

Technical report 07-036

Cascaded parameter estimation for a water treatment plant using particle filters*

Zs. Lendek, K.M. van Schagen, R. Babuška, A. Veersma, and
B. De Schutter

If you want to cite this report, please use the following reference instead:

Zs. Lendek, K.M. van Schagen, R. Babuška, A. Veersma, and B. De Schutter, “Cascaded parameter estimation for a water treatment plant using particle filters,” *Proceedings of the 17th IFAC World Congress*, Seoul, Korea, pp. 10857–10862, July 2008.

Delft Center for Systems and Control
Delft University of Technology
Mekelweg 2, 2628 CD Delft
The Netherlands
phone: +31-15-278.24.73 (secretary)
URL: <https://www.dcsc.tudelft.nl>

*This report can also be downloaded via https://pub.deschutter.info/abs/07_036.html

Cascaded parameter estimation for a water treatment plant using particle filters [★]

Zs. Lendek ^{*} K.M. van Schagen ^{***} R. Babuška ^{*}
A. Veersma ^{***} B. De Schutter ^{*,****}

^{*} *Delft Center for Systems and Control, Delft University of Technology, Mekelweg 2, 2628 CD Delft, The Netherlands (email: {z.lendek, r.babuska}@tudelft.nl)*

^{**} *DHV BV, PO Box 1132, 3800 BC Amersfoort, The Netherlands (email: kim.vanschagen@dhv.com)*

^{***} *Waternet, PO Box 8169, 1005 AD Amsterdam, The Netherlands (email: alex.veersma@waternet.nl)*

^{****} *Marine and Transport Technology Department, Delft University of Technology, Mekelweg 2, 2628 CD Delft, The Netherlands (email: b@deschutter.info)*

Abstract: Advanced online control of drinking water treatment plants requires reliable models. These models in general involve temperature-dependent, uncertain parameters, which can only be measured in laboratory conditions. We propose to estimate these parameters online, using the available pH quality measurements. Since the pH measurements are a nonlinear combination of the system's states, a particle filter is used. Thanks to the cascaded nature of the plant, the estimation is also performed in a cascaded setting. The performance is evaluated both for simulated and real-world data. Results indicate that the filter can be effectively used to improve the model accuracy.

1. INTRODUCTION

In drinking water production and distribution, there is an increased interest in advanced control using online flow and level measurements. However, in the current practice advanced control methods are mainly used for water quantity control. The increased use of these measurements has led to the optimization of the quantitative aspect of production and distribution [Bakker et al., 2003; Hill et al., 2005]. For instance, in the Netherlands, about 25% of the drinking water production and distribution locations use advanced control to optimize the production capacity and storage.

Advanced semi-online water quality measurements have also become more common, although, currently they are predominantly used for monitoring. Such measurements include for instance pH and UV spectra measurements.

Previous research [van Schagen et al., 2006] has shown that before implementing advanced control strategies in drinking water production, it is recommended to investigate the trade-off between the number of necessary measurements, the accuracy of the measurements, and the effort for maintaining the measurement devices. The sensors must be able to perform under industrial circumstances, with very small

variations in water quality. If new measurements that do not meet the above conditions are used in online control, they can worsen the performance of the process.

The measure to describe the super-saturation of calcium carbonate in the water is the Saturation Index (SI), which is defined as the pH offset at which the actual calcium concentration is in equilibrium with the carbonate concentration (see van Schagen et al. [2007] and the references therein). An SI below zero will cause the concrete of the installation to dissolve. An SI above 0.3 will cause scaling on the equipment of the installation, causing malfunction of the valves and dosing units. The difference between a high and low SI therefore about 0.3. The accuracy of the pH measurement is restricted and the process conditions change in time.

The research reported here involves one of the water treatment plants (WTP) of Amsterdam, WTP Weesperkarspel. Together with the WTP Leiduin it produces all the drinking water for the city of Amsterdam (400Ml/day). The Weesperkarspel plant treats seepage water from a polder with eight process steps (stages) in cascade: coagulation, 100 days retention in a lake water reservoir, acid dosage (HCl), rapid sand filtration (RSF), ozone, pellet softening, biologically activated carbon, and slow sand filtration.

In the current situation, the SI is controlled using only pH measurements. Due to the setup of the treatment plant (with a long retention time in the lake), the water quality parameters change slowly. However, the reaction dynamics and model characteristics in each process step depend non-linearly on the temperature, which is reflected in uncertain model parameters. To improve SI control a

[★] This research is part of the PROMICIT project of Waternet, Delft University of Technology, DHV B.V. and ABB B.V. The project is sponsored by SenterNovem, Ministry of Economic Affairs of the Netherlands. Part of the research is also sponsored by Senter, Ministry of Economic Affairs of the Netherlands within the project Interactive Collaborative Information Systems (grant no. BSIK03024).

reliable dynamic model should be available. This research aims to improve the model characteristics based on the pH measurements available in the process. In order to improve the model, we consider the online estimation of these uncertain parameters for two stages: acid dosage and RSF.

The most well-known and widely used probabilistic estimation methods are the Kalman filter and its extensions to nonlinear systems [Kalman, 1960; Welch and Bishop, 2002]. However, these filters have severe limitations and may become unstable even for linear processes. Over the last years, particle filters [Doucet et al., 2000; Arulampalam et al., 2002] have been extensively studied. These filters have been successfully applied to state-estimation problems [Nait-Charif and McKenna, 2004; Li et al., 2007], and allow to handle nonlinear, non-Gaussian dynamic systems.

We propose a method for the online estimation of the uncertain reaction constants of the acid dosage and the RSF stage, in order to improve the model characteristics. Due to the nonlinear and uncertain nature of the process, we use particle filters. Since the process steps are in a cascade, the parameter estimation is performed in a cascaded manner.

In Section 2, the cascaded particle filter methodology is reviewed. Section 3 presents the water treatment plant and the dynamic model used to represent the stages of interest. Section 4 gives the results for simulated and experimental data, respectively. Finally, Section 5 concludes the paper.

2. CASCADED PARTICLE FILTERS

For arbitrary distributions or nonlinear processes, representing the random variables by their mean and covariance only, as Kalman filters do, is not sufficient for a reliable estimation. Moreover, there is no general method to compute the resulting distribution analytically. Particle filters approximate the distributions by samples, rather than by a compact parametric form, which leads to better performance. The price that has to be paid are the higher computational costs.

2.1 Particle Filters

The particle filter (PF) uses probabilistic models for the state transition function and the measurement function, respectively [Doucet et al., 2000]:

$$p(x_k|x_{k-1}), \quad p(y_k|x_k).$$

The objective is to recursively construct the posterior probability density function (PDF) $p(x_k|y_k)$ of the state, given the measured output y_k and assuming conditional independence of the measurement sequence, given the states. The PF works in two stages:

- (1) The *prediction stage* uses the state-transition model to predict the state PDF one step ahead. The PDF obtained is called the *prior*.
- (2) The *update stage* uses the current measurement to correct the prior via the Bayes rule. The PDF obtained after the update is called the *posterior*.

Particle filters represent the PDF by N random samples (particles) x_k^i with their associated weights w_k^i , normalized

so that $\sum_{i=1}^N w_k^i = 1$. At time instant k , the posterior obtained in the previous step, $p(x_{k-1}|y_{k-1})$, is represented by N samples x_{k-1}^i and the corresponding weights w_{k-1}^i . To approximate the posterior $p(x_k|y_k)$, new samples x_k^i and weights w_k^i are generated. Samples x_k^i are drawn from a (chosen) *importance density function* $q(x_k^i|x_{k-1}^i, y_k)$, and the weights are updated, using the current measurement y_k

$$\tilde{w}_k^i = w_{k-1}^i \frac{p(y_k|x_k^i) p(x_k^i|x_{k-1}^i)}{q(x_k^i|x_{k-1}^i, y_k)} \quad (1)$$

and normalized

$$w_k^i = \frac{\tilde{w}_k^i}{\sum_{j=1}^N \tilde{w}_k^j}.$$

If the importance density $q(x_k|x_{k-1}, y_k)$ is chosen equal to the state-transition PDF $p(x_k|x_{k-1})$, the weight update equation (1) becomes:

$$\tilde{w}_k^i = w_{k-1}^i p(y_k|x_k^i).$$

The use of the transition prior as the importance density is a common choice [Arulampalam et al., 2002] and it has the advantage that it can be easily sampled and the weights are easily evaluated. In practice, this effectively means generating samples of the noise and using these samples in the state transition functions to generate the state samples.

The posterior PDF is represented by the set of weighted samples, conventionally denoted by:

$$p(x_k|y_k) \approx \sum_{i=1}^N w_k^i \delta(x_k - x_k^i)$$

where δ is the Dirac delta measure.

The PF algorithm is summarized in Algorithm 1. A common problem of PF is the particle degeneracy: after several iterations, all but one particle will have negligible weights. Therefore, particles must be resampled. A standard measure of the degeneracy is the effective sample size:

$$N_{\text{eff}} = \frac{1}{\sum_{i=1}^N (w_k^i)^2}$$

If N_{eff} drops below a specified threshold $N_T \in [1, N]$, particles are resampled by using Algorithm 2.

Algorithm 1 Particle filter

Input: $p(x_k|x_{k-1})$, $p(y_k|x_k)$, $p(x_0)$, N , N_T

Initialize:

for $i = 1, 2, \dots, N$ **do**

Draw a new particle: $x_1^i \sim p(x_0)$

Assign the weight: $w_1^i = \frac{1}{N}$

end for

At every time step $k = 2, 3, \dots$

for $i = 1, 2, \dots, N$ **do**

Draw a particle from importance distribution:

$x_k^i \sim p(x_k^i|x_{k-1}^i)$

Use the measured y_k to update the weight:

$\tilde{w}_k^i = w_{k-1}^i p(y_k|x_k^i)$

end for

Normalize the weights: $w_k^i = \frac{\tilde{w}_k^i}{\sum_{j=1}^N \tilde{w}_k^j}$

if $\frac{1}{\sum_{i=1}^N (w_k^i)^2} < N_T$ **then**

Resample using Algorithm 2.

end if

Algorithm 2 Resampling

Input: $\{(x^i, w^i)\}_{i=1}^N$
Output: $\{(x_{\text{new}}^i, w_{\text{new}}^i)\}_{i=1}^N$
for $i = 1, 2, \dots, N$ **do**
 Compute cumulative sum of weights: $w_c^i = \sum_{j=1}^i w_k^j$
end for
 Draw u_1 from the uniform distribution $\mathcal{U}(0, \frac{1}{N})$
for $i = 1, 2, \dots, N$ **do**
 Find x^{+i} , the first sample for which $w_c^i \geq u_i$.
 Replace particle i : $x_{\text{new}}^i = x^{+i}$, $w_{\text{new}}^i = \frac{1}{N}$
 $u_{i+1} = u_i + \frac{1}{N}$
end for

The state estimate is computed as the weighted mean of the particles:

$$\hat{x}_k = \sum_{i=1}^N w_k^i x_k^i.$$

For more details on particle filters, the interested reader is referred to [Doucet et al., 2000; Arulampalam et al., 2002].

2.2 Cascaded State Estimation

An important class of nonlinear systems can be represented as cascaded subsystems (e.g., material processing systems, chemical processes, flow processes). For such systems it is also possible to design cascaded observers, which makes their tuning easier.

Consider a general, nonlinear system, for which an observer has to be designed. For a large number of states, and highly nonlinear equations, the design of an observer is clearly problematic. If the states and/or measurements are also corrupted by noise and one uses a particle filter, then a very large sample set is needed, in which case the computational costs may render the observer unusable for online estimation.

A solution is to decompose such a system, and to design separate observers for each subsystem, allowing that one subsystem may use the estimation results of another one. Such a decomposition is presented in Figure 1.




Fig. 1. Cascaded observers.

3. THE WATER TREATMENT PLANT


Most models developed for water treatment plants are steady-state models, and they are used for design purposes. In this paper we consider a dynamic model, that describes the effect of chemical dosing and reactions through the m-alkalinity (M) and p-alkalinity (P), related to the pH. Neither these alkalinities, nor the reaction rates or disturbances can be measured directly. The available measurement is that of the pH, which is a nonlinear function of M and P .

In the current situation, the pH at different stages is kept at fixed values to achieve a desired SI. The desired

pH is based on laboratory measurements of M and P and the current temperature. Therefore the true SI is only controlled at the sampling rate of the laboratory measurements, with a delay of several days. In reality, the process parameters vary with the temperature and other, non-measured disturbances. Moreover, the concentrations of interests cannot be directly measured. This effectively means that changes in process performance/process input can only be detected with a delay of several days. In order to determine the current state of the process, online estimation is needed. This would also give direct feedback to the operators, instead of a delayed evaluation.

So, the goal is to determine M and P online, based on the measured pH and dosing at different steps in the treatment process together with the reaction rates at the treatment steps.

In drinking water production, the pH is mainly determined by the carbonic equilibrium [Wiechers et al., 1975]:



where K_1 , K_2 and K_w are temperature-dependent reaction constants.

The reaction rates of these equilibria are high, and it is therefore assumed that the carbonic fractions are always in equilibrium. The pH (H_3O^+ activity) changes when one of the other concentrations changes due to a reaction. To model the equilibrium, the alkalinities M and P are used. These numbers are ‘‘conservative’’, as they follow the normal rules of mixing (the pH does not). The actual concentrations can be found by solving the following system of equations [van Schagen et al., 2007]:

$$\begin{aligned}
 M &= 2[CO_3^{2-}] + [HCO_3^-] + [OH^-] - [H_3O^+] \\
 P &= [CO_3^{2-}] - [CO_2] + [OH^-] - [H_3O^+] \\
 K_1 &= f^2[HCO_3^-][H_3O^+][CO_2]^{-1} \\
 K_2 &= f^4[CO_3^{2-}][H_3O^+][HCO_3^-]^{-1} \\
 K_w &= f^2[H_3O^+][OH^-]
 \end{aligned} \tag{3}$$

where f is a function of the ionic strength of the water and $[\cdot]$ denotes concentration. This system of five equations can be solved as soon as two variables are known. It can also be used to make a simple model of the treatment plant with respect to the pH, by describing the effect of dosage and reactions on M and P and deducing the (nonlinear) relation to the pH. For the relevant range between 7 and 11 the relation is plotted in Figure 2. The measured pH values are normally between 7.6 and 8.6.

The dynamic model for M and P can be expressed as:

$$\begin{aligned}
 \dot{M} &= \frac{F}{V}(M_{\text{prev}} - M) + \frac{F}{V}f_M(r_{\text{in}}) - R_M(M, P, r, \tau) \\
 \dot{P} &= \frac{F}{V}(P_{\text{prev}} - P) + \frac{F}{V}f_P(r_{\text{in}}) - R_P(M, P, r, \tau) \\
 \dot{r} &= \frac{F}{V}(r_{\text{in}} - r) - R_r(M, P, r, \tau)
 \end{aligned}$$

where F is the flow, V is the water volume in the corresponding process step, r is the concentration of the reactant in the water, r_{in} is the concentration of the




Fig. 2. Dependence of the pH on M and P at 15°C .

reactant added to the water, R_M , R_P and R_r describe the influence of the temperature on the reactions in the treatment step and depend on the temperature τ . The functions f_M and f_P are the instantaneous changes in M and P due to the dosage of chemicals and M_{prev} and P_{prev} are the M and P from the previous treatment stage [van Schagen et al., 2006].

We consider the estimation of the parameters for the *HCl* dosage stage and the RSF stage, for which the functions f_M , f_P , R_M , R_P and R_r for the treatment stages of interest are presented in Table 1. In the *HCl* dosage stage, reactions take place due to the added *HCl*, with the reaction “constant” k_{12} depending heavily on the temperature. During the RSF stage, NH_4 is biologically degraded. The biology uses PO_4 already in the water. The reaction constants again depend on the temperature.

Table 1. Functions used in the estimation model.

Stage	f_M	f_P	R_M	R_P	R_r
<i>HCl</i> dosage	$-r_{\text{in}}$	$-r_{\text{in}}$	0	$k_{12}(P + 0.05)$	0
RSF	0	0	$k_3 r$	$k_3 r$	r
Transportation	0	0	0	0	—

The functions f_M and f_P are considered to be linear in r_{in} , and R_M , R_P and R_r are approximated by linear combinations of M , P and r . The parameters k_i change in each treatment step, and in general they depend on the temperature. The measurements are the pH values after each stage.

For each treatment stage, the input is considered to be the states M and P from the previous stage and the dosing r_{in} . The concentration in each stage refers to the different chemical added. The change during the stage in the M and P numbers depend on the dosing, the reactant already in the water and the temperature. The measurement is the pH after the treatment stage.

Since measurements are only available after the water resulting from RSF has been transported, we need to consider the RSF and transportation models as one stage. This is why Table 1 also contains a “Transportation” stage. It is assumed that the reactions due to the reactant added in a stage are restricted to the corresponding stage.

Due to the nonlinear nature of both the state transition and the measurement model, a nonlinear estimator is needed. Trials have shown that neither the Extended Kalman Filter, nor the Unscented Kalman Filter are able to handle the estimation when measured data is used. Therefore, particle filters were applied. Thanks to the cascaded form of the system, the filters were also applied in cascades.

4. ESTIMATION RESULTS

The particle filter was first tested for simulated data and afterwards applied to measured data. In all cases, the estimation is cascaded and particle filters are used.

4.1 Results Using Simulation Data

The model was simulated using noisy measured flow (F) and reactant (r) data, and both the states and the measurements were corrupted by zero-mean, Gaussian noises.

The initial values were randomly generated. For each process step, both the states and measurements are assumed to be corrupted by a zero-mean Gaussian noise, with known noise covariances. The state transition noise covariance in each stage is 0.005 for M , P , and r and 0.01 for the unknown parameter. The measurement noise covariance is 0.05, which is approximately the measurement error. To simulate the continuous model, a fourth-order Runge-Kutta numerical integration method was used.

Since the water quality parameters in the lake are practically constant, the M_{prev} and P_{prev} for the first *HCl* dosage are known and constant. The models for each stage and the corresponding results are presented below. For each stage, 50 particles were used, with resampling at $N_T = 45$. To estimate the unknown parameters, a random walk model is used.

HCl dosage stage The model used for both simulation and estimation purposes is:


$$\begin{aligned}\dot{M} &= \frac{F}{V}(M_{\text{prev}} - M) - \frac{F}{V}r_{\text{in}} \\ \dot{P} &= \frac{F}{V}(P_{\text{prev}} - P) - \frac{F}{V}r_{\text{in}} - k_{12}(P + 0.05) \\ \dot{r} &= \frac{F}{V}(r_{\text{in}} - r)\end{aligned}\quad (4)$$

When generating the data, $k_{12} = 0.1$ was used. The estimation result for the unknown parameter is presented in Figure 3(a).


Rapid Sand Filtration stage The model used for both simulation and estimation purposes is:

$$\begin{aligned}\dot{M} &= \frac{F}{V}(M_{\text{prev}} - M) - k_3 r \\ \dot{P} &= \frac{F}{V}(P_{\text{prev}} - P) - k_3 r \\ \dot{r} &= \frac{F}{V}(r_{\text{in}} - r) - r\end{aligned}\quad (5)$$

When generating the data, $k_3 = 3$ was used. The estimation result for the unknown parameter is presented in Figure 3(b).



(a) Parameter convergence for acid dosage.



(b) Parameter convergence for RSF.

Fig. 3. *HCl* and RSF estimation results.

As one can see from the presented results, the estimates converge to the true values. The maximum standard deviation computed point-wise over the 30 simulations is below 0.01 for all the estimated states and below 0.001 for the estimated reaction constants.

4.2 Results Using Experiment Data

The simulation results give confidence that using a particle filter the process parameters can be identified based on the measured pH and the measured dosing.

To verify this assumption unfiltered real-world data from the full-scale plant are used. The reactant dosage in the *HCl* dosage stage, the pH after the dosage stage, the incoming NH_4 concentration in the RSF stage, the pH after transportation and the temperature are measured every minute. The results are obtained for a situation where the reactant and flow through the treatment plant are changed, due to operational changes.

HCl dosage stage Since before the *HCl* dosage, the water is in equilibrium with the CO_2 in the air, the M and P numbers stabilize at $M = 3.7$ and $P = -0.05$, respectively. These are the inputs M_{prev} and P_{prev} for the *HCl* dosage stage. The added reactant is measured, and so is the pH at the end of the stage. The states M , P , r and the reaction constant k_{12} need to be estimated. The particle filter uses 100 particles, with resampling at $N_T = 90$. The state transition noise covariance for M , P and r was 0.00005, while for the random walk model of k_{12} it was 0.002. The measurement noise covariance used was 0.05. The results are presented in Figure 4.

With these tuning parameters the prediction of the pH after dosing is very accurate, but the reaction rate varies too quickly. This variation cannot be explained by physical phenomena and therefore it is concluded that the error in the pH measurement must be larger. In a second run the accuracy of the pH measurement was changed to 0.2 and the covariance of the random walk model of k_{12} was changed to 0.0001. The results of this run are shown in Figure 5.




Fig. 4. Estimation results for the first acid dosage - experimental data (solid line - measured, dashed - predicted).




Fig. 5. Estimation results for the first acid dosage - experimental data (solid line - measured, dashed - predicted).

In Figure 5, the pH is within the expected deviation from the measured pH, and the process parameters stabilize at a constant value. However, in simulation, the pH varies more rapidly than in the real-world measurements. An explanation is a possible extra mixing in the *HCl* dosage stage that is not modeled in the current model.

RSF stage and transportation. Since a measurement of the pH is only available after transportation of the water obtained in the RSF stage, the two stages have to be taken together. Reactant is added only in the RSF stage, where the inputs are the M and P numbers at the end of the *HCl* dosage stage. The transportation only delays the M and P s obtained at the end of the RSF stage. The model is given by:

$$\begin{aligned}
\dot{M}_2 &= \frac{F}{V_2}(M_1 - M_2) \\
\dot{P}_2 &= \frac{F}{V_2}(P_1 - P_2) \\
\dot{M}_1 &= \frac{F}{V_1}(M_{\text{prev}} - M_1) - k_3 r \\
\dot{P}_1 &= \frac{F}{V_1}(P_{\text{prev}} - P_1) - k_3 r \\
\dot{r} &= \frac{F}{V_1}(r_{\text{in}} - r) - r
\end{aligned} \tag{6}$$

The measured pH is a nonlinear combination of M_2 and P_2 . Besides the states, the parameter k_3 also has to be estimated. The values used for M_{prev} and P_{prev} are those estimated in the previous stage. A part of the estimated results is presented in Figure 6. The particle filter uses 100 particles, with resampling at $N_T = 90$. The state transition noise covariances are 0.00005 for M_1, M_2, P_1, P_2 , and r and 0.0002 for the random walk model of k_3 . The measurement noise covariance is again increased to 0.2.




Fig. 6. Estimation results for RSF and transportation - experimental data (solid line - measured, dash-dotted - predicted).

The estimation results for this case (Figure 6) show that the convergence to the correct pH value is slower than in the simulated case, but again the process parameter converges to a constant parameter. There still remain changes in the pH measurements that are not explained by the model.

5. CONCLUSIONS

In drinking water production and distribution there is an increased interest in advanced control using online data. However, in the current situation only the pH at different stages is measured. Based on these measurements, a particle filter was applied to the estimation of reaction constants for two stages of a water treatment plant.

Estimation was performed both on simulated data and using data from the full-scale installation. It has been shown that the theoretical accuracy of the pH measurements does not hold in practice, and that an accuracy of 0.2 for the

pH measurement must be used. If that accuracy is taken into account, “constant” process parameters are changing gradually and the pH is within the defined range.

The results are encouraging, but, at the same time, they show that SI cannot be controlled solely by one pH measurement. The actual accuracy of a pH measurement (of 0.2) is within the allowed SI range (of 0.3). To guarantee that the desired SI is kept, it is necessary to implement redundant pH measurements (to increase accuracy) or to use our model-based approach to identify erroneous measurement data.

REFERENCES

- S. Arulampalam, S. Maskell, N. J. Gordon, and T. Clapp. A tutorial on particle filters for online nonlinear/non-Gaussian Bayesian tracking. *IEEE Transactions on Signal Processing*, 50:174–188, 2002.
- M. Bakker, K.M. van Schagen, and J. Timmer. Flow control by prediction of water demand. *Journal of Water Supply: Research and Technology-AQUA*, Vol 52(6):417–424, 2003.
- A. Doucet, S.J. Godsill, and C. Andrieu. On sequential Monte Carlo sampling methods for Bayesian filtering. *Statistics and Computing*, 10:197–208, 2000.
- B. Hill, S. Conrad, H. Kidder, and R. Riddle. Optimal control of water distribution. In *Proceedings of the 2nd International Conference on Instrumentation, Control and Automation*, Busan, South Korea, May 2005.
- R. E. Kalman. A new approach to linear filtering and prediction problems. *Transactions of the ASME—Journal of Basic Engineering*, 82:35–45, 1960.
- P. Li, R. Goodall, P. Weston, C. S. Ling, C. Goodman, and C. Roberts. Estimation of railway vehicle suspension parameters for condition monitoring. *Control Engineering Practice*, 15:43–55, 2007.
- H. Nait-Charif and S. J. McKenna. Tracking poorly modelled motion using particle filters with iterated likelihood weighting. In *The Asian Conference on Computer Vision Systems (ACCV04)*, pages 156–161, 2004.
- K.M. van Schagen, M. Bakker, L.C. Rietveld, A.M.J. Veersma, and R. Babuška. Using online quality measurements in drinking water process control. In *AWWA WQTC Conference*, Denver, USA, November 2006.
- K.M. van Schagen, L.C. Rietveld, and R. Babuška. Dynamic modelling for optimisation of pellet softening. *Journal of Water Supply: Research and Technology-AQUA*, accepted for publication, 2007.
- R. van der Merwe, J. F. G. de Freitas, and A. Doucet. The unscented particle filter. In *Advances in Neural Information Processing Systems*, Cambridge, MA, 2001. MIT Press.
- G. Welch and G. Bishop. An introduction to the Kalman filter. Technical Report TR 95-041, Department of Computer Science, University of North Carolina, NC, USA, 2002.
- H. Wiechers, P. Sturrock, and G. Marais. Calcium carbonate crystallization kinetics. *Water Research* 15: 835-845, 1975.