

Fault detection for control of wastewater treatment plants

O. Schraa, B. Tole and J.B. Copp

Hydromantis Inc., Consulting Engineers, 1685 Main St. West, Suite 302, Hamilton, Ontario, L8S 1G5, Canada (E-mail: copp@hydromantis.com)

Abstract Interest in real-time model-based control is increasing as more and more facilities are being asked to meet stricter effluent requirements while at the same time minimizing costs. It has been identified that biological process models and automated process control technologies are being used at wastewater treatment plants throughout the world and that great potential for optimising biotreatment may exist with the integration of these two technology areas. According to our experience, wastewater treatment plants are indeed looking for ways to successfully integrate their modelling knowledge into their process control structure; however, there are practical aspects of this integration that must be addressed if the benefits of this integration are to be realised. This paper discusses the practical aspects of monitoring, filtering and analysing real sensor data with an aim to produce a reliable real-time data stream that might be used within a model-based control structure. Several real case study examples are briefly discussed in this paper.

Keywords Control; fault detection; modelling; sensors

Introduction

Several things have contributed to the limited application of advanced model-based control at wastewater treatment facilities. First and foremost is reliability (Hill *et al.*, 2002). Operations staff must have confidence in the system and the system must be both reliable and robust in that alarm signals must be kept to a minimum, and those alarms must be useful. Historically, a lack of confidence may have been the result of sensor reliability and, maintenance requirements and as a result, operator reluctance to trust sensors. Sensor technology has improved significantly making the timing of on-line implementation pertinent. Operators are often interested in advanced control systems, but demonstrating the benefits and minimising the problems is crucial. Simulations have been used to predict the benefits of proposed control strategies using off-line historical data. This has worked well and gives a relative indication of the benefits if the plant had been operating with the strategy; however, because the simulations are based on historical data, there is no guarantee that things have not changed.

Although process control is widely implemented without process models, the use of process models for optimisation is becoming increasingly popular. That is, a calibrated process model can be used to identify process inefficiencies including the impact of over-aeration, sub-optimal pump rates and the impact of specific control loops on the process behaviour. This ability to predict the performance of a system results in better process understanding and enables a more accurate accounting of possible control options. The need for a process model is particularly important if a plant wishes to investigate multi-variable control.

One of the most advanced implementations of on-line model-based control has been the development of the Integrated Computer Control System (IC²S) by Hydromantis Inc. (Belia and Takacs, 2002; Takacs *et al.*, 1998). IC²S was based on the premise that an

integrated computer-based approach to wastewater treatment plant operation and control can have a significant impact on the performance of a plant, including:

- reduction in the duration and frequency of water quality effluent excursions
- reduction in energy costs
- deferred capital expenditures
- optimal use of existing facilities
- ability to cope with unusual plant operation conditions
- reduction in the overall pollutant loadings to receiving water bodies.

Although the theoretical benefits of real-time model-based control are evident, the practical aspects of implementing such a system must be addressed. That is, in an ideal world, the data fed to the model will be correct without error, but in reality this is not possible so care must be taken to ensure that a reliable data stream is used. This paper focuses on the practical aspects of dealing with sensor data and discusses issues that can occur with real sensor data and other control issues that can be overlooked.

Fault detection

Successfully implementing process control requires a reliable stream of data. As a result, sensor and process fault detection algorithms can be an important aspect of a control system. There are up to three elements to such a fault detection module. Preliminary sensor fault detection is performed on monitored signals, intermediate fault detection compares filtered and measured outputs while advanced fault detection tracks process trends.

Preliminary sensor fault detection

Sensor signals are subject to a number of possible errors including, but not limited to: noise, drift, catastrophic failure, power outages and transmission problems. Low-level sensor fault detection is essential so that the data stream fed to the control system is a reasonable estimate of the current state of the plant. The objective is to process the raw signals and this may involve some or all of the following:

- ensuring measurements are within the 4 to 20 mA range
- detecting constant signals
- detecting sudden changes
- dealing with missing measurements (catastrophic sensor failure)
- filtering of the measurements.

Sensor signals operate in the range of 4 to 20 mA, so detection of signals outside this typical range is important. For example, a signal outside this range might indicate a sensor malfunction, a power failure or a transmission failure. Detection of these anomalies is crucial to prevent potential process upsets due to incorrect control actions. The constantly changing conditions in a wastewater treatment plant mean that sensor signals are expected to change continuously. Hence, the detection of a constant signal might also indicate a problem such as a sensor malfunction, a dirty sensor or a sensor that is off-line. Large instantaneous changes in sensor signals can occur and might be real depending on the sensor, but this type of change might also indicate a process disturbance, an unreliable measurement or an abrupt sensor failure.

Another step in the processing involves filtering the raw signals because often measurements contain high-frequency components caused by electrical interference, measurement noise and uncontrollable process disturbances. If the control system acts on these high-frequency components, the performance may be undesirable and therefore these components should be removed from the measured signal.

It is important to examine historical measurement data to determine how much filtering is required. Different measurements can have varying degrees of noise depending on

the variable being measured, the type of sensor being used and the location of the sensor in the process. In addition, filtering should consider the dynamics of the process being controlled. Because the goal of filtering is to remove high frequency noise from the measurements, a first-order digital low pass filter as shown below might be used. In this case, the digital filter calculates an exponentially weighted moving average (EWMA) and might be implemented as follows:

$$X_{filtered,n} = \alpha * X_{filtered,n-1} + (1 - \alpha) * X_n$$

where $X_{filtered,n}$ = filtered sensor value at time, n ; $X_{filtered,n-1}$ = filtered sensor value at time, $n - 1$; α = filter constant; and X_n = sensor value at time, n .

In this case, the degree of filtering increases as α is increased. Sensors that produce inherently noisy signals will require more filtering. The value of the low-pass filter constant should be selected to achieve a balance between noise attenuation and dynamic response. Suitable values of the filter constant(s) can be determined by experimenting with different filter constants and plotting the filtered response, or using more advanced techniques such as frequency response analysis.

As an example of filtering, consider the total suspended solids (TSS) measurements from a wastewater treatment plant as shown in Figure 1a. The filter constant used was 0.9 (in this case corresponding to a filter time constant of 135 min). The filtered response follows the major trends in the data but removes noise including large instantaneous spikes without a large time delay. If these data were being used for mixed liquor suspended solids (MLSS) or sludge retention time (SRT) control, more filtering might be beneficial. A control system based on manipulating the waste flow could not, and typically should not, be used to mitigate the diurnal variations that occur during each day. If a filter constant of 0.99 is used most of the diurnal variations that cannot be controlled are removed as shown in Figure 1b.

In some cases, filtering may cause problems. As shown in Figure 2, with an intermittent flow rate and a moderate filter constant of 0.7, the filtered response ceases to accurately track the signal. In this case, the signal does not contain much noise and a better approach might be to perform a minimal amount of filtering (filter constant of 0.1 for example) or no filtering at all.

The low-pass filter discussed above is one of many possible digital filter equations. Other higher-order filter equations can be generated and their filter constants can be designed appropriately.

Intermediate fault detection

The filtering step can be used for determining other sensor problems. In addition to filtering the raw signals, the filtered and unfiltered sensor signals can be analysed to determine




Figure 1 Filtering of TSS data with filter constant of (a) 0.9 and (b) 0.99




Figure 2 Filtering of flow rate data with filter constant of 0.7

sensor or process faults. This can include checking if signals are outside physically realistic boundaries, tracking changes in the variance of the measurement noise and identifying incompatible sensor results (e.g. ammonia > TN).

Detection of signals outside a physically realistic boundary is relatively straightforward, but there are less obvious signal limits that need to be considered. These bounds should be the extreme operating conditions expected for each sensor signal. In this case, it is best to monitor the filtered signals as opposed to the raw signals as the raw signals (especially noisy signals) might be prone to large instantaneous upward or downward spikes that may extend outside normal operating ranges. Monitoring the filtered signals is a better indication of persistent problems.

The variance in the noise of a signal is expected to be reasonably constant and, as such, if the noise level of a sensor signal changes, it may indicate that the sensor is dirty, malfunctioning or could be an indication of a process related disturbance. This fault can be identified by calculating the variance in the residuals between the filtered and measured values. The calculated variance over a period of time is then compared to the variance calculated over a historical time period and a warning is triggered if the difference is substantial. The variance changes that can occur in a measurement signal are illustrated in [Figure 3a](#), which shows another example of TSS measurements from a wastewater treatment plant. As shown, the variability in the measurements starts to increase steadily after 2 days. If the data is filtered with a filter constant of 0.9 and the filter residuals are plotted, as shown in [Figure 3b](#), it is apparent that the variability of the residuals steadily increases, possibly indicating that the sensor is becoming dirty.

Incompatible sensor signals should be detected where possible. Redundant signals should be similar and related signals should make sense. For example, if ammonia and TN sensors are located in the same location, then the signals from these sensors must show a lower ammonia concentration.




Figure 3 TSS sensor data with increasing variability (a) and the relative filter residuals for the same data (b)

Detection of a problem is essential but it is also necessary to activate a default control action until the problem is investigated. That is, if a fault is suspected, then a reasonable response would be to revert any control action involving the suspect signal to a default or manual operation mode until the source of the detected problem is investigated. A summary of fault detection techniques from this paper is given in [Table 1](#).

The above techniques are simple sensor fault detection techniques that are easy to implement and often found in commercial process controllers. More sophisticated statistical techniques, such as data reconciliation, can be used to identify persistent sensor biases not easily detected by simpler methods. Data reconciliation is a technique used to adjust process measurements so that they are consistent with known process constraints such as mass balances. Detailed discussions of data reconciliation can be found elsewhere ([Crowe 1996](#); [Romagnoli and Sánchez, 2000](#)).

Advanced fault detection

Advanced fault detection can be used to detect process upsets and disturbances. As with sensor fault detection, process fault detection can range from simple tests to sophisticated statistical analyses. A simple fault detection system could involve checking for trends in the data using the Mann–Kendall test ([Olsson and Newell, 1999](#)). If a trend is detected, a curve could be fit to the data and the model parameters could be compared to typical values for the process when it is relatively stationary. More advanced process fault detection techniques can involve the use of the data reconciliation techniques discussed earlier or statistical process control charts. Statistical process control techniques typically involve tracking or monitoring process variables (or parameters derived from process data) over time using statistical control charts. The variables of interest are charted over time and compared to control limits to determine if the process is within control.

Because wastewater treatment plant measurements can exhibit autocorrelation, seasonality and non-constant variance ([Berthouex, 1989](#)), it can be difficult to apply traditional control charts such as Shewhart or cumulative sum (CUSUM) charts. In this context, it is possible to fit a time-series model such as an auto-regressive integrated moving average (ARIMA) model to normal operating data and then use the model as a charting tool ([Berthouex, 1989](#)). The model would be used to continually predict process data, given the previous data, and the difference between these predictions and the actual measurements (i.e. residuals) which can be plotted on a conventional control chart. When the process is operating normally, the residuals will be independent, random and have constant variance.

A simpler alternative is to construct EWMA charts. The EWMA statistic is the optimal one-step-ahead forecast for the ARIMA(0,1,1) model ([Montgomery and Mastrangelo, 1991](#)):

Table 1 Summary of simple sensor fault detection methods

Type of fault	Fault detection method
Measurement outside 4 to 20 mA range	Test whether signal is less than 4 mA or greater than 20 mA
Violation of user-defined normal bounds on measurement	Test whether filtered signal is outside user-defined bounds
Constant measurement	Test whether signal has been constant over an extended period
Large instantaneous change in measurement	Determine if absolute value of residual is reasonable
Change in noise variance	Determine if current variance is significantly different from historical variance
Sensor results incompatible	Comparison of redundant signals, and/or the calculated values of related sensors (i.e. $TN > ammonia$)

$$Y_t = Y_{t-1} + \varepsilon_t - \theta \varepsilon_{t-1}$$

where $\alpha = (1 - \theta)$; Y_t = observation from the modelled process at sample time, t ; ε_t = independently and identically distributed random variable at sample time, t ; θ = moving average parameter.

Therefore, the one-step-ahead prediction errors (i.e. residuals between predictions and actual measurements) for an ARIMA(0,1,1) process, calculated using the EWMA statistic, can be plotted on a traditional control chart. As discussed by [Montgomery and Mastrangelo \(1991\)](#), the EWMA approach can be a reasonable approximation of the ARIMA model approach in many cases. For a suitably selected value of α , the EWMA statistic is an excellent one-step-ahead predictor for processes where the mean does not shift too rapidly and the observations are positively autocorrelated. The EWMA statistic is discussed by [Hunter \(1986\)](#).

The above discussion of control charts considers the univariate case. In the multivariate case, where it is desirable to simultaneously monitor many variables from the same process, the above procedures can be used and Hotelling T^2 charts can be constructed for several related variables. Some researchers have looked at the use of multivariate statistical techniques, such as principal components analysis (PCA) to monitor wastewater treatment data. PCA involves the singular value decomposition of the measurement covariance matrix in order to project the measurement space into a lower dimensional space that is easier to visualise. The variables in the lower dimensional space explain the majority of the variance in the process variables. Although PCA considers static covariance relationships, it can be adapted to the analysis of dynamic data. [Lennox and Rosen \(2002\)](#) used adaptive multiscale PCA for online monitoring of wastewater treatment data.

Practical examples

Sensor faults

As an example of a malfunctioning sensor, consider the TSS data shown in [Figure 3](#). As shown, there are extended periods where the signal is constant indicating a problem such as sensor failure, a dirty sensor or that the sensor has been switched out of data collection mode. There is also a period where the signal becomes extremely noisy with some large instantaneous spikes. Again this suggests that the sensor is malfunctioning or dirty.

Because visual inspection of data is usually not feasible, it is important to provide automated procedures to identify sensor faults. Tests to check for constant signals and variance changes could be used to identify faults such as those shown in [Figure 4](#).




Figure 4 TSS sensor data from a dirty or malfunctioning sensor

Handling of unexpected power interruptions

Power interruptions are another potential fault that can cause problems for controllers. As with sensor faults, they compromise the information used to calculate control actions. The handling of controller startup or initialization after power interruptions is straightforward in cases where the controlled variable is measured directly but needs careful consideration for controllers that use measurements to infer the value of the controlled variable such as in SRT control. If the controller is a remote unit located in the field and an unexpected power interruption occurs, it is necessary to ensure that the inferred variable is properly re-initialized at startup. This may entail archiving past measured values or SRT values so that a reasonable SRT estimate is available for re-initialization of the controller. As an example, consider an averaged SRT tracked at an activated sludge plant as shown in Figure 5. If the averaged SRT is re-initialized using the current information after an unexpected power interruption the calculation will require considerable time to move towards the true SRT. If past SRT values are archived and used for re-initialization, the SRT calculation will take less time to resume tracking the true SRT.




Figure 5 Issue with re-initialising SRT calculation after a power interruption




Figure 6 Integral windup when supervisory control system does not implement control actions from lower-level controller

Controller windup

Integral windup is another potential issue for controllers and requires mitigation in all controllers that have integral action. The issue of controller windup is well understood for the situation where the controller cannot compensate for disturbances because of limitations in the final control element. It can also occur in cases where a lower-level controller sends its output to a supervisory control system. If the supervisory control system does not implement the control action from the lower-level controller, the potential for integral windup also exists, as shown in Figure 6. In this case, when the lower-level controller regains control of the manipulated variable there is a potential for a large instantaneous change to occur in the manipulated variable as it is adjusted to match the controller output. It is important to ensure that the controller does not accumulate differences between the calculated control action and the control action actually implemented. Therefore, measurement and archiving of the manipulated variable is important.

Conclusion

Real-time control is dependent on sensor data and even though significant improvements have been made in wastewater treatment sensor technology, sensor signals are still subject to a number of possible errors including, noise, drift, catastrophic failure, power outages and transmission problems. Hence for model-based on-line control that attempts to deduce many aspects of the process from a minimum number of samples, there still exists a need to ensure that the sensor signals are accurate and reflect the true state of the process so that the control scheme is working on the best possible data. Signal processing is complicated by the various potential problems and, as such, considerable care must be taken when designing appropriate algorithms as each signal will have independent characteristics that have to be taken into consideration.

References

- Belia, E. and Takacs, I. (2002). *Integrated On-Line Systems in Wastewater Treatment*, 2002 Borchard Conference. Ann Arbor, Michigan, USA.
- Berthouex, P.M. (1989). Constructing control charts for wastewater treatment plant operation. *J. WPCF*, **61**(9), 1534–1551.
- Crowe, C.M. (1996). Data reconciliation – progress and challenges. *J. Proc. Control*, **6**, 89–98.
- Hill, R.D., Manross, R.C., Davidson, E., Palmer, T., Ross, M.C. and Nutt, S. (2002). *Sensing and Control Systems: A Review of Municipal and Industrial Experiences*, WERF Report 99WWF4, (175 pages), IWA Publishing, London UK.
- Hunter, J.S. (1986). The exponentially weighted moving average. *J. Qual. Tech.*, **18**, 203–209.
- Lennox, J. and Rosen, C. (2002). Adaptive multiscale principal components analysis for online monitoring of wastewater treatment. *Wat. Sci. Tech.*, **45**(4–5), 227–235.
- Montgomery, D.C. and Mastrangelo, C.M. (1991). Some statistical process control methods for autocorrelated data. *J. Qual. Tech.*, **23**(3), 179–193.
- Olsson, G. and Newell, B. (1999). *Wastewater Treatment Systems – Modelling, Diagnosis, and Control*, IWA Publishing, London, UK.
- Romagnoli, J.A. and Sánchez, M.C. (2000). *Data Processing and Reconciliation for Chemical Process Operations*, Academic Press, London, UK.
- Takacs, I., Patry, G., Gall, B. and Patry, J. (1998). The integrated computer control system: a comprehensive, model-based control technology. *Dynamics and Control of Wastewater Systems*, Barnett, M., Stenstrom, M. and Andrews, J. (eds), Technomic Publishing Co. Inc., Lancaster Pennsylvania.