC) and involve having a student based at United Utilities' headquarter pursuing an Engineering Doctorate (EngD) or Doctor of Philosophy (PhD) degree for industrially relevant research. These programmes are therefore also valuable as they enable the training of people capable of working at the interface of traditionally separate informatics, science and engineering disciplines and who understand both data science and the complexities of water challenges.

This article presents three complementary examples of the research work carried out in collaboration with the University of Exeter and the University of Sheffield to improve event management practices. Specifically, the first example focusses on event detection, the second example focusses on approximate event location and the third example focusses on post event response planning. These examples show how United Utilities is pursuing a fully managed life-cycle of events by taking a holistic approach to addressing the challenge of optimising the decision-making process of different teams in order to achieve the required level of service and the best utilisation of the assets at a minimum cost with an effective response time to all events. Indeed, a comprehensive, efficient and effective event management solution is key to such an optimisation challenge, which encompasses cross-organizational functions and works across different management levels.

Event detection

The first objective of a comprehensive event management solution is to provide near real-time, actionable event alerts such as, pipe bursts/leaks, pressure/flow anomalies, and sensor faults / telemetry problems. This enables water companies to become aware of all the events occurring in a timely fashion and better manage the situation, armed with valuable insights about these events (e.g. type, size, indication of their timing, etc.). This section briefly presents an AI-based system[4,5] that not only detects pipe bursts/leaks but also equipment and other failures in WDSs. This section additionally provides a couple of examples of the significant impact that this system has had on United Utilities' ability to deal with events in its WDS.

The detection system briefly presented here makes synergistic use of several self-learning AI techniques and statistical data analysis tools. In the detection system the automatic processing of pressure and flow data communicated by the DMA sensors in near real-time starts with using advanced techniques for ensuring that the data is cleansed and erroneous/missing data removed and/or infilled (e.g. wavelets are used for removing noise from the measured flow and especially pressure signals). The detection system then makes use of the pre-processed data to forecast the signal values in the near future using Artificial Neural Networks (ANNs). These values are then compared with incoming observations to collect different pieces of evidence about the failure event taking place. Statistical Process Control (SPC) techniques are also used for the analysis of the failure event -induced pressure/flow variations and gather additional pieces of evidence about the event occurring. The evidence collected this way is then processed using Bayesian Networks (BNs). BNs enable reasoning under uncertainty and simultaneously (synergistically) analysing multiple event occurrence evidence and multiple pressure/flow signals at the DMA level to estimate the likelihood of the event occurrence and raise corresponding detection alarms. The system also offers the capability to effectively learn from historical events to improve the detection of the future ones[5] (albeit it does not need information about historical events to start making reliable event detections when first applied to a DMA/WDS). It does not make use of a hydraulic or any other simulation model of the analysed WDS - i.e. it works solely by extracting useful information from sensor signals where bursts and other events leave their imprints (i.e. deviations from normal pressure and flows signals). This fact makes the detection system robust and scalable as it enables data to be processed in near real-time (i.e. within a 15 minute time window).

Elements of the aforementioned detection system, developed initially as part of a research at the University of Exeter, have been built into United Utilities' new Event Recognition in the Water Network (ERWAN) system. The ERWAN system's development carried out in United Utilities also benefitted by the following additional technology enhancements: a) development of a new methodology to add the capability to handle alarms from cascading DMAs[12], b) development of a new methodology to add the capability to rank alarms (based on a risk framework that accounts for factors such as mains length, material, number of industrial and key customers in a particular area of the water network), and c) development of a new methodology to add the capability to determine the likely root cause of an event. These enhancements have provided United Utilities additional, helpful event management tools. The ERWAN system has been used operationally companywide since 2015. It processes data from over 7,500 pressure and flow sensors every 15 minutes and detects events such as pipe bursts and related leaks in

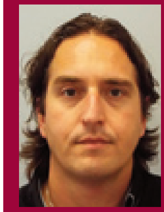

Dr Michele Romano is the Insight and Analytics Manager at United Utilities responsible for developing an advanced analytics strategy and capability (people, processes and technology). He obtained his PhD in Engineering from the Centre for Water Systems at the University of Exeter in the UK and holds a Honorary Research Fellow position at that University. His research interests and expertise are centred around the development and application of novel (geo)statistical and Artificial Intelligence based methodologies aimed at addressing a wide range of issues in urban water and wastewater systems. He has significant water industry experience and has authored a large number of academic publications.

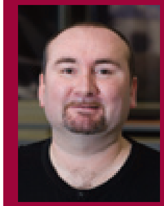

Dr Steve Mounce is a Visiting Research Fellow in Hydroinformatics at the University of Sheffield and the Director of Mounce HydroSmart Ltd. His PhD is in Computer Science and he has over twenty years of experience on research projects with over 100 academic publications. His research combines Artificial Intelligence and Water Engineering for such applications as leakage (including smart meter data mining), CSO analytics, water quality and burst event detection systems and fuzzy RTC.

Professor Zoran Kapelan is a Professor at the Delft University of Technology in the Netherlands where he is leading a research group on urban water infrastructure. He also holds a part-time professorial position at the University of Exeter in the UK. He is an IWA Fellow with 30 years of research and consulting experience in water engineering. His research interests cover a wide range of challenges related to water and wastewater infrastructure including development of various machine learning based technologies. Prof Kapelan pioneered the award winning burst/leak detection technology that is now used companywide in one of the largest UK water utilities resulting in large savings via reduced operational costs. He has published over 150 peer-reviewed journal papers.

Shaun Boatwright is an EngD candidate, based at the University of Sheffield and working with United Utilities, a large water company based in the UK, to develop a technique for localising leak/burst events in water distribution systems. He previously graduated with an MEng in Mechanical Engineering before working for a large engineering design consultancy on a number of potable water and wastewater projects for several UK water companies.

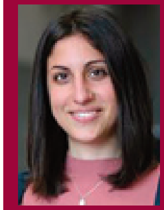

Eirini Nikoloudi is a WISE CDT student at the University of Exeter where she is pursuing a PhD degree. She obtained a MEng in Civil Engineering from National Technical University of Athens (NTUA) in 2015.

a timely and reliable manner - i.e. shortly after their occurrence and with high true and low false alarm rates.

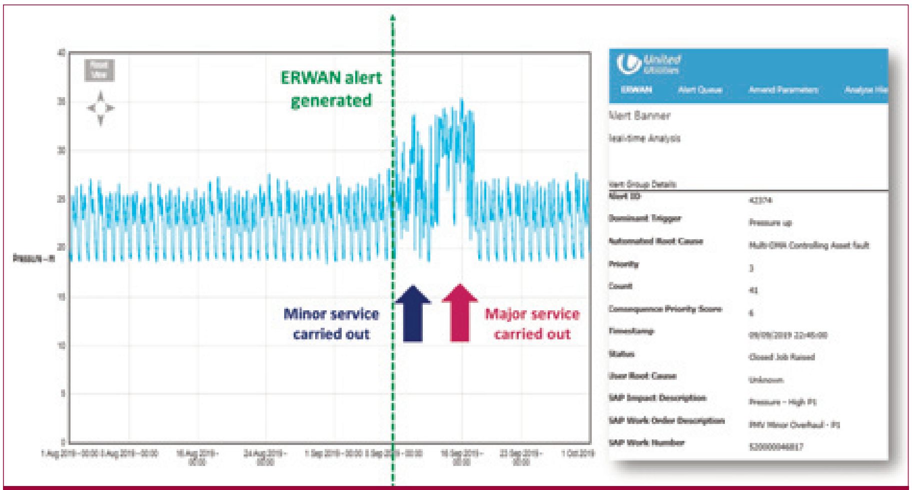

Alert Banner

Real-time Analysis

Alert Group Details

Alert ID	42374
Dominant Trigger	Pressure up
Automated Root Cause	Multi-DMA Controlling Asset fault
Priority	3
Count	41
Consequence Priority Score	6
Timestamp	09/09/2019 22:45:00
Status	Closed Job Raised
User Root Cause	Unknown
SAP Impact Description	Pressure – High P1
SAP Work Order Description	PRV Minor Overhaul - P1
SAP Work Number	520000046817

Figure 1. ERWAN system alert indicating a sudden pressure increase, likely due to a faulty pressure reducing valve.

Compared to previous company practice the ERWAN system has enabled United Utilities to detect pipe bursts and other failure events much more quickly. As an example, on the 31st of May 2016 there was a catastrophic failure of a 450 mm diameter main in the town of Formby which affected 10,600 properties. Using the ERWAN system the burst was identified more than three hours before the customers reported any disruption. This early event detection ensured planned responses were therefore deployed quicker. This also meant that customers were disrupted less as Alternative Supply Vehicles (ASVs – i.e. emergency tankers injecting water into the network) were deployed while the main was repaired. United Utilities estimates that this proactive response reduced interruptions to supply by 42%. Additionally, the ERWAN system has also demonstrated the potential to proactively prevent failures in some cases, e.g. via timely detection of faulty Pressure Reducing Valves (PRVs) often resulting in a follow-on pipe burst event(s). As an example of this situation, Figure 1 shows the ERWAN alert that was generated on the 9th of September 2019, indicating that the pressure had increased in a DMA. This alert prompted the Early Detection Team (EDT) to investigate the issue and immediately schedule a job for a minor PRV service as the automatically identified root cause suggested a fault of that asset. During that site visit, it was identified that the PRV had failed ‘open’. Further work was then scheduled for a network resource to carry out a major service on the asset. Proactively identifying that issue with the PRV may have prevented a pipe burst in the relevant pressure-controlled DMAs (especially considering the observed significant pressure increase). The potential impact of such a failure would have been in excess of £900k in Customer supply Minutes Lost (CML) penalty cost alone, with the ASV fleet and resource utilisation costs and the disruption to the customers adding to that.

The use of the ERWAN system has resulted in major operational cost savings (due to the reduced number of pipe bursts/leaks needed to detect and repair) to date and contributed to United Utilities’ CML, leakage and Customer Measure of Experience (C-Mex) performance (due to the avoidance or reduction in issues such as poor water pressure, no water, or poor water quality - therefore improving the service to over 7 million people and 200,000 business customers). It has also reduced asset maintenance costs by informing the need for maintenance prior to asset failure, and avoiding unneeded maintenance visits. Operational costs are also reduced as it enables problems to be dealt with proactively which is much less expensive than dealing with asset and service failures.

Furthermore, the success of the ERWAN system has been important to influencing change in the ways of working (e.g. making better use of data analytics in the daily operation) and the establishment of the EDT in United Utilities’ Integrated Control Centre (ICC). The ICC is the hub of United Utilities’ operations where a team of highly trained system operators watch over the network 24/7. They use the information and insight provided by ERWAN and other monitoring systems to perform complex event diagnosis and, by making intelligent decisions in the centre, prevent abortive work for field staff and resolve disruption for customers faster. Increasingly, through control and automation, the ICC can intervene remotely to resolve issues faster and more efficiently. This hub is one of the cornerstones of United Utilities’ AMP7 (Asset Management Plan five-year time period used in the English and Welsh water industry) Systems Thinking strategy and will catalyse future benefits.

Event location

After it is established that an event has occurred in a DMA by using automated systems such ERWAN, the next challenge in event management, especially when pipe burst/leak events are considered, is to determine the exact event location. Typically, network resources are deployed to DMAs containing new burst/leak events so that they can be precisely located (or “pinpointed”) and then repaired. There are many cases, such as when the size of a burst/leak event is small, where their location is not readily apparent. In these cases, resource intensive pinpointing activities such as acoustic surveys are carried out so that each of the pipes in a DMA can be examined to find the exact burst/leak location. It can take several days to examine all the pipes in a DMA as, in United Utilities for example, the typical total length of mains is about 13 km. This represents a significant investment of labour, equipment and operational expenditure when this approach is used across an entire WDS. In this scenario, a methodology that enables narrowing down the area that must be searched within a DMA (i.e. approximately locate the event) would be greatly beneficial for water companies.

This section briefly presents the details of a novel methodological framework[9] for the approximate burst/leak location that is being developed as part of a collaboration with the University of Sheffield and one example of its application to a burst event simulated by the controlled opening of a fire hydrant in a United Utilities’ DMA. This framework assumes that an increased number of pressure sensors can be deployed in the DMA being analysed. Due to the financial constraints placed on water companies and the costs of the additional instrumentation required, however, it is desirable to limit the number of additional instruments to be deployed. Therefore, the methodological framework being developed also encompasses a method for selecting the optimal number and location of sensors to be deployed in a particular DMA to achieve a desired level of event location performance. This tight coupling between optimal sensor placement and approximate burst/leak location is of particular importance as an optimal sensor placement strategy depends on the method that is used to locate the potential bursts/leaks and the efficiency of the burst/leak location depends on the sensor placement.

The novel methodological framework is based upon a Spatially Constrained version of the Inverse Distance Weighted (SC-IDW) geospatial interpolation technique[13]. Generally speaking, geostatistical techniques have the potential to limit the number of instruments which are deployed in a DMA as they can estimate the values of parameters at locations which are not measured based on the measurements from nearby sensors and, hence, to enable higher burst/leak location performance to be achieved for a given number of sensors[14]. Bearing this in mind, the use of SC-IDW enables the overcoming of the obvious limitation of traditional geostatistical techniques of using the Euclidean distance instead of the pipe length between the estimation locations and the instrument locations (i.e. not accounting for the actual network layout of a DMA). The framework makes also use of a hydraulic model and of the GALAXY multi-objective evolutionary algorithm[15] (i.e. a nature inspired AI methodology) to identify a Pareto front of optimal sensor configurations which simultaneously minimise the required number of pressure sensors (cost) and the average size of the areas to be searched (best level of burst/leak approximate location accuracy).

The first step for solving the optimal sensor placement problem involves hydraulic modelling of bursts/leaks at all nodes and building a sensitivity matrix. The valid range of burst/leak event sizes to be modelled is determined for each DMA by considering the accuracy of the pressure instruments being used (to find the smallest burst/leak event sizes) and a maximum allowable increase in flow (to determine the largest burst/leak event sizes for each burst/leak event location). The aforementioned sensitivity matrix is based on the changes in pressure for each potential sensor location, which are calculated by comparing the pressure in the hydraulic model with no burst/leak modelled with the pressure in the model with each burst/leak modelled. Additional computations are then conducted aimed at reducing the search space of the optimisation (i.e. grouping together events that cannot be distinguished given the pressure instruments' accuracy). Following this, the values of the pressure changes in the 'grouped' sensitivity matrix are used for building various interpolation surfaces during the optimisation step, which aims at maximising (using an objective function also based on the SC-IDW interpolation technique and a threshold that defines the burst/leak search area on an interpolation surface) the location performance of each configuration of sensors for every burst/leak being modelled. After determining the optimal sensors configuration by looking at the results of the optimisation step (and after deploying the pressure sensors in the field), the SC-IDW interpolation technique can be used operationally to calculate the approximate location of an actual burst/leak occurring in a DMA (once a burst/leak has been detected or is suspected) based on the actual changes (from 'normal') in pressures measured at the sensor locations. The calculated search area is then highlighted on a map of the DMA, which is passed to network resources to aid with pinpointing the burst/leak event.

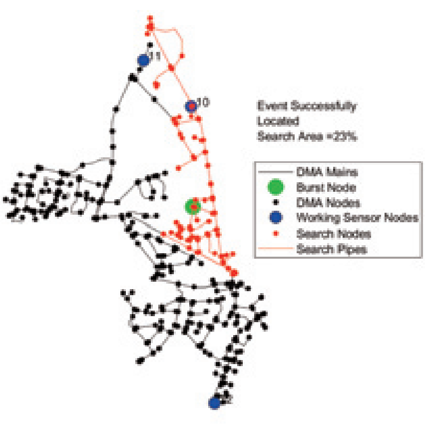

Figure 2. Example of a successful approximate burst location.

Figure 2 shows an example of such a map generated for the approximate location of a burst event simulated on the 14th of February 2020 by the controlled opening of a fire hydrant (so that the exact size and start time are known) in one of United Utilities' DMAs. This DMA contains approximately 2,100 properties and 25 km of mains. A PRV controls the pressure in one section of the DMA because of the highly variable elevation in the area. The fire hydrant opening was adjusted to achieve a flow rate of 0.6 l/s which is equivalent to approximately 6% of the average flow rate into the DMA calculated over a normal week. In Figure 2, the locations of the three optimally placed pressure sensors (determined by considering a total of 934 potential burst/leak event scenarios across 7 burst/leak event sizes) are shown as blue dots. The location of the opened fire hydrant is shown as a green dot. The pipes and nodes within the calculated search area are coloured in red. It can be noticed that this event was successfully approximately located within a search area that is less than a quarter of the total length of mains in the DMA. This example demonstrates the potential of the methodological framework being developed to allow successful approximate location of relatively small burst/leak events by using only a few additional optimally placed pressure sensors. This said, it is expected that the search areas can be further reduced by deploying more sensors. Nevertheless, by reducing the search area to a sub-region within a DMA, significant reductions in the time taken to pinpoint burst/leak events can be achieved (e.g. by ¾ as exemplified here).

Figure 3. New response methodology's flowchart.

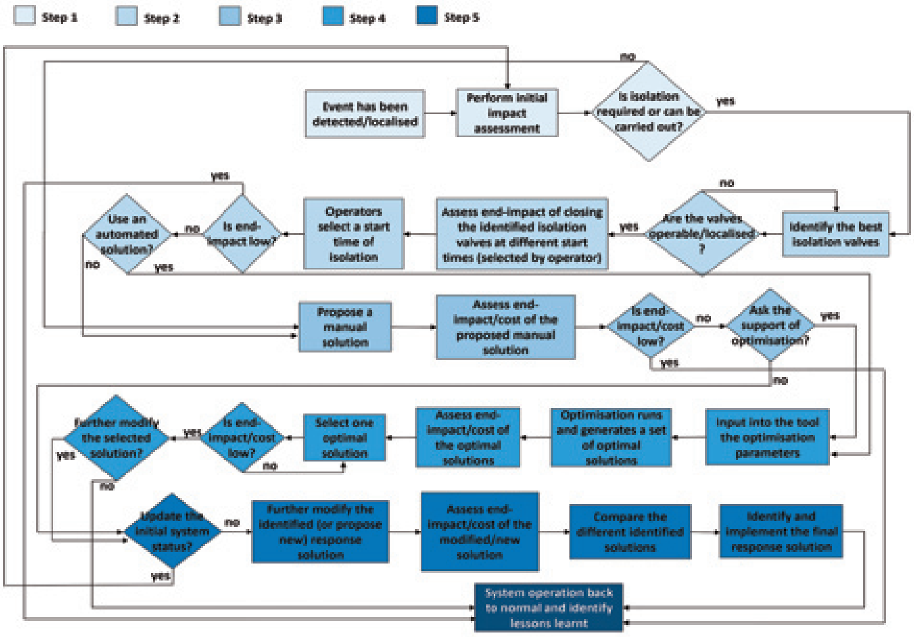

Post event response

After successful detection and location of events, the next considerable challenge for water companies during the event management process is the identification of a suitable strategy to respond to those events in near real-time. This section briefly presents the details of a novel methodology[11] for the response to water network events that is being developed as part of a collaboration with the University of Exeter and the initial, promising results obtained from its application on a semi-real case study.

The novel event response methodology presented here aims at improving United Utilities' current event response practice by supporting/guiding the ICC operators in the identification of low end-impact (i.e. the total impact after implementation of the response solution) and low cost response solutions. It consists of the following main steps: (1) robust initial event impact assessment (over a set horizon), (2) identification of a suitable isolation plan, (3) human-based, but computer-aided, identification of a response solution (i.e. manual solutions proposed by an operator), (4) automatic identification of a response solution generated using Genetic Algorithms (GAs) optimisation, and (5) selection of the response solution to be implemented in the field. Note that these five steps do not need to be necessarily carried out in a sequential manner. The following three-stage routine is performed in each of the aforementioned step: Stage 1) involves obtaining various operators' inputs (e.g. impact horizon, earliest time the repair can be initiated, etc.), Stage 2) involves carrying out hydraulic simulations to assess the end-impact and cost of each solution, and Stage 3) involves visualising the calculated end-impact and computing the cost of each solution. The new response methodology's steps are shown as a flowchart in Figure 3.

The above event response methodology is implemented in the Interactive Response Planning Tool (IRPT), which has been developed in Matlab. In the IRPT, the hydraulic simulations are carried out by using EPANET2[16] and pressure-driven network modelling based on the methodology developed by Paez et al.[17]. The Non-Dominated Sorting Genetic Algorithm II or NSGA II[18] (i.e. another AI tool/technique) is used to solve the multi-objective optimisation problem (albeit work has also been done to develop and use a new heuristic method that offers the advantage of significantly reducing the time taken to find near-optimal response solutions). The IRPT also links to the Quantum Geographic Information System (QGIS) software to visualise the spatial distribution of end-impact on a suitable map of the analysed WDS.

The IRPT facilitates an operator's decision-making by considering/providing: (i) structured yet flexible approach that supports and guides the operator throughout the entire response process, whilst allowing the operator to have a final say, (ii) novel interaction with the operator in near real-time via the simple IRPT graphical user interface (e.g. allowing operators to propose different 'what-if' scenarios without being hydraulic experts), (iii) provision of automatically generated advices (e.g. optimal response solutions and assessed end-impacts/costs), (iv) improved impact assessment using realistic (i.e. based on real-life metrics used by water utilities) impact indicators that cover different aspects of the event, which are consistently calculated for every considered response intervention, (v) capability to select multiple common operational intervention types such as rezoning and water injection (based on operational costs, availability of different types of interventions, etc.), and (vi) capability to easily compare different response solutions by visualising, inter alia, the impact coverage (using maps) and cost of different solutions. As a result, low end-impact and cost solutions can be effectively identified. This has multiple benefits for a water company. The most important benefit is reducing the impact on the customers, which can be costly in many ways (financially but also in terms of reputation, etc.). Other benefits related to costs include: a) operational savings in the long-term as many events may occur each year - although the cost of a single response solution may be small (e.g. hundreds of pounds), and b) less time spent on site for opening valves or injecting water - this could benefit water companies in terms of more efficient scheduling of the network resources' activities.

The IRPT is illustrated here on a semi-real case study to demonstrate the benefit of a response solution identified through interaction with the IRPT (hereafter referred to as the 'new methodology response') by comparing it to a response solution based on typical water companies' current practice (hereafter referred to as the 'current practice response'). Note that the case study under scrutiny is referred to here as "semi-real" because, despite being based on a real system and event, several simplifications were made with regard to the actual response actions taken by the ICC operator in real-life. This is primarily because the IRPT is still in development and did not yet offer the capability of exactly replicating those real-life response actions. Bearing this in mind, note that the used 'current practice response' label should also be construed accordingly. The considered event was a shutdown of a Water Treatment Work (WTW) (serving multiple DMAs and approximately 100,000 customers) due to a burst on a main within the WTW. The shutdown resulted in intermittent supply and low pressure to some customers. The WTW remained shut until the quality of the water leaving the WTW could be assured to meet the required standards. United Utilities mobilised ASVs to the area, which injected water at various points in the affected area and at different times during the incident.

Furthermore, United Utilities implemented a number of network changes (i.e. rezoning) in order to minimize customer end-impact. Bottled water was delivered directly to priority services and sensitive customers. The repair was completed 24 hours after the shutdown. Table 1 summarises the result obtained on this case study in terms of the total end-impact and the cost calculated by the IRPT for the 'new methodology response', 'current practice response' and 'no response' (i.e. initial condition of the system after the event) scenarios. For each of those scenarios, Table 1 also presents the calculated values of the various impact indicators (which make up the total end-impact), namely: a) CML, b) Average Minutes Low Pressure (AMLP), c)

continued on page 119

Table 1. Total end-impact, cost and values of the considered impact indicators for the 'no response', 'current practice response', and 'new methodology response'.

	CML (mins/cust)	*AMLP (mins/cust)*	*UW* (m^3)	*DRI (-)*	Cost (£)	Total end-impact (%)
No response	*4*	*3.6*	*3330*	*14*	**0**	**11.1**
Current practice response	*2.1*	*2*	*1825*	*273*	**894**	**6.5**
New methodology response	*1.6*	*2*	*1475*	*92*	**55**	**5**