Jul 1st, 12:00 AM

Coping with communications delays and different time scales when applying central automatic control

R R P Van Nooijen

A. G. Kolechkina

A. Van Loenen

International Environmental Modelling and Software Society (iEMSs)
2012 International Congress on Environmental Modelling and Software
Managing Resources of a Limited Planet, Sixth Biennial Meeting, Leipzig, Germany
R. Seppelt, A.A. Voinov, S. Lange, D. Bankamp (Eds.)
http://www.iemss.org/society/index.php/iemss-2012-proceedings

Coping with communications delays and different time scales when applying central automatic control

R.R.P. van Nooijen [a], A.G. Kolechkina[a], A. van Loenen[b]

[a] *Faculty of Civil Engineering and Geosciences, Delft University of Technology, Stevinweg 1, 2628 CN Delft, The Netherlands (R.R.P.vanNooyen@tudelft.nl)*

[b] *Operational Water Management, Deltares, Rotterdamseweg 185, 2600 MH Delft, The Netherlands*

Abstract A problem that we encountered during the implementation of a central automatic control system for a sewer system was the variable delay in data and command transmission. Another problem was the discrepancy between safe run times for pumps. Pumps with a large local buffer can be left running longer. In general even the pumps with small local storage can run several minutes on dry weather flow. However, the latter type may necessitate checks on the validity of the current pump settings before the next regular call to the controller, especially when there are variable delays in communication. We describe how these problems affected the final controller design.

Keywords: Central automatic control; sewer system; communication delays.

1 INTRODUCTION

Modern society is more and more concerned with the quality of the environment. At the same time the means we have at our disposal are limited by physical and financial constraints. In the project "Integrale sturing afvalwaterketen en watersysteem Hoeksche Waard" (rough translation: Integral management of drainage system and sewer systems for "Hoeksche Waard") a study was made of the water quality on the island "Hoeksche Waard" located in the Dutch Rhine Meuse Delta near the North Sea. The aim of the study was to find and institute experimental measures to improve water quality. One of the measures considered was the central automatic control of the sewer systems. More information on the design process of the controller can be found in van Nooijen et al. [2011a]. To implement this without incurring excessive costs existing systems were used. This meant a fairly complex chain of communication was needed. This in turn meant we needed to adapt the implementation of the control algorithm. In this paper we explore the theoretical implications of such a communications chain for a controller. The chain of communication is described in section 3. The effects on the controller are discussed in section 5 and 7. Some measures that could be taken to mitigate those effects are described in section 6.

In the Netherlands sewer systems are generally managed by municipalities whereas waste water treatment plants are usually run by a water-board. We were fortunate because the system was inherently stable, but our experiences show that complications should be expected. In fact reuse of communication channels should be done with care as it may limit the effective communication rate, for more on the effects of a limited communication rate see for instance Zhang et al. [2001] and Nair et al. [2007].

2 SYSTEM DETAILS

The sewer systems in this area of the Netherlands are located in flat terrain. The ground water level is relatively high (2 meters or less below street level). Together these two facts limit the amount of head that can be created by sloping pipes and therefore the extent of the area over which gravity driven transport is possible. A typical sewer system in this area therefore consists of sub-networks with gravity driven transport linked to one or two small underground reservoirs with pumps that lift the sewage to a pipe leading to another sub-network. At one location there is a pumping station that links the system to either a Waste Water Treatment Plant (WWTP) or a pressurized pipeline network leading to a WWTP. The pumps are usually under local control. At a pre-set local reservoir level h_1 they switch on and at a lower local reservoir level h_2 they switch off. The small underground reservoir acts as a buffer that prevents rapid switching of the pumps under local control. If the run-off and inflow from other sub-networks exceeds the pump capacity then the pipes and manholes fill up and once a fixed threshold is exceeded the excess water will flow over a weir to surface water. In the case of combined sewers that receive both run-off and household sewage this is called a Combined Sewer Overflow (CSO).

3 THE CHAIN OF COMMUNICATION

There was an administrative boundary in the chain of communication between the municipal pumping stations. This coincided with the boundary between software systems from two different suppliers. It was decided to use two files to exchange data at this boundary. One that was could be written by the municipality and read by the water board and one that could be read by the municipality software and written by the water board software.

In the resulting system there are several key components. At the water board there is a set of software packages that, amongst many other tasks, regulate the data flow to and from the municipalities and the calculation of new pump settings. We will call this set of software packages the *central controller*. At the pumping station there is a computer that we will call the *local controller*, it is in charge of the details of pump switching and acts as a final guard against dangerously low or high water levels. The municipal data centre runs a SCADA (Supervisory Control and Data Acquisition) system. On a computer that shares a file system with the SCADA system there is a *data collector* that acts as the local representative of the central controller. This results in a three link chain: local controller to SCADA, SCADA to data collector and data collector to central controller. For simplicity we treat the case where there is just one municipality and we assume there are no communications failures or time-outs. The central controller contains a *coordinator* that initiates all actions. It also contains a database component. The software that actually calculates the new pump assignments is an external program, we will call it the *calculator*. This program is accessed through a *calculation node*.

Every Δt_{M} seconds the SCADA system starts a loop to collect data from the pumping stations 1 to n and store that data in files $f_1 \texttt{ to } f_n$ on the file system shared with the data collector. The data gathering process may take up to Δt_{G} seconds, see also Fig. 1.

The SCADA system also monitors a file f_C on the file system shared with the data collector. Iff_C changes then it reads its contents and transmits it to the pumping stations, see also Fig. 2.

Every Δt_{W} seconds the coordinator starts the following work flow (see also Fig. 3): first send a request for data to the data collector and wait to receive the data from the data

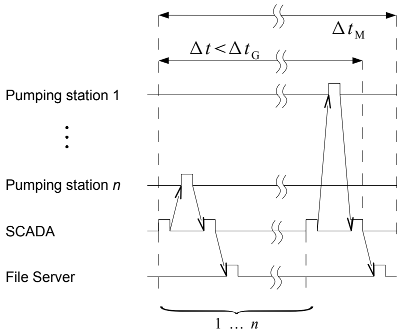

Figure 1: Pump data collection by SCADA

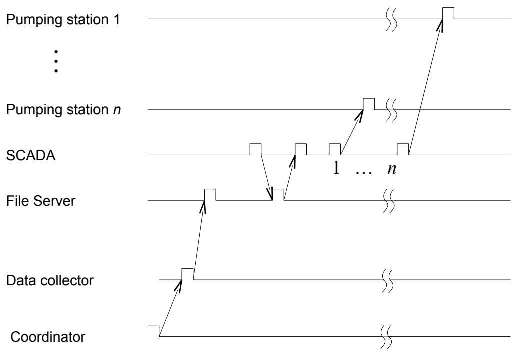

Figure 2: Pump data delivery by SCADA

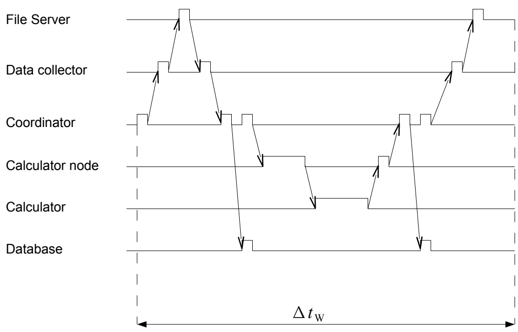

Figure 3: Coordinator work flow

collector then send the data to the database for storage and to the calculation node. Order the node to start the calculator. Wait to receive the results from the calculation node and send the results to the database for storage and to the data collector. The data collector receives the results and writes writes a new f_C , this is assumed to be nearly instantaneous. SCADA detects new f_C and prepares to send it to the pumps.

It is of importance to note that the controller is called every Δt_W seconds, even if the controller time step, that is the period for which the calculated settings are valid, is longer. We can now calculate upper bounds on the age of the data supplied to the controller and the time it takes for the controller to receive information that reflects a change in pump assignment. If we make assumptions on the distribution of response times we can also quantify the probability of exceedance for a certain delay.

4 General information on the properties of the calculator output

The calculator provides an advice that consists of the number of pumps per pumping station to be used during the coming Δt_C seconds. This advice has been verified to be realizable in principle. In other words, given the knowledge of the state of the sewer system available at the time of the call to the calculator and extrapolation over a time step of Δt_{C} based on that state and the proposed pump settings, the selected pumps can run for Δt_{C} without an intervention due to a low water level by the local controller. Moreover, it is necessary to avoid excessive pump switching which would lead to higher maintenance costs. This was implemented by enforcing a minimum time interval Δt_{R} between on-off and off-on state changes. Originally the calculator was written to run as part of and act as controller for simulated sewer system. This meant certain assumptions were embedded in the software:

- It expected to be called at fixed times t_k separated by a fixed time step that was equal to the time step Δt_{C}, so $t_{k+1} - t_k = \Delta t_{\mathrm{C}}$.
- It did not distinguish between the actual pump settings at a given time $t_k + \delta$ with $\delta < \Delta t_{\mathrm{C}}$ and the pump settings that were calculated at t_k.
- In determining the system behaviour it assumed that the delay from data acquisition to control action application would not invalidate the advice.

In moving the controller to a real world application it was decided to consider Δt_{C} as an upper limit on the validity of the pump assignment and schedule a new call to the controller somewhat earlier. In effect a distinction was introduced between the time Δt_{W} between controller calls and the the upper limit Δt_{C} on the validity of the advice.

5 THE EFFECTS OF DELAYS ON THE CONTROL SYSTEM DURING DRY WEATHER

Due to the experimental nature of the project the entire central controller ran on the same Δt_{W} cycle. For wet weather the proper Δt for the system was two to four times Δt_{W} this meant that the controller was asked for a new advice fairly often. For dry weather the same Δt should have been feasible. Certainly the sensor data from the local controller showed that the on and off periods during dry weather could be accommodated by the chosen Δt and Δt_{C}. In practice it turned out that during dry periods there were two major problems. The first was uncertainty about the dry weather flow, this is not constant, so we cannot assume its presence and without it for some of the pumping stations the volume between the level of the bottom of the incoming pipe and the level at which the pump might start to ingest air near the pumps was uncomfortably small given the given pump capacity. An extension is planned to assimilate data on the rate of change of the level into the control system input to generate a short term (Δt_{C}) highly accurate prediction of the inflow which would solve this problem. The other problem was that the cycle from measurement of levels to implementation to the pump command might take more than Δt_{W} which meant an additional loss of precious potential storage space.

6 MITIGATING THE EFFECTS OF SHORT INTERVALS BETWEEN CALLS ON THE CONTROL SYSTEM

The central controller ran at a Δt_{W} of $1/5$ to $1/3$ of Δt_{C} to allow for the occasional malfunction of the communication network. In the experimental system the central controller called the calculator for a new advice with a time step Δt_{W}. Occasionally the delay between data collection and pump state change would exceed Δt_{W}. In that case the controller would be called with incorrect information on the pump state which meant that it might be told that a pump was still running when the local controller had already switched it off. Communication delays were more or less inevitable in the experimental context so we needed to solve this. We decided to let the controller check whether the commands given in the previous call were still usable and in that case simply repeat them with the original time limit on validity. We use the following algorithm. Let A_k be the advice given at time $t_k = t_0 + k\Delta t_{\mathsf{W}}$. The settings given in this advice are intended to be applied from t_k to $t_k + \Delta t_{\mathsf{C}}$. At t_{k+1} the program checks whether $t_{k+1} + \Delta t_{\mathsf{W}} < t_k + \Delta t_{\mathsf{C}}$ and whether these settings can be safely used for the period from t_{k+1} to $t_{k+1} + \Delta t_{\mathsf{C}}$. If this is true, then $A_{k+1} = A_k$ else A_{k+1}will be determined by the controller based on the most recent water levels and pump states. This should prevent short on-off-on cycles due to incorrect data. The controller implementing this scheme will be brought on-line in May of 2012. While this solution will solve the short cycle time, it does nothing to reduce the variable delay between measurement and the request for advice a preliminary evaluation of the effects of that delay is provided in the next section.

7 Quantification of the effects of delays on controller actions

Details on the control algorithm used can be found in Breur et al. [1997]; van Leeuwen [2003]; van Leeuwen and Breur [2005]; van Nooijen et al. [2011b]. We were fortunate that the gravity flow driven sub-networks themselves were stable and that restrictions on pump capacity and on idle time between pump runs effectively eliminated the threat of instability introduced by the control system. The aim of this section is to show that despite this inherent system stability and the measures described earlier communications delays could still cause problems for the controller. For the quantification of the effects of delays we use a controller that is a simplified version of the controller described in those papers.

Let $\widehat{M}$ be the incidence matrix of a connected directed acyclic graph where the gravity driven sub-networks and the WWTP are represented by nodes and the pumping stations are represented by arcs. We assume each node has at most one outgoing arc and that the WWTP has no outgoing arcs. We assume all sub-networks are combined sewer systems. Let n_{n} be the number of nodes and n_{a} the number of arcs. The nodes and arcs are numbered consecutively. Let i_{WWTP} be the node corresponding to the sewage treatment plant, without restriction of generality we can assume $i_{\mathsf{W}} = n_{\mathsf{n}}$. We define the matrix M as the matrix we get by removing row $i_{\mathsf{W}} = n_{\mathsf{n}}$ from $\widehat{M}$. We define the vector $\boldsymbol{m}_{\mathsf{WWTP}}$ to be the transpose of row $i_{\mathsf{W}} = n_{\mathsf{n}}$ of $\widehat{M}$, so

$$\widehat{M} = \begin{pmatrix} M \\ \boldsymbol{m}_{\mathsf{WWTP}}^T \end{pmatrix}$$

where a superscript T denotes the transpose of a matrix or vector. We assumed that $\widehat{M}$ was the incidence matrix of a connected directed acyclic graph, this implies that $M^T M$ is positive definite. Let q_{W} be the limit on the capacity of the sewage treatment plant. Let $\boldsymbol{v}(t)$ be the vector of volumes stored in the sub-networks at time t. We also introduce a vector of volumes $\boldsymbol{v}_{\max}$ that lists the maximum volumes the districts can store without triggering a CSO. Let $\boldsymbol{q}(t)$ be the vector of pump discharge settings at time t and let $\boldsymbol{q}_{\max}$ be the vector of pump capacities. Let $\boldsymbol{r}(t)$ be the vector of inflows into the sub-networks at time t, it represents both precipitation run-off and dry weather flow. For a simple reservoir model without hydrodynamic delays we have conservation of mass

$$\frac{d\boldsymbol{v}}{dt} = \boldsymbol{r} + M\boldsymbol{q}$$

If there were no communication or calculation delays then to optimize use of in system storage at time t_k we would like to solve the following minimization problem

$$\min_{\boldsymbol{u} \in U} \left(\boldsymbol{v}(t_k) + \int\limits_{t=t_k}^{t_{k+1}} \boldsymbol{r}(t)\, dt + M\boldsymbol{u}\Delta t_{\mathsf{C}} - \boldsymbol{v}_{\mathsf{target}}^{(k)} \right)^2$$

where

$$U = \left\{ \boldsymbol{u} \mid 0 \le u_i \le q_{\max,i}, \boldsymbol{m}_{\mathsf{WWTP}}^T \boldsymbol{u} \right\}$$

$$v_{\mathsf{target},i}^{(k)} = \left(\frac{\sum_k v_k^{(k)} - q_{\mathsf{W}} \Delta t_{\mathsf{C}}}{\sum_j v_{\max,j}} \right) v_{\max,i}$$

For future use we introduce the column vector $\boldsymbol{e}$ consisting of $n_{\mathsf{n}} - 1$ ones

$$\boldsymbol{e}^T = \left(\overbrace{1, 1, \ldots, 1}^{n_{\mathsf{n}}-1} \right)$$

and for $\boldsymbol{x} \in \mathbb{R}^n$ the notation diag $(\boldsymbol{x})$ for the $n \times n$ matrix with $\boldsymbol{x}$ on the diagonal and zero off-diagonal elements. We define I as a square matrix with ones on the diagonal and zero off-diagonal elements. Next we define

$$B = \frac{1}{\Delta t_{\mathsf{C}}} \left(I - \frac{\mathsf{diag}\left(\boldsymbol{v}_{\max}\right)}{\boldsymbol{e}^T \boldsymbol{v}_{\max}} \boldsymbol{e}^T \right)$$

$$\boldsymbol{b}^{(k)} = \frac{1}{\Delta t_{\mathsf{C}}} \int\limits_{t=t_k}^{t_{k+1}} \boldsymbol{r}\left(t\right) dt - \frac{q_{\mathsf{W}}}{\boldsymbol{e}^T \boldsymbol{v}_{\max}} \boldsymbol{v}_{\max}$$

It is clear that communication delays can cause significant problems when one or more constraints are active. For the remainder of the analysis we assume we are close enough to the optimal solution that we may disregard the constraints. We take as our starting point

$$\min_{\boldsymbol{u}} \left(M\boldsymbol{u} + B\boldsymbol{v}\left(t_k\right) + \boldsymbol{b}^{(k)} \right)^2 \tag{1}$$

We are primarily interested in the effects of delays, so let us assume that $\boldsymbol{r}$ is piecewise constant and known. Now assume there are known, fixed communications, calculation and processing delays, so we can choose $0 < \lambda < \Delta t_{\mathsf{C}}$ and $0 < \mu < \Delta t_{\mathsf{C}}$ such that at time t_k we are certain that all data for time $t_k - \lambda$ is available and we are certain our commands are available at the pumps and will be executed at time $t_k + \mu$. So, while originally we looked for a solution to Equation 1, now we want to solve

$$\min_{\boldsymbol{u}_\mu} \left(M\boldsymbol{u}_\mu + B\boldsymbol{v}_\mu^{(k)} + \boldsymbol{b}_\mu^{(k)} \right)^2 \tag{2}$$

where

$$\boldsymbol{b}_\mu^{(k)} = \frac{1}{\Delta t_{\mathsf{C}}} \int\limits_{t=t_k+\mu}^{t_{k+1}+\mu} \boldsymbol{r}\left(t\right) dt - \frac{q_{\mathsf{W}}}{\boldsymbol{e}^T \boldsymbol{v}_{\max}} \boldsymbol{v}_{\max}$$

$$\boldsymbol{v}_\mu^{(k)} = \boldsymbol{v}^{(k)} + \int\limits_{t=t_k}^{t_k+\mu} \boldsymbol{r}\left(t\right) dt + \mu M \boldsymbol{u}^{(k-1)}$$

Now suppose that the measurements are not all taken at the same time and the control actions are not all implemented at the same time. We consider only one control time step, so we do not index the time shifts. Let $-1 < \delta_i < 1$ for $i = 1, 2, \ldots, n_{\mathsf{n}-1}$ and $-1 < \epsilon_j < 1$ for $j = 1, 2, \ldots, n_{\mathsf{a}}$ be relative shifts in measurement time and command execution time, so the measurement planned for $t_k - \lambda$ would actually be carried out at $t_k - \lambda - \delta_i \Delta t_{\mathsf{C}}$ and the control action planned for $t_k + \mu$ would actually be carried out at $t_k + \mu + \epsilon_i \Delta t_{\mathsf{C}}$. In that case we would like to solve

$$\min_{\boldsymbol{u}_{\mathsf{s}}} \left(M\left(I - \mathsf{diag}\left(\boldsymbol{\epsilon}\right)\right) \boldsymbol{u}_{\mathsf{s}} + B\boldsymbol{v}_{\mathsf{s}}^{(k)} + \boldsymbol{b}_{\mathsf{s}}^{(k)} \right)^2 \tag{3}$$

where

$$b_{\mathsf{s},i}^{(k)} = \left(M\mathsf{diag}\left(\boldsymbol{\epsilon}\right)\boldsymbol{u}^{(k-1)}\right)_i + \frac{1}{\Delta t_{\mathsf{C}}}\int\limits_{t=t_k+\mu}^{t_{k+1}+\mu} r_i\left(t\right)dt - \frac{q_{\mathsf{W}}}{\boldsymbol{e}^T\boldsymbol{v}_{\max}}v_{\max,i}$$

$$v_{\mathsf{s},i}^{(k)} = v^{(k)} + \int\limits_{t=t_k}^{t_k+\mu} r_i\left(t\right)dt + \mu\left(M\boldsymbol{u}^{(k-1)}\right)_i$$

Now suppose $\boldsymbol{\delta}$ and $\boldsymbol{\epsilon}$ are unknown vectors of elements with absolute value less than one and in stead of problem (3) we solve

$$\min_{\tilde{\boldsymbol{u}}}\left(M\tilde{\boldsymbol{u}} + B\tilde{\boldsymbol{v}}^{(k)} + \tilde{\boldsymbol{b}}^{(k)}\right)^2 \tag{4}$$

with

$$\tilde{b}_i^{(k)} = \frac{1}{\Delta t_{\mathsf{C}}}\int\limits_{t=t_k+\mu}^{t_{k+1}+\mu} r_i\left(t\right)dt - \frac{q_{\mathsf{W}}}{\boldsymbol{e}^T\boldsymbol{v}_{\max}}v_{\max,i}$$

$$\tilde{v}_i^{(k)} = v_i\left(t_k - \lambda - \delta_i\Delta t_{\mathsf{C}}\right) + \int\limits_{t=t_k-\lambda}^{t_k+\mu} r_i\left(t\right)dt + \left(\lambda+\mu\right)\left(M\boldsymbol{u}^{(k-1)}\right)_i$$

The solution $\tilde{\boldsymbol{u}}$ of problem (4) satisfies

$$M^TM\tilde{\boldsymbol{u}} = M^T\left(B\tilde{v}_i^{(k)} + \tilde{b}_i^{(k)}\right)$$

which, for piecewise constant and known $\boldsymbol{r}$, we can write as

$$\begin{aligned} M^TM\tilde{\boldsymbol{u}} &= M^T\left(B\boldsymbol{v}_\mu^{(k)} + \boldsymbol{b}_\mu^{(k)}\right) + \\ &\quad M^TB\left(-\mathsf{diag}\left(\boldsymbol{\delta}\right)\boldsymbol{r}^{(k-1)} - M\mathsf{diag}\left(\boldsymbol{\delta}\right)\boldsymbol{u}^{(k-1)}\right)\end{aligned}$$

The solution $\boldsymbol{u}_{\mathsf{s}}$ of problem (3) satisfies

$$\left(I-\mathsf{diag}\left(\boldsymbol{\epsilon}\right)\right)M^TM\left(I-\mathsf{diag}\left(\boldsymbol{\epsilon}\right)\right)\boldsymbol{u}_{\mathsf{s}} = \left(I-\mathsf{diag}\left(\boldsymbol{\epsilon}\right)\right)M^T\left(B\boldsymbol{v}_{\mathsf{s}}^{(k)} + \boldsymbol{b}_{\mathsf{s}}^{(k)}\right)$$

which, for piecewise constant and known $\boldsymbol{r}$, we can write as

$$\begin{aligned}\left(I-\mathsf{diag}\left(\boldsymbol{\epsilon}\right)\right)M^TM\left(I-\mathsf{diag}\left(\boldsymbol{\epsilon}\right)\right)\boldsymbol{u}_{\mathsf{s}} &= \left(I-\mathsf{diag}\left(\boldsymbol{\epsilon}\right)\right)M^T\left(B\boldsymbol{v}_\mu^{(k)} + \boldsymbol{b}_\mu^{(k)}\right) + \\ &\quad \left(I-\mathsf{diag}\left(\boldsymbol{\epsilon}\right)\right)M^TM\mathsf{diag}\left(\boldsymbol{\epsilon}\right)\boldsymbol{u}^{(k-1)}\end{aligned}$$

For $\boldsymbol{\delta}$ random and uncorrelated we see that the expected value $\mathcal{E}\left[\tilde{\boldsymbol{u}}\right]$ satisfies

$$\mathcal{E}\left[\tilde{\boldsymbol{u}}\right] = \left(M^TM\right)^{-1}M^T\left(B\boldsymbol{v}_\mu^{(k)} + \boldsymbol{b}_\mu^{(k)}\right)$$

so $\mathcal{E}\left[\tilde{\boldsymbol{u}}\right] = \boldsymbol{u}_\mu$. However, to obtain $\mathcal{E}\left[\boldsymbol{u}_{\mathsf{s}}\right]$ we need to do some extra work. We have

$$\begin{aligned} M^TM\left(I-\mathsf{diag}\left(\boldsymbol{\epsilon}\right)\right)\boldsymbol{u}_{\mathsf{s}} &= M^T\left(B\boldsymbol{v}_\mu^{(k)} + \boldsymbol{b}_\mu^{(k)}\right) + \\ &\quad M^TM\mathsf{diag}\left(\boldsymbol{\epsilon}\right)\boldsymbol{u}^{(k-1)}\end{aligned}$$

so

$$\begin{aligned}\boldsymbol{u}_{\mathsf{s}} &= \left(I-\mathsf{diag}\left(\boldsymbol{\epsilon}\right)\right)^{-1}\left(M^TM\right)^{-1}M^T\left(B\boldsymbol{v}_\mu^{(k)} + \boldsymbol{b}_\mu^{(k)}\right) + \\ &\quad \left(I-\mathsf{diag}\left(\boldsymbol{\epsilon}\right)\right)^{-1}\mathsf{diag}\left(\boldsymbol{\epsilon}\right)\boldsymbol{u}^{(k-1)}\end{aligned}$$

If we expand the inverse matrix into a formal power series then we get

$$\begin{aligned} \boldsymbol{u}_{\mathsf{s}} &= \left(I + \mathsf{diag}\left(\boldsymbol{\epsilon}\right) + \left(\mathsf{diag}\left(\boldsymbol{\epsilon}\right)\right)^2 + ...\right)\left(M^T M\right)^{-1} M^T \left(B\boldsymbol{v}_{\mu}^{(k)} + \boldsymbol{b}_{\mu}^{(k)}\right) + \\ &\quad \left(\mathsf{diag}\left(\boldsymbol{\epsilon}\right) + \left(\mathsf{diag}\left(\boldsymbol{\epsilon}\right)\right)^2 + ...\right)\boldsymbol{u}^{(k-1)} \end{aligned}$$

so the expected value of $\boldsymbol{u}_{\mathsf{s}}$ will not be $\boldsymbol{u}_{\mu}$. We see that for relatively large ϵ_i (more than 1/3) we will have a problem because $\mathcal{E}\left[\tilde{\boldsymbol{u}}\right]$ will diverge from $\mathcal{E}\left[\boldsymbol{u}_{\mathsf{s}}\right]$ by more than 10 percent, so we may not be taking the correct action. Evidently this is only a preliminary analysis and a full analysis over multiple time steps is needed to draw definitive conclusions.

8 Summary

During the design of a central controller for a group of sewer systems several problems surfaced. One problem was that for technical reasons the time step for calls to the controller algorithm was shorter than the planned time step between control actions. This was solved by reuse of the advice when suitable. Another problem was the variation in communication delays. Our preliminary analysis indicated that the controller would probably be able to cope with this.

Acknowledgement

This work is an extension of work done for the project "Integrale Sturing Afvalwaterketen en watersysteem Hoeksche Waard" for which AgentschapNL (previously known as SenterNovem) provided a subsidy from the "Innovatieprogramma Kaderrichtlijn Water" of the Dutch Ministry of Transport, Public Works and Water Management.

References

Breur, K.-J., P. E. R. M. van Leeuwen, and N. P. Dellaert. Deriving discharge strategies to reduce CSO in urban drainage systems. *Water Science and Technology*, 36(8-9): 259–263, 1997.

Nair, G. N., F. Fagnani, S. Zampieri, and R. J. Evans. Feedback control under data rate constraints: An overview. *Proceedings of the IEEE*, 95(1):108–137, January 2007.

van Leeuwen, P. E. R. M. *Sturing van Watersystemen*. DUP Science, Delft University Press, Delft, The Netherlands, 2003. In Dutch.

van Leeuwen, P. E. R. M. and K. J. Breur. The modeling policy-maker. On decision support systems in water management. *Integrated Assessment*, 2(2):89–92, 2005.

van Nooijen, R. R. P., A. G. Kolechkina, P. E. R. M. van Leeuwen, and E. van Velzen. Agile design of sewer system control. In Bittanti, S., A. Cenedese, and S. Zampieri, editors, *Proceedings of the 18th IFAC World Congress, Milano (Italy) August 28 - September 2, 2011*, 2011a.

van Nooijen, R., A. Kolechkina, K.-J. van Heeringen, and E. van Velzen. Implementation of central control for multiple sewer systems. In *Proceedings of the 12th International Conference on Urban Drainage, Porto Alegre/Brazil, 10-15 September 2011*, 2011b.

Zhang, W., M. Branicky, and S. Phillips. Stability of networked control systems. *Control Systems Magazine, IEEE*, 21(1):84–99, 2001.