Soft-sensing with Qualitative Trend Analysis for Wastewater Treatment Plant Control

Christian M. Thürlimann[a,b], David J. Dürrenmatt[c], Kris Villez[a,*]

[a] *Swiss Federal Institute of Aquatic Science and Technology (Eawag), 8600 Dübendorf, Switzerland*
[b] *Institute of Environmental Engineering, ETH Zürich, 8093 Zürich, Switzerland*
[c] *Rittmeyer Ltd., 6341 Baar, Switzerland*

Abstract

Ammonia control in municipal wastewater treatment plants typically requires maintenance-intensive instrumentation. A low maintenance alternative is sought for small- to medium-scale applications. To this end, a pH-based soft-sensor is proposed to detect ammonia peak load events. This soft-sensor is based on a newly developed technique for qualitative trend analysis and is combined with a rule-based controller. The use of qualitative trend analysis makes this soft-sensor tolerant towards sensor drifts and thereby reduces the maintenance effort. The method allows controlling any process in which relative changes in the measured output are informative about the system output.

Keywords: activated sludge, aeration, ammonia control, hidden Markov model, qualitative trend analysis, process monitoring

1. Introduction

Ammonia removal from municipal wastewater is conventionally achieved with biological nitrification. This process requires aerobic conditions for the bacteria to oxidise the ammonia to nitrite and further to nitrate. Fine bubble aeration is an efficient way to supply the required oxygen into an activated sludge reactor (Rosso et al., 2008). Ingildsen (2002) reported that on average aeration accounts for up to 60% of the total energy consumption of a wastewater treatment plant (WWTP). Since municipal WWTPs exhibit nutrient and flow dynamics, the control of the plants aeration intensity is essential for operational efficiency. Initial attempts to make the oxygen supply more efficient by automation were pursued in the 1970s (Olsson, 2012). These first automation attempts were aimed at controlling the aeration rate to maintain the dissolved oxygen (DO) concentration in the aeration basins close to a fixed oxygen concentration setpoint. This prevents an under- and, especially, over-supply of oxygen to the bacterial community (e.g. Andersson, 1980). Today this is a standard control strategy for WWTPs. Further aeration efficiency gains can be realised by controlling the ammonia concentration (e.g. Åmand et al., 2013). Practically this is achieved by means of

* Corresponding author, address: Überlandstrasse 133, CH-8600 Dübendorf, Switzerland, email: kris.villez@eawag.ch, fax: +41 58 765 5802

Abbreviations: DO: dissolved oxygen, Eco: Eco controller state, EtN: Eco to Normal rule, HMM: hidden Markov model, Normal: Normal controller state, NtE: Normal to Eco rule, p.e.: population equivalent, $pHd(i)$: Current pH difference value, pHd_{LM}: pH difference value of lowest minimum, ΔpHd_{MAD}: Value of maximal difference parameter, ΔpHd_{MID}: Value of minimal difference parameter, ISE: Ion-selective electrode, QPE: Qualitative Path Estimation, QSE: Qualitative State Estimation, QTA: Qualitative trend analysis, SBR: Sequencing batch reactor, SCADA: Supervisory control and data acquisition, SCS: Shape constrained splines, T_{MIN}: Time value of minimal Normal time parameter, T_{Normal}: Normal controller state timer value, WWTP: wastewater treatment plant, a: Degree of derivative, f: Polynomial function, i: Data point index, q: Primitive, p: Discrete (predecessor) state, p_m: Maximum likelihood predecessor state, s: Discrete state, t: Discrete (target) state, w: Data points index, y: Measurements, κ: Primitive, Λ: Probability/likelihood, λ: Transition likelihood, σ: Variance, $\mathbf{T}(t,p)$: Transition likelihood matrix, τ: Kernel half-width,

cascade control: The master control sets the DO set points in the slave aeration controller according to the ammonia concentration. The slave controller, in turn, adjusts the aeration intensity to reach the given DO setpoint. Reducing the DO concentration in low load situations results in two main aeration efficiency gains: *(i)* it decreases the volume-specific nitrification capacity, which in turn shifts part of the load downstream, thereby involving otherwise unused downstream aerobic reactor volumes in the nitrification process and *(ii)* the larger oxygen concentration gradient between the gas and liquid phase improves oxygen transfer efficiency. In addition, ammonia based aeration control can enhance the nitrogen and phosphorus removal efficiency and can decrease the need for external carbon dosage (Rieger et al., 2014). Today such advanced aeration control is available for many large WWTPs.

Commercially available ammonia control systems rely on in-situ ion selective electrodes or ex-situ analysers for the measurement of the ammonia concentration. However, these instruments are costly to operate as they require frequent maintenance by skilled staff (Kaelin et al., 2008). This is why installation and maintenance costs may not be covered by the energy cost savings on small or remotely controlled WWTPs. To overcome this obstacle soft-sensors offer an interesting alternative to using high-end yet error-prone instruments. Soft-sensors are based on the identification of stable relationships between the variable of interest and the signals of one or multiple cheaper and/or more robust sensors (Haimi et al., 2013). If these relationships are based on empirical black-box modelling instead of first principles, such soft-sensors are often challenged in practice due to their lack of transparency, lack of general validity and/or their (re-)calibration and fine-tuning needs (Dürrenmatt and Gujer, 2012). The use of mechanistic concepts for soft-sensor development - if available - is therefore considered the preferred approach. For biological wastewater treatment in sequencing batch reactors (SBR), soft-sensors based on a mechanistic understanding such as conductivity for ammonia (e.g. Joss et al., 2009) or pH/ORP for control of biological denitrification (e.g. Peng et al., 2002) are described in the literature. Zanetti et al. (2012) derived a comprehensive list of soft-sensors used for the control of wastewater treatment.

For continuously operated WWTPs, Ruano et al. (2009) introduced a mechanistic ammonia soft sensor. They demonstrated that the pH difference over the aerobic zone contains information about the ammonia concentration present in the tanks. Indeed, the pH signal reflects the balance between nitrification and CO_2-stripping. An increasing ammonia load boosts the nitrification and the proton production. Unless an increased aeration rate equalises the proton production by CO_2-stripping, the reactor will experience a net proton production and its pH will drop. The opposite happens if the ammonia load drops: the lower proton production of the reduced nitrification does not outweigh the CO_2-stripping, and the pH rises unless the aeration intensity is decreased.

Ruano et al. (2012) kept the pH difference over the first and last tank of an aerobic tank series at a fixed setpoint by manipulating the airflow in these activated sludge tanks. By maintaining this pH difference close to a given setpoint value, they were able to adequately react to changes in the ammonia load. A given pH difference setpoint corresponds to a certain ammonia effluent concentration in a steady state system (Ruano et al., 2012).

As pH sensors are prone to drift, the controlled pH difference value will tend to deviate from the intended pH difference value. Consequently, the effective ammonia concentration at the WWTP outlet is expected to deviate from the intended setpoint as time progresses. Despite the difference in their time constant, actual process dynamics and sensor drift are hard to distinguish. Thus, a frequent and strict sensor management remains necessary in spite of using state-of-the-art pH sensors.

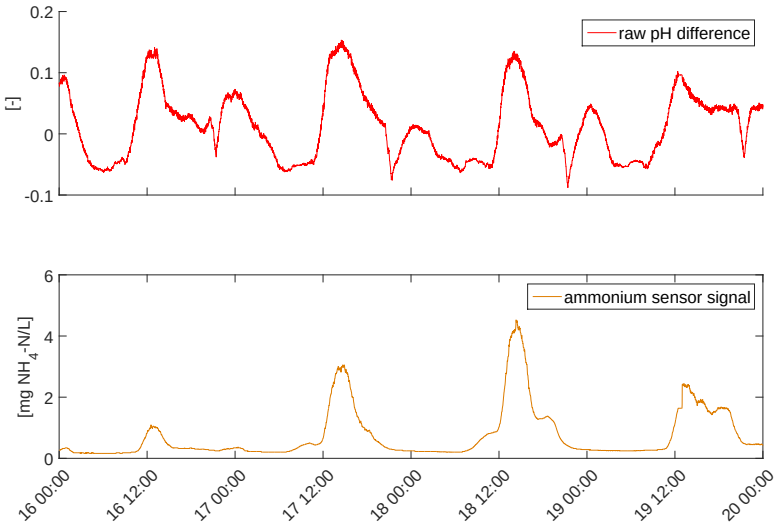

Figure 1: pH difference between first and last tank in a series of three aerated activated sludge tanks (top) and NH_4-N concentration (bottom) measured in the middle tank. Each of the NH_4-N peaks is matched with a peak in the pH difference.

Despite the expected drift of pH sensors, information about the ammonia load can be extracted from the pH difference signal. The pH and ammonium time series in Fig. 1, recorded at the Hard WWTP (details follow), show the synchronous appearance of their extrema. The maxima in particular appear well-synchronised. The pH difference signal (top Fig. 1) exhibits two peaks per day. The one with the higher amplitude exhibits a shape which is similar to the shape of the ammonia concentration peak (bottom Fig. 1). The sharp drop and increase in the pH difference every evening is caused by an air flushing which is used to clean the ceramic diffuser domes. In this study the identified minima and maxima of the pH difference signal (further referred to as qualitative features) are used in a semi-quantitative rule-based control to distinguish between high and low ammonia load situations. Based on this information the controller will determine the DO setpoint in the slave aeration controller.

Qualitative trends and the corresponding features of the signal can be detected by means of qualitative trend analysis (QTA) methods. QTA algorithms split time series into time periods called *episodes* within which the signs of one or more of the signal's derivatives are constant. Cheung and Stephanopoulos (1990) introduced seven primitives which can qualitatively describe most of the trends occurring in a time series. Based on this study, methods were developed to identify these primitives in online signals. A piece-wise fitting method was presented by Konstantinov and Yoshida (1992). Bakshi and Stephanopoulos (1994) used wavelet analysis to extract trends. Another approach was taken by Rengaswamy and Venkatasubramanian (1995), who identified the primitives by means of a neural network. So far, QTA has been used for process monitoring (e.g. Charbonnier and Gentil, 2007), fault diagnosis (e.g. Scali and Ghelardoni, 2008; Maurya et al., 2007) and fault detection (e.g. Villez and Habermacher, 2016). More recently, QTA-like models have also been used for image analysis (Derlon et al., 2017) and for kinetic modelling (Mašić et al., 2017; Thürlimann and Villez, 2017). First ideas to use QTA also in process control were mentioned by Villez et al. (2008). However, QTA has not been applied for any autonomous control purpose yet. The use of this qualitative information in a control system requires an algorithm that detects the qualitative features in a timely and robust manner. In contrast to the mentioned methods, shape constrained splines (SCS) methods are robust to noise yet computationally expensive (Villez and Rengaswamy, 2013; Villez and Habermacher, 2016; Derlon et al., 2017). A recently developed method called Qualitative Path Estimation (QPE) approximates the performance of the SCS in a computationally efficient manner (Villez, 2015). QPE fits a given sequence of shapes, defined by the first and second derivative, to a given signal and optimises the identified change of

the shapes within the sequence (i.e., maxima, minima and inflection point). This means that the sequence of shapes needs to be known beforehand. This is why QPE as well as the SCS methods are, despite their accuracy, simplicity and speed, not suited well for online applications. Therefore, a new algorithm tailored for online application, called Qualitative State Estimation (QSE) based on the QPE algorithm, is proposed in this study.

Tests on six different municipal WWTPs, which differ in size and treatment technologies, demonstrate the long-term validity and ease of deployment of the QSE method in combination with the rule-based controller for ammonia based aeration control. The paper focuses mainly on the results from the WWTP Hard where the QSE technique was developed initially and the longest experimental phase, i.e. 15 months of virtual testing (i.e., post-processing of recorded data without closed loop control) and live testing (i.e., with closed loop control) took place. Experiences gained on other plants are listed and discussed with a focus on the opportunities and limitations of the method. The paper is structured as follows: In the Material and Methods section the theory of the newly developed QSE algorithm is described in detail. Furthermore, a new control strategy for pH based ammonia control is presented. This includes the hardware-setup as well as the rule-based control. Finally, information about the performance evaluation procedure and the characteristics of the WWTPs where the controller was tested are provided. The Results section is used to demonstrate the QSE algorithm and the rule-based controller, and to evaluate the controller performance with regard to various full-scale plants. The implications of the results are discussed at the end of this study.

2. Material and Methods

2.1. Work flow summary

Fig. 2 gives an overview of the data processing steps from the initial pH sensor measurements to the final adaptation of the nitrification process. The first four steps are described in detail in the Material and Method section. First, the signal is processed by the two-step QSE algorithm, then the rule-based control evaluates the QSE result. In parallel to this procedure, the raw signal and some of the plant's process states are analysed to check for anomalous data or events and override the rule-based control decision when necessary. After this final check, the DO setpoints are passed to the DO controller which will adapt the airflow in the biological reactors and subsequently the nitrification rate and the pH in the bulk phase. An implementation of the QSE algorithm in Matlab® is available in the *supplementary information* (*SI*). Additionally a Matlab®/Octave toolbox enables to reproduce our results.

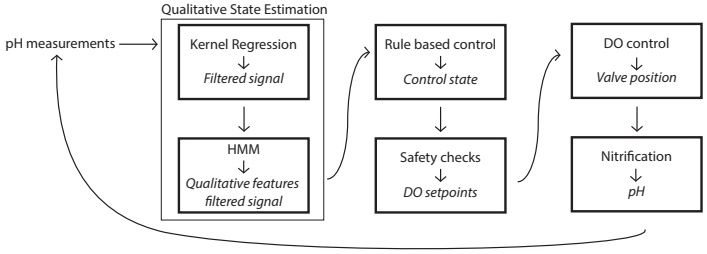

Figure 2: Preprocessing of raw signal with QSE and subsequent control cascade steps leading to an adaptation of the nitrification performance. The normal text names the step, the italic text names the outcome of the step which is passed on to the next step.

2.2. Qualitative State Estimation algorithm

The Qualitative State Estimation (QSE) is an algorithm for the segmentation of time series into *episodes*. Each episode has a distinct set of signs for the signal's derivatives. Each such unique set is referred to as a *primitive* and is reported as an unique alphabetic character. Table 1 shows the thirteen primitives based on the sign of first and/or the second derivative. A character is attributed to each primitive. Some are meant as abbreviations of their form while others are assigned arbitrarily. Depending on the use case, episodes can

either be defined by the first and second derivative (i.e., A (antitonic acceleration), B (boost), C (concave deceleration), D (deceleration), E, F (flat), G) or only by one derivative (i.e., L (lower), N (negative), P (positive), and U (upper)) or by none (i.e., Q (undefined signs)). The end point of one episode is the start point of the next and represents either a maximum, a minimum, or an inflection point. In the application studied in this work, only the primitives defined on the basis of the first derivative sign are used. Moreover, only positive and negative signs are considered while zero-valued or undefined derivative signs are excluded. Therefore, only the primitives U and L are of practical use in this work. The QSE algorithm is however general and allows - in principle - for all primitives in Table 1 to be considered.

Table 1: Thirteen primitives defined by first and second derivative sign of the signal. ? stands for undefined sign. Adapted from (Villez and Rengaswamy, 2013). (*) primitives shared with Cheung and Stephanopoulos (1990). ($^+$) primitives are listed for completeness but are deliberately ignored in this paper.

Derivatives		First			
	Signs	+	0	-	?
Second	+	B*	/	A*	P
	0	G^{*+}	F^{*+}	E^{*+}	O$^+$
	-	C*	/	D*	N
	?	U	/	L	Q

The proposed algorithm is based on Qualitative Path Estimation (QPE) by Villez (2015), which allows fast batch process diagnosis. Both QPE and QSE can be separated into two steps. The first step of both QPE and QSE provides a forward flow of information based on locally weighted least-squares regression. In this step, time-local probabilities for the sign of the derivatives are computed. It is noted that the first filtering step of the QSE algorithm is designed to estimate probabilities for the derivatives' signs to facilitate the further processing.

The second step of QPE is based on the Viterbi algorithm (Viterbi, 1967) for optimal path estimation and is used to convert the continuous probabilities into discrete values for the derivative signs. The Viterbi algorithm includes a forward and a backward pass which renders QPE especially useful for off-line and end-of-batch applications. As the current application is focused on control in continuous-flow processes, we replace this step with a maximum likelihood estimation of the discrete state at every measurement sample. As a result, QSE only requires a forward pass which makes the method especially suitable for on-line application.

Both steps in the QPE and QSE algorithms induce a denoising effect. Importantly however, the second step also allows enforcing a certain order of appearance of the identified inflection points, minima, and maxima. This means that both QPE and QSE allow incorporation of prior knowledge about the shape of the expected signal, similarly to the SCS method and in contrast to conventional QTA methods. The following paragraphs describe the QSE algorithm. The method is first presented in a general form (*General treatment*) and then applied to the pH difference example (*Application*).

Step 1 – Kernel regression to estimate qualitative trend probabilities.

General Treatment. First, kernel regression is used to smoothen the raw signal with a moving window approach. i is the data point at the window centre. The point $i + w$ represents the most recent data point in the window. An equidistant spacing of the data points is assumed. For each window $[i - w\ ,\ i + w]$ around the window centre, a polynomial function with a degree d is estimated by minimising a weighted least square function. The chosen weighting is known as the tri-cube kernel (Hastie et al., 2009). This kernel is parametrised by τ, which is the critical distance from the reference data point i after which all data points have zero weight. The complete vector of measurements included in the window centred around i is given

as $\boldsymbol{y}_i$. The number of data points considered in the kernel regression increases with τ given that $w = \lfloor \tau \rfloor$. This generates two effects which one has to trade off against each other: *(i)* with increasing value for τ the filtered signal is smoother and QSE therefore identifies fewer qualitative features and *(ii)* the delay of the filtered signal increases. Kernel regression enables the computation of the expected distribution for the derivatives of the signal f_i^a as:

$$f_i^a \sim \; N\left(\hat{f}_i^a, \sigma_{f^a,i}\right) \tag{1}$$

with $\sigma_{f^a,i}$ as the variance of the signal ath derivative. For details we refer to Villez (2015).

The probability for a positive sign of the ath derivative at the sampling point i ($\Lambda_i^{(a)}$) is computed by integrating the normal probability density function from infinity to zero:

$$\Lambda_i^{(a)}\Big(+ \mid \boldsymbol{y}_i\Big) = \Lambda\Big(f_i^{(a)} > 0 \mid \boldsymbol{y}_i\Big) = \int_{u=0}^{+\infty} \frac{1}{\sqrt{\sigma_{f^a,i}\Pi}} \exp\left(-\frac{\left(u - \hat{f}_i^a\right)^2}{\sigma_{f^a,i}}\right) du \tag{2}$$

The difference of 1 to the positive sign probability corresponds to the probability for a negative sign:

$$\Lambda_i^{(a)}\Big(- \mid y\Big) = 1 - \Lambda_i^{(a)}\Big(+ \mid y\Big) \tag{3}$$

Note that zero-valued derivatives are deliberately ignored. With Eq. 2 and Eq. 3 one can derive the probability for the primitives L, N, P, Q, and U (Table 1). For the primitives defined by the first and second derivatives one multiplies the probabilities for the sign of the first and the second derivatives to obtain the probability for the primitive. In words, the probability for the primitive (κ_i) to be A is equal to the probability that the first derivative is positive times the probability that the second derivative is negative.

$$\Lambda_i\Big(\kappa_i = A \mid \boldsymbol{y}_i\Big) = \Lambda_i^{(1)}\Big(+ \mid \boldsymbol{y}_i\Big) \cdot \Lambda_i^{(2)}\Big(- \mid \boldsymbol{y}_i\Big) \tag{4}$$

This product ignores the presence of correlation between the estimates for the first and second derivative. However, as pointed out in Villez (2015), no significant loss in accuracy has been demonstrated so far. For each primitive, the products are derived analogously according to Table 1. This concludes the first step.

Application. For the pH difference example, the procedure is limited to Eq. 1-Eq. 3 as only the sign of the first derivative is of interest. Thus, the second derivative is neglected and only the primitives U (positive first derivative) and L (negative first derivative) (cf. Table 1) are used.

Step 2 – Probability estimation by hidden Markov model.

General Treatment. To identify the change of the primitive (derivatives signs), a hidden Markov model (HMM) is used. Every state s in the HMM is associated with a primitive $\boldsymbol{q}(s)$. $\boldsymbol{q}$ is the vector containing all the possible primitives. This HMM model predicts the likelihood of a certain process state t at point i, depending on the likelihoods of all possible states at point $i-1$ as follows:

$$\Lambda\Big(\boldsymbol{s}(i) = t \mid i-1\Big) = \sum_{p=1}^{q} \mathbf{T}(t,p) \cdot \Lambda\Big(\boldsymbol{s}(i-1) = p \mid i-1\Big) \tag{5}$$

At any given time instant, i, the likelihood that the state equals t is a sum of products, where each product consists of the likelihood that the state equalled p at time $i-1$ times the associated transition likelihood, i.e., the likelihood that at point i the process is at state t, conditional to the process state being p at time point $i-1$. The complete set of transition likelihoods is given as a matrix $\mathbf{T}$ with $\mathbf{T}(t,p)$ the transition likelihood

for states p and t. Thus, the prediction of the process states in i are only dependent on the likelihoods at $i-1$ (first-order Markov property).

$$\mathbf{T}(t,p) = \begin{cases} 1, & p = t \\ \lambda_i(t,p) = \lambda, & \text{otherwise} \end{cases} \tag{6}$$

All diagonal elements are equal to one. Each off-diagonal element $\lambda_i(t,p)$ of the transition probability matrix $\mathbf{T}(t,p)$ may take any value between 0 and 1. The lower the value of these off-diagonal elements, the more effective is the filtering by the HMM (i.e, less likely switch from state p to t). Any $\lambda_i(t,p)$ equal to zero prohibits the identification of a state change from p to t within a single time step.

The predictive model is now completed with a sensor model. According to Villez (2015) the HMM sensor model and the underlying assumptions can be written as follows:

$$\Lambda\left(\boldsymbol{y}_i \mid \boldsymbol{s}(i) = t\right) = \Lambda\left(\boldsymbol{\kappa}(i) = \boldsymbol{q}(t) \mid \boldsymbol{y}_i\right) \tag{7}$$

In words, the likelihood of the observed data ($\boldsymbol{y}_i$) conditional to a given state (t) is assumed equal to the probability of the primitive associated with the considered state ($\boldsymbol{q}(t)$) conditional to the observed data ($\boldsymbol{y}_i$). The latter probability, i.e., the right side of Eq. 7, is calculated in the first step of the QSE (cf. Eq. 4). Based on Eq. 5 and Eq. 7, the maximum likelihood state estimate $\hat{s}$ at time i is now given as:

$$\hat{s}_i = \max_t \Lambda\left(\boldsymbol{s}(i) = t \mid i\right) \tag{8}$$

with

$$\Lambda\left(\boldsymbol{s}(i) = t \mid i\right) = \Lambda\left(\boldsymbol{\kappa}(i) = \boldsymbol{q}(t) \mid \boldsymbol{y}_i\right) \cdot \sum_p \mathbf{T}(p,t) \cdot \Lambda\left(\boldsymbol{s}(i-1) = t \mid i-1\right) \tag{9}$$

Each change in the maximum likelihood state corresponds to the identification of a change in primitive and therefore implies the presence of a qualitative feature such as a maximum, a minimum, or an inflection point.

Application. For the pH difference signal the HMM identifies the change in the first derivative's sign (zero-crossing) in the signal resulting in one of the two model states (i.e., upward and downward trend). The vector $\boldsymbol{q}$ has the following form:

$$\boldsymbol{q} = \begin{pmatrix} L \\ U \end{pmatrix} \tag{10}$$

In contrast to the work in Villez (2015) the transition matrix $\mathbf{T}$ is symmetric in this study as both maximum likelihood transitions (i.e., L to U and U to L) have the same λ values. This implies that both state changes are a priori equally likely (cf. section 4.2).

$$\mathbf{T} = \begin{pmatrix} 1 & \lambda \\ \lambda & 1 \end{pmatrix} \tag{11}$$

The time instants where the state estimate $\hat{s}_i$ changes correspond to maxima and minima are used in the proposed rule-based controller.

2.3. Hardware and Control strategy

The next paragraph describes the hardware setup as well as the rule-based controller.

General experimental design. A pH difference signal is computed as the difference between two pH sensor signals. These sensors are placed in the first and the last aerated tank in the biological treatment line as shown in Fig. 3. The difference is calculated as the upstream tank measurement minus downstream tank measurement. It is important to know that this difference is influenced mostly by nitrification and CO_2-stripping (i.e., no denitrification, no control of the pH by any other means). Both pH signals are recorded simultaneously with a high enough sampling interval (e.g., 1 minute) to capture the relevant process dynamics. For development and validation purposes an ion selective electrode (ISE) is installed to record ammonia concentrations. This ammonia signal is only used for performance evaluation and fine-tuning.

Hardware setup, maintenance and signal processing. Fig. 3 shows the chosen hardware setup. The two pH sensors (Orbisint CPS11D, Endress & Hauser, Reinach, Switzerland) are placed in the first and the last aerated tank. Both pH signals are recorded with a sampling interval of one minute. The ion selective electrode (ISE, ISEmax CAS40D, Endress & Hauser AG, Reinach, Switzerland) is placed in the middle tank and is not used for control.

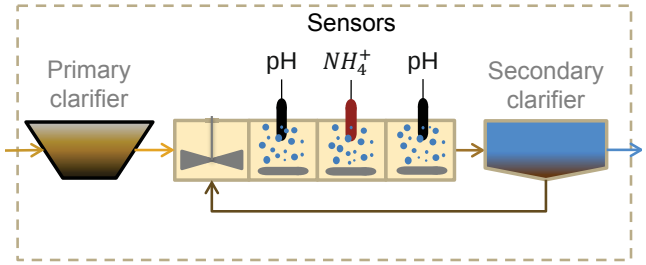

Figure 3: Hardware setup for soft-sensor development and validation on a wastewater treatment plant. pH sensors (black) are placed in front and rear aerated reactor tanks. An ammonia sensor (red) is placed in the middle tank (not used for control).

All sensor signals are converted to 4-20 mA signals (Liquiline CM444, Endress & Hauser AG, Reinach, Switzerland) and fed into the supervisory control and data acquisition system (SCADA, RITOP®, Rittmeyer Ltd., Baar, Switzerland). The final version of the controller is implemented in and fine-tuned with RITUNE® (Rittmeyer Ltd., Baar, Switzerland), a platform for WWTP optimisation. This platform communicates with RITOP® to receive process data and send control tasks. Initial tests and algorithm development are executed in Matlab®. The maintenance of the hardware is limited to a weekly cleaning of the sensors. Recalibrations are executed based on a visual inspection of the signal but not more than once a month. This is applied on all WWTPs considered.

Rule-based control. The extrema resulting from the QSE require further processing for two reasons: First, the QSE step fed with the noisy and dynamic raw signal (Fig. 1 top) results in more detected extrema than one would expect given the rather smooth ammonia concentration time series (Fig. 1 bottom). Second, the variable of interest is the ammonia concentration, which is a continuous signal; the QSE output, however, is binary (i.e., downward trend/maximum, upward trend/minimum) and forces one to distinguish between ammonia peaks that can be handled with low DO concentrations and ammonia peaks that require an enhanced nitrification performance (high DO).
Consequently, additional semi-quantitative rules are added to obtain a meaningful control output. These rules are implemented into a rule-based controller which is the master control in the aeration cascade control. It designed as a gain scheduling controller which switches between two states: An *Eco* controller state which demands a low DO setpoint and a *Normal* controller state which demands a high DO setpoint. These setpoints are passed to the existing DO based airflow controller. These two master control states are further referred to as Eco ($c = 1$) and Normal ($c = 2$). First the general treatment is presented in a formalised way. Then the rule evaluation is explained with an illustrative example. Finally the sought intention for all rules are explained by drawing relationship between the rules, the pH difference signal and the nitrification activity.

General treatment. A set of three rules to trigger controller state change and in turn a DO setpoint adjustment are used in each of the two control states c (c=1,2). Each of them was developed based on an empirical analysis of typical similarities in the pH difference and ammonia concentration signal for the corresponding control state. The applied rules differ depending on the control state. The rules are further referred to as *EtN1*, *EtN2* and *EtN3* for the switch from Eco ($c = 1$) to Normal ($c = 2$) and as *NtE1*, *NtE2* and *NtE3* for the switch from Normal ($c = 2$) to Eco ($c = 1$). The controller only switches the state c if all three rules are simultaneously true.

The rule based control consists of two while-loops. One of these is active when $c = 1$ while the other one is active when $c = 2$. Both loops are described by pseudo-code in Table 2 and Table 3. In the pseudo-code the following states are used: c is the control state (i.e., $c = 1$: Eco, $c = 2$: Normal) and, $\hat{s}$ is the HMM state estimate (i.e., $s = 1$ downward trend (primitive L), $s = 2$ upward trend (primitive U)). pHd_{LM} is the value of the lowest minimum in the Eco and is initialised with the current pH difference $pHd(i)$. Furthermore the following parameters are used: *EtN2* is parametrised by ΔpHd_{MID}. T_{MIN} is a critical wait time to allow a control switch from Eco to Normal (*EtN3*). ΔpHd_{MAD} parametrises the *NtE2* rule.

Table 2: While loop of rule based control used while the controller is in the Eco state with pseudo-code and explanations.

Pseudo code	**Comments**
WHILE $c = 1$	
IF $s(i-1) = 1 \wedge s(i) = 2$	
EtN1 = TRUE	EtN 1: monitors if at least one minimum was identified, updates $\mathrm{pHd_{LM}}$ if minimum is lower than current $\mathrm{pHd_{LM}}$
IF $pHd_{LM} > pHd(i)$	
$\mathrm{pHd_{LM}}$ = pHd(i)	
END	
END	
IF $pHd(i) > pHd_{LM} + \Delta pHd_{MID}$	
EtN2 = TRUE	EtN 2: monitors if current pH difference value pHd(i) is larger than lowest minimum $\mathrm{pHd_{LM}}$ plus a threshold $\Delta\mathrm{pHd_{MID}}$
ELSE	
EtN2 = FALSE	
END	
IF $s(i) = 1$	
EtN3 = TRUE	EtN 3: monitors signal for downward trend
ELSE	
EtN3 = FALSE	
END	
IF $EtN1 = TRUE \wedge EtN2 = TRUE \wedge EtN3 = TRUE$	
EtN1 = FALSE	Checks if all rules are true and if so, resets all rules, resets the Normal timer and switches the controller state
EtN2 = FALSE	
EtN3 = FALSE	
$\mathrm{T_{Normal}} = 0$	
c=2	
END	
END	

Table 3: While loop of rule based control used while the controller is in the Normal state with pseudo-code and explanations.

Pseudo code	**Comments**
WHILE $c = 2$	
IF $s(i-1) = 1 \wedge s(i) = 2$	NtE 1: monitors if at least one minimum was identified
NtE1 = TRUE	
END	
IF $pHd(i) < pHd_{LM} + \Delta pHd_{MAD}$	NtE 2: monitors if current pH difference pHd(i) value is smaller than lowest minimum $\mathrm{pHd_{LM}}$ plus a threshold $\Delta\mathrm{pHd_{MID}}$
NtE2 = TRUE	
ELSE	
NtE2 = FALSE	
END	
IF $T_{Normal} > T_{MIN}$	NtE 3: monitors if Normal control state is active for sufficient time period
NtE3 = TRUE	
END	
IF $NtE1 = TRUE \wedge NtE2 = TRUE \wedge NtE3 = TRUE$	Checks if all rules are true and if so, resets all rules, initialise $\mathrm{pHd_{LM}}$, and switches the controller state
NtE1 = FALSE	
NtE2 = FALSE	
NtE3 = FALSE	
$\mathrm{pHd_{LM}}$ = pHd(i)	
c=1	
END	
END	

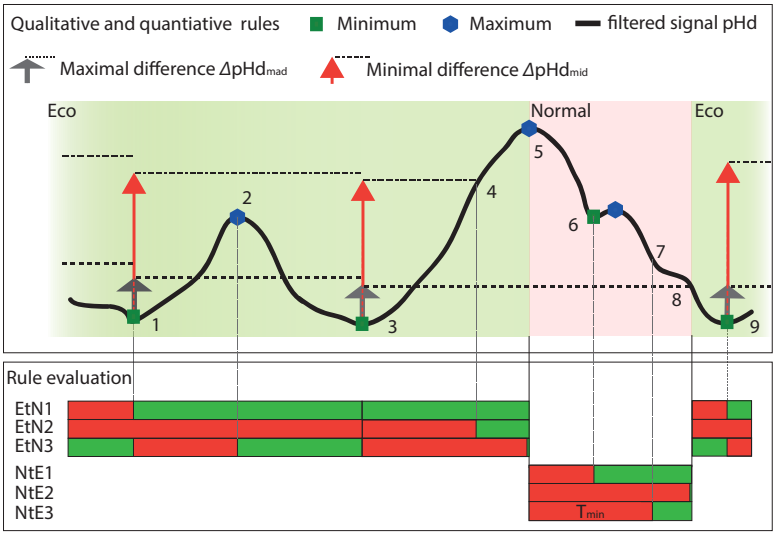

Figure 4: Evaluation procedure of the rule-based control after signal (black full line) processing by QSE. (1) to (9) indicate critical time points in the evaluation of the signal and its maxima (blue hexagon) and minima (green square) which are further explained in the text below. Yellow (grey) background means Eco (i.e., low DO), red (bright grey) background means Normal (i.e., high DO)

Application. Fig. 4 visualises the rule-based controller execution in an exemplary scenario. The signs (1) to (9) indicate important points for the rule evaluations. The controller starts in the Eco state ($c = 1$), which means the while-loop in Table 2 is active. The pH difference exhibits a downward trend. Therefore, *EtN3* is true. At (1) the pH difference starts to increase and in turn the first minimum is identified. This sets *EtN1* to true and *EtN3* is turned back to false. The value at (1) is stored as pHd_{LM} and further used to identify ΔpHd_{MID} (to be used while $c = 1$) and ΔpHd_{MAD} (to be used when $c = 2$). During the following upward trend point (1) to (2), the pH difference does not exceed the current threshold ΔpHd_{MID}. At point (2), the pH difference starts to decrease. This results in an identified maximum (2) setting *EtN3* to true. At point (3), a new minimum is identified. This minimum is lower than the current value of pHd_{LM}, which was last updated at point (1) and set equal to the value of the pH difference at that time. Because of this, pHd_{LM} is updated with the current difference value, i.e., the value at point (3). As the difference signal increases beyond point (3), *EtN3* is evaluated as false. At point (4), the signal increases above the the threshold ΔpHd_{MID} so that *EtN2* is evaluated as true. The rising trend continues until a new maximum is reached at point (5). Consequentially, *EtN3* is evaluated as true. At this point all *EtN* rules are evaluated as true and the controller state is changed from Eco (c=1) to Normal (c=2). This deactivates the while-loop described in Table 2 and activates the while-loop in Table 3. From point (5) onward, the pH difference decreases. At point (6) the trend direction changes from downward to upward to that *NtE1* is evaluated as true. At point (7) T_{MIN} has passed since the time instant where the controller state changed to the Normal state $c = 2$ (point (5)). This means *NtE3* is evaluated as true. The difference signal continues to decrease beyond point (7) and at point (8) the pH difference reaches the value of $pHd_{LM} + \Delta pHd_{MAD}$, with pHd_{LM} lastly updated at point(3). This means all NtE rules are evaluated as true and the controller state is changed to the Eco state (c=1). All *EtN* rules are reset to false. As soon as the while-loop in Table 2 is activated gain, *EtN3* is set to true given that the pH continues to decrease after point (8). The first new minimum in Eco at (9) triggers the same changes as the minimum at point (1). In other words, the controller state has cycled through all possible states.

Empirical relation between nitrification activity and rule evaluation. *EtN1* ensures that the minimal nitrification activity is identified. *EtN2* is used to distinguish between low ammonia load peaks which only elicit a minor pH difference increase and increased ammonia loads. The example shows that relevant peaks (5)

lead to a change in the controller state while irrelevant ones (2) do not. The updating of pHd_{LM} from (1) to (3) ensures that *EtN2* and *NtE2* are evaluated based on the lowest nitrification activity recorded and compensates for any drift between the two time points (1) and (3) and for any drift before point (1). *EtN3* delays the switch to the Normal controller state until the ammonia peak load has broken through and reached the last aerated tank.

NtE1 monitors the identification of a first minimum. This minimum can be an indication that the nitrification activity is stabilising at a low level. However, as such minima can also be identified at high nitrification activities, the condition is only sufficient and therefore paired with *NtE2*. *NtE2* ensures that a switch to a lower oxygen level is only pursued when the observed minimum in nitrification activity is close enough to the last known minimum nitrification activity characterised by the parameter pHd_{LM}. The tolerance ΔpHd_{MAD} is a parameter describing the maximally allowed upward drift since the last known minimum in nitrification activity, i.e., the last time pHd_{LM} was updated. *NtE3* safeguards the hardware such as blowers and valves from unnecessary and excessive wear and tear.

2.4. Controller performance monitoring

Initial tests. On all six tested plants, the hardware (i.e., pH sensors and ammonia ISE) is installed first and the signals are recorded without using them for control. With these initial tests, the QSE and the rule-based control can be tested and tuned off-line first.

Controller tests. In three plants extended tests were executed by activating the control which automatically adapted the DO setpoints. In these cases the controller performances were evaluated in terms of energy efficiency gains and nutrient removal performance of the plants. To this end, a comparison is made between period with and without active rule-based control. The turbo blower energy consumption was normalised by the ammonia load converted in the biological treatment line to enable a fair comparison. This load was derived from the effluent flow from the primary clarifier and both volume proportional 24h composite sample concentrations in the outflow of the primary clarifier and the outflow of the secondary clarifier (cf. *SI* and Fig. 10). The WWTP's laboratories determined the nitrogen species concentrations in the 24h composite samples with LCK 303, 304 and 305 test kits for ammonia, LCK 341 and LCK 342 for nitrite, and LCK 340 for nitrate (Hach-Lange, Berlin, Germany).

2.5. Wastewater treatment plant characteristics

The use of the QSE-based control at the WWTP Hard (Winterthur, Switzerland), which treats wastewater of about 130'000 population equivalent, is studied in greatest detail. The WWTP Hard treatment line consist of a mechanical treatment step for removal of solids, a primary clarifier, a secondary treatment by means of a biological activated sludge process including a series of reactor tanks and settling tanks, and a tertiary sand filtration. After the primary clarifier, the water is separated into two parallel treatment lines. The proposed control is only tested in one of the biological treatment lines. In the biological treatment step the wastewater first enters a stirred anoxic tank for denitrification followed by a series of three aerated tanks for the oxidation of organic compounds and ammonia. Each biological treatment line has its own separated turbo blower group, aeration pipe system, automatic samplers, and secondary clarifier tanks. Thus, the evaluation of the controller is not affected by the line without ammonia based aeration control. The existing aeration control is a cascade controller: First, all three aerated tanks are equipped with an oxygen sensor that controls the airflow coming from the main collector. Then, the line's turbo blower group is controlled by a sliding pressure control in the collector. The diffusers are ceramic domes which require a minimal constant airflow to avoid clogging.

The five other treatment plants differ in size and their biological treatment technology (Table 4). However, all of them are continuously operated, receive typical municipal wastewater and include a predenitrification step. The WWTP Buchs (BCH, Buchs SG, Switzerland) has an additional denitrification step in-between the aerated tanks.

The two other plants in which the control was activated were WWTP Pfäffikon (PFA, Pfäffikon ZH, Switzerland) and WWTP Pfungen (PFU, Pfungen, Switzerland). They have a slightly different aeration setup than

the WWTP Hard: PFU has three aerated compartments in each of the two lines, yet only one airflow valve per line which manipulates the airflow to all three compartments at once. The oxygen is only controlled in the second compartment. In PFA, the biological treatment is separated into two lines with two tanks each. The front tanks are pure activated sludge tanks and the second tanks are designed as hybrid moving bed systems which require a minimum air flow at all time. All tanks have separate DO sensors and airflow valves.

Table 4: WWTP characteristics and performance evaluation of the pH difference based ammonia controller of six different municipal wastewater treatment plants.

WWTP	Altenrhein (ALT AS)	Altenrhein (ALT FB)	Buchs (BCH)	Hard (AS)	Pfäffikon (PFA)	Pfungen (PFU)
Wastewater treatment plant characteristics						
Location	Altenrhein, Switzerland	Altenrhein, Switzerland	Buchs SG, Switzerland	Winterthur, Switzerland	Pfäffikon ZH, Switzerland	Pfungen, Switzerland
Design capacity [p.e.]	120'000[a]	120'000 [a]	45'600	180'000	15'000	15'200
Biological treatment	Activated sludge	Fixed bed	Activated sludge[b]	Activated sludge	Hybrid moving bed	Activated sludge
Controller performance evaluation						
Evaluation type (# days [d])	virtual test (32)	virtual test (18)	virtual test (63)	live test[c] (142)	live test[c] (50)	live test[c] (25)
Active control and reference period (active control / benchmark) [d] [d]	0/0	0/0	0/0	79/23	7/10	7/7
Energy efficiency improvement [%]	N/A	N/A	N/A	-7.2	-13	-8.2[e]
Legal effluent requirements	N/A	N/A	N/A	fulfilled	fulfilled	fulfilled
Control potentially feasible	Yes	No	Yes	Yes	Yes[f]	Yes

[a] Total capacity of WWTP Altenrhein: Activated sludge reactors account for 70% and the fixed bed reactors for 30% of this number.

[b] This plant has an additional anoxic zone in between the aerated tanks. The pH difference was measured over this second anoxic zone.

[c] Preceding virtual tests executed

[d] Numbers do not sum to total days, as benchmarking days are only those at which laboratory nitrogen measurements were available.

[e] This number is not normalised with the ammonia load.

[f] pH difference does not exhibit the same dynamics as other plants. For a more efficient control, adaptations in the rules based control are needed.

2.6. Oxygen slave controller

On all actively tested plants the presented controller was installed as a master controller which sets the DO setpoints in the existing DO based aeration control. The Normal control state DO setpoints were set to the values in use prior to this study (WWTP Hard, 2.0 mgO_2/L). For the Eco control state DO setpoints at WWTP Hard a variability of combinations were tested. First, the setpoint spread was kept conservatively small for safety reasons (2.0 to 1.5 mgO_2/L) as the effects of the controller on the pH signal dynamics were not known exactly. Later this spread was increased to maximise efficiency (2.0 to 0.8 mgO_2/L). The details regarding the different DO setpoints can be found in the *SI*. The combination for the Eco DO setpoints which was used the longest was 1.3 mgO_2/L in the front, 1.1 mgO_2/L in the middle and 1.0 mgO_2/L in the rear tank. The pH difference was amplified with decreasing Eco DO setpoints and urged an adjustment of the parameters in the rule-based control (i.e., *EtN2*, *NtE2*). With an increasing setpoint spread the aeration system's time lag increases and the safety margins decrease. In PFU the DO was controlled at 1.8 mgO_2/L for Normal and 1.3 mgO_2/L for Eco. In PFA the airflow is governed by the hybrid moving bed tank sensor which measures the lower DO concentration of the two lines. This sensor DO setpoint was set to 3 mgO_2/L for Eco and 4.5 mgO_2/L for Normal. The front tank DO setpoints were kept at 1.5 mgO_2/L independent of the controller state.

3. Results

In the next paragraphs the QSE algorithm and the subsequent control strategy are first demonstrated on the basis of a 4-day dataset recorded on the WWTP Hard. After these demonstrations the performance evaluation is presented.

3.1. Demonstration of Qualitative State Estimation algorithm

First, the results from the kernel regression step leading to the trend probabilities are shown. In the kernel regression the moving window filter computes a new filtered signal from the raw signal (cf. Fig A1 in the *SI*). The chosen window length τ is 19 minutes. This corresponds to 39 considered measurements within the moving window and causes a delay of 19 minutes (cf. section 2.2). For this study τ has been selected by trial and error. A good first estimate can be derived from the sensor, process, or actuator time constants. Methods such as generalised cross validation Hastie et al. (2009) also allow tuning of τ. This, requires a cost-function which can be computed with limited uncertainty. For this reason, this is not explored further in this study.

The first derivative and its standard deviation is computed to characterise the distribution of the second derivative of the pH difference signal (Eq. 1). Fig. 5 illustrates the distribution's mean and a $\pm\ 3\sigma$ band around it for the pH difference first derivative over time.

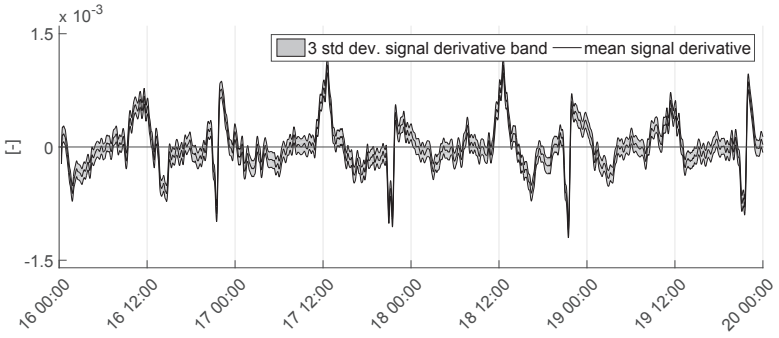

Figure 5: Mean and expected $\pm\ 3\sigma$ band around the mean characterising the calculated signal first derivative's distribution.

The probabilities for a positive first derivative sign are derived by integrating the positive mass of the distribution shown in Fig. 5 (Eq. 2). The corresponding probability for a negative sign is computed with Eq. 3. Fig. 6 shows the resulting positive (upward trend) probability of the pH difference signal. This concludes the kernel regression step of the QSE.

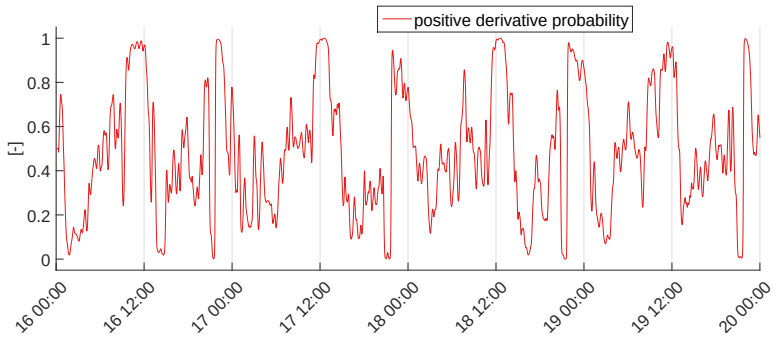

Figure 6: Positive derivative sign (upward trend) probability of the signal. This probability is derived from the point-wise distribution of the signal's first derivative.

In the QSE's second step, the probabilities computed above are used to compute the maximum likelihood HMM state (Eq. 8). Fig. 7 shows the resulting maximum likelihood states based on the pH difference signal. Several values for λ were tested by trial and error. Good results were obtained with λ=0.4. This however means that the filtering effect by means HMM state estimation is marginal. In our case, the HMM therefore primarily functions as a way to transform the continuous probabilities into discrete states. Results obtained with more aggressive filtering $\lambda = 0.05$ are plotted in Figure A2 and Figure A3 (cf. 4.2). The identified trends and extrema are now further processed by the proposed rule-based control.

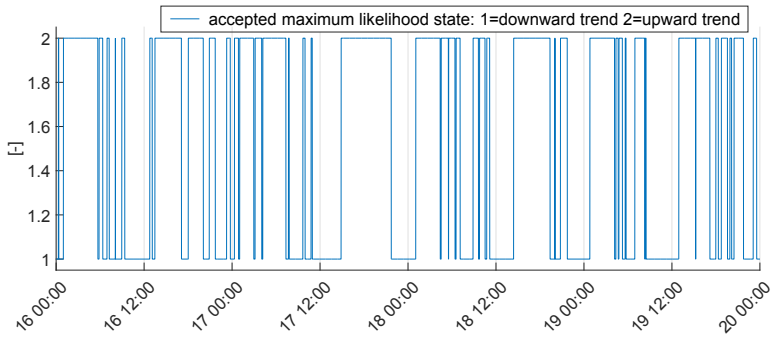

Figure 7: Accepted maximum likelihood state of the signal. 1 = downward trend, 2= upward trend. Changes in the maximum likelihood state indicate the identified maxima and minima.

3.2. Demonstration of the control strategy

Rule-based control. The rule-based control uses the qualitative features identified by the QSE algorithm. It evaluates rules which analyse the filtered pH signal and the appearance of the qualitative features on the signal. Fig. 8 shows the evaluation of these rules in the two different control states based on the example signal from above. Most typically, the *EtN3* is the last rule to become activated before a switch from Eco to Normal. The most influential rule with respect to the aggressiveness of the controller is *EtN2* which defines the length of the Eco to a great extent. For changes from the Normal and Eco state, the *NtE1* and *NtE2*

rule tend to be the last to become evaluated as true. Therefore, we have primarily focused on tuning the controller by manipulating the parameters for the rules *EtN2*, *EtN3*, *NtE1*, and *NtE2*.

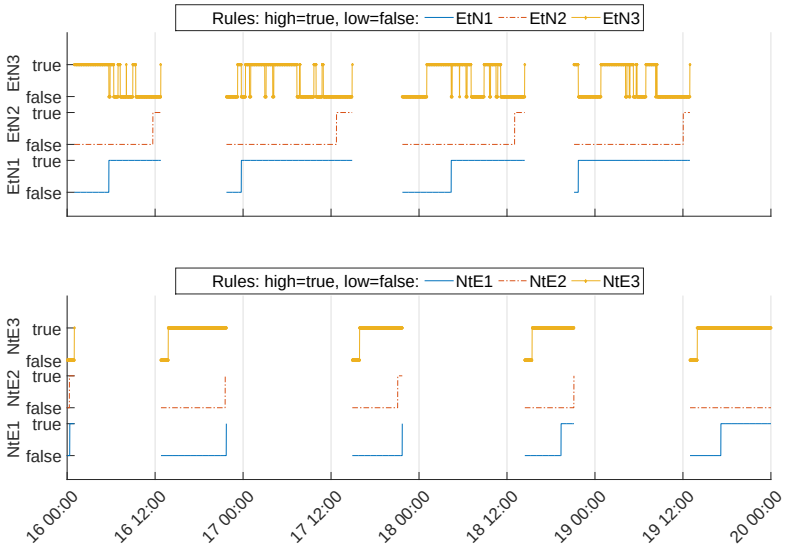

Figure 8: Rule evaluation and pH difference; Top: Eco to Normal (*EtN*) rules; Bottom: Normal to Eco (*NtE*) rules.

The example shows that the four rules *EtN2*, *EtN3*, *NtE1*, and *NtE2* are critical when it comes to controller state switches. In the Eco state (Fig. 8 top) as well as in the Normal state (Fig. 8 bottom) one can see that the rules *EtN2* and *NtE2* typically stay for the longest on false. An exception is found for the third time the controller is in the Normal control state. In this particular case *NtE1* was the last rule to be evaluated as true. This indicates that it is an essential rule for a safe operation. *EtN3* determines the exit from the Eco state. The example further reveals that, as expected, *NtE3* and *EtN1* had no direct influence on the control output since they were always set to true long before the control state switched.

Fig. 9 shows the raw pH difference, the ammonia signal and the final control decision when the pH difference is processed with the QSE algorithm and the rule-based control. The example shows that the algorithm can detect high and low ammonia load situations accurately. Ammonia was measured in the middle tank of the tank series, which explains the slight time shift of the control state switch compared to the ammonia signal, especially given the fact that the controller is designed to delay the increase of the oxygen setpoint until ammonia is present in the last tank. A comparison between the output of the QSE (Fig. 7) and the controller output (Fig. 9) shows the importance of the additional filtering effect induced by the rule-based controller.

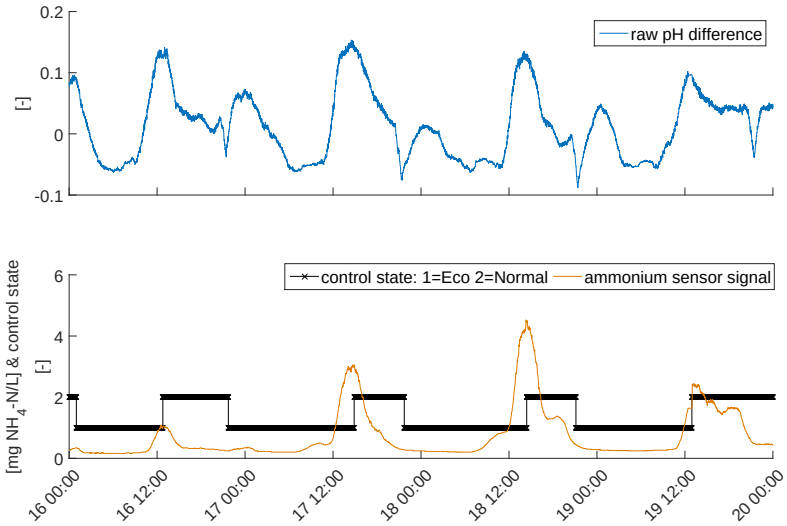

Figure 9: Final resulting control decision with ammonia signal as a reference. Top: Raw pH difference underlying the control decision. Bottom: Reference ammonia signal for validation of the controller and the final control decision 1= Eco control state (low DO setpoints), 2 = Normal control state (high DO setpoints). The control decisions made on the basis of the usage of the pH difference signals match well with the peaks in the ammonia concentration in the central aeration tank.

Internal and external safety checks. For additional operational safety, safety checks were introduced. These safety checks can overrule the rule-based control decision and will override the oxygen setpoints with the Normal DO setpoints in case one or more of the safety checks are violated. Two internal safety checks were defined for the whole cascade control: First, a lower and higher threshold was defined for both pH sensors. Outside of these thresholds, the signals are considered abnormal. Second, RITUNE® assigned an expiration time on the Eco DO setpoints which ensures that the SCADA falls back to Normal DO setpoints in case the communication between RITUNE® and the SCADA fails. Additionally, external safety checks were provided to ensure adequate nitrification capacity at all times. Very common situations in which high loads occur are wet weather events. They increase the load to the plant due to the higher inflow rates and with the first flush a lot of deposits from the sewer system reach the plant, additionally increasing the nutrient load (Larsen et al., 1998). Monitoring the inflow for high flows prepares the biological system for an increased load before it actually attains that state. The pH at the inlet of the plant was also monitored to safeguard against abnormal pH levels. Indeed, unusual pH conditions in the inflow might lead to a faulty control decision. High and low pH values can be linked to toxic wastewater compositions and the abnormal pH values themselves can push the microbiological community to its thermodynamic limits (Park et al., 2007), at which point the control should not impair the bacterial growth conditions further by also lowering DO concentration.

Upon successful evaluation of the safety checks, the desired DO setpoints are sent to the slave aeration control programmed in the basic SCADA system. Quite logically, the external safety checks are not specific to the proposed trend-based controller but can also be applied for any type of cost-reducing ammonia control.

3.3. Performance evaluation

The extensive long-term test in the WWTP Hard is discussed first. The measurement campaign took place from 25th July 2013 until 30th September 2014. The control was tested and benchmarked from 12th March 2014 until 30th September 2014. A list of the set parameters during the test period can be found in the *SI*. The control decision and the corresponding adaptation of the dissolved oxygen setpoints have an impact on the slave oxygen controller as well as the plant performance in general. First, the detailed results

from the longest test on the WWTP Hard, Winterthur are presented. The two other fully tested plants are briefly summarised. Table 4 summarises the results of all tested WWTPs.

Energy. Based on the assumptions stated in the method section, a decrease of 6.6% in the specific energy demand per kilogramme of ammonia degraded (kWh/kg$NH_4^+ - N_{\mathrm{deg}}$) was computed . This efficiency gain was even slightly higher if only dry weather flow days are considered (7.2%). This is partly caused by the decreased efficiency in the ammonia removal due to dilution and partly by the higher proportion of high DO setpoints due to the high inflow safety check during wet weather. Moreover, the two other fully tested plants PFA (13%) and PFU (8.2%) exhibited substantial energy savings (cf. Table 4).

Removal efficiency. The control system reduces the process' safety margins. Therefore, the primary aim of the WWTP - to degrade and remove compounds in the wastewater which have a negative impact on the environment - is more likely to be jeopardised. Hence, the legal effluent limits and removal efficiencies were monitored closely (Fig. 10). Neither the ammonia or nitrite effluent concentrations violated the legal limits. More information regarding the received and discharged nutrient loads can be found in the *SI*. No negative impacts on the operation of the plant were noticed (e.g., foaming due to low DO concentrations (Martins et al., 2003)). The two other plants PFA and PFU were not additionally monitored but the analysis of the regular sampling days did not unveil any legal threshold violations.

Figure 10: Ammonia and nitrite concentration measured in 24h composite samples at the WWTP Hard, Winterthur. Top: Secondary clarifier effluent. Bottom: Primary clarifier effluent. The legal threshold for ammonia in the effluent is 1 mgN/L. For nitrite there is only a non-binding guideline value of 0.3 mgN/L.

4. Discussion

The proposed controller based on the QSE algorithm showed robust behaviour during the tests. Not only were significant energy savings achieved, but all effluent limits were respected as well (cf Fig. 10 and *SI*). The first tests on multiple full-scale plants produced evidence that the control method can be applied successfully on many plants and indicates a general applicability of the control setup. A number of strengths and limitations of the control setup could already be identified.

4.1. Data sources

The kernel regression assumes equidistant spacing between all the data points. This facilitates the implementation of this step as an efficient convolution. If the data are irregularly spaced or when data

points are missing, kernel regression is still feasible yet requires an adaptive solution in which the data point weights are continuously re-evaluated. The QSE method is currently limited to the analysis of univariate signals. However, we have not yet observed an immediate need for multivariate QTA in the context of wastewater treatment automation.

4.2. Kernel regression and hidden Markov model

The two filtering steps in the QSE have different purposes: The kernel regression is the first filter which obtains the qualitative information and smooths the signal symmetrically. This filtering results in probabilities for a set of derivative signs of interest. Note that alternative filtering techniques may be used as well to compute such probabilities. Second, the HMM allows penalising (using small λ values) or even restricting (using zero-valued λ values) certain state sequences according to the prior expert knowledge about the expected signal shape characteristics (e.g., (un)likely sequences of primitives, smoothness condition). During the development and tuning phase of the pH difference based ammonia control we observed that the HMM filtering has limited effects in the studied system. Two reasons are responsible for this effect. Firstly, all λ values are identical. Secondly, all λ values are non-zero (cf. Eq. 11). Under these two conditions, the amount of filtering in the kernel regression step can be traded of against filtering by means of the HMM step. Indeed, decreasing τ and therefore decreasing the smoothing in the kernel regression can be compensated by decreasing λ. In the present study, this meant that the effect of HMM filtering was marginal. The main reason to keep the HMM step within the proposed QSE method is to ensure that incorporation of prior knowledge is feasible with the method, should this be valuable to the tentative user.

4.3. Qualitative information

One of the main advantages of the proposed method is the reduced sensor maintenance (see below). In our experience, the pH sensors that are used for control are far more robust and less prone to malfunctioning compared to ion selective ammonia sensors. Additionally, the use of qualitative information allows for a certain drift in the pH difference signal which reduces the recalibration effort compared to a controller which relies on a drift-free measurement (i.e., absolute values). Sensor drifts are compensated each time the minimal minimum of the pH difference signal is identified (see Fig. 4 point (1) and (3)). As long as the drift of the pH difference within the evaluation period of the parameter is smaller than the set value for the parameter (i.e., ΔpHd_{MID}, ΔpHd_{MAD}), the drift can be compensated. For the downward drift the problem is less severe as only the drift within the time period between the registration of the lowest minimum (pHd_{LM}) and the activation of *EtN2* (ΔpHd_{MID}) has to be compensated. The parameters for the rules *NtE2* and *EtN2*, ΔpHd_{MAD}, and ΔpHd_{MID}, define the accepted upward and downward drift rate of the signal. Their successful use relies on the frequent referencing with pHd_{LM}. Nevertheless, the benefit obtained by using the qualitative information is, theoretically, also one of the largest limitations of this type of control. As the controller is based on the detection of specific features such as minima and maxima, it is important that such dynamic features regularly appear in the analysed signal. More specifically, the relevant features should appear in shorter time-scales than the drift in the analysed signal. This caps the theoretical efficiency gains given that theoretically optimal ammonia control would eliminate all dynamics of the pH difference signal and thereby render the controller useless. Further work is aimed at quantifying to which extent the the efficiency gains are actually limited by the need for a dynamic and information-rich pH difference signal.

4.4. Implementation and fine-tuning

Experience with the applied implementation approach. At each plant the pH difference and the ammonia signal were first recorded without any active ammonia based aeration control. This dataset allowed us to evaluate the applicability of the controller at the specific plant (i.e., virtual testing). Additionally a first set of parameters fitting the controller's sensitivity to the pH difference signal was estimated with which the controller was eventually activated on the plant. The comparison of these first parameter sets showed that they are very similar across the studied plants, despite differences in treatment technology and substantial size differences (cf. *SI*).

Alternative approaches. The experience explained above indicates that one could implement such a control without a reference ammonia signal. This would facilitate the roll-out. A potential strategy could be the activation of the controller with a pre-existing conservative parameter set followed by a fine-tuning of the parameters with grab and composite samples. In this way, the use of the maintenance intensive and drifting online ammonia sensor for the implementation could be avoided.

Fine-tuning. Tuning can be delicate. Firstly, multiple tunable parameters exist within the QSE (i.e., τ and λ) and the rule-based controller (i.e, ΔpHd_{MAD}, DO setpoints, T_{MIN}). Secondly, the tuning parameters have a non-linear effect on the control decision. Hence, it is recommended to tune all parameters with care and with small and incremental changes only. Limiting the fine-tuning to either the QSE or the rule-based control can simplify the tuning. Note that the aeration system's and biological system's dead times (e.g., sliding pressure control in the aeration system) should be considered when choosing the parameters.

4.5. Energy savings and maintenance effort

The energy savings achieved are substantial, even though this study was aimed at proving a concept rather than fine-tuning the controller. The ammonia removal efficiency remains excellent. Consequently, there is potential for increased energy efficiency through even more aggressive parameter tuning.
An average weekly effort of about 10 minutes is expected to maintain a pair of pH sensors, including replacement, cleaning, and calibration. In comparison the average maintenance effort for our ISE ammonia used to validate the method presented took up to 60 minutes per week. These numbers depend on the location of the sensors, the wastewater characteristics, and the possibility of incorporating the maintenance into other regular maintenance cycles of the plant (e.g., DO sensor cleaning). For the control algorithm, an average of 10 minutes effort per week is needed to check the control output. Repetitive manual checks can be automated and linked to the alarm system, which should reduce the effort. For one controlled system this results in a total effort of 18 hours per year. Given a salary of 70 Swiss Francs (CHF)/h the operation costs sum up to roughly 1'350 Swiss Francs CHF (1'215 EUR). At the Hard WWTP (130'000 p.e.) with two separate lanes the maintenance cost is doubled to monitor both lanes (2'700 CHF/y). On the consumption side the total aeration energy saving of 140'000 kWh/y (7% of 2'000'000 kWh/y result in monetary savings of roughly 17'500 CHF/y (0.125 CHF/kWh). For smaller plants (15'000 p.e.) with 25 kWh/p.e./y as the specific aeration energy consumption this leads to savings (7%) of 3'150 CHF/y. This means that depreciation for hardware and software can be maximally 1'800 CHF/y for a smaller plant and 14'800 CHF/y for a larger plant. The current conservative estimation of 7% energy savings, the ignored indirect effects (i.e., improved nutrient removal, less external carbon dosage) and the potential to improve the controller tuning make this controller a promising alternative to the conventional measured ammonia control in particular for smaller WWTPs with limited staff capacity.

4.6. Treatment technologies

The control was initially tested on three different types of biological treatment, of which two exhibited the required qualitative relation between the pH difference and the ammonia concentration. The technology in which the control did not work was a fixed bed reactor in the WWTP Altenrhein (ALT FB, Altenrhein. Switzerland). Based on all tests we assume that the following two factors can impede the controller performance. First, a too short hydraulic residence time between the two pH sensors delivers a signal with a low signal-to-noise ratio, which in turn makes the controller more prone to noise. This was observed on the PFA plant, the ALT FB plant as well as on the conventional activated sludge plants when the hydraulic residence time between the pH sensors was reduced by measuring the difference only over one two instead of three tanks. A second factor is a time-varying relationship between airflow, DO concentration and nitrification activity: This might have caused the failure in the fixed bed system and the weaker pH difference dynamics in the hybrid moving bed system. Both systems work with a biofilm system which is less controllable in terms of nitrification performance given a bulk oxygen concentration and therefore might prevent a clear identification of ammonia peaks by means of the pH difference. Due to the higher volume specific nitrification rate in biofilm systems the hydraulic residence time is also shorter in this case. Consequentially, it

is not clear at this time which of these factors is most important. Overall, the controller showed general applicability for conventional activated sludge systems. Modifications of the control concept are considered to address the observed limitations.

4.7. Future work

The findings of this study are promising. Further research regarding the use of QSE in control loops should be directed towards handling multivariate data sets. Despite the limited applicability in the wastewater field such a development would allow the use of QSE in fields with a higher prevalence of multivariate control loops (e.g., Multiple Input Multiple Output or Multiple Input Single Output). In its current form the QSE algorithm can also be applied in any other field (e.g., biotechnology, chemical engineering) where sensors are prone to drift and the signal represents information in its derivatives. In the wastewater field redox potential sensors are common and exhibit the described sensor characteristics. The presented controller could be improved by the identification of intermediate ammonia load levels - either with quantitative rules or by taking both extrema and inflection points into account.

Another modification of the tested state estimation and control approach is to modify the output of the HMM-based state estimation so that a continuous probability for the state is computed instead of a discrete maximum likelihood estimate. This would allow using continuous controllers (e.g., PID controller, Linear-quadratic regulator) instead of discrete ones (e.g., rule-based). However, the continuous controller output bears two challenges: *(i)* there needs to be a sufficiently high confidence in the continuous information, which is not the case in our study. In the presented case only the confidence in the sign of the first derivative - not its value - is sufficiently high. *(ii)* continuous feedback control may dampen the dynamics of the monitored signal, which is critical for application of QTA (cf. section 4.3).

For a more comprehensive assessment of the rule-based controller some open questions remain: First, a direct comparison with a conventional ammonia sensor-based control would allow the life-cycle costs and benefits of both options to be assessed. Second, the controller was tested only on plants where the nitrification capacity was reduced by decreasing the volume-specific nitrification rate. However, adaptation of the aerobic volume could be considered as well. The constant DO setpoint in such systems and the reduced residence time between the pH sensors presumably require dramatic changes in the control structure. Eventually, a robust method to introduce and auto-tune the controller without the need for an ammonia ISE sensor would make this controller even more attractive.

5. Conclusion

A new algorithm for qualitative trend analysis (QTA) has been proposed and evaluated. In the example presented, a pH difference measured with two pH sensors over an aerobic activated sludge tank series was considered as the raw signal. These are the conclusions:

- The QSE algorithm is a fast and robust QTA method to extract qualitative information from a continuous signal.
- Using QSE for control is especially interesting in cases where sensor signals exhibit well-recognised dynamics but lack accuracy due to drift and other incipient faults. With QSE, aeration control based on an ammonia soft-sensor using only measurement of two pH levels is shown to be feasible. The qualitative analysis is especially valuable as it enables automatic compensation of drift in the pH sensors.
- The automatic drift compensation and the more robust nature of the control concept reduces the sensors' maintenance needs and makes it more attractive for smaller and remote plants to introduce ammonia control.
- The soft-sensor exhibits general applicability mainly because of its mechanistic nature. This makes the controller transparent and allows operators to retrace and inspect the control decision so to critically evaluate the performance of the proposed controller.

- The controller is applicable in continuous activated sludge and continuous hybrid moving bed reactors with sufficient short-term variations in the ammonia influent concentration. So far this kind of control does not appear viable for fix bed reactors. This limitation originate from the need for sufficiently dynamic and thereby information-rich signals. Too short hydraulic residence time and/or weak correlation of airflow, DO concentration and nitrification rate hamper the application of the controller.

6. Acknowledgements

This work was supported by the Commission for Technology and Innovation (CTI) of the Swiss Federal Department of Economic Affairs Education and Research (EAER). (CTI project no. 14351.1 PFIW-IW). The authors thank Prof. Dr. Eberhard Morgenroth (Eawag, Switzerland) and Prof. Dr. Hansruedi Siegrist (Eawag, Switzerland) for their helpful discussions leading to this article and the staff of the participating WWTPs for their support in the field campaigns.

References

Åmand, L., Olsson, G., Carlsson, B., 2013. Aeration control – a review. Water Science & Technology 67, 2374–2398. doi:10.2166/wst.2013.139.

Andersson, L.G., 1980. Energy savings at wastewater treatment plants, in: Strub, A.S., Ehringer, H. (Eds.), New ways to save energy. Springer Netherlands, Dordrecht, pp. 741–749.

Bakshi, B., Stephanopoulos, G., 1994. Representation of process trends—IV. Induction of real-time patterns from operating data for diagnosis and supervisory control. Computers & Chemical Engineering 18, 303–332. doi:10.1016/0098-1354(94)85029-1.

Charbonnier, S., Gentil, S., 2007. A trend-based alarm system to improve patient monitoring in intensive care units. Control Engineering Practice 15, 1039–1050. doi:10.1016/j.conengprac.2006.12.005.

Cheung, J.Y., Stephanopoulos, G., 1990. Representation of process trends Part I. A formal representation framework. Computers & Chemical Engineering 14, 495–510. doi:10.1016/0098-1354(90)87023-I.

Derlon, N., Thürlimann, C., Dürrenmatt, D., Villez, K., 2017. Batch settling curve registration via image data modeling. Water Research 114, 327–337. doi:10.1016/j.watres.2017.01.049.

Dürrenmatt, D., Gujer, W., 2012. Data-driven modeling approaches to support wastewater treatment plant operation. Environmental Modelling & Software 30, 47–56. doi:10.1016/j.envsoft.2011.11.007.

Haimi, H., Mulas, M., Corona, F., Vahala, R., 2013. Data-derived soft-sensors for biological wastewater treatment plants: An overview. Environmental Modelling & Software 47, 88–107. doi:10.1016/j.envsoft.2013.05.009.

Hastie, T., Tibshirani, R., Friedman, J., 2009. The Elements of Statistical Learning: Data Mining, Inference, and Prediction, Second Edition. Springer Science & Business Media.

Ingildsen, P., 2002. Realising full-scale control in wastewater treatment systems using in situ nutrient sensors. Ph.D. thesis. Dept. of Industrial Electrical Engineering and Automation [Institutionen för industriell elektronik och automation], Univ.. Lund.

Joss, A., Salzgeber, D., Eugster, J., König, R., Rottermann, K., Burger, S., Fabijan, P., Leumann, S., Mohn, J., Siegrist, H., 2009. Full-scale nitrogen removal from digester liquid with partial nitritation and anammox in one SBR. Environmental Science & Technology 43, 5301–5306. doi:10.1021/es900107w.

Kaelin, D., Rieger, L., Eugster, J., Rottermann, K., Bänninger, C., Siegrist, H., 2008. Potential of in-situ sensors with ion-selective electrodes for aeration control at wastewater treatment plants. Water Science & Technology 58, 629–637. doi:10.2166/wst.2008.433.

Konstantinov, K.B., Yoshida, T., 1992. Real-time qualitative analysis of the temporal shapes of(bio) process variables. AIChE Journal 38, 1703–1715. doi:10.1002/aic.690381104.

Larsen, T., Broch, K., Andersen, M.R., 1998. First flush effects in an urban catchment area in Aalborg. Water Science & Technology 37, 251–257.

Martins, A.M.P., Heijnen, J.J., van Loosdrecht, M.C.M., 2003. Effect of dissolved oxygen concentration on sludge settleability. Applied Microbiology & Biotechnology 62, 586–593. doi:10.1007/s00253-003-1384-6.

Maurya, M.R., Rengaswamy, R., Venkatasubramanian, V., 2007. Fault diagnosis using dynamic trend analysis: A review and recent developments. Engineering Applications of Artificial Intelligence 20, 133–146. doi:10.1016/j.engappai.2006.06.020.

Mašić, A., Srinivasan, S., Billeter, J., Bonvin, D., Villez, K., 2017. Shape constrained splines as transparent black-box models for bioprocess modeling. Computers & Chemical Engineering 99, 96–105.

Olsson, G., 2012. ICA and me – A subjective review. Water Research 46, 1585–1624. doi:10.1016/j.watres.2011.12.054.

Park, S., Bae, W., Chung, J., Baek, S.C., 2007. Empirical model of the pH dependence of the maximum specific nitrification rate. Process Biochemistry 42, 1671–1676. doi:10.1016/j.procbio.2007.09.010.

Peng, Y.Z., Gao, J.F., Wang, S.Y., Sui, M.H., 2002. Use pH and ORP as fuzzy control parameters of denitrification in SBR process. Water Science & Technology 46, 131–137.

Rengaswamy, R., Venkatasubramanian, V., 1995. A syntactic pattern-recognition approach for process monitoring and fault diagnosis. Engineering Applications of Artificial Intelligence 8, 35–51. doi:10.1016/0952-1976(94)00058-U.

Rieger, L., Jones, R.M., Dold, P.L., Bott, C.B., 2014. Ammonia-based feedforward and feedback aeration control in activated sludge processes. Water Environment Research 86, 63–73. doi:10.2175/106143013X13596524516987.

Rosso, D., Stenstrom, M.K., Larson, L.E., 2008. Aeration of large-scale municipal wastewater treatment plants: state of the art. Water Science & Technology 57, 973–978. doi:10.2166/wst.2008.218.

Ruano, M., Ribes, J., Seco, A., Ferrer, J., 2009. Low cost-sensors as a real alternative to on-line nitrogen analysers in continuous systems. Water Science & Technology 60, 3261–3268. doi:10.2166/wst.2009.607.

Ruano, M., Ribes, J., Seco, A., Ferrer, J., 2012. An advanced control strategy for biological nutrient removal in continuous systems based on pH and ORP sensors. Chemical Engineering Journal 183, 212–221. doi:10.1016/j.cej.2011.12.064.

Scali, C., Ghelardoni, C., 2008. An improved qualitative shape analysis technique for automatic detection of valve stiction in flow control loops. Control Engineering Practice 16, 1501–1508. doi:10.1016/j.conengprac.2008.04.009.

Thürlimann, C.M., Villez, K., 2017. Input estimation as a qualitative trend analysis problem. Computers & Chemical Engineering doi:10.1016/j.compchemeng.2017.04.011. in press.

Villez, K., 2015. Qualitative path estimation: A fast and reliable algorithm for qualitative trend analysis. AIChE Journal 61, 1535–1546. doi:10.1002/aic.14736.

Villez, K., Habermacher, J., 2016. Shape anomaly detection for process monitoring of a sequencing batch reactor. Computers & Chemical Engineering 91, 365–379. doi:10.1016/j.compchemeng.2016.04.012.

Villez, K., Rengaswamy, R., 2013. A generative approach to qualitative trend analysis for batch process fault diagnosis, in: Proceedings of the European Control Conference 2013 (ECC13), pp. 1958–1963.

Villez, K., Rosén, C., Anctil, F., Duchesne, C., Vanrolleghem, P.A., 2008. Qualitative representation of trends: an alternative approach to process diagnosis and control. Water Science & Technology 57, 1525. doi:10.2166/wst.2008.141.

Viterbi, A., 1967. Error bounds for convolutional codes and an asymptotically optimum decoding algorithm. IEEE Transactions on Information Theory 13, 260–269. doi:10.1109/TIT.1967.1054010.

Zanetti, L., Frison, N., Nota, E., Tomizioli, M., Bolzonella, D., Fatone, F., 2012. Progress in real-time control applied to biological nitrogen removal from wastewater. A short-review. Desalination 286, 1–7. doi:10.1016/j.desal.2011.11.056.