Article:

Mounce, S.R. and Boxall, J.B. (2015) Automated Data-Driven Approaches to Evaluating and Interpreting Water Quality Time Series Data from Water Distribution Systems Read More: http://ascelibrary.org/doi/abs/10.1061/(ASCE)WR.1943-5452.0000533. Journal of Water Resources Planning and Management . ISSN 0733-9496

https://doi.org/10.1061/(ASCE)WR.1943-5452.0000533

Automated data driven approaches to evaluating and interpreting water quality time series data from water distribution systems

S. R. Mounce*, J.W. Gaffney**, S. Boult*** and J. B. Boxall****.

* Research Fellow, Pennine Water Group, Department of Civil and Structural Engineering, University of Sheffield, Sheffield S1 3JD, UK.

** Product Manager, Evoqua Water Technologies Ltd, 160 London Road, Sevenoaks, Kent TN13 1BT, UK.

*** Lecturer, School of Earth, Atmospheric and Environmental Sciences, University of Manchester, Williamson Building, Oxford Road, Manchester M13 9PL, UK.

**** Professor of Water Infrastructure Engineering, Pennine Water Group, Department of Civil and Structural Engineering, University of Sheffield, Sheffield S1 3JD, UK.

Corresponding Author: S.R.Mounce@sheffield.ac.uk

Abstract

Water distribution networks are not inert transport systems, the high quality water produced at water treatment works is subject to a variety of complex and interacting physical, chemical and biological interactions within these highly variable, high surface reactors. In particular the ageing and deteriorating asset condition in water distribution systems can result in a degradation of water quality delivered to the customer, often experienced as discolouration due to increasing amounts of fine particulate matter. Here, we propose that by assessing measured turbidity over time, in particular its correlation with local hydraulics, an assessment of change in risk of fouling can be obtained and asset deterioration inferred.

This paper presents a methodology for pairwise monitoring of a hydraulic parameter (flow or pressure) and turbidity using wavelet based semblance analysis – a novel methodology from another domain, which is applied for the first time to water quality data in distribution systems. It is suggested and subsequently explored through case studies that an increasing (anti-) correlation of the turbidity with the (pressure) flow diurnal cycle will be indicative of increasing fouling risk. This can be further supported through evaluation of the rate and magnitude of drift and through assessment of the change in magnitude of the daily turbidity profile. The composite of these approaches is applied to an extensive data set from a UK distribution system revealing the effectiveness of the analysis pre and post flushing (reducing discoloration events by between 64 to 89 percent). With increasing proliferation of monitoring devices and real-time data acquisition the potential for online systems and well informed proactive management is apparent.

Subject headings: Water distribution systems; Water management; Field tests; Data analysis; Quality

INTRODUCTION

Water Distribution Systems (WDS) function to supply treated water safe for human consumption and complying with increasingly stringent quality regulations. Considered primarily an aesthetic issue, discolouration is the largest cause of customer dissatisfaction associated with distribution system water quality. Of the 154,985 customer complaints about drinking water quality in 2007 for England and Wales, 124,671 (~80%) were about discoloured water (Husband and Boxall, 2011). By 2012, the number of customer complaints about drinking water quality for England and Wales had reduced to 106,612, of which 50,456 (~47%) were due to discolouration (DWI 2013). The cost of a water quality incident to a water company is high. When these result in exposure of customers to health risks the cost becomes very substantial both in financial terms and in terms of a significant loss of confidence by both consumers and regulators. In addition, costs rapidly escalate when standards are not maintained and there is a significant cost of failure and low quality. Proactive approaches are therefore required for water quality monitoring to identify, monitor and forecast water quality risks. Water quality sensors potentially facilitate data to inform system assessment, operational monitoring and overall management of the network, enabling forward-looking analysis for capital and operational maintenance planning. Therefore there is increasing interest in strategic deployment of online water quality monitors in the network, enabled by recent advances in sensor technology and remote data transfer options (UKWIR, 2013). While such work concludes that the cost benefit case for water quality monitoring is not clear, this is likely due to the “expert” nature of the review and the uncertainty of what value or information can be derived from water quality data. Despite this uncertainty temporal and spatial relationships can be inferred through improved network model and geographic information system (GIS) integration utilising this data and customer water quality complaints (for example Furnass et al. 2013). And some water companies have been making use of water quality information in real time during the operation of trunk main valves, as a means of monitoring and controlling particle mobilisation (Husband et al. 2011). Blokker et al. (2011) showed how existing networks can be made self-cleaning through adjusting valve positions to increase daily peak velocities. However at a more strategic level water companies need to understand and minimise the impacts of deterioration of the network, safeguarding water quality. The potential of online water quality monitoring in WDS reported here is seen as two-fold: first to identify contamination and second to detect deterioration in water quality due to the interactions between the water and the network: particularly to infer risk of fouling.

BACKGROUND

Discolouration

Discolouration is the result of fine particulates, of biological (from the biofilm) and/or inorganic (e.g. iron, manganese, chalk etc.) origin. Particulate matter in the distribution system may be present in the source water, be due to carry-over from the treatment works or, be produced by corrosion of iron pipes or by ingress and other physical, chemical and biological processes. The hydraulics of the system determines where this material is accumulated and where it is not. They similarly determine when accumulated particles are mobilised and may cause a discolouration problem. Figure 1 summarises these complex and interacting processes.

{Figure 1 approximately here}

Current modelling research suggests that discolouration can be conceptualised as: material accumulation over the entire circumference of distribution mains in cohesive layers, with the daily demand patterns and resultant shear stress governing how the material accumulates in these layers. Mobilisation of these layers occurs due to an increase in the shear stress above that found in the daily pattern (Boxall and Saul 2005, Vreeburg and Boxall 2007, Husband et al. 2008).

Biofilms contribute to fouling (defined as deposit formation on the internal surface of pipeline) of the WDS, including the accumulation of particles of iron and manganese – the main component of most discoloured water samples. Consequently, over time, certain pipes may develop fouling due to biofilms; this undesirable process - often affecting taste, odour and discolouration of water (Dotuerelo et al. 2013) - is known as biofouling. Biofilms are the predominant site of microbial activity in a WDS although growth in the bulk water can occur under certain conditions. The micro-organisms that grow in a biofilm do not represent a serious threat to health. All surfaces in contact with water can support a biofilm including pipes, walls and particulate matter.

Discolouration events often occur due to an increase in shear stress such as caused by a burst or valve/ hydrant operation that may be some distance upstream of where the customer complaint occurs. Historically, water companies in the UK have flushed water mains in response to customer complaints or following maintenance work. Predictive modelling tools such as Predicting Discolouration in Distribution Systems (PODDS), Boxall and Saul (2005), can be used to assess those pipes which have the greatest increase in shear stress and therefore discolouration potential due to increases in demand in other pipes. Such tools can be applied by water supply companies to improve the flushing process, assess required flushing velocities and focus on specific lengths of main to improve the effectiveness of the flushing operations. A new improved model has been proposed to track the relative amount of discolouration material that is bound to the pipe wall over time at each of a number of shear strengths (Furnass et al. 2014).

Since discolouration material can be conceptualised as accumulating in cohesive layers on the inside of pipe walls it can be assumed that anything inserted in the mains or exposed to flow for prolonged periods is likely to be similarly fouled over time. Therefore material is likely to accumulate on the optical surfaces of turbidity instruments which could adversely cause light scatter or absorption during the measurement process. A drift is often observed over time due to this process of material depositing (or fouling) the lens. The characteristic of this type of fouling manifests as a ‘drift’ where the turbidity recorded steadily increases over time. It has been observed in field deployment that the gain term is stable and the fouling of instruments only affect the instrument’s offset (Cook 2007). Similar drift was observed in field deployment of other instruments and consequently this drift can be adjusted for relatively easily (Gaffney and Boult 2011). However, the rate and magnitude of this drift will also be intrinsically linked with the amount of material present and the pipeline condition and so could act as an index of asset condition/rate of fouling.

Water distribution system time series

For hydraulic and water quality sensors used in water distribution networks, the sampling period can vary from 1 to 60 minutes (less than one minute is unusual). A 15-minute interval is widely adopted for strategic flow and pressure measurement. In the United Kingdom, and increasingly in other parts of the world, WDS are subdivided into district metered areas or distribution management areas (DMAs). Depending upon the sensor, and routinely for DMA flow meters, measurements may be averaged over a time period (for example 15 minutes) to produce the next measurement (i.e. they are not instantaneous). This averaging process can help to reduce the effect of spurious signals. Time series produced by sensors deployed in WDS are, in general, non-stationary and manifest significant noise (both observational and measurement) because of changing network characteristics (for example, a valve in the network which is closed may result in a new flow profile) as well as consumption patterns altering over longer periods. A global diurnal cycle generally manifested in hydraulic parameters is a reflection of the dominating residential consumption pattern. This is nearly always present in DMA level data although not necessarily in larger bulk water transfer data such as trunk main monitoring. A time series is sometimes described as having period s, in that it repeats after s time periods: $x(t+s) \cong x(t)$ for time t. In the case of a diurnal cycle and, say, a fifteen minute reading interval then s = 96. Secondly, a weekly cycle may also be present in WDS data with the main difference being between weekdays and the weekend and whose strength is most likely a function of the amount of residential and industrial demand present in the network area under consideration. Less is known about patterns or the periodicity of water quality data.

The coefficient of correlation is a measure of the degree of linear association between two variables, the degree can be either positive or negative ($-1 \leq \rho \leq 1$). Covariance is a similar measure, but in contrast it does depend on the variance of the two variables Correlation and covariance statistical measures can be used for example to monitor the movement of two financial investments or indices (Hamoa et al. 2000). Correlations can change over time and in different economic conditions in this domain. Detection of structural changes in the relationship between two time series has been an issue in econometrics for many years. Moving windows have proved useful in identifying where a break point may occur (Hansen 2001) but are limited by the window length chosen. Since the data is continuously arriving a time series is often termed time-evolving implying the presence of non-stationarity. A single, static correlation score for an entire time series is thus not particularly useful.

Over time, the material that accumulates on pipeline walls may be responsible for increasing incidence of discolouration (Cook and Boxall 2011). There may be an underlying correlation between flow (or inversely with pressure used as a surrogate) and turbidity which might be capable of being exploited for drawing conclusions about system fouling in water distribution systems. There is a certain level of correlation of the diurnal flow cycle with turbidity (Vreeburg 2007) and laboratory scale studies have supported this (Sharpe et al. 2010). Lehtola et al. (2004) showed diurnal variation in turbidity and other water quality parameters for a cast iron pipe with which there was a known discolouration risk thus suggesting that a mechanism occurs whereby cohesive material layers are continually eroding and regenerating at the pipe wall. A daily turbidity cycle seen in the data, with turbidity at its highest at peak daily flow in the morning, has the potential to be useful in indicating how ‘clean’ the network being monitored is. ‘Clean’ could refer to the amount and mobility of discolouration material at a given point within the network. Locations with a higher turbidity correlation with the flow diurnal change (or higher anti-correlation with pressure diurnal change) and a larger turbidity amplitude change per day could have a higher fouling risk than those locations without, since they are fully fouled. As well being observed in turbidity and hydraulic data, such daily patterns and correlations have been observed in WDS microbial data (Sekar et al. 2012). The rate of drift of turbidity measurements could also be an indicator for this. Based on this the hypothesis is that if the pipe condition is clean, there is little correlation between turbidity and hydraulics, whilst for fully ‘fouled’ pipes there would be a strong daily cycle in turbidity following the hydraulic diurnal profile.

The dependency of hydraulics with turbidity has non-linear behaviour. In addition, data sets will manifest non-stationarity, measurement noise and missing data. Various network events and works can impact on measurements so the temporal evolution of flow (or pressure) and turbidity depends on external variables in a complex fashion and their evolution is too complicated to solve by analytical models. Husband and Boxall (2011) note from field results that the relative significance of peak and variable daily hydraulics are not clear in relation to discolouration. The amplitude of the daily turbidity cycle has also been proposed as a simple indicator of the discoloration potential within a network (Cook 2007). Higher amplitudes suggest that there is more discoloration material within a network which is more easily mobilised and hence can help inform on the effectiveness of intervention.

This paper proposes automated data driven strategies for both detecting abnormal events and also for assessing the rate of pipeline fouling and associated discolouration incidents using metrics for evaluating longer term trends and patterns. An analysis technique first developed for the geophysics domain which allows for correlation changes to be viewed with varying window lengths through time in one graph is utilised. Correlations of pairs of time series are analysed with a particular focus on the daily and weekly temporal scale. Note that utilising hydraulic parameters in this way is not a means of predicting turbidity events per se, but that changing correlations between flow/pressure and turbidity can provide information on network operational conditions. Additionally, the rate and magnitude of turbidity drift and its daily amplitude magnitude and development over time are derived. The software analysis tool is developed and tuned by application to a real world WDS case study data set with some known interventions to demonstrate the structural changes in pipeline conditions. If hydraulic and turbidity correlations display periodicities and changes over time, discovering these may reveal valuable information that can be used for network condition assessment and decision making. Importantly, all these techniques and tools should be able to infer system state prior to the occurrence of customer impacting levels of discolouration.

EVENT DETECTION SYSTEMS, SEMBLANCE ANALYSIS AND METRIC METHODOLOGY

Event Detection Systems

The combination of telemetric data and an alert system is a step towards integrated network management allowing proactive response to events. A number of approaches from the fields of Artificial Intelligence and Statistics have been applied for detecting abnormality in water distribution systems from hydraulic time series data. Event Detection Systems (EDS) that convert flow and pressure sensor data into usable information in the form of timely alerts have been developed to help with the issue of leakage reduction (e.g. Mounce et al. 2010). A field of work has been emerging around automated and semi-automated data analysis. Interest is also growing in applying similar EDS to online water quality measurements. The detection of anomalous events (such as those outside normal turbidity range) is of interest for daily operational management (to maintain high water quality) as well as for identification of contamination events (either natural or intentional).

In Gaffney and Boult (2011) turbidity measurements at high temporal resolution (generally 15 minute) from several sites within a distribution network were assembled to form a comprehensive record of the variability of turbidity. The aim was to determine whether more intensive and extensive observation of turbidity is likely to be necessary to improve assessment of risk of fouling and, if so, how it can be used to do so. The data sets were analysed to investigate the correlation of turbidity events with pressure fluctuation. The data was manually mined to identify and extract all changes in turbidity that were greater than 0.5 NTU. The pressure data was then extracted at the same location in the hour prior to the onset of the turbidity change. Specifically, the following approach was used to determine discrete events and associated pressure change:

- A single event was defined as the period of time from when the event is triggered (i.e. greater than a 0.5 NTU increase) to when it returns to its pre-event turbidity value.
- In terms of the pressure data prior to the event, the 4 readings in the hour before the event were taken and the range calculated of those numbers.

This methodology was implemented in an automated manner in MATLAB (Version 7.14; The Mathworks Inc.) by directly processing the hydraulic CSV files for the four case study sites. A difference function was used to find the 0.5NTU change in turbidity (increase) rather than actual value and then the pressure change calculated. The output gives the event timestamp start and defined end as described previously.

CANARY is an open source software platform for EDSs developed by the United States Environmental Protection Agency (US-EPA) which can be used for the analysis of water quality time series data (Hart et al. 2009). CANARY can read in Supervisory Control and Data Acquisition (SCADA) data, perform an analysis in near real-time and then return the evaluated probability of a water quality event occurring at the current time step. CANARY uses statistical and mathematical algorithms to identify the onset of periods of anomalous water quality data, while at the same time, limiting the number of false alarms that occur. A two-step process is adopted: state estimation for future water quality value and a second stage of residual classification for determination of expected or anomalous value (an outlier). Algorithms calculate a background water quality profile for each water quality sensor (for one parameter or across the set of parameters), using some user defined period, and compare each new water quality measurement to that background profile to determine if the new measurement is an outlier or not. The definition of the water quality background is updated continuously as new data become available. Each monitoring station is analysed independently using CANARY. The values of the configuration parameters for each station might vary from one utility to the next and could vary across monitoring stations within a utility (EPA 2010). Mounce et al. (2012) used CANARY for a UK case study in which nine water quality sensors measuring six parameters were deployed in three connected District Meter Areas, fed from a single water source. A one year period of analysis was evaluated using comprehensive water utility records including repair and customer complaint data and which showed 86% of event detection clusters correlated to causes (such as burst repairs).

A linear filter algorithm, LPCF (McKenna et al. 2006), was utilised on two inputs (pressure and turbidity). Some of the parameters are key in the type of events that are detected by CANARY. These include window length (used for the prediction) and outlier threshold (measured in sigma). CANARY was applied to all six sites with the following parameters (after some sensitivity testing): a two day history window, a Binomial Event Discriminator (BED) window of 1 hour, outlier threshold of 2.5 and BED probability of 0.5. Event probability threshold was set at a default of 0.9 which is the value recommended by US-EPA. The automated approach based on differencing is looking for instantaneous changes above a certain threshold. When an event detection window is utilised a larger BED window results in the ability to identify abnormality over a longer time period and to limit the number of ghost alerts i.e. only identify events of interest. The size of the BED window was defined at one hour in order to detect short duration events. To some extent the EDS configuration will depend on the sampling rate and the type and length of events to be detected. Integrating results over greater numbers of time steps prior to increasing the probability of event detection generally results in fewer false positive detections, but at the expense of faster detection time.

Semblance analysis

For comparing the evolution of separate time series it is desirable to have a notion of correlation that evolves with time, is able to track changing relationships and which is not overly sensitive to transients depending on scale. Semblance filtering compares two data sets on the basis of their phase, as a function of frequency. It splits each input dataset into two output datasets consisting of the portion of the input datasets that is correlated to a given degree, and the portion that is not. The original formulation of semblance analysis uses the Fourier transform and analyzes the differences in the phase angles for each frequency (von Frese et al. 1997). This approach uses the entire data set to determine the relationship between the phases but does not allow for the frequency content of the time series data changing with time. If during a time series sequence there is a local oscillation representing a particular feature its position on the time axis will be lost, although it still contributes to the calculated Fourier transform. There is no way of knowing whether the value of the calculated Fourier Transform $F(w)$ at a particular window size w derives from frequencies present throughout the length of $f(t)$ or during just one or a few particular periods.

Wavelet analysis provides an alternative way of breaking down a signal into constituent parts. It does this by analysing at different scales (or resolutions) of data. Wavelet analysis has found a wide range of applications in speech, image processing and compression and signal processing generally. Masters (1995) defines a wavelet as "a function that defines a set of filters obtained by scaling this one function." This function is called an analysing or mother wavelet. Haar (1910) proposed the Haar basis function – in effect the first wavelet mother function. Daubechies (1988) constructed a set of wavelet orthonormal basis functions that have been widely adopted.

Temporal analysis is performed with a contracted, high frequency version of the mother wavelet, while frequency analysis is performed with a dilated, low frequency version of the same wavelet. These dilations and translations of the mother wavelet $\Phi(x)$ define an orthogonal basis termed the wavelet basis (for the discrete wavelet transform (DWT) which is the analogous technique to the discrete Fourier transform (DFT) for discrete time series data):

$$\Phi\left(_{s,l}\right)(x) = 2^{-\frac{s}{2}}\Phi\left(2^{-s}x - l\right) \qquad (1)$$

where s and l are integers (scale index and location index respectively) that scale and dilate the mother function to generate wavelets. s indicates the wavelet's width and is analogous to frequency and l gives its position. In order to span the data domain at different resolutions, the analysing wavelet is used in a scaling equation:

$$W(x) = \sum_{k=-1}^{N-2} (-1)^k c_{k+1} \Phi(2x + k) \qquad (2)$$

where $W(x)$ is the scaling function for the mother function, and the c_k are the wavelet coefficients which must satisfy the constraints:

$$\sum_{k=0}^{N-1} c_k = 2, \qquad \sum_{k=0}^{N-1} c_k c_l = 2\delta_{l,0} \qquad (3)$$

where δ is the delta function and l the location index. The coefficients $\{c_0, ..., c_n\}$ can be considered as a filter that is placed in a transformation matrix and applied to a raw data vector. The subsequent ordering of the coefficients results in two patterns: one that works as a smoothing filter and another that brings out the data's detail. The original time series signal can be represented in terms of a wavelet expansion by using linear combinations of coefficients.

The continuous wavelet transform (CWT) of a dataset $x(t)$ used to divide a continuous-time function into wavelets is given by (Mallat, 1998):

$$\mathrm{CWT}(u,s)=\int_{-\infty}^{\infty} x(t)\frac{1}{|s|^{1/2}}\Psi*\left(\frac{t-u}{s}\right)dt \tag{4}$$

Where s is scale, u is displacement, ψ is the mother wavelet used and * denotes the complex conjugate. Thus the CWT is a convolution of the data with the scaled version of the mother wavelet. Cooper and Cowan (2008) proposed a semblance analysis based on using the continuous wavelet transform (CWT) for initial transformation of the data. The wavelet based transform is superior to the Fourier transform as it does not assume that the frequency content is constant through time and allows for the analysis of changes in the frequency content within the data. They note that the time co-ordinate in Eq. (4) can be the spatial co-ordinate y and use the complex Morlet wavelet for this (Teolis, 1998):

$$\Psi(y)=\frac{1}{\pi f_b}e^{2\pi i f_c y}e^{-y^2/f_b} \tag{5}$$

where f_b controls the wavelet bandwidth and f_c is the wavelet center frequency. Cooper and Cowan (2008) set f_c = 1 so that scale becomes equivalent to wavelength. They discovered that the cross-wavelet transform for comparing two different time series (Torrence and Campo, 2008) did not give good results, and instead devised a wavelet equivalent of differencing the phase angles at each frequency for the Fourier transforms of two data sets. This resulting measure, the semblance S does not require the two series to be in the same units of measure and can take values from -1 (inversely or anti-correlated) through zero (uncorrelated) to +1 (perfect phase correlation).

Figure 2 provides an illustration of how semblance analysis can be presented for two turbidity instruments (sampling every 5 minutes consequently 288 readings per day) located in the same section of a DMA hence for which high correlation is expected (Khan et al. 2005). This reference data set manifests a diurnal cycle, occasional spikes and continual drift; which is more pronounced in instrument 1. Note that, calculated over the full data set of 35 days the Pearson coefficient was 0.22. The semblance graph denotes areas of correlation by red and inverse correlation in blue. This type of exploration (it is suggested that daily and weekly scales are most relevant for water distribution datasets) can indicate developing changes which occur in the relationship between the two series. The CWT graphs pick out three abnormal peaks in the turbidity. Over a scale of one day (288 on the wavelength axis), the data is highly correlated as evidenced in red on the semblance graph and in the final plot of the correlation coefficient at scale 288. The semblance graph indicates some areas on the edges where this relationship does not hold which occurs due to the boundary effect where the wavelet transform values close to the edge of the data are tainted by the discontinuous nature of the series edge (Addison, 2002). One mechanism for avoiding this effect is by truncation. Since both sensors are in quite close proximity within one DMA here, the index is highly correlated.

{Figure 2 approximately here}

Metrics

As well as seeking to examine the correlation with hydraulics, two other calculated metrics, and in particular their possible change over time, are expected to be indicative of fouling: drift and amplitude of turbidity.

Turbidity sensors are typically calibrated prior to deployment, and can be recalibrated following removal from the network – this allows calculation of the drift over time such that data can be subsequently corrected according to the observed fouling rate during the deployment period (Gaffney and Boult 2011). This drift effect and offset instability cause problems comparing information between different loggers, because over longer time periods, loggers tend to drift at differing rates and amounts and this can reduce confidence in absolute value. A linear trend might indicate a systematic increase in the observed values over time, such as sensor drift or material depositing (or fouling) on the lens.

In order to estimate a trend component without making parametric assumptions, a filter can be used to transform the time series. One popular example of a linear filter is the moving average which can be used to estimate a slow-moving trend. A symmetric (centered) moving average filter, adapted to deal with reduced elements at the ends and not alter the extremes when applied offline thus having zero phase, was applied with a weekly window. A best fit line was then calculated (using MATLAB detrend) per week and a gradient calculated for this line to give an NTU/day or NTU/week figure. Figure 3 shows the result of applying this approach for instruments 1 and 2 with NTU per week drift calculation – note that drift is cumulative per week and not a total figure. Hence, as supported by the raw data plot (see figure 2), we can see that instrument 2 is relatively stable compared to instrument 1, which shows increasing drift particularly in weeks 3 and 4.

{Figure 3 approximately here}

The final metric for assessing the evolving turbidity value is derived from the amplitude. A daily turbidity cycle may be more or less evident in the data depending on degrees of the corrosion and mobilisation processes present (Cook 2007). It has the potential to be useful in indicating how ‘clean’ (based on the amount and mobility of discolouration material) the network is. Networks with a higher change in turbidity per day/week could have a higher incidence of discolouration than networks with a lower change in turbidity per day/week. In order to cope with drift, and any spikes from discolouration events, an appropriate metric is to examine the (unbiased) standard deviation in turbidity recorded per day (or averaged per week as shown in figure 4).

{Figure 4 approximately here}

CASE STUDY

Description

Four Hydraclam instruments (Siemens) were deployed in a single DMA of a UK water company and data logged (at fifteen minute period) for a period of approximately four to six months. The total pipe length of this area of the network is approximately 26km and with a population of 6829 served in the DMA. The source water type is mostly ground source (greensand boreholes), with some surface water in a blend at times to supplement the supply. For this DMA, material types were composed as follows: Cast iron (CI) (55%), Ductile iron (DI) (20%), Asbestos cement (AC) (18%) and other materials (mainly plastics) being the remaining 7%. Table 1 provides the site details (site A is the inlet with flow and pressure data only) and figure 5 a schematic of the DMA. In Table 1, the Type of the site refers to a subjective classification based on the flow route - with main and dead-end being self-explanatory, and ‘loop’ referring to instrument location in an area likely to be effected by flow reversals or so call 'tidal points' due to multiple potential flow paths.

{Table 1 approximately here}

{Figure 5 approximately here}

Data sets and programme of activities

A time line indicating the sequence of activities in the DMA over the period of monitoring is shown in Figure 6 (letters on the x axis indicate sampling occurred weekly in the months indicated). During the period late November to early December operations were undertaken to systematically flush the DMA in a sequential order. This followed a large mains burst upstream of the DMA in an external trunk main in early November. Note the distribution of discolouration customer contacts clustered around the burst event prior to the flushing operations.

{Figure 6 approximately here}

RESULTS

Event detection using automated thresholding analysis and CANARY

Automated thresholding analysis resulted in 166 turbidity change events in total. When examining the relationship between turbidity increases and pressure changes (as a surrogate for hydraulic disturbance), with all events from the four sites amalgamated, as in Gaffney and Boult (2011) this relationship was found to be weak.

An Event Detection System such as CANARY can be used to perform similar analysis for detecting abnormal turbidity and pressure events. It was applied here in an offline manner (historic analysis, but simulated online by processing in time series order one step at a time). Figure 7 shows an example of automated detection of a large event resulting in detection across both pressure and turbidity (for site B). Figure 7b shows another event cluster with a turbidity only change. In figure 7b, the dashed box indicates the only turbidity change detection during the period. Two later CANARY detections in this week are not picked up by the threshold based method due to the turbidity differencing not quite exceeding 0.5. This illustrates that EDS such as CANARY can pick up more subtle events than threshold based methods. Thus an online CANARY system has the potential to allow real-time detection of abnormal turbidity and hydraulic conditions. It should be noted that pressure is a more local measurement, so a pressure and turbidity correlation is probably due to local change. But turbidity changes could originate from further afield, even outside the DMA, in which case we would expect no pressure change. Conversely we could have a pressure event, but the local pipe has experienced this recently so no new corresponding turbidity event is experience. I.e. there are three possible combinations telling you different things about the event and the network state.

{Figure 7a and 7b approximately here}

As will be described in the next section, it was of particular interest to assess how the frequency of turbidity events occurred pre and post the flushing operations conducted in the DMA (see figure 6 for timeline). This information is summarised in figure 8 for both automated thresholding analysis and CANARY detections (averaged to number of events per week to take account of variable logging periods). It is apparent that there was a significant improvement in the frequency of turbidity events across all sites after the flushing i.e. the flushing was successful in improving water quality in the short term. This was observed across both the automated thresholding and CANARY detections. Thus the benefit of flushing (or other interventions) can be quantified.

{Figure 8 approximately here}

Strategic level analysis

Earlier, three possible metrics were presented for evaluating the developing status of a network condition as regards fouling using turbidity (and a flow / pressure if present): hydraulic correlation (using semblance analysis), drift and standard deviation of the turbidity calculated per day. The case study data is now used to illustrate how these metrics could be utilised for assessing pipeline condition potentially providing strategic level evaluation for a WDS. Calculations of the drift and standard deviation metrics have already been outlined. The semblance correlation coefficient at the daily scale can be averaged in order to provide a weekly index summarising how correlated (or anti-correlated) the turbidity parameter is with another parameter such as a hydraulic variable. Truncation of some data points can be applied to avoid boundary effects. From figure 2 (bottom plot) it is apparent that two highly correlated sensors would have a near positive unity weekly score.

Automated MATLAB code was developed for Semblance, drift and fouling analysis which incorporates checking routines for any missing date time stamps, filling minor amounts of missing data (using seasonal averaged filling) and calculating the metrics. Full results for the case study data sets are provided in figure 9 (data availability periods are as given in figure 6). Note that the y axis range are not identical between graphs, however it is the evolution of the individual metrics at a particular site which are most significant.

{Figure 9 approximately here}

DISCUSSION

Figure 8 is showing the efficacy of the flushing (DMA wide) in reducing the number of turbidity events at all sites before/after flushing (discoloration events were reduced by between 64 to 89 percent.

Figure 9 is showing the three metrics over the same total period. In Figure 9 it can be observed that the calculated metrics provide longer term information (than EDS alerts) about how turbidity measurements are varying at sensor locations within the DMA. The first (leftmost) plot show how the baseline turbidity is drifting over time and the next indicates daily variation about the mean turbidity. The two plots on the right show how the correlation, at a diurnal scale, varies over time between the turbidity and i) pressure measured at the same point (purely local conditions) ii) flow at the DMA inlet (local turbidity correlated with DMA level). These latter two are inversely related for obvious reasons. It has been postulated that parts of a network with higher turbidity correlation with the flow diurnal change (or higher anti-correlation with pressure diurnal change) could have a higher incidence of discolouration than those parts of the network without. In Figure 9 the start of the two week flushing period is indicated. Consequently we can assess how the flushing affected the semblance correlation per site. We can conclude a benefit was seen at all sites, although the magnitude of this benefit varied.

For sites B and D (located on central mains) we can see a positive effect on the semblance correlation i.e. after the flushing we see a reduced positive correlation between turbidity and inlet flow (mirrored in the local pressure semblance). This indicates that these sites benefitted strongly from flushing. In contrast, two sites had little beneficial change on their positive (negative) turbidity to flow (pressure) correlations after flushing (C and E). In Table 1 we see that site C has low/possibly reverse flow and is located in a looped area of the network. Site E in the northern dead end of the network is likely to suffer from stagnation upon inspection of their minimum and maximum flows from Table 1. These results support the concept that turbidity is very pipe segment specific and is closely related to the type of pipe material, condition of the pipe and local site characteristics. These factors result in significant variation in the behaviour of discoloration events within the same network. The techniques presented here if implemented as automated systems for condition monitoring have the potential to allow inference of asset degradation over time.

CONCLUSIONS

This paper has explored automated techniques for evaluating and interpreting water quality time series data from water distribution systems. Through advanced automated analysis it has been shown that high resolution turbidity data can provide a data source for proactively assessing fouling risk, incidence of discolouration and potentially asset deterioration: all prior to customer impacting discolouration. A novel semblance methodology, originating in another field, was applied for the first time to WDS water quality data. Its use for assessing measured turbidity over time along with its correlation with local hydraulics, in particular the diurnal profile, was presented. Other key techniques found to be valuable include assessment of drift and evaluation of change in the diurnal pattern over time. The software tools were applied to a DMA data set for four instruments from a UK distribution system revealing the effectiveness of the analysis pre and post flushing. Findings included:

- EDS can be used advantageously to combine hydraulic and water quality data and produce improved detection results (for example with systems such as CANARY).
- A significant improvement in the frequency of turbidity events across all sites after the flushing was observed i.e. the flushing was successful in improving water quality in the short term, thus the benefit of flushing (or other interventions) can be quantified. For the case study, the flushing programme reduced discoloration events by between 64 to 89 percent.
- Semblance analysis indicates the changes in the correlation coefficient as a function of time and window length allowing further analysis to identify the timing and the duration of shifts in pipeline condition. One does not need to know a priori the fluctuation at different scales or the exact magnitude of those changes. Changes in the local correlation were observed strongly at some of the case study sites after flushing and can provide information on network operational conditions.
- The results support the concept that turbidity is very pipe segment specific and is closely related to pipe material, condition and local site characteristics. These factors result in significant variation in the behaviour of discoloration events within the same network.

Benefits for these approaches include:

- Near real-time network water quality information linked to online sensors and data acquisition systems
- Semblance analysis enables sophisticated analysis of water quality data paired with a hydraulic or operational signal within water networks. The methodology allows analysis which is sensitive

enough to infer system state prior to the occurrence of customer impacting levels of discolouration.

- The ensemble of metrics presented here allows an assessment of flushing efficacy at pipe level
- Potential for a strategic online operational tool - particularly incorporating the daily scale for metric based analysis. Analysing how these metrics change over a weekly/monthly period could allow the monitoring of asset deterioration/degradation and proactive evaluation and selection of intervention options.

ACKNOWLEDGEMENTS

This work was supported by the Pipe Dreams project (EP/G029946/1) funded by the U.K. Science and Engineering Research Council. The authors would like to thank South East Water - particularly Neil Hudson and Benjamin Smith - and Evoqua Water Technologies Ltd for field access and data provision.

REFERENCES

Addison, P.S. (2002). “The Illustrated Wavelet Transform Handbook, Introductory Theory and Applications in Science, Engineering, Medicine and Finance”, Institute of Physics Publishing, London.

Blokker, M. J., Schaap, P. G. and Vreeburg, J. H. G. (2011). “Comparing the fouling rate of a drinking water distribution system in two different configurations”. In Urban Water Management: Challenges and Opportunities, Exeter, UK, 2011.

Boxall, J.B., and Saul, A.J. (2005). “Modelling discolouration in potable water distribution systems”. Journal Environmental Engineering ASCE, 131(5), 716-725.

Cook, D. (2007). “Field investigation of discolouration material accumulation rates in live drinking water distribution systems”. PhD Thesis, University of Sheffield.

Cook, D. and Boxall, J. (2011). “Discoloration Material Accumulation in Water Distribution Systems”. J. Pipeline Syst. Eng. Pract., 2(4), 113–122.

Cooper, G. R. J., Cowan, D. R. (2008). “Comparing time series using wavelet-based semblance analysis”. Computer and Geosciences, 34, 95-102.

Daubechies, I. (1988). “Orthonormal bases of compactly supported wavelets”. Comm. Pure Appl. Math., 41, 906-966.

Douterelo, I., Sharpe, R. L. and Boxall, J. B. (2013). “Influence of hydraulic regimes on bacterial community structure and composition in an experimental drinking water distribution system”. Water Research, 47, 503-516.

DWI. Annual report 2012: Letter to minister - England. Technical report, Drinking Water Inspectorate, 2013b. URL http://dwi.defra.gov.uk/about/annual-report/2012/letter-english.pdf.

EPA (2010). “Water Quality Event Detection Systems for Drinking Water Contamination Warning Systems: Development, Testing and Application of CANARY.” EPA/600/R-10/036, Washington, DC. http://www.epa.gov/ord.

Furnass, W. R., Mounce, S. R. and Boxall, J. B (2013). “Linking distribution system water quality issues to possible causes via hydraulic pathways”. Journal of Environmental Modelling and Software, 40, 78-87.

Furnass, W. R., Collins, R. P., Husband, P. S., Mounce, S. R. and Boxall, J. B. (2014). “Modelling both the continual erosion and regeneration of discolouration material in drinking water distribution systems”. IWA Water Science and Technology: Water Supply, 14 (1), 81-90.

Gaffney, J. W., Boult, S. (2011). “The Need for and Use of High Resolution Turbidity Monitoring in Managing Discolouration in Distribution”. Journal of Environmental Engineering, 1, 416–416.

Haar, A. (1910). “Zur Theorie der orthogonalen Funktionensysteme”. Math. Ann., 69, 331-371.

Hamao, Y., Masulis, R. W. and Ng, V. (1990). “Correlations in Price Changes and Volatility across International Stock Markets”, The Review of Financial Studies, 3 (2).

Hansen, Bruce, E., (2001). “The New Economics of Structural Change: Dating Breaks in U.S. Labor Productivity”, The Journal of Economic Perspectives, 15(4), 117-128.

Hart, D. B., and McKenna, S. A. (2009). CANARY user's manual, version 4.1, EPA/600/R-08/040A, U.S. Environmental Protection Agency, Office of Research and Development, National Homeland Security Research Center, Cincinnati, OH.

Husband P. S., Boxall J. B. and Saul A.J. (2008) "Laboratory studies investigating the processes leading to discolouration in water distribution networks". Water Research, 42(16), 4309-4318.

Husband, P. S. and Boxall, J. B. (2011). "Asset deterioration and discolouration in water distribution systems". Water Research, 45, 113-124.

Husband, P. S., Jackson, M. and Boxall, J. (2011). "Trunk Main Discolouration Trials and Strategic Planning". Proceedings of Computing and Control for the Water Industry (CCWI), Exeter, 5/7th September.

Khan, A., Widdop P. D., Day, A. J., Wood, A. S., Mounce, S. R., Machell. J. (2005). "Performance assessment of leak detection failure sensors used in a water distribution system". Journal of Water Supply: Research and Technology – AQUA, 54 (1), 25-36.

Lehtola, M. J., Nissinen, T. K., Miettinen, I. T., Martikainen, P. J. and Vartiainen, T. (2004). "Removal of soft deposits from the distribution system improves the drinking water quality". Water Research, 38(3), 601–610.

Mallat, S. (1998). "A Wavelet Tour of Signal Processing", Academic Press, New York.

Masters, T. (1995). “Neural, novel and hybrid algorithms for time series prediction”. John Wiley & Sons, New York.

McKenna, S. A., Klise, K. A., and Wilson, M. P. (2006). “Testing water quality change detection algorithms.” Proc., 8th Annual Water Distribution Systems Analysis Symposium, ASCE, Reston, VA.

Mounce, S. R., Boxall, J.B. and Machell, J. (2010). “Development and Verification of an Online Artificial Intelligence System for Burst Detection in Water Distribution Systems”. ASCE Water Resources Planning and Management, 136(3), 309-318.

Mounce, S. R., Machell, J. and Boxall, J. B. (2012). “Water quality event detection and customer complaint clustering analysis in distribution systems”. IWA Journal of Water Science and Technology: Water Supply, 12 (5), 580-587.

Sekar R., Deines P., Machell J., Boxall J. B., Biggs C. A. and Osborn A. M. (2012). “Bacterial water quality and network hydraulic characteristics: A field study of a small, looped water distribution system using culture-independent molecular methods”. Journal of Applied Microbiology, 112(6), 1220-1234.

Sharpe, R. L., Smith, C. J., Biggs, C. A. and Boxall, J. B., (2010). “Pilot scale laboratory investigations into the impact of steady state conditioning hydraulics on potable water discolouration”. Proceedings of 11th Annual International Symposium on Water Distribution Systems Analysis, Arizona, USA, September.

Teolis, A., 1998. “Computational Signal Processing with Wavelets”. Birkhauser Boston Inc., Boston, MA, 352pp.

Torrence, C. and Compo, G.P. (1998). “A practical guide to wavelet analysis”. Bulletin of the American Meteorological Society, 79, 61–78.

von Frese, R.R.B., Jones, M.B., Kim, J.W. and Kim, J.H. (1997). “Analysis of anomaly correlations”. Geophysics, 62 (1), 342–351.

UKWIR (2013). “Cost Benefit Analysis of Ubiquitous Data Collection in Water Distribution - CBA Scenarios”. 13/DW/12/2 - ISBN: 1 84057 692 8.

Vreeburg, J. (2007). “Discolouration in drinking water systems: a particular approach”. Technical University Delft. PhD: 183.

Vreeburg, J. and Boxall, J.B. (2007). “Discoloration in potable water distribution systems: a review”. Water Research, 41(3), 519-529.

Table captions:

Table 1: Site characteristics for case study

Tables

Table 1: Site characteristics for case study

Site	Distance from Inlet (km)	Min Flow (l/s)	Max Flow (l/s)	Pipe diameter (mm)	Pipe material at site	Pipe materials through system	Type
B	1.18	0.28	1.71	150	CI	Predominantly CI	Main
C	1.90	-0.03	0.16	150	CI	Predominantly CI, some DI	Loop
D	1.96	0.27	1.48	150	CI	Predominantly CI, some DI	Main
E	2.74	0.08	0.4	100	AC	Predominantly CI, some AC	Dead end