Proceedings of the 17th World Congress
The International Federation of Automatic Control
Seoul, Korea, July 6-11, 2008

Distributed controller design for open water channels [1]

Yupinig Li * **Michael Cantoni** *,[2]

* *The University of Melbourne, Department of Electrical and Electronic Engineering, Parkville, VIC 3010, Australia*

Abstract: In the design of an automatic controller to achieve water-level set-point regulation and off-take load-disturbance rejection for an open water channel, a key concern is an inherent trade-off between local performance and the way water-level errors propagate due to control action. Here a structured optimal controller synthesis problem is formulated to systematically manage this trade-off, using $\mathcal{H}_\infty$ loop-shaping ideas. The loop-shaping weights can be scalably designed and the imposed structure ensures the controller only involves local information exchange. Importantly, the distributed control structure we consider confines water-level error propagation to upstream pools, with corresponding benefits in terms of water distribution efficiency. Moreover, it coincides with the interconnection structure of a channel, and so the corresponding optimal synthesis problem has a convex characterisation; detailed state-space formulae are provided. Field test data are presented to illustrate overall performance.

1. INTRODUCTION

In large-scale irrigation networks, water is often distributed via open water channels under the power of gravity alone (i.e. there is no pumping). The flow of water through the network is regulated by automated gates positioned along the channels (Mareels et al., 2005). To give some idea of scale, there are about 7000 kilometers of irrigation channels in the Goulburn-Murray Irrigation District, the largest in Australia. These annually supply more than $2\mathrm{Gm}^3$ of fresh water for agricultural purposes.

The stretch of a channel between two gates is commonly called a pool. Water off-take points to farms and secondary channels are distributed along the pools, typically with a degree of concentration just upstream of the gates (i.e. at the downstream end of pools). As such, an important control objective is set-point regulation of the water-levels immediately upstream of each gate, which enables flow demand at the (often gravity-powered) off-take points to be met without over-supplying. Water off-takes into secondary channels and onto farms act as load disturbances that are to be rejected by the controller, with the assistance of disturbance feedforward/preview when possible. The control of open water channels to achieve such objectives has been of interest for some time; see (Clemmens et al., 1998; Mareels et al., 2005) and the references therein.

When the number of pools to be controlled is large and the gates widely dispersed, centralised feedback control (Weyer, 2003; Li et al., 2004; Weyer, 2007) can incur a heavy communication overhead. This can make it difficult to ensure timely exchange of water-level sensor measurements and gate actuator commands between all gates and potentially a central host. A natural way to overcome this is to employ a distributed control structure, requiring only local exchange of information between neighbouring gates (Li, 2006; Cantoni et al., 2007). In this paper, we present the details of a systematic approach to feedback compensator synthesis, subject to a particular distributed control information structure. The distributed structure accommodates the exchange of control information, akin to a feed-forward action for de-coupling, between neighbouring local compensators. This facilitates the management of a key trade-off between the local performance and the way water-level errors propagate due to control action (Li et al., 2005). Importantly, the structure also confines water-level error propagation to the pools upstream of a load-disturbance, which is important in terms of water distribution efficiency and quality of service (Litrico et al., 2003; Mareels et al., 2005). This is not the case for centralised feedback control in general.

The paper is structured as follows. An open water channel model is briefly developed in Section 2 for the purpose of feedback controller design. A structured optimal controller synthesis problem is then formulated, based on $\mathcal{H}_\infty$ loop-shaping ideas, to deal with the design trade-off described above. This includes a discussion of loop-shaping weight design, which is ultimately achieved in a decentralised/scalable fashion. Section 4 presents a convex characterisation of the optimal synthesis problem, using state-space realisations and a specialisation of Thm. 4 in (Langbort et al., 2004), which accommodates heterogeneity in the pool dynamics and finiteness of the channel extent. This yields a distributed controller which is different to the one suggested in (Li et al., 2005), where each local compensator is obtained using a spatially-invariant, infinite-extent relaxation of the problem. To illustrate the performance of a distributed controller designed as described herein, simulation and field test results are discussed in Section 5. The field test were made possible by Rubicon Systems Australia Pty. Ltd., whose support is gratefully acknowledged.

[1] Work supported in part by the Australian Research Council (Linkage Grant LP0349134 and Discovery Grant DP0664225).
[2] Corresponding author (email: cantoni@unimelb.edu.au).

2. CHANNEL MODELLING FOR CONTROL

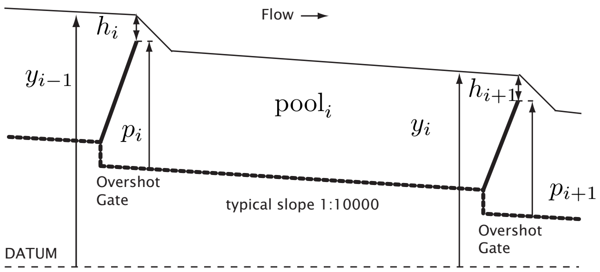

Fig. 1. Stretch of a channel with over-shot gates. Although we consider over-shot gates, our approach to design approach is applicable with other gate structures.

A side view of an open water channel is shown in Fig. 1: y_i denotes the downstream water level in pool$_i$; p_i the position of the i-th gate; and h_i the "head" over the i-th gate, which yields a flow $u_i = \gamma_i h_i^{3/2}$ (Bos, 1978). To meet off-take flow demand, without over-supplying, it is important to regulate the downstream water-levels y_i to calculated set-points r_i. Towards this end, a simple frequency-domain model of the water level in pool$_i$ can be obtained by conservation of mass (Weyer, 2001):

$$P_i \;:\; y_i(s) = \frac{1}{s\,\alpha_i}\Big(\exp(-s\tau_i)u_i(s) - v_i(s) - d_i(s)\Big), \quad (1)$$

where d_i models off-take load-disturbances in pool$_i$, u_i is the flow over the i-th gate, τ_i is a transport delay, α_i is a measure of the pool surface area and $v_i = u_{i+1}$, which characterises the interconnection of neighbouring pools. The model parameters α_i and τ_i are readily obtained via system identification (Weyer, 2001; Ooi et al., 2003). While (1) does not capture wave dynamics, it is suitable for feedback control design since it represents a pool from the perspective of desired closed-loop behaviour (Cantoni et al., 2007). Indeed, it is required that automatic control not excite the wave dynamics; this can lead to unacceptable fluctuation in flows at off-take points, flooding and damage of the channel. To summarise, an open water channel can be thought of as a string of pools, each with dynamics of the form (1), as shown in Fig. 2.

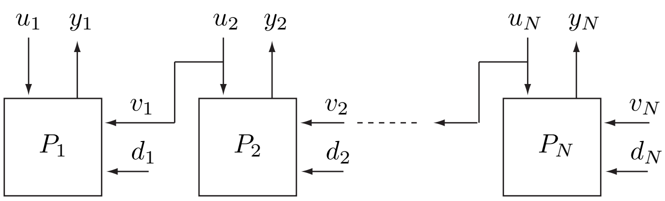

Fig. 2. A string of N pools, each with dynamics (1)

By adjusting the position of the i-th gate relative to the water-level upstream of it (i.e. by adjusting the "head"), the flow u_i into pool$_i$ can be controlled. Exploiting this, the objectives of water-level set-point regulation and (unmeasured) off-take load-disturbance rejection can be achieved with a feedback controller which sets the flows u_i on the basis of water-level measurements y_i and set-points r_i. A centralised controller (Li et al., 2004; Weyer, 2007) can exploit all measured water-levels to set each gate flow. This can yield good water-level regulation performance, but in general, it leads to both upstream and downstream propagation of water-level errors. Downstream propagation is undesirable from a quality of service perspective and in terms of distribution efficiency (Litrico et al., 2003; Mareels et al., 2005). Moreover, the communication overhead for centralised control can be prohibitive.

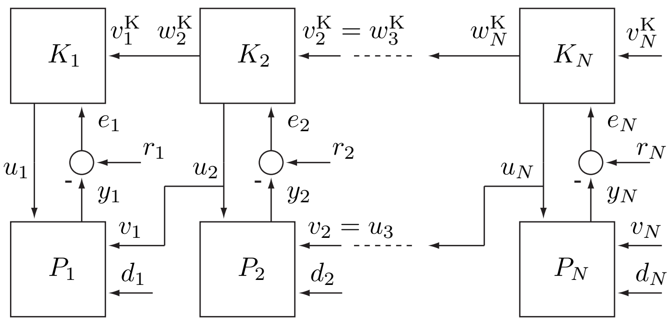

Fig. 3. Distant-downstream distributed control

A distributed control structure (Li, 2006; Cantoni et al., 2007), on the other hand, only requires local exchange of information. Moreover, with the so-called distant-downstream control structure shown in Fig. 3, whereby the downstream water-level in pool$_i$ is effectively controlled via the flow over the upstream (i.e. i-th) gate, the propagation of water-level errors is confined to upstream pools. In this way, it is possible to separately control the flow over the last gate to be zero (or some other constant), with benefits in terms of distribution efficiency.

While the distributed distant-downstream control information structure exhibits the favourable properties just described, there exists an inherent design trade-off between local performance and the way errors propagate (Li, 2006; Cantoni et al., 2007). Indeed, it is possible for a component of the error to be amplified as it propagates upstream (Li et al., 2005), particulary when there is worst-case coupling between pools (i.e. homogeneity). This can lead to gate saturation and unacceptable deviations from set-point; e.g. failure to meet off-take flow demand or flooding.

Since water-level error propagation arises from control action, which adjusts the flow into one pool and hence out of another, the information exchanged between the local compensators can facilitate the management of the design trade-off described above. Indeed, one possibility corresponds to the so-called decentralised feedback with feed-forward structure described in (Cantoni et al., 2007; Weyer, 2007); specifically,

$$K_i = \begin{pmatrix} v_i^K = u_{i+1} \\ e_i \end{pmatrix} \mapsto \begin{pmatrix} w_i^K = u_i \\ u_i \end{pmatrix} = \begin{pmatrix} F_i & C_i \\ F_i & C_i \end{pmatrix},$$

where F_i provides a de-coupling feed-forward action, on the basis of the flow u_{i+1} out of pool$_i$, and C_i is a decentralised feedback compensator. Systematically exploiting the compensators F_i to manage error propagation, however, is difficult (Li, 2006). Below, we formulate a structured optimal control problem, based on loop-shaping ideas, to design a distributed distant-downstream controller. The global performance index captures the trade-off between local performance and error propagation, among other things, and the design of the loop-shaping weights is quite systematic in that it is achieved on pool-by-pool basis, focusing on local performance objectives.

3. $\mathcal{H}_\infty$ LOOP-SHAPING BASED DISTRIBUTED CONTROLLER DESIGN

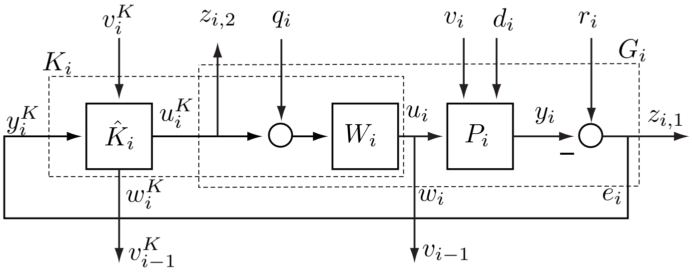

Fig. 4. Localised portion of a controlled channel

For the purpose of design, consider the closed-loop system shown in Fig. 4; this represents a localised portion of a channel under distributed distant-downstream feedback control. As before, P_i is the nominal model (1) for pool$_i$. The local component K_i of the distributed controller as it is shown in Fig. 3, on the other hand, is now effectively split into a loop-shaping weight W_i and a compensator $\hat{K}_i$; this is consistent with the well-known $\mathcal{H}_\infty$ loop-shaping approach to feedback controller design (Skogestad et al., 1996). The role of the weight W_i is to classically shape the nominal plant model dynamics to obtain a local loop-gain that is consistent with the local performance and robustness objectives across frequency. To this end, the flow coupling between pools that arises from control action in downstream pools, can be treated as an additional unknown load-disturbance to be rejected. That this component of the load-disturbance is actually known can be exploited by the additional compensators $\hat{K}_i$, which each pass control information to the corresponding compensator upstream. In particular, for a given collection of loop-shaping weights, the extra design freedom can be used to optimise a measure of global performance, accounting for the trade-off between local performance objectives and the way water-level errors propagate.

With reference to Fig. 4, the transfer function of the i-th weighted generalised pool

$$\begin{pmatrix} v_i \\ n_i \\ u_i^K \end{pmatrix} \mapsto \begin{pmatrix} w_i = u_i \\ z_i \\ y_i^K = e_i \end{pmatrix},$$

takes the form

$$G_i(s)=\begin{pmatrix} 0 & \begin{pmatrix}0 & 0 & W_i(s)\end{pmatrix} & W_i(s) \\ \begin{pmatrix} \frac{1}{s\alpha_i} \\ 0 \end{pmatrix} & \begin{pmatrix} 1 & \frac{1}{s\alpha_i} & \frac{\exp(-s\tau_i)}{-s\alpha_i}W_i(s) \\ 0 & 0 & 0 \end{pmatrix} & \begin{pmatrix} \frac{\exp(-s\tau_i)}{-s\alpha_i}W_i(s) \\ 1 \end{pmatrix} \\ \frac{1}{s\alpha_i} & \begin{pmatrix} 1 & \frac{1}{s\alpha_i} & \frac{\exp(-s\tau_i)}{-s\alpha_i}W_i(s) \end{pmatrix} & \frac{\exp(-s\tau_i)}{-s\alpha_i}W_i(s) \end{pmatrix}, \quad (2)$$

where the set-point and off-take load-disturbance have been lumped into the signal $n_i := (r_i, d_i, q_i)^{\mathrm{T}}$, together with a signal q_i which is used to model additional uncertainty in the flow over gate$_i$. Moreover, the pool interconnection is characterised by $v_i = w_{i+1}$ with boundary condition $v_N = 0$; recall that under distant-downstream control setting $v_N = 0$ indeed possible. Within this context, the aforementioned control design trade-off corresponds to managing the collective effect of the exogenous signals $n_i := (r_i, d_i, q_i)^T$ on the signals

$$z_i := \begin{pmatrix} y_i^K = e_i \\ u_i^K \end{pmatrix}.$$

This is achieved by judicious choice of the interconnection of distributed compensators, denoted $\hat{K} = (\hat{K}_1, \ldots, \hat{K}_N)$, where

$$\hat{K}_i = \begin{pmatrix} v_i^K \\ y_i^K \end{pmatrix} \mapsto \begin{pmatrix} w_i^K \\ u_i^K \end{pmatrix}$$

and $v_i^K = w_{i+1}^K$ with boundary condition $v_N^K = 0$. In particular, using $G := (G_1, \ldots, G_N)$ to denote the interconnection of the generalised pool models (2), and $H(G, \hat{K})$ to denote the overall closed-loop transfer function from the vector of exogenous signals $n := (n_1^T, \ldots, n_N^T)^T$ to the vector of signals $z := (z_1^T, \ldots, z_N^T)^T$ shown in Fig. 5, the trade-off can be handled by selecting the $\hat{K}$ to stabilise $H(G, \hat{K})$ and minimise

$$\|H(G, \hat{K})\|_\infty = \sup_{\omega \in \mathbb{R}} \bar{\sigma}(H(G, \hat{K})(j\omega)),$$

where $\bar{\sigma}(\cdot)$ denotes the maximum singular value of a matrix. This structured $\mathcal{H}_\infty$ synthesis problem corresponds to optimising a weighted worst-case measure of global performance, accounting for local objectives and the mechanism underpinning water-level error propagation.

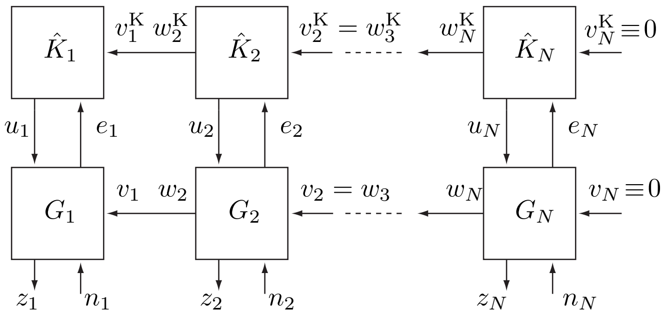

Fig. 5. Generalised distributed structure for synthesis

In Section 4 we present a convex state-space characterisation of distributed compensators $\hat{K}$ that achieve closed-loop stability and a specified bound on the global performance index $\|H(G, \hat{K})\|_\infty$. Given a stabilising $\hat{K}$ that achieves a minimal (or satisfactory) performance bound, the i-th component of the distributed distant-downstream controller, shown in Fig. 3, is recovered as

$$K_i = \begin{pmatrix} v_i^K \\ e_i \end{pmatrix} \mapsto \begin{pmatrix} w_i^K \\ u_i \end{pmatrix} = \begin{pmatrix} 1 & 0 \\ 0 & W_i \end{pmatrix} \hat{K}_i.$$

3.1 Loop-shaping weight design

The design of the loop-shaping weights can be achieved on a pool-by-pool basis. The weight W_i should be designed to classically shape the local loop-gain $L_i(s) = W_i(s)/s\,\alpha_i$ across frequency, so that the local closed-loop transfer functions

$$T_{r_i \to e_i}(s) := \frac{1}{1 + L_i(s)\exp(-s\tau_i)},$$

$$T_{d_i, v_i \to e_i}(s) := \frac{1}{1 + L_i(s)\exp(-s\tau_i)} \left(\frac{1}{s\,\alpha_i} \right)$$

and

$$T_{d_i \to u_i}(s) := \frac{L_i(s)}{1 + L_i(s)\exp(-s\tau_i)},$$

i	τ_i	α_i	φ_i
1	4 mins	6492m^2	0.48 rad/min
2	2 mins	2478m^2	1.05 rad/min
3	4 mins	6084m^2	0.48 rad/min
4	4 mins	5658m^2	0.48 rad/min
5	6 mins	7650m^2	0.42 rad/min

Table 1. Pool model parameters: delay (τ_i), surface area (α_i) and wave frequency (φ_i)

are consistent with the local performance and robustness objectives; i.e. water-level set-point regulation, load-disturbance rejection and the requirement that control action should not excite the dominant wave dynamics. In particular, to achieve zero steady-state water-level error for step load disturbances, it follows by the final value theorem that $W_i(s)$ should have at least one pole at $s = 0$. Furthermore, in line with classical design ideas (Skogestad et al., 1996), good set-point regulation and load disturbance rejection is achieved by ensuring that the local loop-gain $|L_i(j\omega)|$ is large at frequencies where the set-point reference r_i and the load disturbance d_i are significant (i.e. low frequency). Moreover, the bandwidth of the loop-gain $|L_i(j\omega)|$ must be set to lie below the frequency of the dominant local wave dynamics, which are not captured by the first-order model (1). Indeed, these local objectives can be achieved with an essentially PI compensator

$$W_i(s) = \frac{\kappa_i(1 + s\phi_i)}{s\,(1 + s\rho_i)}. \quad (3)$$

Specifically: κ_i is used to set the loop-gain bandwidth – this should also sit below $(1/\tau_i)$ rad/min, because of the delay which is not reflected in L_i; ϕ_i is used to introduce phase lead in the cross-over region to reduce the roll-off rate for stability and robustness; ρ_i is used to provide additional roll-off beyond the loop-gain bandwidth to ensure sufficiently low gain at the (un-modelled) dominant wave frequency. In fact, the W_i in (3) has the form of the completely decentralised compensators described in (Cantoni et al., 2007; Weyer, 2007), and it corresponds to the basic form of the weight used for optimal centralised controller synthesis in (Li et al., 2004).

Figure 6 shows the shaped and un-shaped local loop-gains used for the five pools of Coleambally Channel, Number 6, NSW, where the field tests reported in Section 5 were conducted. The identified model parameters for these pools are summarised in Table 1, where the wave frequency provided is representative of the dominant wave dynamics not modelled in (1). These loop-shapes correspond to the weight parameters κ_i, ϕ_i, ρ_i summarised in Table 2. The un-shaped loop-gains are also scaled by a constant η_i, to facilitate use of the same plot; these are ultimately part of the parameter κ_i. It can be seen that compared to the un-shaped loop-gains, the shaped gains are significantly higher at low frequencies, and somewhat lower at the wave frequencies, with a smooth transition in between. Also note that the bandwidth of each shaped loop-gain is much less than the corresponding value of $1/\tau_i$.

4. A STATE-SPACE CHARACTERISATION OF THE $\mathcal{H}_\infty$ SYNTHESIS PROBLEM

The existence of an interconnection $\hat{K} = (\hat{K}_1, \ldots, \hat{K}_N)$ of distributed compensators that stabilises the interconnection $G = (G_1, \ldots, G_N)$ of weighted generalised pool models and achieves a specified bound on the performance index $\|H(G, \hat{K})\|_\infty$, is characterised below in terms of convex conditions involving a finite-dimensional state-space model for G. This state-space model is constructed by employing a Padé approximation of the delay associated with each pool. In particular, each weighted generalised pool model G_i, originally defined in (2), is approximately realised with the finite-dimensional state-space model

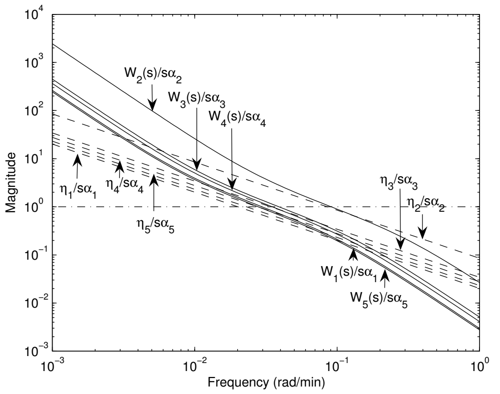

Fig. 6. Shaped and un-shaped local loop-gains

i	κ_i	ϕ_i	ρ_i	η_i
1	1.69	113.64	9.97	130
2	6.47	37.17	3.26	223
3	2.37	86.96	7.60	183
4	2.21	96.15	8.47	170
5	1.68	113.64	9.97	153

Table 2. Parameters for loop-shaping weights

$$\begin{pmatrix} \dot{x}_i \\ w_i \\ z_i \\ y_i^K \end{pmatrix} = \begin{pmatrix} A_i^{tt} & A_i^{ts} & B_i^{tn} & B_i^{tu} \\ A_i^{st} & A_i^{ss} & B_i^{sn} & B_i^{su} \\ C_i^{tz} & C_i^{sz} & D_i^{zn} & D_i^{zu} \\ C_i^{ty} & C_i^{sy} & D_i^{yn} & D_i^{yu} \end{pmatrix} \begin{pmatrix} x_i \\ v_i \\ n_i \\ u_i^K \end{pmatrix}$$

$$:= \left(\begin{array}{cccc|c|cc|c|c} 0 & \frac{1}{\alpha_i} & \frac{-1}{\alpha_i} & 0 & \frac{-1}{\alpha_i} & 0 & \frac{-1}{\alpha_i} & 0 & 0 \\ 0 & \frac{-2}{\tau_i} & \frac{4}{\tau_i} & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & \frac{\kappa_i\phi_i}{\rho_i} & \frac{\kappa_i\phi_i}{\rho_i} \\ 0 & 0 & 0 & \frac{-1}{\rho_i} & 0 & 0 & 0 & \frac{\kappa_i(\rho_i-\phi_i)}{\rho_i^2} & \frac{\kappa_i(\rho_i-\phi_i)}{\rho_i^2} \\ \hline 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ \hline -1 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \\ \hline -1 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \end{array}\right) \begin{pmatrix} x_i \\ v_i \\ n_i \\ u_i^K \end{pmatrix}, \quad (4)$$

where the local state $x_i = (y_i, \Delta_i, u_i, \Omega_i)^T$, the sub-state Ω_i corresponds to the loop-shaping pole at $s = -1/\rho_i$ and the sub-state Δ_i corresponds to the pole in the Padé approximation of the delay. For $i = N$, the input $v_N = 0$, and so, the model can be simplified, in this case, by removing the second block column (i.e. replacing it with an empty matrix with zero column dimension). Similarly, for $i = 1$, there is no need to retain the output w_1, from the perspective of the synthesis problem, and so, the corresponding model can be simplified by removing the second block row. Exploiting this yields the following specialisation of Thm. 4 in (Langbort et al., 2004), which leads to a state-space characterisation of the structured $\mathcal{H}_\infty$ synthesis problem introduced above; see Remark 2.

Theorem 1. There exists an interconnection $\hat{K}$ of distributed compensators $\hat{K}_i$ each with a state-space model

$$\begin{pmatrix} \dot{x}_i^K \\ w_i^K \\ u_i^K \end{pmatrix} = \mathcal{S}_i \begin{pmatrix} x_i^K \\ v_i^K \\ y_i^K \end{pmatrix}, \qquad (5)$$

where the dimension of x_i^K is equal to that of x_i in (4) and the interconnection corresponds to $v_i^K = w_{i+1}^K$ with the boundary condition $v_N^K = 0$, that

(1) stabilises the interconnection $G = (G_1, \ldots, G_N)$ of weighted generalised pool models (4), subject to $v_i = w_{i+1}$ with the boundary condition $v_N = 0$, and
(2) achieves $\|H(G, \hat{K})\|_\infty < \gamma$, where $H(G, \hat{K})$ denotes the closed-loop transfer function from the vector of signals $(n_1^T, \ldots, n_N^T)^T$ to the vector of signals $(z_1^T, \ldots, z_N^T)^T$ shown in Fig. 5,

if, and only if, there exist symmetric matrices $X_i^{tt} > 0$ and $Y_i^{tt} > 0$, for $i = 1, \ldots, N$, and symmetric matrices $X_{i,i-1} < 0$ and $Y_{i,i-1} < 0$, for $i = 2, \ldots, N$, such that for $i = 1, \ldots, N$,

$$\begin{pmatrix} X_i^{tt} & I \\ I & Y_i^{tt} \end{pmatrix} \geq 0, \quad \begin{pmatrix} X_{i,i-1} & -I \\ -I & Y_{i,i-1} \end{pmatrix} \leq 0,$$

$$(\Pi_i^X)^T \begin{pmatrix} 0 & X_i^{tt} & 0 & 0 & 0 & 0 \\ X_i^{tt} & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & -X_{i,i-1} & 0 & 0 & 0 \\ 0 & 0 & 0 & X_{i+1,i} & 0 & 0 \\ 0 & 0 & 0 & 0 & I & 0 \\ 0 & 0 & 0 & 0 & 0 & -\gamma^2 I \end{pmatrix} \Pi_i^X < 0$$

$$(\Pi_i^Y)^T \begin{pmatrix} 0 & Y_i^{tt} & 0 & 0 & 0 & 0 \\ Y_i^{tt} & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & -Y_{i,i-1} & 0 & 0 & 0 \\ 0 & 0 & 0 & Y_{i+1,i} & 0 & 0 \\ 0 & 0 & 0 & 0 & I & 0 \\ 0 & 0 & 0 & 0 & 0 & -\gamma^{-2} I \end{pmatrix} \Pi_i^Y > 0,$$

where

$$\Pi_i^X = \begin{pmatrix} I & 0 & 0 \\ A_i^{tt} & A_i^{ts} & B_i^{tn} \\ A_i^{st} & A_i^{ss} & B_i^{sn} \\ 0 & I & 0 \\ C_i^{tz} & C_i^{sz} & D_i^{zn} \\ 0 & 0 & I \end{pmatrix} \mathcal{N}_i^X,$$

$$\Pi_i^Y = \begin{pmatrix} (A_i^{tt})^T & (A_i^{st})^T & (C_i^{tz})^T \\ -I & 0 & 0 \\ 0 & -I & 0 \\ (A_i^{ts})^T & (A_i^{ss})^T & (C_i^{sz})^T \\ 0 & 0 & -I \\ (B_i^{tn})^T & (B_i^{sn})^T & (D_i^{zn})^T \end{pmatrix} \mathcal{N}_i^Y,$$

$\mathcal{N}_i^X$ is a matrix with columns that span the null-space of $\begin{pmatrix} C_i^{ty} & C_i^{sy} & D_i^{yn} \end{pmatrix}$, $\mathcal{N}_i^Y$ with columns that span the null-space of $\begin{pmatrix} (B_i^{tu})^T & (B_i^{su})^T & (D_i^{zu})^T \end{pmatrix}$, $A_1^{st} = []$, $A_1^{ss} = []$, $B_1^{sn} = []$, $A_N^{ts} = []$, $A_N^{ss}[]$, $C_N^{sz} = []$, $X_{1,0} = []$, $Y_{1,0} = []$, $X_{N+1,N} = []$, $Y_{N+1,N} = []$, and $[]$ denotes an empty matrix with the row or column dimension appropriately set to zero; this is on the understanding that, when an empty matrix appears as a block in a larger matrix, the remaining blocks in the same row/column take on empty status and zero dimension as required.

Remark 1. The convex conditions in Theorem 1 are coupled through $X_{i,i-1}$ and $Y_{i,i-1}$ in a fashion paralleling the directed interconnection structure of the weighted generalised pool models G_i and distributed compensators $\hat{K}_i$. Nevertheless, minimising γ, subject to such convex constraints is numerically tractable via standard and efficient software tools (Gahinet et al., 1995), even for fairly large problems.

Remark 2. Given matrices X_i^{tt}, Y_i^{tt}, $X_{i,i-1}$ and $Y_{i,i-1}$, which satisfy the collection of LMIs in Thm. 1 for a specific value of $\gamma > 0$, the matrices

$$\mathcal{S}_i = -\epsilon_i^{-1} \mathcal{U}_i \Phi_i \mathcal{V}_i^T (\mathcal{V}_i \Phi_i \mathcal{V}_i^T)^{-1}, \quad i = 1, \ldots, N, \qquad (6)$$

are such that the interconnection $\hat{K} = (\hat{K}_1, \ldots, \hat{K}_N)$ of the distributed compensators with state-space models (5), satisfies

$$\|H(G, \hat{K})\|_\infty < \gamma,$$

where $\Phi_i = (\epsilon_i^{-1} \mathcal{U}_i^T \mathcal{U}_i - \mathcal{G}_i)^{-1}$, $\mathcal{U}_i = \begin{pmatrix} (\mathcal{T}_{i,1}^u)^T \mathcal{X}_i^{tt} & 0 & (\mathcal{T}_{i,2}^u)^T \end{pmatrix}$, $\mathcal{V}_i = \begin{pmatrix} \mathcal{T}_i^y & 0 \end{pmatrix}$,

$$\mathcal{G}_i = \begin{pmatrix} (\mathcal{T}_i^{11})^T \mathcal{X}_i^{tt} + (\mathcal{X}_i^{tt})^T \mathcal{T}_i^{11} & (\mathcal{X}_i^{tt})^T \mathcal{T}_i^{12} & (\mathcal{T}_i^{21})^T \\ * & \begin{pmatrix} \mathcal{X}_{i+1,i} & 0 \\ 0 & -\gamma^2 I \end{pmatrix} & (\mathcal{T}_i^{22})^T \\ * & * & \begin{pmatrix} \mathcal{X}_{i,i-1}^{-1} & 0 \\ 0 & -I \end{pmatrix} \end{pmatrix},$$

$$\mathcal{X}_i^{tt} = \begin{pmatrix} X_i^{tt} & Z_i^{tt} \\ * & I \end{pmatrix}, \quad Z_i^{tt} (Z_i^{tt})^T = X_i^{tt} - (Y_i^{tt})^{-1},$$

$$\mathcal{X}_{i,i-1} = \begin{pmatrix} X_{i,i-1} & I \\ I & Z_{i,i-1} \end{pmatrix}, \quad Z_{i,i-1}^{-1} = X_{i,i-1} - (Y_{i,i-1})^{-1},$$

$$\begin{pmatrix} \mathcal{T}_i^{11} & \mathcal{T}_i^{12} \\ \mathcal{T}_i^{21} & \mathcal{T}_i^{22} \end{pmatrix} = \begin{pmatrix} \begin{pmatrix} A_i^{tt} & 0 \\ 0 & 0 \end{pmatrix} & \begin{pmatrix} A_i^{ts} & 0 & B_i^{tn} \\ 0 & 0 & 0 \end{pmatrix} \\ \begin{pmatrix} A_i^{st} & 0 \\ 0 & 0 \\ C_i^{tz} & 0 \end{pmatrix} & \begin{pmatrix} A_i^{ss} & 0 & B_i^{sn} \\ 0 & 0 & 0 \\ C_i^{sz} & 0 & D_i^{zn} \end{pmatrix} \end{pmatrix},$$

$$\begin{pmatrix} \mathcal{T}_{i,1}^u \\ \mathcal{T}_{i,2}^u \end{pmatrix} = \begin{pmatrix} \begin{pmatrix} 0 & 0 & B_i^{tu} \\ I & 0 & 0 \end{pmatrix} \\ \begin{pmatrix} 0 & 0 & B_i^{su} \\ 0 & I & 0 \\ 0 & 0 & D_i^{zu} \end{pmatrix} \end{pmatrix}, \quad \mathcal{T}_i^y = \begin{pmatrix} 0 & I & 0 & 0 & 0 \\ 0 & 0 & 0 & I & 0 \\ C_i^{ty} & 0 & C_i^{sy} & 0 & D_i^{yn} \end{pmatrix},$$

ϵ_i is a positive scalar such that $\epsilon_i \ll 1/\mu_i$, and

$$\mu_i = \lambda_{\max}((\mathcal{U}_i^+)^T (\mathcal{G}_i - \mathcal{G}_i \mathcal{U}_i^\perp ((\mathcal{U}_i^\perp)^T \mathcal{G}_i \mathcal{U}_i^\perp)^{-1} (\mathcal{U}_i^\perp) \mathcal{G}_i) \mathcal{U}_i^+),$$

with $+$ denoting the Moore-Penrose inverse, $\perp$ the ortho-completion and $\lambda_{\max}$ the largest eigenvalue of a matrix.

Remark 3. The construction of $\mathcal{S}_i$ in (6) is a direct application of Theorem 1 in (Iwasaki et al., 1994), exploiting the full-rank properties of $\mathcal{T}_i^u$ and $\mathcal{T}_i^y$ for the particular problem considered here. The construction from (Langbort et al., 2004) proves more complex, and unfortunately, leads to numerical difficulties.

5. SIMULATION AND FIELD TEST RESULTS

Figure 7 shows the simulated water-level errors for the five pools of Coleambally Channel, Number 6, NSW, in response to a large step change in the flow out of pool$_5$, when

- operating under a distributed distant downstream controller that was designed using the models, loop-shaping weights and optimal synthesis technique described in the preceding two sections (achieved $\gamma_{\min} = 2.88$), and when

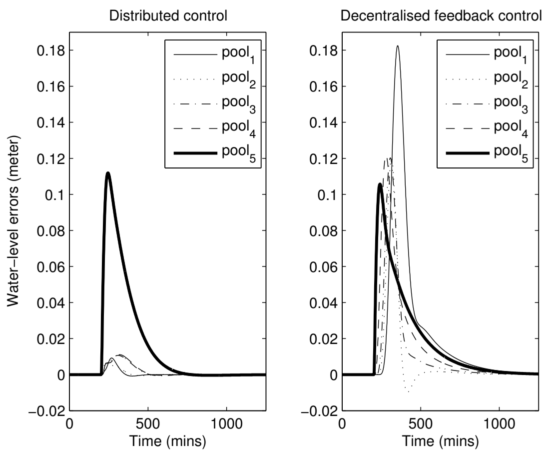

Fig. 7. Simulated water-level errors

- operating under a decentralised controller (Cantoni et al., 2007), where the isolated local compensator for pool_i is taken to be the loop-shaping weight W_i.

The optimal controller clearly achieves significantly improved performance in terms water-level error propagation, while yielding very similar local load-disturbance rejection performance (see pool_5). This comes at the very modest expense of four extra scalar states in each local compensator and the local information exchange overhead.

The distributed controller has also been tested in the field; specifically on gates Coly6-1 through Coly6-5. The test data are shown in Fig. 8, which provides both water levels and gate flows over a period of many hours. The water-level set-points were set to $r_1 = 1.450$m, $r_2 = 1.510$m, $r_3 = 1.554$m, $r_4 = 1.500$m and $r_5 = 1.520$m. Before 398 min (thick solid lines), gates Coly6-1 to Coly6-5 are controlled by the distributed controller, designed as described above. At 398 min, control is transferred, with the assistance of anti-windup/bumpless-transfer compensation, to a decentralised controller with additional decentralised feed-forward action – see the end of Section 2 and (Weyer, 2007). To perturb the system, the flow over Coly6-6 is manually increased from 20ML/day to 45ML/day at 159 min. This off-take load-disturbance is ceased at 462 min, which perturbs the system again. The performance of the optimal distributed controller is clearly superior. Indeed, when the off-take is removed the decentralised controller causes gate saturation.

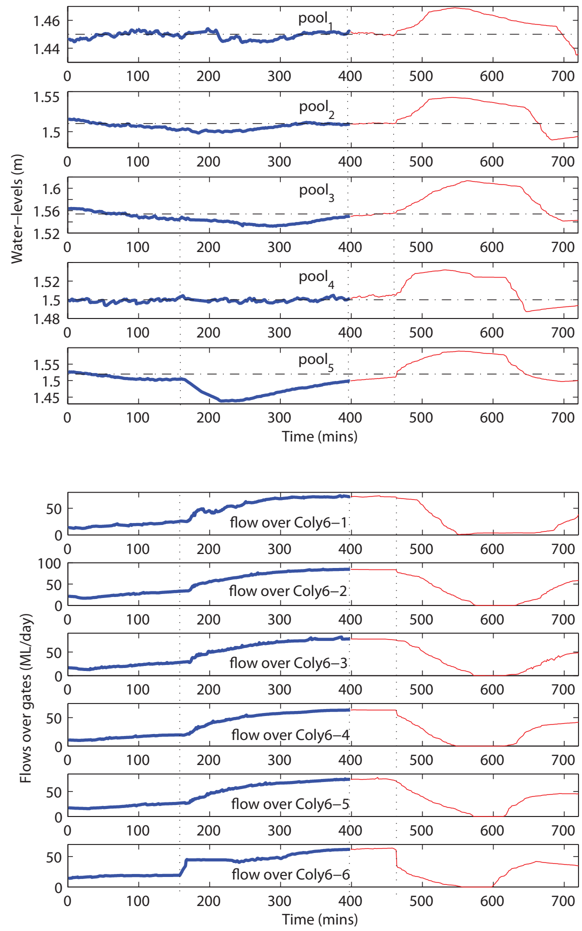

Fig. 8. Field test data

REFERENCES

Bos, M.G. (1978). *Discharge measurement structures.* International Institute for Land Reclamation and Improvement/ILRI, Waageningen, Netherlands.

Cantoni, M., Weyer, E., Li, Y., Ooi, S.K., Mareels, I., and Ryan, M. (2007). "Control of large-scale irrigation networks." Special Issue on the Emerging Technology of Networked Control Systems, *Proceedings of the IEEE*, vol. 95(1), pp. 75–91.

Clemmens, A.J., et al. (1998). Special issue on control and automation of irrigation canals. *J. Irrigation and Drainage Engineering*, vol. 124(1), pp. 1-62.

Gahinet, P., Nemirovski, A., Laub, A., and Chilali, M. (1995). *LMI Control Toolbox User's Guide.* The Mathworks, Natick, MA.

Iwasaki, T., and Skelton, R. (1994). "All controllers for the general $\mathcal{H}_\infty$ control problem: LMI existence conditions and state space formulas." *Automatica*, 1994, vol. 30(8), pp. 1307–1317.

Langbort, C., Chandra, R., and D'Andrea, R. (2004). "Distributed control design for systems interconnected over an arbitrary graph." *IEEE Transactions on Automatic Control*, vol. 49(9), pp. 1502–1519.

Li, Y., Cantoni, M., and Weyer, E. (2004). "Design of centralised controller for an irrigation channel using $\mathcal{H}_\infty$ loop-shaping." *Proceedings of Control 2004*, Bath, UK.

Li, Y., Cantoni, M., and Weyer, E. (2005). "On water-level error propagation in controlled irrigation channels." *Proceedings of the combined 44th IEEE CDC and ECC 2005*, Seville, Spain.

Li, Y. (2006). "Robust control of open water channels." PhD Dissertation, The University of Melbourne.

Litrico, X., and Fromion, V. (2003). "Advanced control politics and optimal performance for an irrigation canal." *Proceedings of the 2003 ECC*, Cambridge, UK.

Mareels, I., Weyer, E., Ooi, S.K., Cantoni, M., Li, Y., and Nair, G. (2005). "Systems engineering for irrigation systems: Successes and challenges." *Annual Reviews in Control (IFAC)*, vol. 29, pp. 191–204.

Ooi, S.K., Krutzen, M., and Weyer, E. (2003). "On physical and data driven modeling of irrigation channels." *Control Engineering Practice*, vol. 13(4), pp. 461–471.

Skogestad, S., and Postlethwaite, I. (1996). *Multivariable feedback control analysis and design.* Wiley, Chichester.

Weyer, E. (2001). "System identification of an open water channel." *Control Engineering Practice (IFAC)*, vol. 9, pp. 1289-1299.

Weyer, E. (2003). "LQ control of an irrigation channel." *Proceedings of the 42nd IEEE CDC*, pp. 750–755, Hawaii, USA.

Weyer, E. (2007). "Control of open water channels." *IEEE Transaction on Control Systems Technology*, to appear.