INTERACTIVE CONTROL ENGINEERING
Principles and Applications

Volume 9 Model Predictive Control and Economic Optimization

CHAPTER 5

Irrigation Canal Simulator for Distant Downstream Control, Resonance and Releasing Water Ahead of Demand

Pages 98–124

Edited by John D. Hedengren

How to cite this chapter. "Irrigation Canal Simulator for Distant Downstream Control, Resonance and Releasing Water Ahead of Demand," in *Model Predictive Control and Economic Optimization*, John D. Hedengren, Ed., vol. 9 of *Interactive Control Engineering: Principles and Applications*. Advanced Process Solutions, LLC, 2026, ch. 5, pp. 98–124.

The page numbers in this file are those of the printed volume.

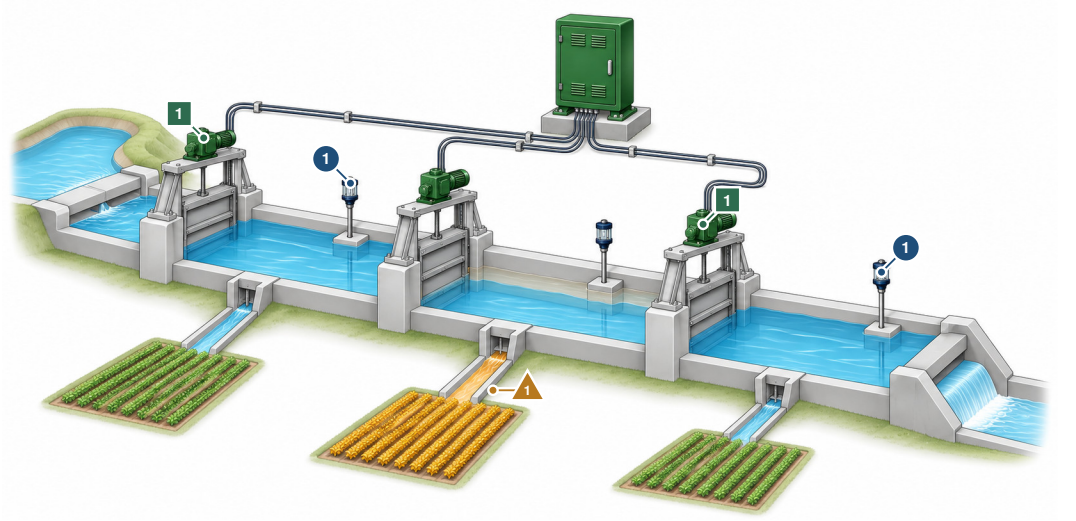
The app for this chapter
https://apmonitor.com/apps/irrigation-canal/

The volume online: contents, other chapters and citation records
https://apmonitor.com/apps/books/v09-mpc-optimization/#irrigation-canal

Copyright © 2026 Advanced Process Solutions, LLC. All rights reserved. Produced by Advanced Process Solutions, LLC.

CHAPTER 5

Irrigation Canal Simulator for Distant Downstream Control, Resonance and Releasing Water Ahead of Demand



The controller moves 1 head gate and two cross-regulators

The controller reads 1 water level at each reach end

What disturbs it 1 farm orders, booked or unannounced

Each gate answers to the level at the far end of the reach below it, so the water it releases is seen there only after a wave has crossed the reach. Illustration, not to scale.

IN PLAIN TERMS Sending Water Down a Canal on Time

The main idea. An irrigation canal carries water from a reservoir to farms through a chain of gates. Water let through a gate needs many minutes to reach the end of the stretch below, so each gate must act early.

The challenge. Each stretch also sloshes when its gate is pushed too hard. A controller that plans all the gates at once can share a surprise between the stretches, but it needs the farmers' booked orders.

Why it matters. Most of the fresh water that people take goes to crops, so water lost at the end of a canal is a real loss.

THREE WAYS TO READ THIS CHAPTER

Explore	12 min	the first page, the fact in Section 5.1, the history in Section 5.2, the try-it segments in Sections 5.5.1 and 5.5.6, the second fact, and the key points
Study	55 min	the whole chapter with its three pause points and the twelve exercises at its end
Practice	25 min	the first page, the results (Section 5.5), Technology in Practice (Section 5.6), Limitations (Section 5.7) and the open challenges

The control question Each reach of an irrigation canal stores water, delays every change by a gravity wave's travel time and rings at the wave's round trip. Which gate should answer to which level, how hard can its loop be tuned, and what is a published order, or one plan for every gate, worth against feedback?

Learning objectives

- Compute the wave travel time and round trip of a reach, and the lead at which each gate releases a published order.
- Compare local upstream and distant downstream control in water spilled against level deviation, and explain where each pays.
- Explain how a level filter placed between the crossover and the resonance lets a faster gain survive.
- Explain why one plan for three gates shares an unannounced offtake over the reaches, and why it trails gate stroking on published orders.
- Tell what exchanged plans recover over decentralized ones, and what is left to cooperation.

Prerequisites Sections 1.6.2 (velocity form), 1.4.2 and 1.4.3 (prediction model and quadratic program), 1.4.5 (input-disturbance estimate), 1.4.11 (distributed plans) and 1.6.4 (how an exercise names a run), and Chapter 2 for an integrating plant behind a dead time.

The app https://apmonitor.com/apps/irrigation-canal/

See also Volume 5, Chapter 2 for fuel that burns after a transport delay and Volume 5, Chapter 4 for a process where every useful move is a prediction. Volume 3, Chapter 6 for feedforward measured ahead of a long delay, and Volume 3, Chapter 2 for a loop structure that decides before any tuning.

5.1 Introduction

A canal reach looks like a tank, and a tank is the first process a control course treats. Water enters through a gate at the head, leaves through a turnout and a gate at the tail, and the level integrates the difference. The reach differs from a tank in two ways that decide its control. A change at the head reaches the tail only after a gravity wave has crossed the reach, which on the canal modeled here takes 11 to 19 min, and a reach closed at both ends by flow-controlled gates reflects that wave and resonates at its round trip [1, 2]. A controller designed for the tank is either too slow for the delay or fast enough to feed the resonance.

> ✦ DID YOU KNOW? **Two Hundred Times Slower Than a Pipe**
>
> A flow change travels along a canal as a translatory wave, typically at 3 to 6 m/s, about 200 times slower than a pressure wave in a pipe, so a canal answers a change in demand more slowly than a pipeline [3]. The gravity wave of this chapter's canal runs at about 4.4 m/s (Section 5.3).

The choice that matters most is made before any tuning, as it is for the structure of the level loop of Volume 3, Chapter 2. In local upstream control each gate holds the level just upstream of itself and the head releases a fixed supply, which is easy and keeps every loop local, but whatever the farms do not take spills at the tail. In distant downstream control each gate holds the level at the far end of the reach below it, so the canal releases only what is drawn, but every correction must cross a reach before it can be seen [4]. Published orders offer a way out of the delay. A farmer who books water in advance lets the operator release it early, and doing so gate by gate, one travel time ahead of each order, is the scheduling method known as gate stroking [5, 6].

This chapter is built against the misconception that a canal reach is a tank. The broader class is transport with storage and reflection controlled from its boundaries, which includes rivers regulated by weirs, drainage networks, district heating lines and gas pipelines, and it shares three lessons: the philosophy of control decides whether errors are paid in waste or in delay, the tuning is capped by both a delay and a resonance, and known future demand belongs in a schedule rather than a feedback

loop. Irrigation takes most of the fresh water that people withdraw, which makes operational losses in canals a question of real consequence [7, 8].

This chapter works with a three-reach canal simulator in which the control philosophy, the tuning, the level filter and the use of published orders can be changed while the canal runs, and the outcome is read in water released, delivered and spilled and in level deviations. The hydraulic model, downstream control and gate stroking are established methods. A model predictive controller is offered as a third philosophy, with decentralized and distributed variants, and the chapter compares it with the loops and the variants with one another on the same information about the orders. Section 5.2 traces the automation of canals. Section 5.3 sets out a compact local-inertia model that keeps the wave, the delay and the resonance, with gates that behave as flow sources. Section 5.4 describes the level loops, the order feedforward and the predictive controller with its three architectures. Section 5.5 reports what was measured, starting with the wave timed by hand in Section 5.5.1 and ending with the graded challenge in Section 5.5.7. Section 5.6 describes the gauges, gates and control platforms of real canals, Section 5.7 states where the model stops, and the exercises in Section 5.9 run the same cases in the application.

5.2 Automation of Irrigation Canals

Most large canals were built for supply-driven operation. The operator sets a release at the head from the expected demand, gatekeepers adjust the cross-regulators by hand to hold levels upstream of the turnouts, and water that is released but not taken runs to the tail. Hydraulic automatic gates, which hold an upstream level with a float and a counterweight and need no power, made local upstream control automatic, and it remains the most common arrangement in service [9]. The equations that describe the flow, the one-dimensional Saint-Venant equations, were set out for engineering use and solved numerically by the end of the 1960s [10], and methods to compute gate movements that move a canal from one steady state to another without transients followed at the same time [5, 6].

> ⧖ HISTORY **1952: A Little Man at the Check Gate**
>
> On one long supply canal in California, holding the level above a check gate once needed an operator stationed at the gate, and even then the level wandered. Agency staff built a simple device that moved the gate whenever the upstream level left a preset deadband around its target. A canal superintendent called it the best little man on his staff, and the Little-Man went on trial in 1952 as an immediate success. Its successors filtered the level, first through a capillary tube into a second well, which dust and dead insects plugged, then with an electronic circuit, and added a reset that removed the offset of proportional action. That electronic controller held levels downstream of the gates on two canals, with 12 and 3 check structures, from 1974 [3].

Downstream control, in which each gate answers to a level it cannot see, needed remote measurement and a way to live with the delay. A classification of canal control algorithms by controlled variable, manipulated variable, logic and design method, published together with a set of standard test canals, gave the field a common vocabulary in 1998 [4]. Simple controllers based on an integrator-delay model of each reach were developed and tested [2, 11], simplified models that retain the first resonance were derived for controller design [1], and robust tuning of distant downstream PI controllers was worked out for a single reach [12]. The modeling and control of hydrosystems has since been collected in a monograph [13].

Field programs moved the ideas onto canals. Downstream feedback control was analyzed and simulated for branching networks [14], system identification from operating data produced models accurate enough to design controllers without detailed hydraulics [15], and programmable logic controller automation was applied to working canals in modernization projects [8]. Large Australian networks were automated end to end with a combination of these ideas, and the experience has been reviewed as a systems engineering success with continuing challenges [7]. Centralized model

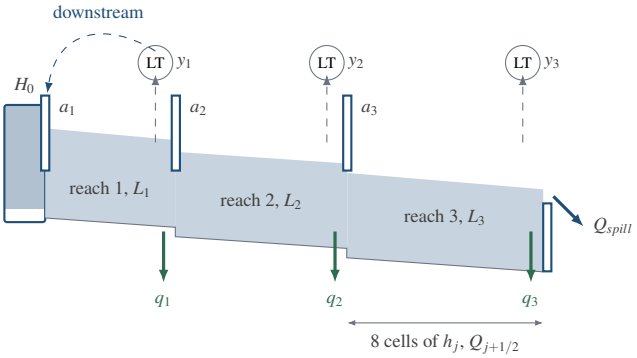


Figure 5.1 Three-reach irrigation canal. A reservoir at level H_0 feeds the canal through the head gate of opening a_1, and cross-regulators of openings a_2 and a_3 divide it into reaches of lengths L_1, L_2 and L_3, each solved on eight cells of depth h_j and face flow $Q_{j+1/2}$. Farms draw q_1, q_2 and q_3 at the reach ends, where the level transmitters read y_1, y_2 and y_3, and the escape weir spills Q_{spill} at the tail. In distant downstream control, as drawn for reach 1, each gate answers to the level at the far end of the reach below it.

Table 5.1 Variables at the Design Operating Point

Symbol	Role	Nominal	Range
y_1, y_2, y_3	CV, reach end depths	2.00 m	SP 1.6–2.4
Q_1, Q_2, Q_3	MV, gate flow setpoints	5.0, 3.5, 2.0 m^3/s	0–15
a_1, a_2, a_3	states, gate openings	—	0–1.6 m
$h_j, Q_{j+1/2}$	states, 24 depths, 21 flows	—	$h \geq 0.05$ m
q_1, q_2, q_3	DV, farm orders	1.5, 1.5, 2.0 m^3/s	0–4
—	DV, published schedule	per scenario	—
n_y	DV, level chop (1σ)	4 mm	0–30
Q_{sup}	MV, head supply (upstream)	5.0 m^3/s	0–12

predictive control, which sets every gate from one optimization over a horizon, was evaluated on ASCE Test Canal 1, one of the standard test canals published with that 1998 classification [16]. Distributed model predictive control, in which the controller of each pool solves its own problem and exchanges information with its neighbors, has been applied to the pools of irrigation canals [17], within a wider taxonomy of decentralized, distributed and hierarchical architectures [18, 19].

5.3 Process Model

A reservoir at a fixed level feeds a rectangular canal 6 m wide through a head gate. Two gated cross-regulators split the canal into reaches of 3, 4 and 5 km, each ending in a farm turnout, and the last reach ends in an escape weir. At design the canal carries 5 m^3/s at about 2 m depth: 1.5 to the first farm, 1.5 to the second and 2.0 to the tail. The mean velocity is about 0.42 m/s and the Froude number about 0.1. The arrangement is shown in Figure 5.1.

Table 5.1 classifies the variables. Forty-eight states make up the model: eight depths and seven internal flows in each of the three reaches, and the three gate openings.

Table 5.2 Canal Parameters, All Editable in Show Model

Symbol	Meaning	Value
B	bottom width	6.0 m
S_0	bed slope	1.0×10^{-4}
n	Manning roughness	0.018 s/m$^{1/3}$
L_1, L_2, L_3	reach lengths	3, 4, 5 km
H_0	reservoir level	3.0 m
C_d	gate coefficient	0.60
a_{max}	maximum opening	1.6 m
—	gate motor speed	0.24 m/min
—	weir crest (depth)	2.3 m
h_t, h_{min}	turnout invert, full-order depth	1.2, 1.8 m
g	gravity (fixed)	9.81 m/s^2

5.3.1 Local-Inertia Saint-Venant Equations

Each reach is divided into $N_c = 8$ cells, with depths at the cell centers and flows at the faces between them. Mass is balanced on each cell and momentum on each internal face, Equations (5.1) and (5.2).

$$B\,\Delta x\,\frac{\mathrm{d}h_j}{\mathrm{d}t} = Q_{j-\frac{1}{2}} - Q_{j+\frac{1}{2}} - q_r\,\delta_{j,N_c} \tag{5.1}$$

$$\frac{\mathrm{d}Q_{j+\frac{1}{2}}}{\mathrm{d}t} = -gA\left(\frac{z_{j+1}-z_j}{\Delta x} + S_f\right), \quad S_f = \frac{n^2 Q|Q|}{A^2 R^{4/3}} \tag{5.2}$$

where B is the width, $\Delta x = L_r/N_c$ the cell length, $z = z_b + h$ the water surface, $A = Bh$ the area and $R = Bh/(B+2h)$ the hydraulic radius at the face, and $n = 0.018$ the Manning roughness. The advective term of the full momentum equation is dropped. This local-inertia form is accurate at the low Froude numbers of an irrigation canal and is the formulation used for efficient inundation modeling [20]. It retains gravity waves at $\sqrt{gh} \approx 4.4\,\mathrm{m/s}$ and their reflections, so the travel times of the three reaches at 2 m depth are $\tau_r = L_r/\sqrt{2g}$, about 11, 15 and 19 min. The bed falls at $S_0 = 10^{-4}$ and drops a further 0.25 m at each cross-regulator.

5.3.2 Gates, Turnouts and the Weir

Each gate is a submerged orifice across the full width, driven by the surface difference across it, Equation (5.3).

$$Q_g = C_d\,B\,a\,\sqrt{2g\,|z_u - z_d|}\;\mathrm{sgn}(z_u - z_d) \tag{5.3}$$

with $C_d = 0.6$. The head gate draws from the reservoir at $H_0 = 3.0\,\mathrm{m}$ above the head bed. Each gate's motor follows its commanded opening with a 10 s lag limited to 0.24 m/min, between zero and 1.6 m. A turnout takes its full order while the depth at the reach end is at least 1.8 m, and less below that, Equation (5.4).

$$q = q_{order}\sqrt{\frac{h - h_t}{h_{min} - h_t}}, \quad h_t < h < h_{min} \tag{5.4}$$

with the turnout invert at $h_t = 1.2\,\mathrm{m}$. The escape weir at the tail is sharp-crested across the full width with its crest at 2.3 m depth, $Q_{spill} = 1.7\,B\,(h - 2.3)^{3/2}$. Freeboard is flagged when a reach end exceeds 2.7 m.

With the values of Table 5.2, every case opens on a steady backwater profile marched upstream from the setpoint depth in every reach, with each gate at the opening that passes its reach's flow.

5.3.3 Assumptions

The reaches are prismatic rectangular channels with uniform slope and roughness. Real canals are trapezoidal and carry structures such as siphons and drops that change the wave behavior [13].

Eight cells per reach reproduce the integrator, the delay and the first resonances, and smear the wave front over several hundred meters. Resonant periods are approximate.

Gates are submerged orifices with a constant coefficient and a smooth motor. Gate motors on real canals often move in discrete steps with deadbands, which adds its own limit cycling.

The reservoir level is constant and there is no seepage or evaporation, so water released is delivered, stored or spilled.

Orders are either published exactly or not at all. Real schedules are uncertain in timing and amount.

5.4 Implemented Control

The controllers run every 30 s and the hydraulics advance by fourth-order Runge-Kutta steps of 5 s, on the engine of Section 1.6.1.

5.4.1 Level Loops in Flow Units

Every level loop is a PI controller in the velocity form of Section 1.6.2, and it computes a gate flow setpoint Q_k every $\Delta t_c = 30\,\mathrm{s}$, as shown in Equation (5.5).

$$\begin{aligned} e_k &= \gamma\left(y_{SP} - y_k^f\right), \qquad P_k = K_c e_k \\ s_k &= s_{k-1} + P_k - P_{k-1} \\ c_k &= Q_{ff,k} + s_k + I_{k-1} \\ Q_k &= \mathrm{clamp}\left(c_k, 0, Q_{max}\right) \\ I_k &= I_{k-1} + \frac{K_c}{\tau_I} e_k \Delta t_c + \frac{\Delta t_c}{T_t}\left(Q_k - c_k\right) \end{aligned} \tag{5.5}$$

where γ is the controller direction, y^f the filtered level, Q_{ff} the feedforward bias described below, $Q_{max} = 15\,\mathrm{m^3/s}$ and $T_t = \tau_I$ the tracking time of the back-calculation. The recursion is that of Equations (1.14) to (1.17) without the derivative, with γ for the action a, s_k for S_k, the demanded flow c_k for v_k and the feedforward Q_{ff} as the bias u_b, and the integral used at scan k was accumulated through scan $k-1$. The gate opening follows from inverting Equation (5.3) with the water surfaces on either side and the declared gate coefficient and width, $a = Q/(C_{d,0} B_0 \sqrt{2g\,\Delta z})$ with $\Delta z \geq 5\,\mathrm{mm}$, which turns each gate into a flow source at its declared rating and keeps the controller gain in flow units. A gate whose coefficient or width is edited passes a different flow, and the level loop corrects it. Every controller reads the declared gate rating, the declared reach lengths and the turnout meters, which measure what each turnout delivers, so an edit of the canal is model mismatch for the loops as it is for the MPC. The opening is limited to a_{max} after the PI, which does not see that limit, so the integral of a loop whose gate is fully open keeps integrating. In the cases of Section 5.5 only the unstable loops reach either limit. Each level measurement passes a first-order filter τ_f before the controller. The defaults are $K_c = 15\,\mathrm{m^3/s}$ per meter, $\tau_I = 3600\,\mathrm{s}$ and $\tau_f = 200\,\mathrm{s}$, one tuning shared by the three reaches.

In distant downstream control, gate r holds y_r at the far end of reach r with direct action ($\gamma = +1$), so a falling level raises the flow into the reach, and the bias is the feedforward described below. In local upstream control, the head releases the fixed supply Q_{sup} and gate $r+1$ holds y_r just upstream of itself with reverse action ($\gamma = -1$) and no bias, so a rising level opens it. In manual mode the operator sets the head release, the cross-regulators hold their openings so their flows follow the levels, and each integral is set every scan so that Equation (5.5) reproduces the actual gate flow, which makes the return to automatic bumpless.

5.4.2 Order Feedforward

In downstream control each gate can add the orders of every farm below it. With the feedforward set to *now*, gate r adds $\sum_{k \geq r} q_k(t)$, which decouples the reaches but ignores the delay. With it set to

Table 5.3 Reach Models of the MPC Against the Nominal Values

Reach	BL_r (m^2)	$A_{s,r}$ (m^2)	τ_r (s)	θ_r (s)	First rise (s)	RMS (mm)
1	18000	16217	677	215	580	16.1 / 2.5
2	24000	21941	903	295	780	11.8 / 2.7
3	30000	28361	1129	330	980	10.0 / 2.0

First rise is the time for the reach end to move 1 mm after its own gate steps by 0.3 m^3/s. RMS is the error over 3 h in predicting that response, nominal against identified.

ahead, gate r adds each order as it will be one travel time later per reach between that gate and that farm, the travel times of the declared reach lengths, Equation (5.6).

$$Q_{ff,r}(t) = \sum_{k=r}^{3} q_k \left(t + k_{lead} \sum_{m=r}^{k} \tau_m \right) \tag{5.6}$$

with $k_{lead} = 1$ by default. The head gate therefore releases for the tail farm three reaches early and the last gate one reach early, which is gate stroking in its simplest form [5, 6]. An unpublished offtake enters through the turnout meters, which read the flow each turnout delivers, only when it happens. A change of feedforward setting is absorbed into each integral, which keeps the gate flows from jumping.

5.4.3 Model Predictive Control

A model predictive controller (MPC) is offered as a third control philosophy. Distant downstream control remains the default, both when the application opens and in every scenario other than scenarios 10 and 11, and its three PI loops are the fallback. Every $T_s = 300$ s, or every tenth scan, one quadratic program sets the three gate flow setpoints together over $N = 36$ steps, or 3 h, in 11 move blocks whose lengths in steps are 1 four times, then 2 twice, then 3, 4, 6, 9 and 6. The plan's first move alone goes to the gates and stays until the next decision, the receding horizon of Section 1.4.3 [21, 22]. Each planned flow passes through the same inversion of Equation (5.3), opening limit and 0 to 15 m^3/s clamp as a PI output, so every gate remains a flow source under all three philosophies.

The controller's model is one integrator behind a dead time per reach, the structure long used for canal controller design [2], Equation (5.7). Chapter 2 identifies the same structure as a step response for a balloon's altitude.

$$A_{s,r} \frac{\mathrm{d}y_r}{\mathrm{d}t} = Q_r(t - \theta_r) - Q_{r+1}(t) - q_r(t) + d_r, \quad Q_4 \equiv 0 \tag{5.7}$$

Here $A_{s,r}$ is the reach's effective surface, θ_r its dead time and d_r an unexplained inflow supplied by the estimator, which also absorbs the tail spill and the model's error. A turnout delivering less than its order enters through the turnout meters the preview reads. The numbers come from step tests run on the simulator at its declared design point, with every parameter at its declared value, orders of 1.5, 1.5 and 2.0 m^3/s, levels at 2.0 m and no level chop. The three gate flows were held as flow sources, one gate flow or one order was stepped by ± 0.3 m^3/s after 30 min, and an integrator with a dead time was fitted over the following 3 h in both directions. Each surface is the mean of the fits for every flow into and out of its reach, so balanced flows predict no drift, and each dead time is that of the reach's own gate. The values, frozen in the controller so that an edit of the canal becomes model mismatch, are listed in Table 5.3. An independent refit agrees within 2 m^2 and 2 s. Over 3 h after a step the identified model predicts the eight input and output pairs within 2.0 to 3.7 mm RMS, against 5.6 to 16.1 mm for the nominal surfaces BL_r and wave travel times τ_r.

The dead times are the best straight-line fit over the 3 h horizon, not the wave's arrival. A reach end first moves 1 mm after 580, 780 and 980 s, and a fit over 1 h returns 266, 404 and 478 s, so the

Table 5.4 What Each Structure Is Told About the Farm Orders

MPC preview	Information	Same as
Off	the levels only	feedback only
Now	the turnout meters now, held over 3 h	feedforward *now*
Schedule	the published orders over 3 h, as 5 min means	feedforward *ahead*

model has no waves, no resonance and no reflection. With $\theta_r = (n_r + f_r)T_s$, $0 \le f_r < 1$, a zero-order hold gives the exact sampled model of Equation (5.8),

$$\begin{aligned} y_{r,k+1} = y_{r,k} + \frac{T_s}{A_{s,r}}\Big[&(1-f_r)Q_{r,k-n_r} + f_r Q_{r,k-n_r-1} \\ &- Q_{r+1,k} - q_{r,k} + d_{r,k}\Big], \qquad d_{r,k+1} = d_{r,k} \end{aligned} \tag{5.8}$$

where $n_r = 0$, 0 and 1 and $f_r = 0.717$, 0.983 and 0.100, so the inflow of a period comes partly from the decision before.

The decisions are the moves of the three gate flows from the flows applied now, constant within each block. Each plan minimizes Equation (5.9), in which $\hat{y}_{r,k}$ is a predicted level, j counts the blocks and $\Delta Q_{r,0}$ is measured from the applied flow.

$$\begin{aligned} \mathscr{J} = &\sum_{k=1}^{N}\sum_{r=1}^{3}\left(\frac{\hat{y}_{r,k} - y_{SP,r}}{0.01\,\mathrm{m}}\right)^2 + \lambda \sum_{j=0}^{10}\sum_{r=1}^{3}\left(\frac{\Delta Q_{r,j}}{0.1\,\mathrm{m^3/s}}\right)^2 \\ &+ \sum_{r=1}^{3}\left(\kappa\,\varepsilon_r^2 + \nu\,\varepsilon_r\right) \\ \text{s.t. } 0 \le Q_{r,j} \le \bar{Q}_r, &\quad |\Delta Q_{r,j}| \le C_d B\sqrt{2g\,\Delta z_r}\,\dot{a}_{max} T_s \\ \underline{y}_{r,k} - \varepsilon_r \le \hat{y}_{r,k} &\le \bar{y}_{r,k} + \varepsilon_r, \quad \varepsilon_r \ge 0 \end{aligned} \tag{5.9}$$

With $\lambda = 3$ a move of 0.1 $\mathrm{m^3/s}$ costs as much as a 1.7 cm level error held for one step. The gate limits are hard. The capacity is $\bar{Q}_r = \min\left(15, C_d B a_{max}\sqrt{2g\,\Delta z_r}\right)$, where Δz_r is the water-surface difference across gate r now, as the inversion reads it, C_d and B are the declared values, and a_{max} is the opening limit read as the inversion clips it. The move limit uses the declared motor speed $\dot{a}_{max} = 0.24\,\mathrm{m/min}$. The level bands are soft, with one slack ε_r per reach priced at $\kappa = 10^7$ per $\mathrm{m^2}$ and $\nu = 10^5$ per m, as in Equation (1.7). Each band, from $\underline{y}_{r,k}$ to $\bar{y}_{r,k}$, is the loops' ± 5 cm about the setpoint, raised where necessary to 2 cm above the full-order depth h_{min} and, at the tail, lowered to 2 cm under the weir crest, so the plan holds water back rather than spill it. When a level already lies outside its band, the band starts at that level and closes back with a 30 min time constant. With the band prices set to zero, every compared case changes its summed IAE by at most 0.56 % and its worst error by at most 0.011 cm, so the bands decide nothing in the comparison below. Each program goes to the dual active-set solver of Goldfarb and Idnani [23], and its result is discarded when the solver does not report it optimal, when a value is not finite or when a flow leaves its bounds.

The orders over the horizon come from a preview with three settings that mirror the feedforward, as listed in Table 5.4. Each reads the same function of the turnout meters and the published schedule that the matching feedforward reads, so each knows exactly what that feedforward knows. The feedforward *ahead* looks one travel time per reach ahead, at most 45 min at $k_{lead} = 1$, and the Schedule preview looks 3 h ahead at the same published orders. An unannounced offtake reaches every structure only when the turnout begins to carry it. Replay tests in which two canals differ only in orders not yet due find identical plans until each preview may see the difference.

A steady-state Kalman filter on (y_r, d_r) runs every 300 s under every philosophy. It predicts with the mean applied gate flows and the mean orders the preview allows over the last period, and corrects with the three level transmitters ahead of the level filter. Its variances per decision are 9×10^{-6} m^2 for each level, standing for the waves the model leaves out, $2.5 \times 10^{-3} \rho$ (m^3/s)^2 for each inflow and 2.5×10^{-5} m^2 for each transmitter, with $\rho = 0.01$. The random-walk inflow is the input-disturbance estimate of Section 1.4.5, and it makes the plan offset-free [24], slowly. The filter's time constants are 162, 219 and 283 min.

The horizon was fixed before tuning. With $\lambda = 1$ and $\rho = 1$, a horizon of 36 steps gave a summed IAE within 1.2 % of 120 steps on four offtake cases and 3.1 % lower on the challenge. One pass over nine values of λ from 0.01 to 100 and five of ρ from 0.01 to 100 then scored scenario 10 from 0 to 10 h by its summed IAE relative to distant downstream control with feedforward *ahead*, admitting only settings that move the gates no more than that structure, 3.13 m. Of 45 settings 23 were admitted, and the best, $\lambda = 3$ and $\rho = 0.01$, scored 0.79. Settings with $\rho \geq 10$ and small λ ring the reaches by up to 64 cm, most likely because the estimate then follows the waves. Nothing was retuned afterwards for any case.

The three PI loops track the applied flows every scan while the MPC acts, so a rejected plan hands the gates to them without a bump until a later plan solves and resumes control. Three forced rejections with the chop off move the gates by 1.31×10^{-3} m^3/s at the hand-over, and none of the measured runs had a rejected plan. A switch between distant downstream control and the MPC moves a gate flow by at most 0.019 m^3/s. From local upstream control the flows are held until the first plan, which then moves each gate by 3.6 to 6.7 m^3/s, its answer to the levels the upstream loops left. The opening limit, the inversion, the weir, the turnout cut-back and the freeboard alarm act under the MPC as under the loops. The controller reads the three transmitters, the applied flows, the head across each gate, the orders its preview allows, the setpoints, the opening limit and the declared values, and replay tests over canals with other widths, lengths, slopes, roughness, gate coefficients, motors, crests and turnouts find identical plans. A plan takes 0.83 to 1.05 ms in Chrome, 0.33 to 0.42 % of one core at 1200 times real time, while the page draws 59.9 to 60 frames per second.

5.4.4 Central, Decentralized and Distributed Plans

A second MPC setting, the architecture, decides who plans which gate. Central, the default, is the single program above. Decentralized and Distributed divide it among three agents, one per gate, each with the horizon, blocks, weights, bands, gate limits and preview of the central program. Agent r plans only Q_r and prices only the end of reach r, on the rows of reach r in Equation (5.8) with the flow through the gate below replaced by a trajectory $\hat{Q}_{r+1,k}$ that the agent is told, as shown in Equation (5.10).

$$\begin{aligned} \mathscr{J}_r = &\sum_{k=1}^{N} \left(\frac{\hat{y}_{r,k} - y_{SP,r}}{0.01\,\mathrm{m}} \right)^2 + \lambda \sum_{j=0}^{10} \left(\frac{\Delta Q_{r,j}}{0.1\,\mathrm{m^3/s}} \right)^2 \\ &+ \kappa\, \varepsilon_r^2 + \nu\, \varepsilon_r \\ \hat{y}_{r,k+1} = &\hat{y}_{r,k} + \frac{T_s}{A_{s,r}} \Big[(1 - f_r) Q_{r,k-n_r} + f_r Q_{r,k-n_r-1} \\ &- \hat{Q}_{r+1,k} - q_{r,k} + d_{r,k} \Big], \qquad \hat{Q}_{4,k} = 0 \end{aligned} \tag{5.10}$$

Agent r minimizes $\mathscr{J}_r$ over its own 11 moves, subject to the limits of gate r and the band of reach r from Equation (5.9). When every $\hat{Q}_{r+1}$ equals the plan of the gate below, the local objectives add up to the central one, $\mathscr{J} = \mathscr{J}_1 + \mathscr{J}_2 + \mathscr{J}_3$, which the central program minimizes over all 33 moves at once.

The architectures differ only in $\hat{Q}_{r+1}$, as listed in Table 5.5. A decentralized agent is told the flow that the gate below passes now, held over the horizon, and nothing is exchanged. A distributed agent is told the latest plan of the gate below. The three agents solve at once, exchange their plans and solve

Table 5.5 What Each Gate's Plan Is Told About the Gate Below

Architecture	$\hat{Q}_{r+1}$ over the horizon	Programs
Central	none needed, one plan sets all three flows	1 of 33 moves
Decentralized	the flow it passes now, held	3 of 11 moves
Distributed	its latest plan, exchanged until settled	8.86 on average

Programs per decision of scenario 11. The tail agent is told nothing, since no gate lies below it.

again, round p taking the plans of round $p-1$, until no plan an agent was told moves by more than δ, as shown in Equation (5.11),

$$\hat{Q}_{r+1}^{(p)} = Q_{r+1}^{(p-1)}, \qquad \max_{r,k} \left| Q_{r+1,k}^{(p)} - \hat{Q}_{r+1,k}^{(p)} \right| \le \delta \tag{5.11}$$

with $\delta = 1 \times 10^{-3}\,\mathrm{m^3/s}$ and at most 10 rounds. The first round starts from the plans accepted at the previous decision, shifted one period, or from the held flows after a change of architecture, a standby or a refused plan, so one round without earlier plans is the decentralized decision itself. This is the iterative, non-cooperative form of distributed MPC [18, 19]. The agents exchange whole plans and no price, unlike the dual decomposition of Section 1.4.11, which Chapter 6 applies to a heat network.

Because each gate is a flow source in Equation (5.10), the tail agent's problem depends on no other plan, gate 2's only on the tail's, and the head's only on gate 2's. The exchange therefore settles exactly in three rounds, with the tail's plan final after the first, gate 2's after the second and the head's after the third. In scenario 11, 137 of the 144 decisions took three rounds and seven stopped after two, their warm start already within δ. A zero tolerance gives three rounds at every decision and the same run, and the cap of 10 is never reached.

Pause. Try Exercise 5.3 before reading on.

The settled plans form a Nash point: each is its agent's best response to the plan of the gate below. No agent prices the reach above its own gate, although its flow leaves that reach at once as the term $\hat{Q}_{r,k}$ in the prediction of agent $r-1$. The central program prices that coupling and the agents do not, which is why a non-cooperative scheme need not reach the central optimum [25]. Computed independently for every decision of scenario 11, by back-substitution from the tail with a separate interior-point solver, the Nash point lies within $9.4 \times 10^{-4}\,\mathrm{m^3/s}$ of the accepted distributed plans over the whole horizon and within $3.4 \times 10^{-4}\,\mathrm{m^3/s}$ in the first move, inside δ. It differs from the central plan by up to 0.128, 0.175 and $0.146\,\mathrm{m^3/s}$ at gates 1, 2 and 3, and judged by the central objective it costs 1.13 times the optimum on average, from 1.004 to 1.32 times. Agents that also price the reach above their gate, minimizing $\mathscr{J}_r + \mathscr{J}_{r-1}$ in turns from the tail, reach the central optimum, within $1.2 \times 10^{-3}\,\mathrm{m^3/s}$ after 20 turns each and $7 \times 10^{-5}\,\mathrm{m^3/s}$ after 40 at the three decisions examined. The whole gap between the Nash point and the central plan is therefore the cross-reach term that no agent prices. At 6.33 h the agents' own costs are 52, 48 and 26 at the Nash point against 27, 33 and 40 under the central plan, which buys reaches 1 and 2 with reach 3.

All three architectures use the Kalman filter described above. Its steady-state gain is exactly block diagonal, the twelve entries that would correct one reach's level or inflow from another reach's transmitter being zero, so the shared filter is already three filters, one per reach, and agent r starts from its own reach's estimate.

A local plan reaches the gates only after a full acceptance check. The solver must report it optimal, the first move must be finite and equal the plan's first step, all 36 planned flows must be finite and inside the gate's limits, and the predictions, the cost and the slack must be finite, the slack non-negative. One refused plan hands all three gates to the tracking PI loops and marks the historian, and the next accepted plans are blended in over ten scans. Seven faults injected into each agent under both

Table 5.6 Measured Results

Case and condition	Result
1, head +0.5 m³/s	reach ends up 2 mm after 10, 29, 92 min
2, upstream, farm 2 closes	9900 m³ spilled in 6 h, reach 3 +56 cm
7 on upstream, tail +1 m³/s	tail farm 5300 m³ short in 6 h
3, downstream, farm 2 closes	no spill, upper reaches +10 cm
3, same, release over 6 h	28 300 m³ less than case 2
4, $K_c = 3$, tail +1 m³/s	tail −14 cm, reach 1 −23 cm
4, $K_c = 15$	worst 6.7 cm
5, $K_c = 20$, no filter	ring 1–3 cm, gate travel 10.6 m in 5 h
5, $K_c = 25$, no filter	grows to 34 cm, gates shut to open
6, $K_c = 20$, $\tau_f = 400$ s	ring <1 cm, gate travel 1.2 m
7, feedback only	reach ends 7, 6, 6 cm
7, feedforward now	reaches 1 and 2 about 10 cm
8, tail order +1.5 m³/s, feedback	10, 9, 9 cm
8, feedforward now	16, 15, 8 cm
8, feedforward ahead	2.3, 1.9, 2.0 cm
8, lead 0.5 / 2 travel times	about 6 / 10 cm
9, three overlapping orders, ahead	about 2 cm
3, MPC, preview now	no spill, 4.8, 5.8, 4.1 cm
7, MPC, preview now	3.2, 4.2, 5.6 cm
8, MPC, preview schedule	1.6, 2.0, 2.9 cm
9, MPC, preview schedule	2.0, 2.3, 2.9 cm
10, downstream, ahead	5.4, 5.7, 3.0 cm
10, MPC, preview schedule	2.0, 2.3, 2.8 cm

architectures, 42 runs with the chop off, were each refused with its reason. The hand-over moved a gate by 1.58×10^{-2} m³/s under Decentralized, against the loops' own next move of 1.70×10^{-2} m³/s, and by 5.7×10^{-4} m³/s under Distributed, and a take-back of 0.218 m³/s arrived as ten steps of at most 0.0218 m³/s. Central keeps its acceptance as described above.

A decision costs one program under Central, three under Decentralized and 8.86 on average under Distributed in scenario 11, at 2.95 rounds. In Chrome on the built application at 1200 times real time, over three sessions on a shared host, the whole controller takes 0.39 to 0.59 % of one core under Central, 0.27 to 0.31 % under Decentralized and 0.35 to 0.59 % under Distributed, at 60.1 frames per second, and the page's run equals a headless run at every one of 602 logged instants.

5.5 Results

Table 5.6 gathers what cases 1 to 10 produced on the engine at their shipped seed. Each case carries the briefing and completion rule of Section 1.6.4, and the challenge is judged on the case as given, so a changed model parameter or disturbance voids the attempt. Every case except 1 and 4 carries the 4 mm level chop, and over five further seeds no comparison below changes order, except the six named in Section 5.5.5, while Section 5.5.6 reports its own counts over ten seeds. Level deviations are peaks at the 5 s physics step. The 30 s record that the application displays matches them to 0.02 cm, except for reach 1 in case 8, where it reads 2.2 cm against 2.3 cm.

5.5.1 Timing the Wave

In manual, a 0.5 m³/s step in the head release raises the end of reach 1 by 2 mm after about 10 min, a little under its 11 min wave travel time, the end of reach 2 after about 29 min and the end of reach 3 after about 92 min, counted from the instant the head gate is commanded. The delay grows faster than the distance because each fixed gate passes only the part of the change that its rising upstream level can push through, so each reach has to fill a little before the next one hears about it. Two hours after the step, gate 3 passes only 0.05 m³/s of the extra release and reach 3 has risen 0.4 cm.

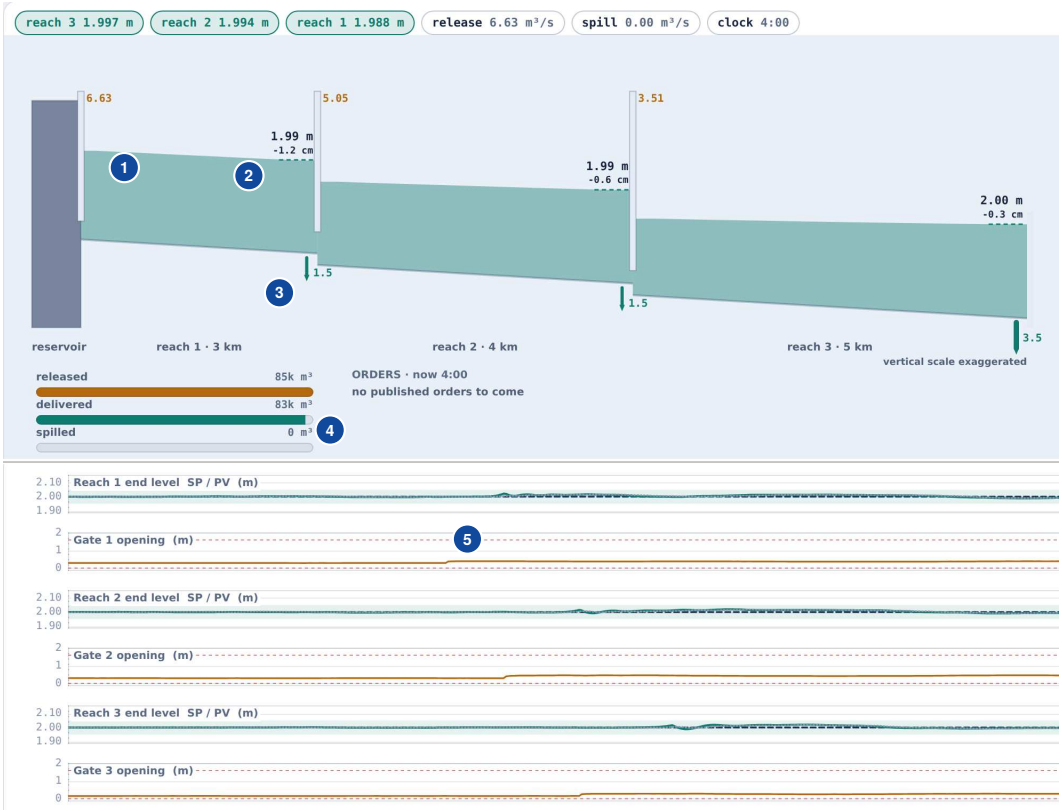


Figure 5.2 The canal running scenario 8, distant downstream control with a 1.5 m^3/s order for the tail farm published to start at two hours and released one travel time ahead at each gate, four hours into the run. The figure joins two parts of the application's screen, the long section above and the trends below, and leaves out the control and Learn panels. The long section is drawn with an exaggerated vertical scale and shows the reservoir, the three reaches with their water surfaces, the gates with their flows (1), the level at each reach end with its deviation from setpoint (2), the farm takes (3) and the escape weir, over the water ledger of released, delivered and spilled volumes (4). The trends show the three reach end levels held within about 2 cm of setpoint (a worst level error of 2.2 cm in the 30 s record, Section 5.5) while the three gate openings step up one after another (5), the head gate first, as the order is released ahead of its start (by the record at 1.26 h, gate 2 at 1.45 and gate 3 at 1.70 h). At the instant shown the gates pass 6.63, 5.05 and 3.51 m^3/s and nothing has spilled.

> **TRY IT IN THE APP** **1 · Open the Head Gate and Follow the Wave**
>
> Load the scenario, press Run and let half an hour pass, then raise Head gate release from 5.0 to 5.5 m^3/s, as the instructions ask. Watch the three level trends. The end of reach 1 first rises 2 mm after about 10 min, reach 2 after about 29 min and reach 3 only after about 92 min.

5.5.2 Two Philosophies

Local upstream control holds the upper reaches close to setpoint, and the cost appears at the tail. When the reach 2 farm closes its turnout, reach 1 stays within 0.2 cm, reach 2 peaks 7.1 cm high 35 min later and is back within 1 cm by 3.3 h, while reach 3 rises about 56 cm and about 9900 m^3 spills in the first six hours. When the tail farm takes 1 m^3/s more without ordering it instead (case 7 run on local upstream control), the tail farm goes about 5300 m^3 short. Distant downstream control

spills nothing when the reach 2 farm closes, lets the upper reaches ride about 10 cm high for a while, and releases about 28 300 m^3 less in six hours. The trade is waste against delay, and on a canal the waste is water.

🕮 **Pause.** Try Exercise 5.1 before reading on.

5.5.3 Tuning Against the Delay and the Resonance

A gain of 3 m^3/s per meter is too timid for the delay. After an unannounced 1 m^3/s tail offtake the tail falls about 14 cm, reach 1 falls about 23 cm, below the 1.8 m at which its turnout delivers the full order, so the reach 1 farm goes about 150 m^3 short, and the levels are still rolling by up to 19 cm five hours later. At the default gain of 15 the worst excursion is about 6.7 cm.

Raising the gain toward what the delay alone would allow finds the resonance. Driven at the head by a sinusoidal flow while gate 2 passes a constant flow, reach 1 answers 3.7 to 6.8 times more than a pure integrator of the same surface at periods between 10 and 26 min, with peaks near 12 and 22 min, the half and the whole of the 22.6 min round trip of a wave [1]. At a gain of 20 with no filter the levels ring by 1 to 3 cm from two to five hours after the offtake, and the gates travel about 10.6 m in the five hours after it. The ring decays by about a quarter each hour without the level chop, but the 4 mm chop keeps reach 1 ringing by about 1 to 2 cm through a 24 h run. Started at 25, the loop is unstable, the reach 1 ring grows from 15 cm in the hour after the offtake to 34 cm four hours later, and gates 2 and 3 swing between shut and fully open. A 400 s level filter, placed between the loop crossover and the resonance, reduces the ring below 1 cm and the gate travel to about 1.2 m. Most of the unfiltered travel is the level chop passed through the gain, and without the chop the gates travel 3.9 m unfiltered and 0.8 m filtered. At 25 with the filter the loop no longer runs away. Started at 25, the reach 1 ring, about 12 cm in the hour after the offtake, decays to about 1.5 cm four hours later, and raised to 25 at three hours, as scenario 6 asks, every reach end stays within about 1 cm. Without the chop, the gain at which the ring stops decaying rises from about 22 to about 30 with the filter. A 600 s filter rings more than a 400 s one, 1.7 against 0.8 cm in reach 1 from two to five hours after the offtake, most likely because its phase lag at the crossover begins to cost more margin than it saves at the resonance.

5.5.4 Feedforward Early and on Time

An unannounced offtake can only be answered after it is seen. Volume 3, Chapter 6 measures its disturbance where it enters, ahead of the delay. A turnout meter reads an offtake where the water leaves, behind every gate, so it gives the gates no lead. Feedback alone holds the three reach ends to about 7, 6 and 6 cm, and adding feedforward from the measured offtakes, with every gate moving at once, improves reach 3 slightly and makes reaches 1 and 2 worse, about 10 cm each. Every gate moves together, so each upper reach gives up a travel time of water to the reach below before its own extra supply arrives.

🕮 **Pause.** Try Exercise 5.2 before reading on.

A published order changes the picture, as shown in Figure 5.2. For a 1.5 m^3/s order at the tail, feedback alone moves the reach ends by about 10, 9 and 9 cm, and releasing the order at every gate when it starts makes the upper two worse, 16 and 15 cm, with 8 cm at the tail. Released one travel time ahead per reach, the same order moves the levels by only 2.3, 1.9 and 2.0 cm. Half the lead or twice the lead gives about 6 and 10 cm, so the lead that works is the wave travel time and not a tuning parameter to be searched. Three overlapping orders scheduled the same way hold all three reach ends to about 2 cm.

5.5.5 One Plan or Three Loops

Scenario 10 publishes two order changes, the tail farm up 1 m^3/s at 2 h and, at 5 h, the reach 1 farm down 1 m^3/s while the reach 2 farm goes up by as much, and at 7 h the tail farm takes 0.5 m^3/s more without asking. The MPC with the Schedule preview, as the scenario loads it, is compared with distant downstream control and the feedforward *ahead*, which receives the same published orders, over

Table 5.7 Scenario 10 in Two Windows. IAE in m s, Worst Level Error in cm, Gate Opening Travel in m Over 12 h

	0–12 h		0–7 h		7–10 h		
Structure	IAE	Worst	IAE	Worst	IAE	Worst	Travel
MPC, schedule	581	2.84	269	2.33	275	2.84	1.20
Each reach alone[a]	1065	5.37	649	5.37	385	3.90	1.51
Downstream, ahead	720	5.66	267	1.56	421	5.66	3.51
Same, $\tau_I = 7200$ s[b]	620	5.65	244	1.54	337	5.65	3.47
MPC, off	5084	9.80	2692	9.80	1756	9.76	1.00
Downstream, feedback	1664	6.93	1182	6.93	432	4.12	2.85

[a] A diagnosis copy of the MPC, not a setting of the application, in which each reach's model replaces the flow through the gate below by the orders below it, so no plan for one gate accounts for the others. It is not the Decentralized architecture of Section 5.4.4, which holds the gate below at its present flow. [b] The best of 225 fixed loop settings by 12 h IAE among the 197 that neither spill nor leave an order short, chosen after the run (K_c from 5 to 25, τ_I of 1800, 3600 or 7200 s, τ_f of 0, 200 or 400 s, feedforward off, *now* or *ahead* at k_{lead} of 0.75, 1 or 1.25). It ties at 620 with $K_c = 10$, $\tau_I = 7200$ s, $\tau_f = 0$ and $k_{lead} = 0.75$, which reaches 5.13 cm with 12.70 m of travel.

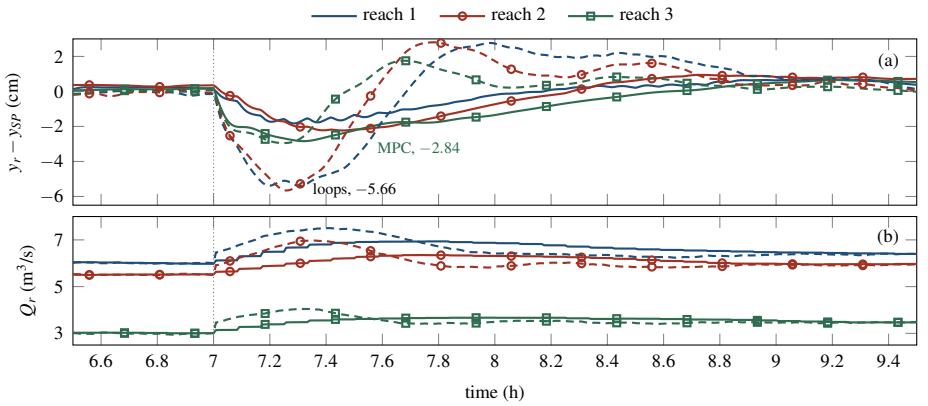


Figure 5.3 Scenario 10 from 6.5 to 9.5 h under the MPC with the Schedule preview (solid) and distant downstream control with the feedforward *ahead* (dashed). (a) Level error at each reach end and (b) gate flow, reach 1 without markers, reach 2 with circles and reach 3 with squares, with a dotted mark at the unannounced 0.5 m^3/s tail offtake at 7 h. The loops add the offtake at every gate from the next scan and reaches 1 and 2 fall 5.41 and 5.66 cm. The plan raises the three gates together over about 40 min and reach 3 falls furthest, 2.84 cm.

12 h at the same seed and chop, with the loops at their default tuning and the MPC at the scenario's settings. In this subsection the worst error is read from the 30 s record, as the scorecard reads it, and IAE is the scorecard's integral of absolute level error summed over the three reaches. The MPC rows of Table 5.6 agree with the physics step to 0.01 cm. The MPC made 144 plans in each 12 h run and none was rejected. Table 5.7 splits the run at the surprise.

Before the surprise the two structures tie on IAE, 269 against 267, and the plan's worst error is the larger, 2.33 against 1.56 cm. The whole lead of the MPC over 12 h, 581 against 720, comes after 7 h, where it holds 275 against 421 and 2.84 against 5.66 cm. The turnout meters tell both structures of the offtake at once. Every loop then adds the whole 0.5 m^3/s to its own gate, so each upper reach gives up water before its own extra supply arrives, the mechanism of case 7, while the plan ramps the head gate from 5.98 to 6.93 m^3/s by 7.67 h and spreads the shortfall over the reaches, as shown in Figure 5.3. The plan also moves the gate openings 1.20 against 3.51 m. Over ten seeds the MPC has the lower IAE and the lower worst error over 12 h at all ten, 585 against 728 on average and 2.8 against 5.6 cm. Before 7 h its worst error is the larger at all ten, 2.3 against 1.8 cm on average, although its IAE is the lower at eight of ten, so the IAE tie at the default seed is not general. After 7 h

Table 5.8 Every Offtake Case Over 12 h: IAE in m s / Worst Level Error in cm, From the 30 s Record

Case (MPC preview)	Loops, same information	MPC	Each reach alone[a]	Best fixed setting[b]	MPC, off	Loops, feedback
3, farm 2 closes, unannounced (now)	1194 / 11.46	406 / 5.84	388 / 8.10	459 / 9.36	3099 / 12.91	1141 / 9.73
7, tail +1 m^3/s, unannounced (now)	797 / 10.32	543 / 5.59	652 / 5.67	698 / 8.64	2500 / 9.87	1064 / 7.32
8, tail order +1.5 m^3/s (schedule)	383 / 2.15	396 / 2.88	824 / 8.37	278 / 5.11	3644 / 15.13	1499 / 10.01
9, three orders (schedule)	412 / 2.15	501 / 2.88	985 / 8.37	279 / 5.13	5475 / 15.13	1713 / 10.01
10, one plan or three loops (schedule)	720 / 5.66	581 / 2.84	1065 / 5.37	620 / 5.65	5084 / 9.80	1664 / 6.93
12, challenge, MPC as loaded (schedule)	741 / 6.07	668 / 3.28	1027 / 4.27	579 / 5.39	4198 / 11.49	2862 / 7.32

The loops with the same information use the feedforward *now* where the MPC preview is Now and *ahead* where it is Schedule, at the default tuning, which is also the challenge's suggested configuration. [a] The diagnosis copy of Table 5.7, not a setting of the application. [b] Chosen per case after the run from the same 225 settings, among those that neither spill nor leave an order short, a bound rather than a design. In cases 9 and 12 it moves the gate openings 6.74 and 12.87 m, against 1.16 and 1.46 m for the MPC.

both are lower at all ten.

The lead after 7 h is partly the comparators'. The best fixed loop setting, the default tuning with τ_I doubled, holds 337 after the surprise and beats the MPC before it, 244 against 269 and 1.54 against 2.33 cm, although its IAE there is the higher at one of five further seeds. Over 12 h it ties with a second setting at 620 against the MPC's 581. None of the 225 settings beats the MPC on both IAE and worst error, and none that spills nothing moves the gates as little. What only the joint plan delivers is the shared worst error after the surprise, 2.84 cm against 3.90 cm for the same model planned reach by reach and 5.65 cm for the best tuning, and that order holds at all five further seeds. The 3.90 cm belongs to the diagnosis copy, whose gates pass the orders below them on. The Decentralized architecture that the application offers, whose gates instead hold the flow of the gate below, lets reach 1 fall 7.01 cm after 7 h in the same scenario, and the Distributed architecture lets reach 2 fall 3.36 cm (Section 5.5.6).

The other unannounced cases, in Table 5.8, show the same split. On the reach 2 farm's closure the MPC with the Now preview scores 406 against 1194 for the loops with the feedforward *now*, but most of that margin is the comparators'. The same model planned reach by reach scores 388, and the order of those two changes over five further seeds, while the best fixed setting scores 459. The joint plan's own contribution is the shared worst error, 5.84 against 8.10 cm per reach and 9.36 cm for the best setting, lower at all five further seeds. On the unannounced 1 m^3/s tail offtake the joint plan keeps part of the total, 543 against 652 per reach and 698 for the best setting, while the worst errors tie, 5.59 against 5.67 cm, an order that changes at one of five further seeds. With the same Now information in scenario 10 the MPC scores 937 against 1361 and 5.63 against 10.53 cm.

On published orders alone the MPC loses to gate stroking. Case 8 gives 396 against 383 and 2.88 against 2.15 cm, and case 9 gives 501 against 412 and 2.88 against 2.15 cm. The worst error is the larger at all five further seeds in both cases, and so is the IAE of case 9, while the IAE order of case 8 changes at three of five. The cause is the model. Its reaches are integrators behind dead times of 215 to 330 s with no waves, while the wave first arrives 580 to 980 s after a gate moves, so the plan meets a stepped order with a ramp. For the 1.5 m^3/s tail order of case 8 the plan's gate flows pass half their change 19.5, 14.5 and 4.5 min before the order starts, against 44.5, 33.0 and 18.0 min for gate stroking, which steps each gate once at the wave's lead. Reach 3 therefore fills ahead of the turnout, 2.88 cm high as the order starts and 2.76 cm low five minutes later, and reaches 1 and 2 fall 1.6 and 2.0 cm about 23 min after it. Nothing tried removes the peak. A move price λ of 0.1, 0.3, 1 or 10 raises case 8's worst error to 4.98, 4.57, 3.91 or 4.47 cm, a decision every 150 s to 6.15 cm, every 60 s makes the plan unstable, and model dead times halfway to the wave travel times or equal to them

raise it to 18.3 and 76.7 cm. None of these was adopted, and the MPC carries the loss as measured.

Told nothing about the orders, the MPC loses every case to the feedback-only loops, by a ratio of IAE from 2.35 to 3.20 on cases 3, 7, 8, 9 and 10 and 1.47 on the challenge. The preview Off has only its estimate of the unexplained inflow, and with $\rho = 0.01$ that estimate is slow. It is offset-free. Over 48 h every reach end settles within 0.03 cm of its setpoint and each estimate within 0.007 m^3/s of minus its order, but in scenario 10 the worst error falls below 1 cm only after about 20 h. The Off loss is a tuning trade, made on a case with published orders.

> ✦ DID YOU KNOW? **A Planner Told Nothing Does Worse Than Feedback**
>
> From the simulator: in scenario 10 the MPC with its preview Off accumulates an IAE of 5084 m s and a worst level error of 9.80 cm over 12 h, against 1664 and 6.93 cm for three feedback-only loops (Table 5.7). A plan is only as good as what it is told about the orders.

Reach 3 has the larger IAE under the MPC than under the loops with the same information in every case, 132 against 51 on the reach 2 farm's closure, 236 against 162 on the tail offtake, 175 and 195 against 91 in cases 8 and 9, 225 against 140 in scenario 10 and 265 against 147 on the challenge. The tail is the joint plan's buffer. The plan prices every reach end alike, so a volume error shared by the canal spreads roughly in proportion to the reach surfaces, 1 to 1.35 to 1.75, and on the tail offtake the three reach ends fall 3.24, 4.16 and 5.59 cm. Planned reach by reach, reach 3 returns toward the loops, 34 on the closure, 147 on the tail offtake and 169 in scenario 10.

Water is not the price of any of this. Every structure that releases on demand releases the same water within 0.1 %, 264628 against 264 617 m^3 in scenario 10, and neither the MPC nor any distant downstream row spills or leaves an order short in any case. Local upstream control with its fixed supply spills or leaves an order short in every case, from 24 664 m^3 short on the tail offtake to 41 810 m^3 spilled on the closure. Asked to hold the tail at 2.4 m, above the 2.3 m weir crest, in case 3, the plan's tail band keeps the level at 2.288 m and spills 362 m^3 in 12 h, against 6614 m^3 in the first six hours for the feedback-only loops, which hold 2.40 m.

5.5.6 One Plan per Gate

Scenario 11 publishes one order change, the tail farm up 1 m^3/s at 2 h, and at 6 h the same farm takes 0.5 m^3/s more without asking. It loads the MPC with the Schedule preview and the Distributed architecture. Each architecture was run for 12 h at the scenario's settings, with the same seed, chop, orders and gates as the loops at their default tuning, the feedforward *ahead* receiving the same published orders. Table 5.9 splits the run at the surprise, with IAE and the worst error read from the 30 s record as in Section 5.5.5, and none of the 144 decisions of any MPC run was refused. The level errors under the three architectures are shown in Figure 5.4.

On the published order Distributed holds every reach end within 2.18 cm and Decentralized within 4.99 cm, against 1.93 cm for Central. The difference is how early each gate releases. Told only the flow of the gate below now, a decentralized gate follows each move of that gate after it is made, and its own water reaches its reach end a dead time after the reach has given water up. Between 1 and 2 h the decentralized head gate releases up to 0.11 m^3/s less than gate 2 and the reach 1 farm draw, and reaches 1, 2 and 3 fall 4.99, 4.12 and 2.18 cm. Told the plan of the gate below, each distributed gate above the tail releases 0.03 to 0.16 m^3/s more than the gate below and its own farm draw in the same hour, which is gate stroking done by messages, and the reaches fall 1.46, 1.66 and 2.18 cm. After the surprise the exchange carries the tail gate's answer up the canal within the decision, 3.10 against 6.68 cm. What remains is cooperation. At 6.6 h the distributed tail gate draws 3.769 m^3/s, what its own reach needs, and reaches 1 and 2 fall 3.10 and 3.09 cm, while the central plan draws 3.631 m^3/s, lets reach 3 take the largest error, 2.90 cm, and spares reaches 1 and 2, 2.01 and 2.28 cm. The worst errors of every structure in both windows are redrawn in Figure 5.5.

Table 5.9 Scenario 11 in Two Windows. IAE in m s, Worst Level Error in cm, Gate Opening Travel in m Over 12 h

	0–6 h		6–10 h		0–12 h	
Structure	IAE	Worst	IAE	Worst	IAE	Travel
MPC, central	183	1.93	281	2.90	490	0.94
MPC, distributed	243	2.18	281	3.10	550	1.26
MPC, decentralized[a]	563	4.99	473	6.68	1067	1.78
Distributed, one warm round[b]	271	2.18	341	4.88	641	1.40
Central, 1 h horizon	327	2.65	324	3.04	677	1.09
Distributed, 1 h	257	2.08	284	3.08	568	1.28
Decentralized, 1 h	526	5.85	475	6.60	1030	1.72
Downstream, ahead	230	1.52	390	5.31	647	3.01
Downstream, now	718	10.53	398	5.67	1143	3.93
Downstream, feedback	944	6.93	480	3.47	1458	2.58
Local upstream[c]	3529	41.01	7643	58.48	15452	1.38

[a] Also the Distributed architecture with one round and no earlier plans, identical in every logged column. [b] One round per decision, starting from the plans of the decision before. [c] The tail farm goes 29 989 m^3 short.

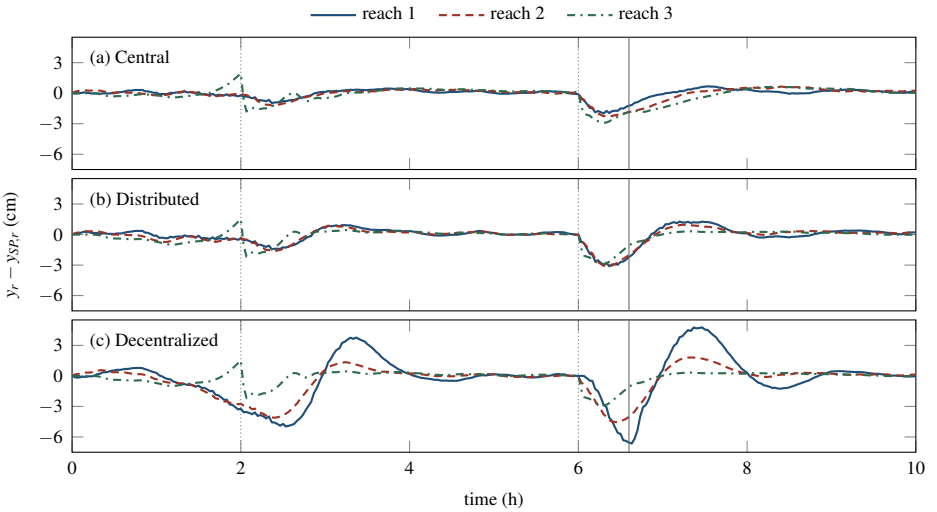


Figure 5.4 Scenario 11 from 0 to 10 h under (a) Central, (b) Distributed and (c) Decentralized, level error at each reach end, reach 1 solid, reach 2 dashed and reach 3 dash-dotted. Dotted lines mark the published tail order at 2 h and the unannounced offtake at 6 h. At the thin solid line, 6.6 h (23 760 s), reach 1 stands at 1.988, 1.978 and 1.934 m, and each plan predicts a worst deviation over its next 3 h of 2.2, 2.1 and 4.9 cm. Drawn from the engine's 30 s record every 2 min, since three captures of the application would be illegible at this size.

Over ten seeds Distributed beats Decentralized at all ten in both windows, and Central beats Distributed at nine of ten before the surprise and eight of ten after it. An independent set of ten seeds gives 10, 10, 8 and 8 of ten. The loops with the feedforward *ahead*, given the same published orders, keep the smallest peak before the surprise at every seed of both sets, 1.65 cm on average in the first set against 2.03 for Central and 2.38 cm for Distributed, and lose to Central and Distributed after it at every seed, 5.13 against 2.95 and 3.10 cm. On the published-order peak no architecture of the MPC beats gate stroking, as the central plan does not in Section 5.5.5.

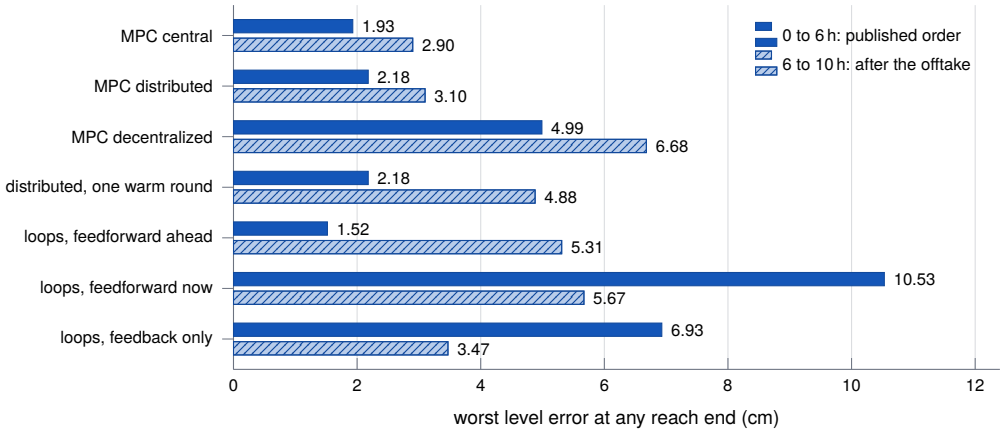


Figure 5.5 Scenario 11 redrawn from Table 5.9: the worst level error at any reach end on the published order (dark bars) and after the unannounced offtake (hatched bars). The 1 h horizon variants and local upstream control, 41.01 and 58.48 cm, are left out. Notice that gate stroking by the loops holds the smallest error on the published order, that Central and Distributed hold the smallest after the surprise, and that the decentralized agents, which exchange nothing, trail both in each window.

TRY IT IN THE APP **11 · Central, Decentralized or Distributed MPC**

Load the scenario and run it nine hours with MPC architecture on Distributed. Every reach end stays within 2.18 cm on the published order and within 3.10 cm after the unannounced offtake at 6 h. Reset, set MPC architecture to Decentralized before Run and repeat: 4.99 cm and 6.68 cm.

The counterfactuals attribute the gain. Removing the exchange is the Decentralized architecture, and one round per decision without earlier plans reproduces Decentralized in all 78 logged columns other than the architecture's own and the wall clock, in scenarios 10, 11 and the challenge. One warm round per decision keeps the published-order result, 2.18 cm, but gives up 4.88 cm after the surprise, since news of the offtake then climbs one gate per decision, and capping the exchange at two rounds leaves 137 decisions unsettled and 4.11 cm. The agents plan over the central horizon, so the central controller at the local horizon is Central itself. With all three shortened to 1 h, Distributed still beats Decentralized in both windows, 2.08 against 5.85 and 3.08 against 6.60 cm, so the horizon does not cause the exchange's gain. The shorter horizon reverses the order of Central and Distributed before the surprise, 2.65 against 2.08 cm, as Central loses 0.72 cm and Distributed gains 0.10 cm, while after it Central still leads, 3.04 against 3.08 cm.

The other cases tell the same story. On the reach 2 farm's closure, the unannounced tail offtake and the published orders of cases 8 and 9, with the previews of Table 5.8, the worst error over 12 h is 5.84, 5.59, 2.88 and 2.88 cm under Central, 8.14, 5.68, 3.38 and 3.38 cm under Distributed and 11.00, 10.71, 8.42 and 8.37 cm under Decentralized. In scenario 10 after the surprise the diagnosis copy of Table 5.7, whose gates pass the orders below them on, reaches 3.90 cm, while the shipped Decentralized reaches 7.01 cm. With the loaded challenge switched to the MPC, Central and Distributed meet it, 3.28 and 3.35 cm all day and 1.37 and 1.46 cm over the last four hours, and Decentralized misses at 8.56 and 4.73 cm, as listed in Table 5.11.

The safety comparison covers the twelve scenarios and thirteen stress cases under eight structures, with the worst rows in Table 5.10. No architecture overtops the bank, and neither local architecture spills where Central does not, 345 against 362 m^3 with the tail asked above the crest. Under the MPC, orders go short only in the stress cases. When every farm asks for 4 m^3/s at once, unannounced at 1 h, Central leaves 293 m^3 short, Distributed 1131, Decentralized 1964 and the loops with the feedforward

Table 5.10 Worst Cases Under Every Structure

Case	Cen.	Dec.	Dist.	FB	Ahead	Upstr.
All shut, overtopping (s)	0	0	0	0	0	33850
Tail asked 2.4 m, spill	362	345	345	13589	13590	53564
All at 4, 1 h, short	293	1964	1131	3583	1534	233466
All at 4, 6 h, short	301	2968	2231	5780	2289	143376
All at 4, 1 h, slow motor, short	693	2170	1638	3723	2137	233466
All at 4, published, short	0	112	0	3517	0	208266
$h_{min} = 2.0$ m, short	387	701	381	1508	365	32026
$B = 4.5$ m, $C_d = 0.5$, short	9	0	0	1	1	106743
All at 4, 6 h, lowest (m)	1.692	1.553	1.585	1.551	1.568	1.183

Volumes in m^3 over 12 h. Central, Decentralized and Distributed MPC, distant downstream control with feedback only and with the feedforward *ahead*, and local upstream control. All at 4: every farm asks for 4 m^3/s, unannounced at 1 h on case 3, at 6 h on scenario 11, or published to start at 2 h.

Table 5.11 Challenge Outcomes Over Twelve Hours

Configuration	Worst, day	Worst, 8–12 h	Short
as loaded, local upstream	68.0 cm	68.0 cm	26 400 m^3
downstream, feedback	7.3 cm	6.7 cm	0
downstream, now	12.2 cm	4.1 cm	0
downstream, ahead	6.1 cm	1.2 cm	0
MPC, schedule	3.3 cm	1.4 cm	0
MPC, schedule, distributed	3.3 cm	1.5 cm	0
MPC, schedule, decentralized	8.6 cm	4.7 cm	0
MPC, now	6.1 cm	3.7 cm	0
MPC, off	11.5 cm	4.6 cm	0

ahead 1534 m^3, so Decentralized is then less safe than the loops it would replace. Published two hours ahead, the same demand leaves Decentralized 112 m^3 short and every other structure except feedback alone and local upstream control nothing. With the full-order depth raised to the setpoint Decentralized leaves 701 m^3 short against 387 for Central and 381 for Distributed. On a canal with narrower gates of lower coefficient, which every controller now sees as mismatch, Central is the one MPC that leaves an order short, 9 m^3.

5.5.7 The Graded Challenge

The graded case covers twelve hours with five published order changes and a tail farm that takes 0.5 m^3/s more without asking at six hours, starting on local upstream control with a 5 m^3/s supply. It is met with less than 1 m^3 spilled, less than 1 m^3 of any order left short, every reach end within 8 cm all day and within 3 cm over the last four hours. Nine configurations were run over the full twelve hours, with the outcomes in Table 5.11.

As loaded, local upstream control with a fixed supply of 5 m^3/s leaves the tail farm 26 368 m^3 short, about 26 400 m^3, and reach 3 ends 68.0 cm low, so the case fails. Distant downstream control delivers every order but misses 3 cm late in the day, releasing at the moment of each order is worse over the whole day than feedback alone, and releasing ahead meets every condition with the last four hours held to 1.2 cm. That suggested configuration, distant downstream control at the default tuning with the feedforward released *ahead* at $k_{lead} = 1$, spills nothing, leaves no order short and stays inside the two bands by 1.9 and 1.8 cm. The unannounced offtake at six hours is the largest remaining excursion, 6.1 cm in reach 1 against 1.8 cm before it, because it is the only change left to feedback alone. Over five further seeds the suggested configuration holds the last four hours to 1.2 to 1.6 cm and feedback alone misses at 6.6 to 6.9 cm.

Changing only the control philosophy of the loaded case to the MPC, with its Schedule preview, also meets the target. Nothing spills, no order goes short, and the worst deviations are 3.3 cm all day

and 1.4 cm over the last four hours, with the gate openings moving 1.46 m against 3.78 m for the suggested configuration. Over the last four hours that is a loss at the shipped seed, 1.4 against 1.2 cm, and over five further seeds the order changes, 1.22 to 1.48 cm against 1.15 to 1.59 cm, with the MPC lower at three. The MPC reads no loop tuning, so the suggested settings with the MPC give the same run. With the Now preview it misses late in the day, 3.7 cm, and with Off it misses at 11.5 cm. The best fixed loop setting chosen after the run, $K_c = 10$, $\tau_I = 7200$ s, no filter and $k_{lead} = 0.75$, has the lower IAE, 579 against 668, with a worst error of 5.39 cm and 12.87 m of gate travel. The challenge, its target and its suggested configuration are unchanged.

5.6 Technology in Practice

What Measures It

An automated canal measures water levels, flows and gate positions at its control structures, and its control equipment must detect an abnormal sensor reading before the controller acts on it [3]. Distant downstream control needs the level at the far end of each pool, away from the gate that answers to it. The three level transmitters of Figure 5.1 stand for these gauges, and the turnout meters for the flow measured at each farm offtake.

> **TECHNOLOGY Sensor: The Pool Level Gauge**
>
> An automated canal reads the level in each pool with a float, pulley and counterweight that turn an absolute shaft encoder, with a bubbler that feeds gas slowly through a tube to an orifice below the lowest water and reads the back pressure, or with a submerged pressure transmitter. Float and pressure gauges reach 0.1 percent of range or better, and a float system spanning 15 m errs by about 0.01 m in its mechanics. The float sits in a stilling well whose small inlet damps wind waves and surges [26, 27]. A silted intake gives erroneous readings, a clogged bubbler tube or a fallen object on the float upsets the gauge, and frost damages a submerged transmitter. The chapter's transmitters carry a 4 mm chop and no such fault.

What Moves It

In the classic electronic downstream controller a float-driven potentiometer reads the level in the pool below, a second potentiometer driven by the hoist reads the gate position, and raise and lower relays run the gate motor until the gate reaches the position the algorithm has computed [3]. A set-operate-time actuator moves the gate in increments, and a 30 mm step at 7.7 mm/s takes 4 s. Check structures also carry bypass weirs sized for the canal's largest flow in case the gates close while the canal runs full [26]. The chapter's gates are flow sources: each opening follows from inverting Equation (5.3), and only the tail has a weir.

> **TECHNOLOGY Final Element: The Motorized Check Gate**
>
> A check structure holds the level above it with radial gates on larger canals and vertical lift gates on small ones, raised by electric, hydraulic or manual hoists with constant or variable speed motors [3]. Automated hoists are designed to set a gate within about 6 to 15 mm, typically move 0.3 to 0.6 m/min and carry a high duty cycle [26]. The flow under a radial gate follows from both water levels and the gate position, and silt or subsidence shifts that rating until the gate is recalibrated [27]. On a power failure a gate closes or stays put, as designed [3]. The chapter's gates move at 0.24 m/min and invert a fixed declared rating.

Where the Controller Runs

Local automatic control runs on microprocessor equipment at each check structure. Supervisory control places a remote terminal unit at each site, which reports levels, gate positions and equipment status to a master station and moves its gate on the instructions it receives over a two-way link [3]. Typical scan times of such water systems range from 2 s to 10 min [26]. Gate stroking needs the central system, because every gate must move at a precise time and by the correct amount, and the schedule can be stored in each remote unit for execution [3]. The simulator's loops scan every 30 s

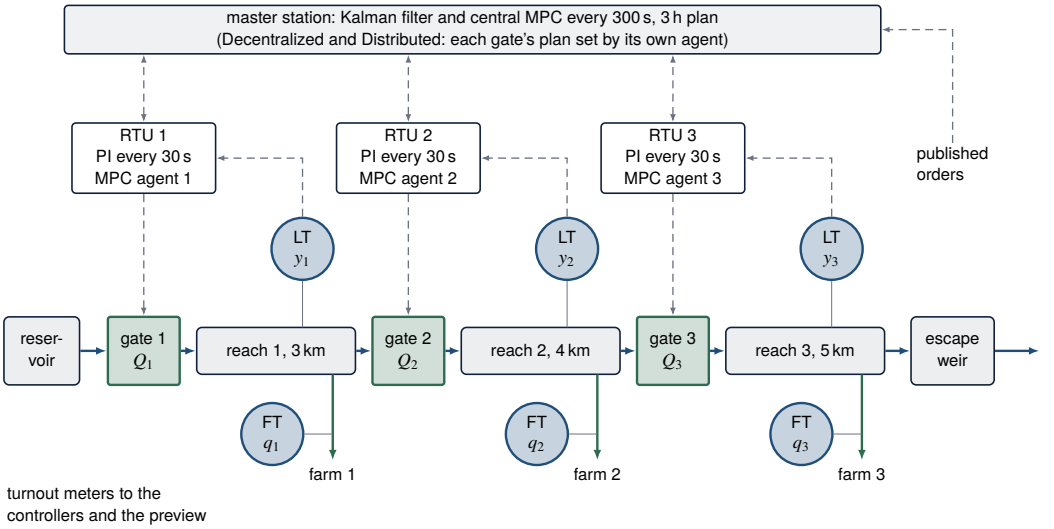


Figure 5.6 The canal's loops on field hardware, drawn from Sections 5.3 and 5.4. Green boxes are the three gates, blue circles the level transmitters at the reach ends and the turnout meters, dashed arrows signals. In a field system a remote terminal unit at each gate reports to a master station. Notice that each gate answers to the level at the far end of the reach below it, so its water is seen only after a wave has crossed that reach.

and its MPC decides every 300 s, inside that range. Its Central architecture plays the master station, and its Decentralized and Distributed agents play controllers at the gates, as drawn in Figure 5.6.

> **IN PRACTICE** **Filter and Deadband Before Gain**
>
> Canal practice filters the level before the controller to remove wave action and cut the number of gate moves, and the filter keeps the controller from amplifying critical frequencies. Oscillation follows when the pool's natural frequency matches the frequency of control actions, or when each correction is too large for the error. A deadband around the target stops the gate while the level stays inside it [3]. For a proportional-plus-reset controller a filter time constant of 100 s is usually enough [26]. The chapter's loops use 200 s, 400 s against the resonance, and no deadband.

What Is Modernizing It

Water users have long ordered water 24 to 48 hours ahead and received it on a schedule [3]. On networks fitted with sensors, actuators and a supervisory network, system identification and control have given near on-demand supply, with a distribution efficiency better than 90 percent against 70 percent under manual operation and better than 90 percent of orders met on demand. In simulations of larger systems, decentralized control with feedforward could not be told apart from centralized control in most scenarios, and the centralized control could not be implemented there [28]. This chapter puts the same question to one canal with its three architectures and its Schedule preview, and finds that exchanged plans keep most of what the single plan delivers.

5.7 Limitations

The channel is prismatic and rectangular, which puts the resonant periods near their textbook values. Real canals with varying cross-sections and structures resonate at less regular periods [13].

The local-inertia form omits advection. It is accurate at Froude numbers near 0.1 and would not be at the steep, fast flows of mountain canals.

The coarse grid smears fronts and approximates the resonance. A controller designed on this model and deployed on a real reach would need identification of the actual reach, as field work has done [15].

Gate motors are smooth and fast relative to real stepping motors with deadbands.

The MPC's model leaves out the gravity waves and the resonance, the change of the friction slope with flow, the gate motor's lag, the turnout cut-back except at the delivery its meter reads now, and the weir except through the estimated inflow. Its dead times are the best fit over its 3 h horizon rather than the wave's arrival, and the plan is sensitive to them, since case 8's worst error grows from 2.88 to 18.3 cm with dead times halfway to the wave travel times. The model is frozen at the design point, so an edit of the canal is model mismatch, and with the preview Off its estimate needs about 20 h to bring scenario 10 within 1 cm. The preview takes the published orders as exact.

The decentralized and distributed architectures share the central model, estimate, weights and horizon, so they differ only in who plans which gate and what each agent is told. The exchange of plans is instantaneous and never loses a message, while a field network between cross-regulators has delays and dropouts. Only the non-cooperative exchange is offered, and a decision that reaches the cap of ten rounds would be accepted unsettled without a mark in the historian, which the shipped tolerance never reaches because the exchange settles in three.

> **OPEN CHALLENGE** **Plans Shared Over a Patchy Network**
>
> On one automated network each site samples regularly, but the center hears only of events, such as a level that has moved 2 cm up or down. A local controller does not suffer from this. A controller for the whole network must cope with information that arrives at varying rates and delays, and its designers judged that substantial further research was needed [28]. This chapter's agents exchange plans at once and never lose one. Progress would be plans that stay safe and close to the central optimum when messages arrive late or not at all.

The farms either publish their orders exactly or not at all, and the canal has no branching, seepage or evaporation.

> **OPEN CHALLENGE** **Which Gate Should Listen to Which Level**
>
> Which inputs a decentralized canal controller should pair with which outputs has had too little attention, and the possible structures are combinatorially many, so heuristics are needed. How a disturbance travels along a string of pools is open too, in particular how much feedforward speeds the response before it amplifies the disturbance [28]. This chapter compares one plan, three silent plans and three exchanged plans on three pools. Progress would be rules that pick the structure for long chains of pools.

5.8 Conclusion

A canal reach is a tank with a delay in front of it and a resonance inside it, and the simulator of this chapter lets each of the three be measured. A 0.5 m^3/s step at the head reaches the three reach ends after about 10, 29 and 92 min. Local upstream control spills about 9900 m^3 in six hours when a farm closes, while distant downstream control spills nothing and lets the upper reaches ride about 10 cm high. A gain tuned to the delay rings at the reach resonance and runs away at a slightly higher gain, and a 400 s level filter tames it. Feedforward that releases a published order at every gate at once makes the upper reaches worse, 16 against 10 cm, while releasing it one wave travel time ahead per reach moves the levels only about 2 cm, and over a twelve-hour challenge the same schedule holds the last four hours to 1.2 cm. A model predictive controller, offered as an option, that sets the three gates in one plan spreads an unannounced offtake over the reaches, 2.84 cm against 5.66 cm for the loops and 3.90 cm for the same model planned reach by reach with the orders below each gate passed on, and trails gate stroking on published orders because its model has no waves. Divided into one plan per gate, the same controller keeps most of that result only when the gates exchange their plans,

3.10 cm against 6.68 cm without the exchange and 2.90 cm for the single plan after the offtake of scenario 11, and the remainder is the cooperation that no gate's own objective prices.

The model uses prismatic reaches, a coarse grid and perfect published orders. A fuller model would add trapezoidal sections, gate deadbands, uncertain orders and a branching network. The canal remains a compact case of controlling water that arrives late, stays, and echoes.

KEY POINTS

- A canal reach is a tank with a delay in front of it and a resonance inside it.
- When a farm closes, local upstream control spills about 9900 m^3 in six hours and distant downstream control spills nothing.
- Released one wave travel time ahead per reach, a published order moves the levels about 2 cm, and the MPC, whose model has no waves, trails this gate stroking.
- One plan for three gates holds an unannounced offtake to 2.84 cm, against 5.66 cm for the loops. Split per gate, the plans keep most of that only when exchanged, 3.10 cm against 6.68 cm.

Where else this applies. Transport with storage and reflection also runs district heating lines (Chapter 6) and gas pipelines (Volume 10, Chapter 12).

References

[1] X. Litrico and V. Fromion, "Simplified Modeling of Irrigation Canals for Controller Design," *Journal of Irrigation and Drainage Engineering*, vol. 130, no. 5, pp. 373–383, 2004.

[2] J. Schuurmans, A. J. Clemmens, S. Dijkstra, A. Hof, and R. Brouwer, "Modeling of Irrigation and Drainage Canals for Controller Design," *Journal of Irrigation and Drainage Engineering*, vol. 125, no. 6, pp. 338–344, 1999.

[3] U.S. Bureau of Reclamation, *Canal Systems Automation Manual, Volume 1*, 1st ed. Denver, CO: U.S. Department of the Interior, Bureau of Reclamation, 1991. [Online]. Available: https://www.usbr.gov/tsc/techreferences/mands/mands-pdfs/CanalSysAuto1.pdf

[4] P.-O. Malaterre, D. C. Rogers, and J. Schuurmans, "Classification of Canal Control Algorithms," *Journal of Irrigation and Drainage Engineering*, vol. 124, no. 1, pp. 3–10, 1998.

[5] E. B. Wylie, "Control of Transient Free - Surface Flow," *Journal of the Hydraulics Division*, vol. 95, no. 1, pp. 347–362, 1969.

[6] W. E. Bodley and E. B. Wylie, "Control of Transients in Series Channel with Gates," *Journal of the Hydraulics Division*, vol. 104, no. 10, pp. 1395–1407, 1978.

[7] I. Mareels, E. Weyer, S. K. Ooi, M. Cantoni, Y. Li, and G. Nair, "Systems engineering for irrigation systems: Successes and challenges," *Annual Reviews in Control*, vol. 29, no. 2, pp. 191–204, 2005.

[8] C. M. Burt and X. Piao, "Advances in PLC-based irrigation canal automation," *Irrigation and Drainage*, vol. 53, no. 1, pp. 29–37, 2004.

[9] D. C. Rogers and J. Goussard, "Canal Control Algorithms Currently in Use," *Journal of Irrigation and Drainage Engineering*, vol. 124, no. 1, pp. 11–15, 1998.

[10] T. S. Strelkoff, "One-Dimensional Equations of Open-Channel Flow," *Journal of the Hydraulics Division*, vol. 95, no. 3, pp. 861–876, 1969.

[11] J. Schuurmans, A. Hof, S. Dijkstra, O. H. Bosgra, and R. Brouwer, "Simple Water Level Controller for Irrigation and Drainage Canals," *Journal of Irrigation and Drainage Engineering*, vol. 125, no. 4, pp. 189–195, 1999.

[12] X. Litrico and V. Fromion, "Tuning of Robust Distant Downstream PI Controllers for an Irrigation Canal Pool. I: Theory," *Journal of Irrigation and Drainage Engineering*, vol. 132, no. 4, pp. 359–368, 2006.

[13] ——, *Modeling and Control of Hydrosystems*. Springer London, 2009.

[14] B. T. Wahlin and A. J. Clemmens, "Automatic Downstream Water-Level Feedback Control of Branching Canal Networks: Theory," *Journal of Irrigation and Drainage Engineering*, vol. 132, no. 3, pp. 198–207, 2006.

[15] E. Weyer, "System identification of an open water channel," *Control Engineering Practice*, vol. 9, no. 12, pp. 1289–1299, 2001.

[16] B. T. Wahlin, "Performance of Model Predictive Control on ASCE Test Canal 1," *Journal of Irrigation and Drainage Engineering*, vol. 130, no. 3, pp. 227–238, 2004.

[17] R. R. Negenborn, P.-J. van Overloop, T. Keviczky, and B. De Schutter, "Distributed model predictive control of irrigation canals," *Networks & Heterogeneous Media*, vol. 4, no. 2, pp. 359–380, 2009.

[18] R. Scattolini, "Architectures for distributed and hierarchical Model Predictive Control – a review," *Journal of Process Control*, vol. 19, no. 5, pp. 723–731, 2009.

[19] P. D. Christofides, R. Scattolini, D. Muñoz de la Peña, and J. Liu, "Distributed model predictive control: A tutorial review and future research directions," *Computers & Chemical Engineering*, vol. 51, pp. 21–41, 2013.

[20] P. D. Bates, M. S. Horritt, and T. J. Fewtrell, "A simple inertial formulation of the shallow water equations for efficient two-dimensional flood inundation modelling," *Journal of Hydrology*, vol. 387, no. 1-2, pp. 33–45, 2010.

[21] D. Q. Mayne, J. B. Rawlings, C. V. Rao, and P. O. M. Scokaert, "Constrained model predictive control: Stability and optimality," *Automatica*, vol. 36, no. 6, pp. 789–814, 2000.

[22] S. J. Qin and T. A. Badgwell, "A survey of industrial model predictive control technology," *Control Engineering Practice*, vol. 11, no. 7, pp. 733–764, 2003.

[23] D. Goldfarb and A. Idnani, "A numerically stable dual method for solving strictly convex quadratic programs," *Mathematical Programming*, vol. 27, no. 1, pp. 1–33, 1983.

[24] G. Pannocchia and J. B. Rawlings, "Disturbance models for offset-free model-predictive control," *AIChE Journal*, vol. 49, no. 2, pp. 426–437, 2003.

[25] J. B. Rawlings and B. T. Stewart, "Coordinating multiple optimization-based controllers: New opportunities and challenges," *Journal of Process Control*, vol. 18, no. 9, pp. 839–845, 2008.

[26] U.S. Bureau of Reclamation, *Canal Systems Automation Manual, Volume 2*, 1st ed. Denver, CO: U.S. Department of the Interior, Bureau of Reclamation, 1995. [Online]. Available: https://www.usbr.gov/tsc/techreferences/mands/mands-pdfs/CanalSysAuto2.pdf

[27] ——, *Water Measurement Manual*, third, revised reprint ed. U.S. Department of the Interior, Bureau of Reclamation, Water Resources Research Laboratory, 2001. [Online]. Available: https://www.usbr.gov/tsc/techreferences/mands/wmm/

[28] I. Mareels, E. Weyer, S. K. Ooi, M. Cantoni, Y. Li, and G. Nair, "Systems engineering for irrigation systems: Successes and challenges," in *IFAC Proceedings Volumes (16th IFAC World Congress)*, vol. 38, no. 1, 2005, pp. 1–16.

5.9 Exercises

5.1 Concept check · 2 min

Under local upstream control the head releases a fixed supply and each cross-regulator holds the level just upstream of itself. At one hour the reach 2 farm closes its turnout. Where does the water it no longer takes go?

(a) It stays in reach 2, whose level rises until the farm opens its turnout again

(b) It is passed from gate to gate down to the tail and spills over the escape weir

(c) The head gate cuts its release, so the extra water never enters the canal

(d) It backs up into reach 1, because gate 2 cannot pass more than its design flow

5.2 Concept check · 2 min

In case 7 the tail farm takes 1 m^3/s more without ordering it. With the order feedforward on Now, every gate adds the metered offtake to its flow at the next scan. Feedback alone lets the reach ends fall about 7, 6 and 6 cm. With the feedforward, reach 3 improves slightly and reaches 1 and 2 fall about 10 cm each. Why do the upper reaches get worse?

(a) The turnout meters report the offtake late, so the feedforward acts after the feedback has

(b) The feedforward adds to the integral term, so each loop overshoots after the offtake

(c) Each gate draws its extra flow out of the reach above at once, and that reach's own extra supply arrives a travel time later

(d) The level filter slows the feedback, so the feedforward acts alone through the first hour

5.3 Concept check · 2 min

In the Distributed architecture the three gate agents exchange plans until no plan an agent was told moves by more than δ. On this canal the exchange settles in three rounds. Why three?

(a) The number of rounds is capped at three to fit inside the five-minute decision period

(b) The Kalman filter's gain is block diagonal, so each agent needs only its own reach's estimate

(c) Each agent also prices the reach above its gate, so the three objectives agree after three swaps

(d) Each agent's problem depends only on the plan of the gate below, so the plans become final one gate at a time from the tail

5.4 Theory and calculation · 6 min

The three reaches are 3, 4 and 5 km long and run at 2 m depth.

(a) Compute the gravity-wave speed $\sqrt{gh}$ and the wave travel time of each reach in minutes. *Checkpoint:* 4.43 m/s, then 11.3, 15.1 and 18.8 min.

(b) A 1.5 m^3/s tail order is published to start at 2 h, as in scenario 8. With $k_{lead} = 1$, at what time does each gate begin to release it according to Equation (5.6)? Give the times in hours. *Checkpoint:* Head gate at 1.25 h, gate 2 at 1.44 h, gate 3 at 1.69 h.

(c) Reach 1 resonates at the round trip of a wave between its two gates. Compute that period and compare it with the periods at which reach 1 answers most strongly. *Checkpoint:* 22.6 min. The peaks near 12 and 22 min are its half and its whole.

5.5 Theory and calculation · 6 min

The MPC models reach r as an integrator of surface $A_{s,r}$ behind a dead time θ_r, sampled every $T_s = 300$ s (Equation (5.8) and Table 5.3).

(a) Write each dead time as $\theta_r = (n_r + f_r)T_s$ with $0 \le f_r < 1$. *Checkpoint:* $n_r = 0, 0, 1$ and $f_r = 0.717$, 0.983, 0.100.

(b) All flows are balanced when the head gate flow steps up by 0.3 m^3/s at decision k and holds. How much does the model raise the end of reach 1 over the period from k to $k+1$, and over each period after that? Compare with the first rise of the real reach in Table 5.3. *Checkpoint:* 1.57 mm, then 5.55 mm per period.

(c) The plan prices the three reach ends alike, so a volume error shared by the canal tends to spread into equal level errors. What share of the volume does each reach take, and how far do all three fall if an unannounced 0.5 m^3/s offtake goes unanswered for 30 min? *Checkpoint:* 24.4 %, 33.0 % and 42.6 %, and about 1.35 cm.

5.6 Theory and calculation · 6 min

Three boundary laws decide who goes short and where water leaves the canal: the turnout cut-back of Equation (5.4) with $h_t = 1.2$ m and $h_{min} = 1.8$ m, and the escape weir $Q_{spill} = 1.7\,B\,(h - 2.3)^{3/2}$ with $B = 6$ m.

(a) At the timid gain of case 4, reach 1 falls about 23 cm from its 2.00 m setpoint. What does the reach 1 turnout deliver at 1.77 m against its 1.5 m^3/s order? *Checkpoint:* 1.462 m^3/s, 0.038 m^3/s or about 2.5 % short.

(b) Feedback-only loops asked to hold the tail at 2.40 m in case 3 hold it there. What does the weir spill at that level, and how much in six hours? The chapter measures 6614 m^3 in the first six hours. Why is the measurement lower? *Checkpoint:* 0.323 m^3/s and about 6970 m^3.

(c) Under local upstream control with every farm shut, the head keeps releasing 5 m^3/s. At what tail depth would the weir pass all of it? Compare with the 2.7 m freeboard flag. *Checkpoint:* About 2.92 m, above the flag.

5.7 App challenge · 6 min

Load scenario 1, *Open the head gate and follow the wave (MANUAL)*, default seed. It starts in Manual with the head release at 5.0 m^3/s and no level chop. Press Run, pause at about 0.5 h (the exact instant does not matter), raise Head gate release to 5.5 m^3/s, and run at speed 300 until at least two hours after the step, about 2.6 h. Export the CSV, whose time column is in hours.

(a) Take the step as the last row before `gate1_flow` rises. For each reach end, find the first row in which `y1_m`, `y2_m` or `y3_m` stands 2 mm above its value at the step, and give the delay in minutes. *Checkpoint:* About 10, 29.5 and 92.5 min.

(b) Read how much `gate3_flow` and `y3_m` have risen two hours after the step, and compare the delays of (a) with the wave travel times of reaches 1, 2 and 3 taken in turn. Explain the difference. *Checkpoint:* About 0.05 m^3/s and 0.4 cm.

Give the delays to 1.5 min.

5.8 App challenge · 8 min

Load scenario 8, *A published order, released one travel time ahead*, default seed. It runs distant downstream control with Order feedforward on Ahead and Lead, in travel times at 1, with a 1.5 m^3/s tail order published to start at 2 h. Run each part to 4 h at speed 1200 and read the Worst level error tile, or the per-reach maxima in the last rows of the CSV columns `metric_maxdev1_m` to `metric_maxdev3_m`.

(a) Run the scenario as loaded. *Checkpoint:* 2.2 cm.

(b) Reset, set Order feedforward to Now before Run, and repeat. *Checkpoint:* 16.2 cm.

(c) Press the reload button to restore the scenario, set Lead to 0.5 before Run, and repeat. Then reload again and repeat with Lead at 2. *Checkpoint:* 6.5 cm and 10.6 cm.

(d) Say why the lead of one travel time wins, and why part (c) needs a reload rather than a Reset.

Give the tile readings to 0.2 cm.

5.9 App challenge · 9 min

Load scenario 11, *Central, decentralized or distributed MPC*, default seed. It runs the MPC with the Schedule preview and MPC architecture on Distributed (advanced view), with a 1 m^3/s tail order published for 2 h and an unannounced 0.5 m^3/s tail offtake at 6 h. Run each part to 10 h at speed 1200 and export the CSV, whose time column is in hours. Read the largest of `metric_maxdev1_m` to `metric_maxdev3_m` in the row at 6 h (the worst before the surprise), and the three columns in the last row.

(a) Run the scenario as loaded. *Checkpoint:* 2.18 cm before, then 3.10, 3.09 and 2.89 cm at the end.

(b) Reset, set MPC architecture to Decentralized before Run, and repeat. *Checkpoint:* 4.99 cm before, then 6.68, 4.55 and 2.89 cm at the end.

(c) Reset, set MPC architecture to Central before Run, and repeat. *Checkpoint:* 1.93 cm before, then 2.01, 2.28 and 2.90 cm at the end.

(d) Explain what the exchange of plans buys on the published order, and what only the central plan buys after the surprise.

Give the values to 0.05 cm.

5.10 Application area · 3 min

A district manager prefers local upstream control because the trends of the upper reaches stay flat and every gate acts on a level it can see. Using the results of cases 2 and 3, state what each philosophy costs on this canal and who pays it.

5.11 Application area · 4 min

A canal authority plans to run the Distributed architecture in the field, with one controller at each cross-regulator and a radio link between neighbors. Which results of this chapter rest on assumptions the field will not meet, and what would you test before commissioning?

5.12 Review the agent · 5 min

An AI assistant reviews case 8 under the MPC, which reaches a worst level error of 2.88 cm against 2.15 cm for gate stroking, and writes: “The plan is too sluggish. The move price $\lambda = 3$ penalizes gate moves, so the plan ramps the gates instead of stepping them. Lower λ to 0.3 and the MPC will step the gates like gate stroking and beat it.” Find the flaw, and say how you would check the claim in the application.