INTERACTIVE CONTROL ENGINEERING
Principles and Applications

Volume 6 Multivariable and Plantwide Control

CHAPTER 2

Seawater Reverse Osmosis Virtual Laboratory to Teach Loop Pairing and Decoupling

Pages 22–48

Edited by John D. Hedengren

How to cite this chapter. "Seawater Reverse Osmosis Virtual Laboratory to Teach Loop Pairing and Decoupling," in *Multivariable and Plantwide Control*, John D. Hedengren, Ed., vol. 6 of *Interactive Control Engineering: Principles and Applications*. Advanced Process Solutions, LLC, 2026, ch. 2, pp. 22–48.

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

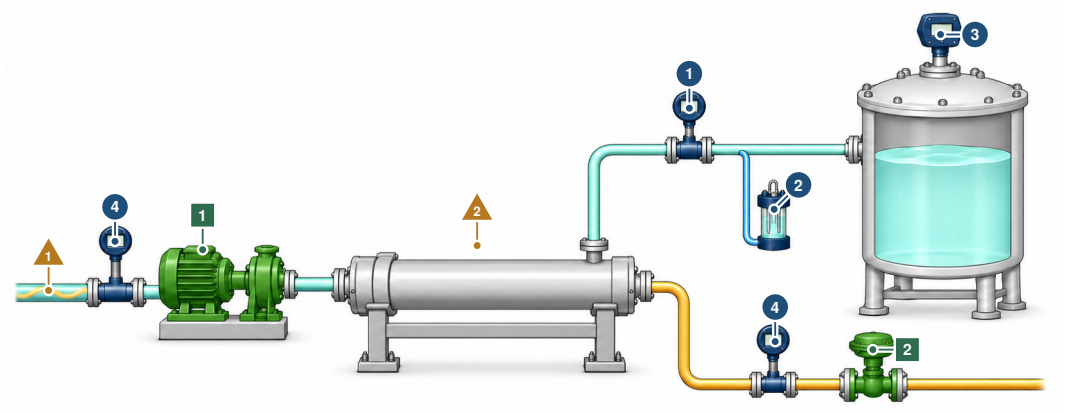
The app for this chapter
https://apmonitor.com/apps/ro-desalination/

The volume online: contents, other chapters and citation records
https://apmonitor.com/apps/books/v06-multivariable-plantwide/#ro-desalination

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

CHAPTER 2

Seawater Reverse Osmosis Virtual Laboratory to Teach Loop Pairing and Decoupling



The controller moves 1 high-pressure pump, whose speed sets the push on the membrane
2 concentrate valve, which lets the salty water that is left flow away

The controller reads 1 permeate flow meter: the production 2 conductivity cell: how salty the permeate is
3 level of the permeate tank 4 feed and concentrate flow meters: the share of seawater recovered

What disturbs it 1 seawater whose saltiness and temperature change 2 fouling that builds up on the membrane

Notice that one pump and one valve set every flow, salinity and pressure of the train, so pairing each with a target is the first design decision. Illustration, not to scale.

IN PLAIN TERMS **Two Handles on a Plant That Makes Fresh Water**

The main idea. A seawater plant pushes seawater hard against a thin film that lets water through and holds most of the salt back. A pump sets the push, and a valve lets the salty water that is left flow away.

The challenge. The usual test for whether two control loops will fight says these two will not. Yet each pump move still changes how much seawater turns into fresh water, and giving production to the wrong handle sends the plant far off target.

Why it matters. Which handle holds which target is decided before any tuning, and no tuning can repair a wrong choice.

THREE WAYS TO READ THIS CHAPTER

Explore	12 min	the first page, the fact on energy (Section 2.1), the history in Section 2.2, the two try-its and the fact on fouling (Section 2.6) 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 2.6), Technology in Practice (Section 2.7), Limitations (Section 2.8) and the two open challenges

The control question A seawater reverse osmosis train has two actuators, a high-pressure pump and a concentrate valve, and its relative gain reads 1.00005, the value that means no interaction. Which actuator should hold production, why do the loops still interact, and what keeps the membrane under its pressure rating as it fouls?

Learning objectives

- Derive the permeate flow and salinity at the design point from the solution-diffusion and polarization equations.
- Explain why the valve gain on permeate flow passes through zero and changes sign at the design point.
- Compute the relative gain and the Niederlinski index of both pairings from the gain matrix, and tell a triangular plant from a diagonal one.
- Design the one-way static decoupler from two gains and judge its calibration from a sweep.
- Compare the decoupled loops with the MPC on a pressure limit and name the mechanism that holds it.

Prerequisites Sections 1.4.5 (relative gain and pairing), 1.4.6 (simplified decoupling), 1.4.7 (directionality), 1.6.2 (velocity-form PID) and 1.6.4 (scenarios and data). Section 1.6.3 is needed only for the predictive controller.

The app https://apmonitor.com/apps/ro-desalination/

See also Volume 7, Chapter 4 for a setpoint at the maximum of the plant characteristic, where the process gain reverses sign. Volume 9, Chapter 4 for offset-free predictive control on a two-by-two plant with a limit, and Volume 9, Chapter 8 for a cost optimum that sits on a moving constraint.

2.1 Introduction

Seawater reverse osmosis recovers fresh water from a saline feed by forcing it through a semipermeable membrane at a pressure above the osmotic pressure of the concentrate, and it now supplies a large share of the municipal water produced from the sea [1]. A train has two manipulated variables. The high-pressure pump sets the pressure applied at the membrane, and the concentrate valve sets the fraction of the feed that leaves without passing through it. Permeate flow, recovery, permeate salinity, feed pressure, specific energy and the margin to scaling are all functions of that pair, so two of the six receive setpoints and four follow. Which two receive them is a design decision taken before any controller is tuned, and it is the one decision that tuning cannot repair.

Sweeping the concentrate valve at a fixed pump speed produces an interior maximum in permeate flow rather than a monotone curve. Permeate flow rises from 438.71 m^3/h at 40 % valve travel to 450.00 m^3/h at 50 % and falls to 436.51 m^3/h at 60 %, and the maximum is the design point. A production controller paired with that valve has no steady-state gain to work against at the operating point and a gain of the opposite sign a few percent to either side. The pump carries no such defect, so production belongs on the pump and recovery on the concentrate valve.

The scenarios of the app cover loop pairing and the relative gain array, one-way static decoupling, gain sign reversal inside the normal operating envelope, override control through a low selector with external reset, cascade control from a surge tank onto a production setpoint, and the drift of every process gain as an asset ages [2, 3]. The physical fundamentals behind the model are a total and a component mass balance over one membrane array, a driving force written as the applied pressure less the osmotic pressure at the membrane wall, and a separation of time scales that runs from a hydraulic node with a three second time constant to a fouling state that moves over weeks [4, 5]. Specific energy is the largest operating cost a desalination plant carries, and the recovery that minimizes it lies close to the concentration at which the sparingly soluble salts crystallize [6], so the distance between an optimum and a constraint is what the control system is defending.

DID YOU KNOW? The Energy Price of Water From the Sea

A research roadmap of the United States Department of Energy gives the energy of seawater desalination in that country as 3.0 to 3.8 kWh/m^3. Supplying potable water in large cities of the country takes 0.34 kWh/m^3 on average, in a study the same roadmap cites [7].

Assessments of process control education call for problems in which the control structure is decided rather than supplied [8, 9]. Interaction is usually demonstrated on plants whose relative gain array announces it, which omits the case in which the standard tool reports no interaction and the loops interact anyway.

This chapter works with a seawater reverse osmosis train on which the pairing decision, its failure mode and its remedy are all measurable. Section 2.2 reviews desalination practice and reverse osmosis control. Section 2.3 gives the model and every constant it uses. Section 2.4 derives the gain matrix, its relative gain and the one-way coupling that the relative gain cannot show. Section 2.5 describes the two regulatory loops, the static decoupler, the supervisory layer and an optional model predictive controller (MPC) of both actuators. Section 2.6 reports the scenario results, from the production loop alone and the reversed pairing through decoupling, warm feed, the energy optimum, fouling and the graded shift, and Section 2.6.8 measures which mechanism of the MPC keeps the membrane under its pressure rating. Section 2.7 describes the real instruments, drives and controllers. Sections 2.8 to 2.10 state the limitations, the three results that carry past this plant and the conclusions.

2.2 Background

Reverse osmosis displaced thermal distillation as the dominant seawater desalination technology over three decades, and energy rather than capacity drove the change. A thermal plant pays for a phase change, while a membrane plant pays only for the pressure applied to the water it processes. Energy recovery closed most of the remaining gap. A pressure exchanger returns the concentrate pressure to the feed at high efficiency, so the high-pressure pump raises only the water that becomes permeate [10], and the specific energy of seawater trains has fallen toward the thermodynamic limit as a result [6]. Desalination now supplies water to a large fraction of the world’s coastal population and consumes electricity in proportion [1].

HISTORY 1952: A Federal Program for Salt Water

Federal funding of desalination research in the United States began in 1952, when the Saline Water Act was passed. An Office of Saline Water was created in 1954, and an Office of Water Research and Technology followed in 1974. A federal handbook for planners later judged these programs primarily responsible for developing reverse osmosis, at a time when few water supply problems on land had even been identified in the country, and noted that the same research gave other fields kidney dialysis with membranes. The Bureau of Reclamation served as a testing laboratory for the office, with membrane pilot plants in the field, and federal funding for desalination largely stopped in the early 1980s [11].

Control of reverse osmosis trains has been studied mainly as a multivariable regulation problem. System identification and decentralized control of a small train date to the late nineteen eighties [12], and reviews of the subject group the published work into empirical transfer function models closed with PID loops and nonlinear model-based designs [13]. Nonlinear model-based control has been demonstrated on an experimental unit operating against pressure and flow constraints [14]. The steady-state trade-off that fixes the operating point, between the feed that has to be pressurized and the osmotic pressure of the concentrate, produces an interior minimum in specific energy that is well established for spiral-wound seawater trains [15].

Membrane performance is not stationary. Permeability declines as foulant accumulates, and on a seawater train the decline is fast enough to move every process gain within one cleaning interval [16]. Sparingly soluble salts crystallize on the membrane once the concentrate passes saturation, which

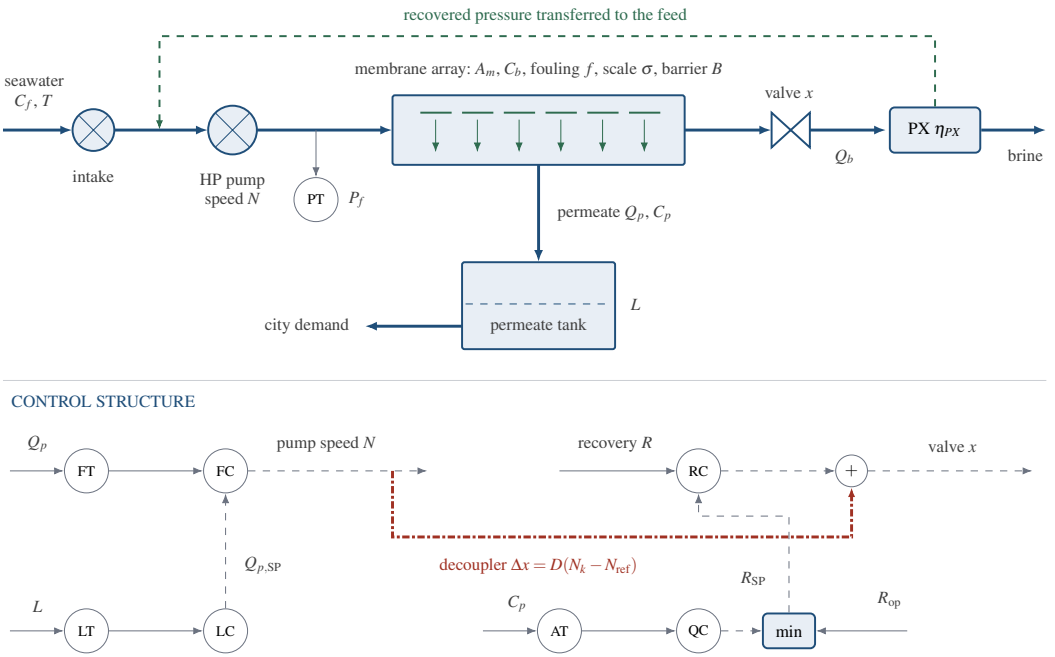


Figure 2.1 The simulated train. Production is controlled on the high-pressure pump (FC) and recovery on the concentrate valve (RC), which is the pairing established in Section 2.4. The permeate tank level cascades onto the production setpoint (LC) and the permeate conductivity low-selects onto the recovery setpoint (QC), so the recovery target R_{SP} is the lower of the operator value and the quality demand and the production target $Q_{p,\mathrm{SP}}$ comes from the tank. The thick dash-dotted path is the one-way static decoupler that carries the pump command directly to the valve, which is the compensation the triangular gain matrix of (2.13) makes possible.

no rinse reverses and which practice controls through antiscalant dosing and through the recovery setpoint [17]. Aging therefore makes the plant a slowly varying system, and a controller commissioned on a clean membrane is running a different plant a month later.

2.3 Process Model

The plant is a single-stage seawater train with a pressure exchanger, sized at 33 000 m^2 of membrane area in one lumped array. Figure 2.1 shows the flowsheet, the instruments and the four controllers.

2.3.1 Working Units and the Design Point

The clock is in minutes, flows in m^3/h, pressures in bar, salinities in g/L and flux in LMH, which is liters per square meter per hour and the unit used on membrane data sheets. At the design point the feed is 35.0 g/L at 25.0 °C, the pump runs at 88 % and the concentrate valve sits at 50 %. That gives 1000 m^3/h of feed, 450.0 m^3/h of permeate at 45.00 % recovery, 62.00 bar of feed pressure, 13.636 LMH of flux, 250.0 mg/L of product and 2.4627 kWh/m^3.

Table 2.1 lists every constant the model uses, so the simulator is reproducible from this chapter alone.

2.3.2 Solution-Diffusion Transport

Water and salt cross the membrane by independent mechanisms [4]. Water crosses in proportion to the applied pressure remaining after the osmotic pressure is subtracted. Salt crosses in proportion to the concentration difference across the membrane and does not depend on pressure,

Table 2.1 Model constants. Every value is the application default.

Symbol	Meaning	Value
A_m	active membrane area	33000 m^2
A_{25}	water permeability at 25 °C	0.6922 LMH/bar
B_{25}	salt permeability at 25 °C	0.06513 LMH
k	polarization mass transfer coefficient	143.11 LMH
$E_{A,w}$	water permeability Arrhenius coefficient	2700 K
$E_{A,s}$	salt permeability Arrhenius coefficient	5100 K
k_π	osmotic coefficient of seawater	0.771 bar/(g/L)
P_p	permeate back pressure	0.5 bar
V_{mem}	feed-channel volume of the array	40 m^3
k_p	pump shut-off head at full speed	90.393 bar
k_Q	pump curve droop	8.0×10^{-6} bar/(m^3/h)^2
C_v	concentrate valve size	393.86 (m^3/h)/bar$^{1/2}$
R_v	concentrate valve rangeability	30
P_{back}	concentrate discharge pressure	1.0 bar
ΔP_d	array pressure drop at design flow	2.50 bar
$n_{\Delta P}$	pressure-drop flow exponent	1.7
$k_{\Delta P}$	pressure-drop increase per unit fouling	3.0
C_{hp}	high-pressure loop capacitance	0.075 m^3/bar
τ_N	pump drive lag	0.10 min
τ_x	concentrate valve travel time	0.15 min
η_{HP}	high-pressure pump efficiency	0.82
η_B	booster pump efficiency	0.78
η_{PX}	pressure-exchanger efficiency	0.96
η_I	intake pump efficiency	0.80
P_{in}	intake and pre-treatment pressure	2.5 bar
k_f	fouling rate at design flux	7.0×10^{-4} min^{-1}
J_0	reference flux for the fouling law	13.636 LMH
Q_{cross}	crossflow below which fouling accelerates	300 m^3/h
k_ϕ	low-crossflow fouling multiplier	2.5
$C_{\text{sat},0}$	saturation salinity, no antiscalant	57.4 g/L
a_s	antiscalant effectiveness	0.45
$D_{1/2}$	antiscalant half-saturation dose	1.5 mg/L
k_σ	scaling rate above saturation	0.15 min^{-1}
k_{ox}	oxidant damage rate	0.020 min^{-1}/(mg/L)
A_t, H_t	permeate tank area and height	60, 6.0 m^2, m
C_{div}, w_{div}	diversion salinity, switch width	600, 25 mg/L
$P_{\max}$	membrane pressure rating	70 bar
$J_{\max}$	design flux limit (alarm only)	17.0 LMH
θ_C	permeate conductivity transport delay	0.45 min
τ_C	permeate conductivity cell response	0.35 min
τ_m	flow, recovery and level transmitter filters	0.05 min
τ_T	intake temperature first-order lag	8.0 min
Δt, Δt_c	integration step, controller scan	0.005, 0.02 min

$$J_w = A\,(\bar{P} - P_p - \pi_w), \qquad J_s = B\,(C_w - C_p). \tag{2.1}$$

That asymmetry governs reverse osmosis quality control. Raising the applied pressure raises the water flux without raising the salt flux, so the permeate becomes fresher, and the fastest correction for an off-specification permeate is more flux rather than less. The permeate composition follows from equating the two fluxes,

$$C_p = \frac{B C_w}{J_w + B}, \qquad Q_p = \frac{J_w A_m}{1000}, \qquad R = \frac{Q_p}{Q_p + Q_b}. \tag{2.2}$$

At design $A = 0.6922\,\text{LMH/bar}$ and $B = 0.065\,13\,\text{LMH}$, which returns 250.0 mg/L of product and 99.286 % rejection.

2.3.3 Concentration Polarization and Temperature

The osmotic pressure that opposes the flux is the one at the membrane wall rather than the one in the bulk. Permeate removal concentrates the layer against the surface, and the two together give

$$C_w = \bar{C}\exp\!\left(\frac{J_w}{k}\right), \qquad \bar{C} = \frac{C_b - C_f}{\ln(C_b/C_f)}, \tag{2.3}$$

with the log-mean standing in for the bulk profile along an array whose ends are C_f and C_b. At design $\bar{C} = 47.82\,\mathrm{g/L}$, the polarization factor is 1.100 and the wall sits at 52.60 g/L, so the wall osmotic pressure is 40.55 bar against a net driving pressure of 19.70 bar. The polarization layer is thickest when the flux is highest, which is a negative feedback that makes the train self-limiting. The layer is also where scale nucleates [17].

Osmotic pressure and both permeabilities carry temperature,

$$\pi = k_\pi C\,\frac{T + 273.15}{298.15}, \quad A(T) = A_{25}\,\mathrm{e}^{E_A\left(\frac{1}{298.15} - \frac{1}{T+273.15}\right)} \tag{2.4}$$

with $E_A = 2700\,\mathrm{K}$ for water and 5100 K for salt. The two activation energies differ, and the salt one is larger, so a warm feed is pressurized more cheaply and produces a saltier permeate while a cold feed costs more pressure and produces a fresher one. Holding both actuators fixed and moving the intake from 16 °C to 32 °C takes the permeate from $400.5\,\mathrm{m^3/h}$ at 159 mg/L to $482.4\,\mathrm{m^3/h}$ at 353 mg/L.

2.3.4 The Hydraulic Node

Feed pressure is the value at which the pump curve, the membrane and the concentrate valve balance,

$$P_{\mathrm{pump}} = k_p\left(\frac{N}{100}\right)^2 - k_Q Q_f^2, \tag{2.5}$$

$$Q_b = C_v R_v^{x/100-1}\sqrt{P_b - P_{\mathrm{back}}} \tag{2.6}$$

with an equal-percentage characteristic of rangeability 30 on the concentrate valve. The node is integrated rather than solved algebraically so that its transient is represented,

$$\frac{\mathrm{d}P_f}{\mathrm{d}t} = \frac{Q_{\mathrm{pump}}(N, P_f) - Q_f}{60\,C_{hp}}, \tag{2.7}$$

$$\frac{\mathrm{d}C_b}{\mathrm{d}t} = \frac{Q_f C_f - Q_b C_b - Q_p C_p}{60\,V_{\mathrm{mem}}}. \tag{2.8}$$

With $C_{hp} = 0.075\,\mathrm{m^3/bar}$ the pressure state has a time constant near three seconds, which is the fastest state in the application and four orders of magnitude faster than the fouling state of (2.10). The array pressure drop is 2.50 bar at design, scales with feed flow to the power 1.7 and is multiplied by $(1 + k_{\Delta P} f)$ as the membrane fouls.

🕮 **Pause.** Try Exercise 2.1 before reading on.

2.3.5 Energy Recovery and Specific Energy

A pressure exchanger returns the concentrate pressure to most of the feed, so the high-pressure pump raises only the water that becomes permeate [10],

$$W = \frac{Q_p(P_f - P_{in})}{36\,\eta_{HP}} + \frac{Q_b\left(P_f - \eta_{PX}(P_b - P_{\mathrm{back}})\right)}{36\,\eta_B} + \frac{Q_f P_{in}}{36\,\eta_I}. \tag{2.9}$$

At design the three terms are 907.0 kW, 114.4 kW and 86.8 kW for 2.4627 kWh/m^3. Setting $\eta_{PX} = 0$ in the same expression gives 4.907 kWh/m^3, which is the difference the device makes and is consistent with the reported industrial trend [6].

2.3.6 Fouling, Scaling and Oxidation

Fouling, scaling and oxidation are separate states with separate rate laws,

$$\dot{f} = r_f \, k_f \, \phi_x \left(\frac{J_w}{J_0} \right)^{1.5} (1 - f), \tag{2.10}$$

$$\dot{\sigma} = k_\sigma \left(\frac{C_b}{C_{\text{sat}}} - 1 \right)_+^{2} (1 - \sigma), \tag{2.11}$$

$$\dot{B} = k_{ox} \, c_{\text{ox}}. \tag{2.12}$$

Fouling grows with flux to the three-halves power and is multiplied by $\phi_x = 1 + k_\phi (1 - Q_b / Q_{\text{cross}})$ when the concentrate flow falls below 300 m^3/h, so a starved crossflow ages the membrane faster than a high flux does [16]. The fouling clock r_f is zero unless a scenario starts it. Scaling is exactly zero below saturation and quadratic above it, and no rinse reverses it. Oxidation raises the salt permeability alone, leaves the water permeability untouched and does not reverse. The antiscalant dose D raises the saturation salinity as $C_{\text{sat}} = C_{\text{sat},0} \left[1 + a_s D / (D_{1/2} + D) \right] / h$, where h is the scaling tendency of the feed, so the benefit saturates in dose, from 57.40 g/L untreated to 74.62 g/L at 3 mg/L and only 79.15 g/L at 8 mg/L.

2.3.7 Permeate Tank and Distribution Demand

Permeate above a diversion salinity of 600 mg/L is sent to waste rather than to the tank through a logistic switch of width 25 mg/L, narrow enough to act as a switch and smooth enough for a fixed-step integrator. The tank is a volume and a salt inventory driven by the on-specification fraction and the distribution demand, and it converts a demand set outside the plant into a production setpoint the train can follow.

2.3.8 Initialization

Initialization solves the equilibrium rather than approximating it. Given a pump speed and a valve position, the feed pressure is bisected until the pump curve and the plant agree, and the concentrate salinity is closed by a damped fixed point on the array salt balance. Both converge to well below display resolution, so the initial condition is a true equilibrium and every scenario starts at steady state. The same solver backs the two manual-mode calculators, which return the pump speed for a chosen permeate flow and the valve position for a chosen recovery.

2.3.9 Assumptions

The array is one well-mixed feed channel with a log-mean bulk salinity rather than a discretized series of pressure vessels. Concentration polarization uses a fixed mass-transfer coefficient. Fouling and scaling are scalar states with empirical laws. Salinity is one equivalent salt. Integration is fixed-step RK4 at 0.005 min with a 1.2 s controller scan, so the trends and the exported CSV reproduce exactly on any machine (Section 1.6.1).

2.4 Gain Structure and Loop Pairing

The steady-state gains of the model are obtained by re-solving the equilibrium of Section 2.3 on either side of the design point with a central difference of 0.5 % of actuator travel, a step small enough that halving it changes no tabulated digit. Table 2.2 lists the gain of the concentrate valve on permeate flow across the maximum.

Table 2.2 Gain of the concentrate valve on permeate flow, pump held at 88 %, 35.0 g/L feed at 25.0 °C. The sign reverses at the design point.

Valve x (%)	Recovery (%)	$\partial Q_p/\partial x$ (m^3/h/%)
42	50.9	+1.772
46	48.0	+0.938
50	45.0	−0.001
54	41.9	−1.045
58	38.6	−2.192

Two effects cancel at the design point. Closing the valve raises the feed pressure, which raises the flux through (2.1), and it raises the concentrate salinity, which raises the osmotic pressure at the membrane wall through (2.3) and lowers the flux. At 45.0 % recovery on this train the two are equal to within the resolution of the numerical derivative. A production controller wired to that valve has no gain to work against at the operating point and a reversed plant one step away from it, which is a defect that no tuning repairs [2]. Volume 7, Chapter 4 meets the same defect at the friction peak of a braking tire, where the slip setpoint sits on the maximum of the plant characteristic. The pump gain on permeate flow is 17.147 m^3/h/% at the design point and does not change sign anywhere in the operating envelope.

The relative gain array of Section 1.4.5 is the standard tool for the pairing decision [2, 18]. For the correct pairing on this train the four steady-state gains at the design point are

$$K = \begin{bmatrix} \partial Q_p/\partial N & \partial Q_p/\partial x \\ \partial R/\partial N & \partial R/\partial x \end{bmatrix} = \begin{bmatrix} 17.147 & -0.001 \\ 0.7029 & -0.7687 \end{bmatrix} \tag{2.13}$$

in units of m^3/h/% and %/%, and the array follows from

$$\lambda_{11} = \left(1 - \frac{K_{12}K_{21}}{K_{11}K_{22}}\right)^{-1} = 1.00005. \tag{2.14}$$

A relative gain of one is the textbook signature of a plant with no interaction, and it is returned here by a plant that interacts strongly. Every term in (2.14) is a product of the two off-diagonal gains, so setting either one to zero returns $\lambda_{11} = 1$ no matter how large the other is. The matrix in (2.13) is triangular rather than diagonal, and a triangular plant has coupling that runs in one direction only.

The direction is the useful part. The valve does not move the permeate flow, so closing the recovery loop leaves the production loop undisturbed. The pump moves the recovery by 0.703 %/% of speed, so closing the production loop disturbs the recovery loop every time production changes. Section 2.6 measures both halves of that statement.

Triangularity also carries its own remedy. A disturbance that arrives entirely from one known signal is cancelled from that signal rather than corrected after it is measured, and one static feedforward is sufficient. In the terms of Section 1.4.6 it is the second term of the simplified decoupler alone, and the first term vanishes with K_{12}. A fully interacting plant requires two decouplers that interact with each other, which is a large part of why plants with genuine two-way coupling are given a model-based controller instead [13, 14]. Chapter 3 has such a plant, with a relative gain above one and a decoupler with both terms.

2.5 Controller Design

2.5.1 Velocity Form PID

All four loops share the velocity-form PID of Section 1.6.2, Equations (1.11) to (1.14), with the derivative taken on the measurement [19]. The controller scan interval is $\Delta t_c = 1.2$ s. Anti-windup

is back-calculation applied to the integral alone, with a tracking time T_t equal to τ_I on every loop. Derivative action is provided only on the tank level loop, where the measurement is filtered with a time constant of 0.05 min before it is differenced, and it is zero in every shipped tuning. Tunings follow the SIMC rules [20] applied to gains measured on the app itself.

2.5.2 The Two Regulatory Loops

The production loop takes permeate flow to the high-pressure pump with $K_c = 0.02\,\%/\mathrm{m}^3/\mathrm{h}$ and $\tau_I = 0.4\,\mathrm{min}$. It is reverse acting, because the process gain is positive and the pump command falls as the measured permeate flow rises. The recovery loop takes recovery to the concentrate valve with $K_c = 0.6$ and $\tau_I = 0.6\,\mathrm{min}$. It is direct acting, because opening the valve lowers the recovery and the valve command therefore rises as the measured recovery rises. Neither loop carries appreciable dead time, and the production loop reaches 95 % of a $60\,\mathrm{m}^3/\mathrm{h}$ step in 4.01 min. Each controller is seeded on the actuator it drives in the chosen pairing, and a pairing change in automatic re-seeds it by output tracking, so changing the structure moves no actuator.

2.5.3 Static Decoupler

The pump moves recovery by +0.7029 %/% and the valve by −0.7687 %/%, so a valve move of $0.7029/0.7687 = 0.914$ times the pump move cancels the coupling before the recovery controller sees an error. The same ratio can be measured in manual mode. The calculator that returns the valve position for a chosen recovery, run for 45 % recovery at 87 and 89 % pump speed, gives $-K_{21}/K_{22} = 0.915$ per percent of speed from the two valve positions it reports. The app ships 0.95. The implementation is given by Equation (2.15),

$$\Delta x = D\,(N_k - N_{\mathrm{ref}}), \qquad x_k = \mathrm{clamp}\big(x_k^{\mathrm{PID}} + \Delta x\big) \tag{2.15}$$

where x_k^{PID} is computed with the controller's own output limits shifted by Δx rather than back-calculated against the applied value. Back-calculating against the applied value feeds the feedforward into the integral, which cancels the very correction it makes and leaves a steady offset. The reference speed $N_{\mathrm{ref}} = 88\,\%$ is calibrated on a clean membrane, which is the assumption the decoupler exposes.

2.5.4 Supervisory Layer

The permeate tank level cascades onto the production setpoint with $K_c = 30\,\mathrm{m}^3/\mathrm{h}/\%$ and $\tau_I = 25\,\mathrm{min}$, deliberately slow so that the tank absorbs demand rather than passing it to the membrane. The permeate conductivity drives a second controller whose output is a recovery setpoint, low-selected against the operator target,

$$R_{\mathrm{SP}} = \min\big(R_{\mathrm{op}},\, R_{\mathrm{quality}}\big)\,, \tag{2.16}$$

so the override can only ask for less recovery. The quality controller ($K_c = 0.2\,\%$ of recovery per mg/L, $\tau_I = 12\,\mathrm{min}$) receives an external reset while the operator target wins the selector. Its integral term is set to the applied recovery setpoint in place of its own integration, so its output is that setpoint plus its proportional term. It cannot wind up while it is out of control, and with a constant operator target it takes over without a bump on the first scan at which the measured salinity exceeds the limit [19]. The selected value is never written back into the operator target, so the recovery returns to R_{op} once the quality demand clears.

2.5.5 Model Predictive Control

A separate structure selector can replace the flow and recovery loops, the decoupler and the quality override with one model predictive controller (MPC) for both actuators, the receding-horizon method behind most industrial multivariable control [21], set up as Section 1.6.3 describes in general. Scenario 11 alone loads with the MPC, and the application and its other scenarios start with the loops. Every $\Delta t_m = 0.1\,\mathrm{min}$, which is five scans of Δt_c, the controller solves a quadratic program for the

pump speed and the valve position over $H = 90$ steps (9 min) in twelve move blocks, the shortest a single step near the present and the longest 45 steps at the far end of the horizon, 24 decisions in all. The train receives only the first move, held until the next plan. With the recovery loop open the valve stays where the operator set it and the plan moves the pump alone. The pairing then only chooses which loops would replace a rejected plan, and the level cascade still sets the flow target. Predictive control of reverse osmosis units has been built on identified transfer function models [22], and optimization-based control has been used to minimize the specific energy of a train [23].

The controller's model is the membrane of Section 2.3 rather than an identified two-by-two, because the step responses have two modes. A 1 % pump step raises the permeate flow by 20.55 $\mathrm{m^3/h}$ within half a minute and leaves it 17.16 $\mathrm{m^3/h}$ higher after the concentrate has salted up, while a 1 % valve step moves it by −3.79 $\mathrm{m^3/h}$ at half a minute and by −0.01 $\mathrm{m^3/h}$ after twenty. The near-zero valve gain of Table 2.2 is therefore a fast effect and a slow one cancelling, and first-order fits of the eight responses of permeate flow, recovery, feed pressure and permeate salinity to the two actuators leave an RMS residual of 1.0 to 26.3 % of each response. The plan uses Equations (2.1) to (2.6) and the array pressure drop with the constants of Table 2.1 and a design feed of 35 g/L at 25 °C through a new pump, with the dynamics of Equation (2.17),

$$\begin{aligned} \tau_N \frac{\mathrm{d}\hat{N}}{\mathrm{d}t} &= u_N - \hat{N}, \qquad \tau_x \frac{\mathrm{d}\hat{x}}{\mathrm{d}t} = u_x - \hat{x} \\ \frac{\mathrm{d}\hat{C}_b}{\mathrm{d}t} &= \frac{Q_f C_f^0 - Q_b \hat{C}_b - Q_p C_p}{60\, V_{\mathrm{mem}}} \\ \left(Q_p, Q_b, P_f, C_p\right) &= \mathscr{M}\left(\hat{N}, \hat{x}, \hat{C}_b;\ \hat{a}A_{25},\ \hat{b}B_{25}\right) \end{aligned} \tag{2.17}$$

where u_N and u_x are the pump and valve commands, $\hat{N}$, $\hat{x}$ and $\hat{C}_b$ are the model's pump speed, valve position and concentrate salinity, $C_f^0 = 35\,\mathrm{g/L}$ is the design feed salinity, and $\mathscr{M}$ returns the flows, the feed pressure and the permeate salinity at which the pump curve, the membrane and the valve agree, with the declared water and salt permeabilities scaled by the estimates $\hat{a}$ and $\hat{b}$. The feed pressure is treated as settled, because the node of Equation (2.7) has a time constant of 3.2 s. At every plan, central differences about the current estimate give a linear model, discretized over Δt_m with its free motion kept as an affine term. At the design point its steady-state gains are those of Equation (2.13), 17.1473, −0.0010, 0.7029 and −0.7687, its valve gain on permeate flow is +1.772 at 50.9 % and −2.192 at 38.6 % recovery as in Table 2.2, and its salt balance has a time constant of 3.235 min against 3.207 min in the tail of the plant's own pump step. Against ±1 % step tests on the plant at the design point its RMS residual is 0.31 to 2.06 % of each response's range, most of it in the first 0.2 min, while the plant's pressure node is still moving.

Two multipliers and an output bias carry what the membrane has become and remove steady offset [24]. Volume 9, Chapter 4 removes offset with a Kalman estimator that carries a fuel bias instead. At every scan $\hat{a}$ is the measured flux divided by A_{25} and by the net driving pressure that the measured feed pressure leaves over the estimated wall osmotic pressure, and $\hat{b}$ is the conductivity reading divided by the model's permeate salinity passed through the cell's declared 0.45 min delay and 0.35 min response, each filtered over 1 min. The concentrate salinity is integrated from the measured permeate flow, recovery and conductivity. A bias on the permeate flow, the recovery and the feed pressure, filtered over 0.2 min, compares each measurement with the model's output passed through the declared pressure node and transmitter filters. Fouling, scale, temperature, feed salinity and pump wear therefore reach the plan only through these estimates. At the end of scenario 11 $\hat{a} = 0.771$, where 20.89 % fouling leaves 0.791 of the clean permeability, because $\hat{a}$ also absorbs the larger pressure drop of the fouled array. Over the same run the relinearized valve gain on permeate flow moves from +1.476 to −6.802 $\mathrm{m^3/h/\%}$ as the recovery falls to 33.27 %, crossing the sign change of Table 2.2 without a rejected plan.

Each plan minimizes the objective Φ of Equation (2.18),

$$\begin{aligned}\Phi = &\sum_{k=1}^{H}\left[\left(\hat{Q}_{p,k}-Q_p^*\right)^2+q_R\left(\hat{R}_k-R^*\right)^2\right]\\ &+\mu\sum_{j=1}^{12}\left(s_N\,\Delta u_{N,j}^2+s_x\,\Delta u_{x,j}^2\right)+\sum_{\ell}\left(\rho_\ell\,\varepsilon_\ell^2+\sigma_\ell\,\varepsilon_\ell\right)\\ \text{s.t. }&\hat{P}_{f,k}\le P_{\max}-\delta_P+\varepsilon_P,\quad \hat{C}_{p,k}\le C_{\lim}-\delta_C+\varepsilon_C\\ &\hat{Q}_{p,k}\le J_{\max}A_m/1000-\delta_Q+\varepsilon_Q,\quad \varepsilon_\ell\ge 0\\ &40\le u_{N,j}\le 100,\quad 5\le u_{x,j}\le 100\end{aligned}\tag{2.18}$$

where the hats are predictions at step k, $\Delta u_{N,j}$ and $\Delta u_{x,j}$ are the command changes between consecutive blocks, counted for the first block from the commands now applied, Q_p^* and R^* are the flow and recovery targets after the adjustment described below, $q_R = 4\,w_R$ per %^{2} against 1 per (m^3/h)^2, $s_N = 10$ and $s_x = 1$ per %^{2}, and the move price μ and the recovery weight w_R are both 1 unless the reader changes them. The salinity limit $C_{\lim}$ is the quality limit when level and quality supervision is on and the 500 mg/L specification otherwise, and the flux ceiling $J_{\max}A_m/1000$ is 561 m^3/h. The bands sit $\delta_P = 0.2$ bar, $\delta_C = 2$ mg/L and $\delta_Q = 1$ m^3/h inside the real limits, with prices $\rho_P = 1\times10^5$ per bar^2 and $\sigma_P = 1\times10^4$ per bar on the pressure and 100 and 100 on the salinity and the flow. The shared solver gives each soft limit one slack ε_ℓ for the whole horizon. Where the free response already breaks a limit, that bound starts at the free response and decays to the band with a time constant of 1 min, or 2 min for the salinity, so the program stays feasible. The command bounds are hard and equal the loops' clamps.

The targets are made consistent with the limits before each plan, and this adjustment holds the rating, as Section 2.6.8 measures. At the linear model's steady state the feed pressure and the permeate salinity are linear in the two targets. If holding both would leave less than 0.5 bar under the rating or 5 mg/L under the salinity limit, R^* is lowered until that margin is restored, and Q_p^*, which is also capped 3 m^3/h under the flux ceiling, is lowered only once R^* has reached its 25 % floor or the valve belongs to the operator. Production is what the distribution system draws, and a lower recovery relieves the pressure because a less concentrated channel presents less osmotic pressure at the wall, so recovery is spent first. The steady state used for this step ignores the pump's top speed. The plan carries no price on energy. Two handles holding two targets have no steady freedom left to trade, so a linear price of 1 per kilowatt per step on Equation (2.9), linearized at the design point, only pulls the targets off, and over scenario 8 it raises the flow error from 12.4 to 678.6 m^3/h min for 0.19 % less specific energy.

The controller accepts a plan only if the solver's status is optimal, which is checked even when the solve returns success, every returned array has its shape, the cost, the slacks and the predictions are finite, no slack is negative, all planned commands respect their bounds and the first of them equals the move applied. Otherwise the loops of the chosen pairing take over until a later plan is accepted. Under the MPC those loops follow the commands the plan applies on every scan, without preload, so a rejection forced at 150 min of scenario 11, while the plan holds the pressure under the rating, moves the valve by 0.0007 % on the hand-over scan. No run reported here rejected a plan.

The plant integrates its eleven states with the run's own parameters, the instruments add their own filters, delay and noise, and the plan holds only the declared constants and the design feed, so a change a reader makes to the membrane, the pump or the feed reaches the controller as model mismatch. In a replay test, a fresh controller outside the simulator receives the measurements of 150 min of scenario 11, recorded on a plant with a different hidden membrane, pump, valve and feed, and reproduces all 7500 applied commands bit for bit. In the run measured for this chapter a plan takes 0.403 ms on average in Chrome, a load of 0.40 % of a processor core when the simulation advances one minute per wall-clock second and 8.1 % at its fastest rate of 20 min per second.

2.6 Results

Twelve scenarios run on one plant, each with the briefing, completion rule and data export of Section 1.6.4. Every number below is read from the app at the stated time with the disturbances at their scenario values. The measurement noise generators are off in every shipped scenario, so each run is deterministic and repeats exactly. Each case is selected, any setting named in the text is changed, and the case is run from $t = 0$.

2.6.1 Production Loop Alone

With the recovery loop in manual and the valve pinned at 50 %, a production step from 450 to 510 m^3/h at 10 min in scenario 3 reaches and stays within 5 % of the step after 4.01 min. Recovery, which is uncontrolled in this configuration, drifts from 45.00 to 47.30 % because the pump raises permeate and feed by different fractions. The saturation index rises from 0.850 to 0.887. On this feed the excursion is harmless. On a harder feed or at a lower antiscalant dose it consumes the whole scaling margin, and it arrives as a side effect of a production change that the production loop reports as a success.

2.6.2 The Reversed Pairing

Scenario 5 assigns production to the concentrate valve and recovery to the pump, which is the arrangement a valve-controls-flow reading of the flowsheet suggests. Its production controller is given a gain 15 times larger than the pump loop's, which is the authority the valve has away from the maximum.

The scenario starts in manual and transfers both loops to automatic at 5 min. The transfer is bumpless, because each controller has tracked its own actuator in manual, and nothing moves until the setpoint steps from 450 to 500 m^3/h at 10 min. The flow controller has no gain to work against at the maximum of Table 2.2, so it integrates, and it closes the valve, which is the right direction only on the opening side of the maximum. Moving down the closing side of the curve raises the recovery, the recovery controller answers by slowing the pump, the permeate flow falls further and the two loops drive each other. The valve reaches its 5 % stop 3.9 min after the step, the recovery peaks at 76.3 %, and the concentrate passes saturation (a saturation index of 1.14) long enough to cost 4.07 % of permeability to scale. At 120 min the pump has been dragged from 88 to 68.39 %, feed pressure has collapsed from 62.00 to 42.01 bar and the train is making 82.5 m^3/h of 1241 mg/L permeate against a 500 m^3/h setpoint. The integrated absolute flow error over those two hours is 42 118 m^3/h min. The state at 90 min is shown in Figure 2.3.

Nothing else was at fault. Both loops are stable alone, the tuning is reasonable for the gain the valve has away from the design point, the transfer is bumpless, and neither loop carries dead time. Only the pairing is wrong.

> **TRY IT IN THE APP** **5 · The Pairing You Would Have Guessed**
>
> Load the scenario and press Run. Both loops switch to automatic at 5 min, and the production setpoint steps from 450 to 500 m^3/h at 10 min. Watch the concentrate valve trend. The valve runs to its 5 % stop 3.9 min after the step, the recovery loop slows the pump, and the permeate flow falls instead of rising.

2.6.3 Effect of the Static Decoupler

Table 2.3 carries the central measurement of this chapter. The correct pairing, whose relative gain array reads 1.000 05, deviates 0.937 points in recovery and requires the concentrate valve to travel from 50.00 to 53.44 % to bring it back, taking 8.4 min to return within 0.05 of a point. That is the interaction (2.14) reports as absent. Figure 2.2 draws the table on logarithmic axes.

Adding (2.15) reduces the peak by a factor of 4.4 and the integral error by a factor of 7.2, and it does so with a single multiplication in the control scan. Measured over the 50 min following the step in scenario 6, rather than across the paired step and return of scenario 4, the peak falls from 0.937

Table 2.3 Response of the recovery loop to the 70 m^3/h production step of scenario 4, with the pairing chosen before the run. The peak deviation is taken over the 135 min following the step and the integral error over the whole 180 min scenario. The swapped row uses the scenario 5 flow-loop tuning ($K_c = 0.30\,\%/m^3/h$, $\tau_I = 2.5$ min). It starts without a bump at the design point and leaves it with no step at all, so 885.1 of its 3585.0 % min accrue before the step.

Structure	Peak $\lvert\Delta R\rvert$ (% pts)	Recovery IAE (% min)	Final Q_p (m^3/h)
Direct pairing	0.937	6.87	450.0
Direct plus decoupler	0.213	0.95	450.0
Swapped pairing	24.35	3585.0	450.0

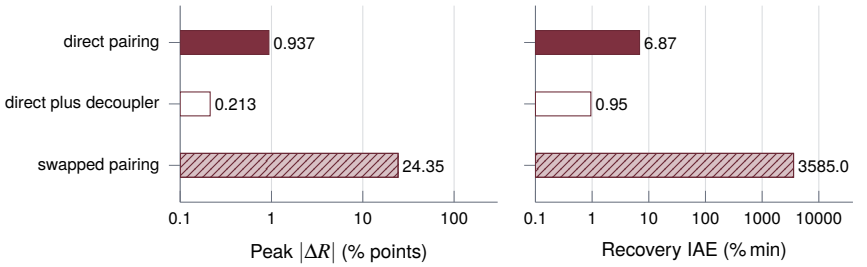


Figure 2.2 The recovery loop on the production step of scenario 4, from Table 2.3, on logarithmic axes: the peak deviation (left) and the integral error (right). The dark bar is the direct pairing, the white bar the direct pairing with the decoupler and the hatched bar the swapped pairing. Notice that the decoupler cuts both measures of an interaction that a relative gain of 1.000 05 reports as absent, while the swapped pairing is worse by orders of magnitude.

Table 2.4 Decoupler gain sweep on the scenario 6 production step, measured over the 50 min following the step. Each gain is set before the run. The design value from the measured gain ratio is 0.914.

Decoupler gain D	0.00	0.50	0.95	1.30	1.60
Peak $\lvert\Delta R\rvert$ (pts)	0.937	0.452	0.160	0.504	0.864
Recovery IAE (% min)	3.44	1.38	0.48	1.92	3.15

to 0.160 points. The swapped row fails without being disturbed. With the recovery loop closed on the pump, opening the valve raises the permeate flow by about $K_{12} - K_{11}K_{22}/K_{21} = 18.75\,m^3/h/\%$, because the pump speeds up to hold the recovery, which is the opposite of the sign the flow controller acts on. Its relative gain is $1 - \lambda_{11} = -0.000\,05$, and the design point is unstable. The recovery leaves it by a full point 4.0 min into the run and settles by 20 min at 22.4 %, with the pump on its 100 % stop and the valve at 83.3 %, while the flow is held on setpoint from the far side of the maximum. With the recovery loop also detuned to $K_{c,R} = 0.05$ and $\tau_{I,R} = 1.5$ min, the drift grows slowly enough (0.1 point at 33 min) for the step to decide the outcome, and the run ends with the valve on its 5 % stop and 82.2 m^3/h of permeate. The row is reported at the same scale because decoupling is a refinement applied to a correct pairing rather than a substitute for one.

The decoupler gain is not a free parameter. Table 2.4 sweeps it across the design value, and over-decoupling at 1.6 is nearly as costly as no decoupling at all. The curve has a minimum rather than a plateau, which is what makes the reference speed N_{ref} worth attention, because it is a static calibration on a plant whose gains drift.

🕮 **Pause.** Try Exercise 2.2 before reading on.

Table 2.5 Specific energy and saturation index against recovery with permeate production held at 450 m^3/h and the antiscalant at 3 mg/L.

R (%)	N (%)	P_f (bar)	SEC (kWh/m^3)	SI
40	88.33	60.41	2.5156	0.780
45	88.00	62.00	2.4627	0.850
49	88.19	63.55	**2.4520**	0.916
52	88.55	64.89	2.4590	0.973
54	88.90	65.88	2.4703	**1.015**
58	89.83	68.13	2.5082	1.111

TRY IT IN THE APP 6 · A One-Way Decoupler

Load the scenario, run it to 60 min and watch the recovery after the production step at 10 min. Then set the Pump-to-valve decoupling gain to 1.6, press Reset and run again. The peak recovery deviation over the 50 min after the step grows from 0.160 points at the shipped 0.95 to 0.864 points (Table 2.4).

2.6.4 Warm Feed and the Quality Override

The intake warms from 25.0 to 32.5 °C over scenario 7. The production loop holds its setpoint throughout, because a warm feed is pressurized more easily. The permeate quality is what changes. With the supervisory layer set to the tank level alone, so that the quality override is out of the loop, the permeate reaches 379 mg/L at the end of five hours against 250 mg/L at design, and the specific energy falls to 2.358 kWh/m^3 because the same production requires only 59.06 bar. A seawater plant therefore produces its cheapest water and approaches its salinity specification most closely in the same season.

With the override active at a 340 mg/L internal limit, well inside the 500 mg/L contractual specification, (2.16) pulls the recovery setpoint from 45.0 to 32.6 % and holds the product at 340 mg/L. It takes the selector at 26.1 min, the first scan at which the conductivity cell reads above the limit, and the permeate peaks at 356.2 mg/L at 37 min before the lower recovery returns it to the limit. The cost is quantified: feed pressure 55.69 bar but specific energy 2.607 kWh/m^3, which is 10.6 % above the unconstrained case. Lower recovery yields less permeate per unit of feed pressurized, which is what the override spends. The trade is favorable only because off-specification water has no value.

2.6.5 Specific Energy Against Recovery

Recovery is the efficiency variable. Recovering more of the feed leaves less seawater to be raised to pressure and discarded, and it also produces a saltier concentrate whose osmotic pressure the pump has to overcome. The two effects produce an interior optimum, which is well established for spiral-wound seawater trains [15] and which the application resolves numerically at fixed production.

The minimum in Table 2.5 sits at 49 % recovery and 2.4520 kWh/m^3, and it is shallow. The whole curve spans 2.6 % between 40 and 58 % recovery, which is why the operating point is a standing argument rather than a calculation. The saturation index crosses one at 53.3 %, so the energy optimum is feasible with 4.3 points of margin. That margin is not a property of the plant. It is a property of an antiscalant dose and of a feed chemistry that is not measured online. At 52 % recovery the index is 1.265 untreated, 1.072 at 1 mg/L and 0.973 at 3 mg/L, for $7.27/h of chemical.

2.6.6 Fouling Drift and Oxidative Damage

Scenario 8 starts the fouling clock at 5 min and runs six hours at constant production. No loop loses control of its variable at any point.

Table 2.6 records a control system holding every setpoint while the plant it holds them on becomes expensive. Specific energy rises from the design value of 2.463 to 2.759 kWh/m^3, 12.0 %, over the

Table 2.6 Fouling with permeate production held at 450 m^3/h, scenario 8 run on past its six hours to 600 min. The element rating is 70 bar and the pump gain $\partial Q_p/\partial N$, by central difference on the fouled steady state, is in m^3/h/%.

t (min)	f (%)	P_f (bar)	N (%)	SEC (kWh/m^3)	$\partial Q_p/\partial N$
10	0.3	62.07	88.05	2.466	17.12
120	7.7	63.94	89.21	2.556	16.61
240	15.2	66.09	90.54	2.656	16.06
360	22.0	68.38	91.93	2.759	15.53
480	28.3	70.83	93.39	2.866	15.01
600	34.1	73.45	94.93	2.978	14.50

six hours of the scenario and by 20.9 % at 600 min, on the largest operating cost a desalination plant carries. Feed pressure crosses the 70 bar element rating at 440.2 min and 26.3 % fouling. The production controller reports no error at any point, because production is not what is degrading. Every gain in the plant moves at the same time. The production loop's own plant gain falls from 17.12 to 14.50 m^3/h/%, a loss of 15.3 %, which detunes the loop in the safe direction here and in the unsafe direction immediately after a clean.

> **DID YOU KNOW? Every Setpoint Held, Every Gain on the Move**
>
> From the simulator: in scenario 8 the membrane fouls while every loop holds its setpoint. By 600 min the feed pressure has risen from 62.07 to 73.45 bar, past the 70 bar rating, and the pump's own gain on permeate flow has fallen from 17.12 to 14.50 m^3/h/% (Table 2.6).

Scenario 10 is a disturbance of a different kind. Twenty-five minutes of dechlorination failure at 0.6 mg/L raises the salt permeability by 30.0 % through (2.12), and it remains raised. The product moves from 250 to 324 mg/L and rejection from 99.286 to 99.073 %, permanently, on a train whose permeate flow holds 450.0 m^3/h throughout the event and afterward. Just under a third of the distance to the 500 mg/L specification is consumed in less than half an hour and no loop recovers it. The production measurement alone records nothing at all.

2.6.7 Graded Shift

The graded scenario is five hours on a train that is already fouling, against a feed whose salinity and scaling tendency rise, an intake that warms and a distribution demand that rises from 450 to 510 m^3/h and then falls to 420 m^3/h. The target is 2300 m^3 delivered, production within 8 m^3/h of its setpoint over the last third of the shift, product under 500 mg/L throughout, no scale, no time above 70 bar and under 2.60 kWh/m^3 averaged over the delivered water. It loads with weak tuning, the recovery loop open, both supervisors off and the antiscalant at 2 mg/L. The recommended configuration closes both loops with the tuning of scenario 4 and the decoupler, turns on level and quality supervision with a 380 mg/L limit, and sets 46 % recovery and 5 mg/L of antiscalant.

The as-loaded run fails on volume alone, at 2238 m^3, 62 m^3 short of the target, because with both supervisors off the production setpoint carries no information about the distribution demand. It spends only 69.9 % of the shift inside the flow band, although its worst deviation over the judged last third, 7.5 m^3/h, passes. The tank drains from 60 to 17.3 %, short of the 12 % low-level alarm. Closing the level cascade converts the tank inventory into a production setpoint, and the recommended configuration passes with a margin of 43 m^3. Among the structural elements the cascade alone accounts for the delivered volume and for the flow integral error falling from 1821 to 270 m^3/h min. It needs the tuned flow loop to do so. With the shipped weak tuning, the cascade alone delivers 2426 m^3 but lets the production wander 117 m^3/h from its setpoint over the last third.

Table 2.7 The graded five hour run. Each row after the second changes one element of the recommended configuration before the run. In band is the share of the whole shift with production within 8 m^3/h of its setpoint.

Configuration	Volume (m^3)	In band (%)	SEC (kWh/m^3)	Scale (%)	Max SI
As loaded	2238	69.9	2.480	0.000	1.000
Recommended	2343	98.1	2.537	0.000	0.958
dose 2 mg/L	2343	98.1	2.538	0.677	1.025
no decoupler	2343	98.0	2.537	0.000	0.958
no override	2343	98.1	2.537	0.000	0.958

The remaining rows are informative for what they do not show. Returning the antiscalant to 2 mg/L scales the membrane, with 0.677 % of permeability lost irreversibly and a peak saturation index of 1.025, and every other column unchanged. Removing the decoupler or demoting the override changes the table only in the in-band share, by 0.1 point without the decoupler. The decoupler still cuts the recovery integral error from 12.56 to 6.78 % min, by 46 %, but it leaves the peak deviation at 0.79 points (0.78 without it), because the level cascade moves the production setpoint far more slowly than the setpoint steps of scenarios 4 and 6. The override binds once, from 193.5 to 213.1 min, after the conductivity cell crosses the 380 mg/L limit, and lowers the recovery target to 45.35 %. The product then spends 14.2 min above the limit with a peak of 382.6 mg/L, against 19.1 min and 383.4 mg/L with the override demoted.

Two of the four structural elements carry this particular shift and two are insurance against disturbances larger than the shift contains. Structure that returns nothing on the scorecard of a favorable shift is not thereby unnecessary, and a scorecard that cannot separate the two is a poor guide to what to build [8].

2.6.8 Predictive Controller Against the Decoupled Loops

Table 2.8 compares the MPC with the decoupled loops on the cases in which a limit binds or production steps. Each case runs for its own duration with the noise off, as it ships. The loops are the decoupled pairing with the recovery loop closed and the case's supervisory layer, which is level and quality supervision in scenarios 7, 8, 11 and 12 and none in scenarios 4 and 9, whose events are operator setpoint steps that the level cascade would overwrite. The MPC carries the settings of scenario 11 into every case without retuning. Three comparators separate what the plan contributes. The pressure override is a PI controller on the measured feed pressure, with its setpoint 0.5 bar under the rating, that low-selects the recovery target with external reset ($K_c = 4\,\%/\text{bar}$, $\tau_I = 8$ min, chosen on scenario 11 after the fact), and it exists only in the measurement scripts because the application ships none. Hindsight is the best fixed tuning of the loops for that case on a grid searched after the fact, and the last comparator hands the loops the plan's own targets every 0.1 min. Flow error integrates the distance between the permeate flow and the flow that the tank or the operator requests.

Scenario 11 runs a membrane that fouls at its normal rate from $t = 0$ at 50 % recovery with level and quality supervision, and the distribution demand rises from 450 to 540 m^3/h between 90 and 200 min. Both structures deliver 2415 m^3. The loops hold the recovery at 50 % and let the feed pressure follow the fouling and the demand, so it spends 108.6 min over the 70 bar rating and reaches 74.00 bar. The MPC takes its recovery target below 49.9 % at 97.5 min and keeps the pressure at or below 69.81 bar, as shown in Figure 2.4. At 200 min the recovery is 33.6 % with the pump at full speed, while the target, blind to that speed limit, has fallen to 31.6 %, and by 300 min both are back at 49.21 %. Holding the rating costs 2.7569 against 2.6755 kWh/m^3 over the run, 3.0 % more, and 3.009 against 2.811 kWh/m^3 over the peak, because a lower recovery raises more seawater to pressure for the same permeate.

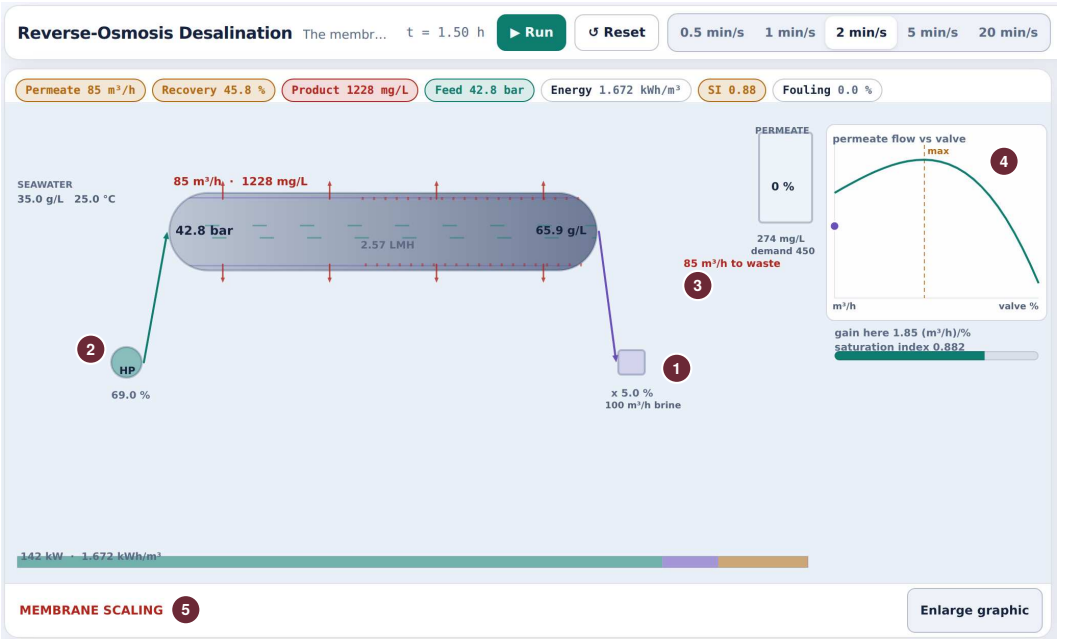


Figure 2.3 The reversed pairing of scenario 5 at 1.50 h, eighty minutes after a production setpoint step from 450 to 500 m^3/h was applied to a controller wired to the concentrate valve, cropped to the application header and the plant view. (1) The concentrate valve has integrated to its 5 % stop. (2) The recovery loop has dragged the pump down to 69.0 %. (3) The permeate, at 1228 mg/L, is diverted to waste, and the permeate tank has emptied. (4) The operating map plots permeate flow against valve position and marks the maximum on which the controller began. (5) The scaling alarm records the recovery excursion that followed the step.

The ablations place the gain in the target adjustment rather than in the optimization. Without it the same plan bends its pressure slack and spends 96.1 min over the rating, because one slack for the whole horizon costs less than a recovery target held over 90 steps. The model frozen at its first plan still holds 69.80 bar, so relinearizing matters little on this case. The classical override also holds the rating, at 69.87 bar, with a flow error of 500.8 at the shipped tuning and 69.1 at the hindsight tuning against the plan's 107.3 m^3/h min. What the optimization adds is coordination while the limit binds. The loops handed the plan's targets every 0.1 min swing the valve through 134.8 % of travel, take the recovery down to 25.07 % and still spend 11.7 min over the rating with a flow error of 1410.4 m^3/h min, while the plan, moving both actuators together, holds the rating with a flow error of 107.3. With instrument noise on the flow meter (2 m^3/h), the conductivity cell (5 mg/L) and the tank level (0.3 %) at four seeds, the MPC spends 2.0 to 4.0 min over the rating against 109.6 to 113.9 min for the loops, and its highest pressure, 70.30 bar, lies 0.5 bar above its band. Halving the integration step moves only the flow errors of the loops, the MPC and the MPC without its target step, by at most 0.3.

In the other cases the plan holds the rating wherever the loops break it. With the recovery raised in steps to 58 % in scenario 9 the loops spend 73.5 min over the rating while the MPC and the override hold it, and the lower recovery they choose also cuts the permeability lost to scale from 20.354 to 13.163 and 12.852 %, although the plan carries no saturation limit. Without its target step the plan still holds the rating in scenario 9, at 69.80 bar. Run on to 600 min, scenario 8 needs more than the pump can give. The MPC holds 69.88 bar with the recovery down to 26.78 % and the pump at full

Table 2.8 The MPC Against the Decoupled Loops, With the Classical Pressure Override and the Plan's Information Handed to the Loops

Case	Structure	$P_{f,\max}$ (bar)	Over rating min (bar min)	Over $C_{\lim}$ min (mg/L min)	Flow error	SEC
11	loops	74.00	108.6 (302.7)	0	301.1	2.6755
	MPC	69.81	0	0	107.3	2.7569
	MPC, no target step	70.75	96.1 (35.8)	0	106.8	2.7129
	MPC, model frozen	69.80	0	0	147.9	2.7535
	loops + override	69.87	0	0	500.8	2.7549
	hindsight + override	69.87	0	0	69.1	2.7506
	hindsight	74.00	108.3 (298.9)	0	39.3	2.6752
	loops, plan's targets	70.47	11.7 (3.4)	0	1410.4	2.7840
9	loops	73.11	73.5 (111.8)	0	135.3	2.5256
	MPC	69.72	0	0	8.9	2.5130
	loops + override	69.57	0	0	96.3	2.5099
8[a]	loops	73.45	159.6 (271.1)	0	138.5	2.7109
	MPC	69.88	0	0	96.0	2.8441
	hindsight + override	70.19	7.7 (0.65)	0	2248.6[b]	2.8295
7	loops	62.00	0	274.7 (742.6)	72.4	2.5629
	MPC	62.00	0	14.9 (42.7)	44.1	2.6506
	loops, plan's target	62.00	0	19.3 (122.7)	88.6	2.6468
4	loops	66.03	0	0	157.4	2.5307
	MPC	66.13	0	0	15.4	2.5308
12[c]	loops, met	66.58	0	14.2 (21.2)	270.5	2.5372
	MPC, met	66.59	0	0	27.9	2.5374

Flow error in m^3/h min, SEC in kWh/m^3 over the run. The salinity limit $C_{\lim}$ is 340 mg/L in scenario 7 and 380 mg/L in scenarios 11 and 12, and the permeate stays under 291 mg/L in the others. Every MPC run planned every 0.1 min (3000 plans on scenario 11) and rejected none, at 0.39 to 0.44 ms per plan in Node in the measured runs. [a]Run on to 600 min. [b]The level cascade's request winding up against a pump at full speed, not lost water (4491.9 against 4490.4 m^3 for the MPC). [c]The challenge with its suggested settings.

speed for the last 61.1 min, at 4.9 % more specific energy than the loops, which spend 159.6 min over the rating. Without its target step the plan spends 105.7 min over it there. In scenario 7 the plan treats the 340 mg/L quality limit as its own constraint and is over it for 14.9 min against 274.7 min under the shipped quality override, at 3.4 % more energy, and the loops handed the plan's recovery target come to 19.3 min.

Two results are losses. Flow tracking is mostly tuning. The hindsight loops of Table 2.8 reach a flow error of 39.3 on scenario 11 while spending 108.3 min over the rating, and a wider grid finds loops that reach 4.6 there, still 108.2 min over the rating, and 14.0 against the plan's 27.9 on the challenge, with 41.7 against 13.8 % of pump travel. On the production step of scenario 4 the plan lets the recovery swing 1.83 points, where the decoupler of Equation (2.15) holds it to 0.21. The plan weighs 1 per $(m^3/h)^2$ of flow error against 4 per $\%^2$ of recovery error, so the 70 m^3/h step outweighs the recovery in its first decisions, and with the recovery weight raised to 4 the swing falls to 0.72 points.

The challenge is met under both structures with its suggested settings, at 2343 m^3 with production within 0.1 m^3/h of its setpoint over the last third under the MPC and 0.7 m^3/h under the loops. The MPC also keeps the permeate under the unscored 380 mg/L quality limit, peaking at 378.2 mg/L, where the loops exceed it for 14.2 min with a peak of 382.6 mg/L. As loaded, with the weak tuning, the recovery loop open and both supervisors off, neither structure meets it, at 2238 and 2250 m^3, and the MPC loses 0.025 % of the permeability to scale.

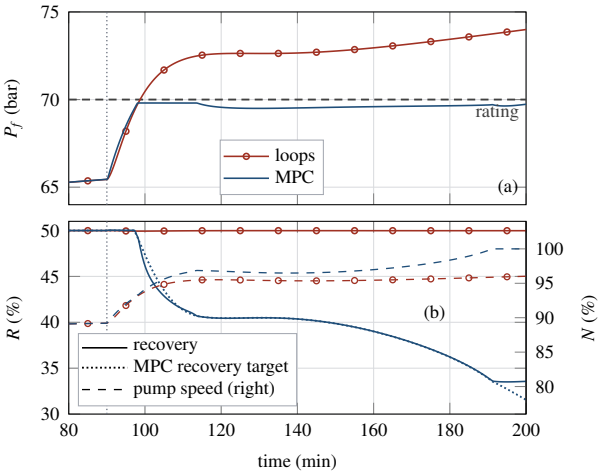


Figure 2.4 Scenario 11 from 80 to 200 min with the MPC and with the decoupled loops. Curves with open circles are the loops and curves without markers the MPC. (a) Feed pressure (solid) against the 70 bar element rating (dashed). (b) Recovery (solid), the recovery target the MPC holds (dotted) and the pump speed command on the right axis (dashed). The demand steps from 450 to 540 m^3/h at 90 min. At 200 min the loops are at 74.00 bar and the MPC at 69.73 bar with 33.6 % recovery and the pump at 100 %.

2.7 Technology in Practice

What Measures It

Recovery is not read by one instrument. It is the permeate flow divided by the feed flow [11], so the recovery loop rests on two flow measurements. A magnetic flowmeter suits such lines: a conductive liquid crossing a magnetic field induces a voltage across the pipe that is linear in the volumetric flow, provided the liquid conducts well enough, the pipe runs full and the flowtube is grounded [25]. Electrodes are set across the pipe horizontally so that a passing gas bubble cannot break their contact with the liquid, and an ungrounded flowtube reads stray currents as error [25]. The chapter gives its flow, recovery and level transmitters a 0.05 min filter and no other error, and the feed pressure transmitter is what the predictive controller's bias compares with its model.

> **TECHNOLOGY** **Sensor: The Permeate Conductivity Cell**
>
> Permeate quality is read as electrical conductivity, an approximate measure of the salt in the water, because dissolved salts carry current and pure water hardly does. The reading is taken in a fraction of a minute and reported at 25 °C, and it is turned into dissolved solids by a relation derived for each water, since conductivity changes with temperature and with the mix of ions [11]. A cell of two electrodes fouls as ions bond to the metal and then reads the water as less conductive than it is, so the permeate looks fresher on the trend than it is. A four-electrode cell is immune to that error, and an electrodeless toroidal cell resists fouling [25]. The chapter's cell adds a 0.45 min transport delay and a 0.35 min response.

What Moves It

The production loop moves the pump through its drive, and the recovery loop moves the concentrate valve. An equal-percentage trim is decidedly nonlinear in its flow capacity, and its formula can never reach zero, so it cannot be trusted at small stem positions [25]. The chapter's valve has a rangeability of 30 and a 0.15 min travel time, and its lower clamp of 5 % is the stop that the reversed pairing of scenario 5 reaches.

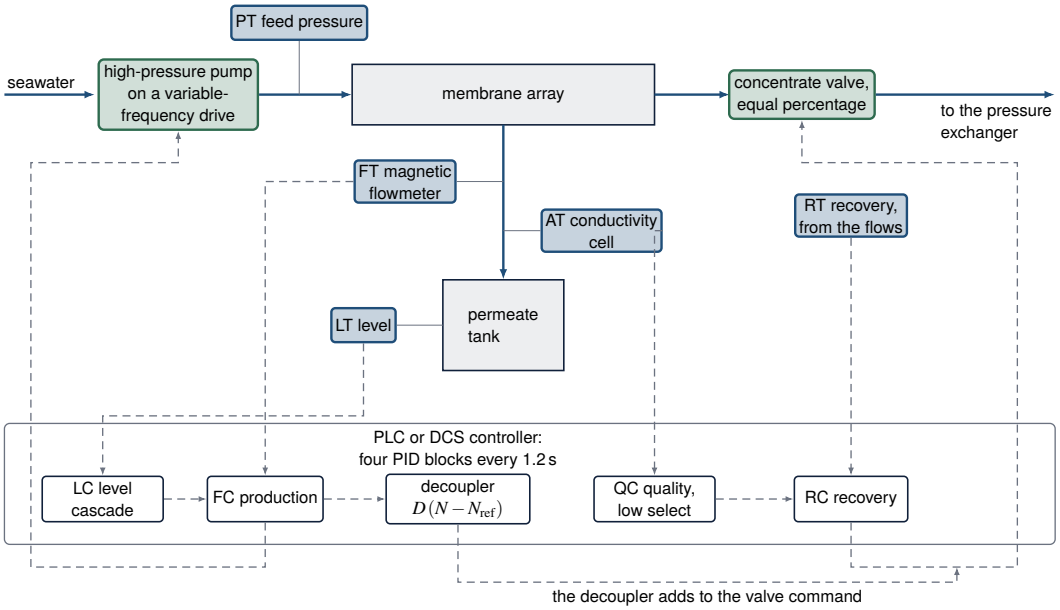


Figure 2.5 The two regulatory loops on plant hardware. The magnetic flowmeter on the permeate line feeds the production controller, which drives the pump through its variable-frequency drive. The recovery, taken from the flows, feeds the recovery controller, which moves the equal-percentage concentrate valve. The level cascade writes the production setpoint, and the conductivity cell's quality controller can only lower the recovery setpoint. Notice the decoupler path: the production controller's own output reaches the valve before the recovery controller sees an error.

> **TECHNOLOGY** **Final Element: The Pump Drive**
>
> The high-pressure pump is a centrifugal machine, and its discharge pressure falls as its flow rises. A variable-frequency drive sets its speed: a rectifier turns line power into direct current, a filter smooths it, and an inverter rebuilds alternating current at the frequency asked for. An induction motor turns slightly slower than the field that frequency sets. Running only as fast as needed saves power, cuts vibration and lets the drive ramp the speed to spare the machine [25]. Such drives were already commonplace on brackish-water trains when a federal handbook for planners was written [11]. In the chapter the drive is a 0.10 min lag on the speed command, and the pump curve droops with flow as Equation (2.5) gives.

Where the Controller Runs

Distributed control systems serve water and wastewater treatment, with programmable logic controllers as local controllers under a supervisory scheme, and supervisory control and data acquisition systems monitor and control water distribution. Loops at both levels run continuously with cycle times from milliseconds to minutes, and some are nested so that one loop's setpoint comes from another loop's process variable [26]. In a research roadmap for water treatment these systems are described as the result of piecemeal development over decades [27]. The chapter's structure maps onto standard function blocks: four PID blocks scanned every 1.2 s, a low selector with external reset on the recovery setpoint, a cascade from the tank level onto the flow setpoint, and one multiplication and one addition for the decoupler. The optional predictive controller plans every 0.1 min, a supervisory period. Figure 2.5 draws the two regulatory loops on their devices.

IN PRACTICE **Recovery Comes First**

A federal handbook for planners calls recovery the most important process variable of a reverse osmosis train, the one the others largely depend on, because it sets the concentration factor $1/(1-R)$. Most trains dose a scale inhibitor, sometimes with acid, because at least one salt is likely to be supersaturated in the concentrate. The handbook lists the ASTM practices for membrane systems, among them D4516 for standardizing performance data, D4472 for record keeping and D3739 for the Langelier saturation index [11]. A later research roadmap asks for inline sensors that would trigger cleaning instead of delaying it until a reduction of water permeability is observed [7].

What Is Modernizing It

Varying the speed of the machine that drives a liquid is more energy-efficient than running it at full speed and throttling its discharge, and it removes a control valve from the line [25]. Energy recovery had become fairly commonplace on seawater trains when a federal handbook put the recovered energy at 25 to 35 % of the input energy and gave the pressure exchanger, then used on smaller systems, an efficiency above 90 % [11]. The chapter's exchanger is a constant efficiency of 0.96. A national research roadmap asks for model-based control and data-driven models that optimize set-points for energy, quality and recovery, and for control systems secured against manipulation by unauthorized users [27]. The chapter's predictive controller, which lowers the recovery target to keep the membrane under its rating, is one small instance of that program.

2.8 Limitations

The lumped array is the largest simplification. A real train is six or seven elements in series in each vessel, the lead element produces two to three times the flux of the tail, and it fouls and fails first. A single well-mixed channel with a log-mean bulk salinity reproduces the concentration factor and the average driving force and cannot show that distribution, so nothing here speaks to element-by-element flux balancing or to the staging decisions that follow from it.

Concentration polarization uses a fixed mass-transfer coefficient. In practice that coefficient depends on the crossflow velocity, which falls along the array as permeate is removed, so polarization is worst at the tail where the salinity is already highest. The application therefore understates the gradient in the place where scale actually starts.

Fouling is one scalar with one rate law. Real fouling is biological, colloidal, organic and inorganic at once, with different kinetics, different symptoms and different cleaning chemistries, and identifying which mechanism is active is a substantial part of the operating skill. The fouling clock here is also compressed by roughly one hundred and seventy times so that six weeks fits in six hours, and the oxidation rate is compressed similarly, so the exposure in scenario 10 is far smaller than the exposure that would do that much damage to a real polyamide film.

OPEN CHALLENGE **No Sensor Yet Reports Fouling Early**

A Department of Energy research roadmap calls membrane fouling the weak point of membrane technology. It lowers productivity, degrades water quality and raises the energy and chemical cost of cleaning, and frequent cleaning shortens membrane life. Practical sensors for fouling have not been adopted commercially by the desalination industry. The roadmap asks for inline sensors that measure several scale-forming solutes at once, and others that report the organic carbon and biomass that feed biofilms, so that pretreatment can change before severe fouling [7]. In this chapter the scaling margin rests on a dose and a feed chemistry that no instrument measures online. Progress would be a measured scaling margin that the recovery loop could defend.

Pre-treatment, intake and post-treatment are boundary conditions. There is no cartridge filter differential pressure, no coagulant dose, no silt density index and no remineralization, so the application cannot show the most common industrial cause of a fouling event, which is a failure in the

pre-treatment train. Salinity is a single equivalent salt, which removes boron, the ion seawater plants most often struggle to reject and the one that decides whether a second pass is needed. The pressure exchanger is a constant efficiency rather than a device with its own mixing, leakage and speed. The high-pressure trip is an alarm and the train keeps running, where a real interlock stops the pump and the lost production usually exceeds the excursion that caused it.

OPEN CHALLENGE **Control That Sees the Whole Train**

A national water research roadmap notes that most treatment systems still need operators to adapt to changing water quality and to worn components, and that the data of their many sensors serve mainly to detect failures. It asks for control networks coupled with models of the interdependencies between unit processes, and for supervisory systems that reach the data of pretreatment and of grab samples, which existing ones often cannot. It adds that the savings of predictive control are hard to quantify without grounding in actual facilities [27]. This chapter's train has no pretreatment, so its supervisory layer sees only the membrane. Progress would be a plantwide structure that carries intake and pretreatment measurements forward to the recovery and production targets.

The predictive controller's model is the plant's own structure at the design feed. Fouling, scale, temperature, feed salinity and pump wear reach it only through two permeability multipliers and an output bias, the feed pressure is taken as settled, and the membrane is held as it is now over the whole 9 min horizon. Each soft limit has one slack for the whole horizon, so the target adjustment rather than the slack keeps the plan off the rating, and that adjustment ignores the pump's top speed, which leaves the recovery at 33.6 % against a target of 31.6 % at the end of the peak in scenario 11. The 0.2 bar band is sized for quiet instruments, and the plan carries neither the saturation index nor a price on energy. Switching from the loops to the MPC at 150 min of scenario 11, with the pressure 2.77 bar over the rating, opens the valve by 21.95 points in one decision and costs up to 63.6 m^3/h of production for 4.56 min. The same switch at 110 min, with the feed pressure at 72.28 bar, opens the valve from 48.12 to 67.53 % in the first plan and brings the pressure under the rating within 0.40 min, at the cost of up to 57.4 m^3/h of production for 3.19 min. The plan serves the limit before either target.

Pause. Try Exercise 2.3 before reading on.

2.9 Discussion

Three results generalize past this plant.

The first is that a relative gain of one is a statement about a product of two numbers. It cannot distinguish a diagonal plant from a triangular one, and triangular is the common case whenever one actuator is much more influential than the other. Chapter 10 has another triangular plant, in which added oxygen raises the total pressure and the nitrogen does not move the oxygen. A large pump against a trim valve, a main heater against a side stream and a feed rate against a reflux ratio all produce nearly triangular gain matrices and a pairing tool that reports no interaction. The remedy is not to abandon the array but to read the gain matrix that produced it, which the single summary number makes easy to skip [2, 18].

The second is that one-way coupling admits a one-way remedy at low cost. Equation (2.15) is one multiplication and one addition per scan, it requires no model of the dynamics, and Table 2.3 shows it removing most of an interaction that survives good tuning. It does require a calibration, and the calibration ages. That is a smaller liability than a full decoupling matrix, which is why triangular plants are worth identifying as such rather than escalating to a multivariable controller [12].

The third result concerns constraints and optima. Table 2.5 has a genuine interior energy minimum and a scaling limit 4.3 points above it, and the distance between them is set by a chemical dose rather than by anything the control system measures. Moving recovery toward the constraint buys a fraction of a percent of energy and risks a permeability loss that no rinse recovers, a trade that is

clear when it is written down and less clear during operation, when the specific energy is displayed continuously and the saturation index is not. Volume 9, Chapter 8 has the same arrangement, a lowest-cost operating point that sits on a constraint that moves with the weather. Desalination supplies water to a large fraction of the world’s coastal population and consumes electricity in proportion [1], so the distance between an optimum and a limit is worth stating precisely [5].

2.10 Conclusion

This chapter studied a seawater reverse osmosis train, a two-by-two plant whose pairing tool reports no interaction and whose loops interact anyway. The concentrate valve has a gain of −0.001 m^3/h/% on permeate flow at the design point and gains of opposite sign 8 % to either side, so the reversed pairing drives the valve to its stop and leaves the plant at 82.5 m^3/h against a 500 m^3/h setpoint. The correct pairing has a relative gain of 1.000 05 and still deviates 0.937 points in recovery on a production step, because the gain matrix is triangular rather than diagonal and the array is structurally unable to report that. One static feedforward from the pump command to the valve, calibrated from a gain ratio of 0.914, reduces the peak to 0.160 points. Behind the pairing question sit a shallow energy optimum at 49 % recovery, a crystallization limit 4.3 points above it, a fouling drift that adds 12.0 % to the specific energy in six hours without moving a single controlled variable, and an oxidation event that permanently consumes just under a third of the product quality margin while the permeate flow reads exactly on setpoint. The optional model predictive controller holds the feed pressure of scenario 11 at 69.81 bar where the decoupled loops spend 108.6 min above the 70 bar rating, at 3.0 % more specific energy, and a classical pressure override on the recovery target also holds it, at 69.87 bar, so the rating is held by giving up recovery at the limit rather than by the optimization itself. The app lets a reader measure loop pairing, decoupling and constraint-aware operation on one plant.

KEY POINTS

- Which handle holds which target is decided before any tuning: production belongs on the pump and recovery on the concentrate valve.
- The valve’s gain on permeate flow passes through zero at the design point, so the reversed pairing drives the valve to its stop and leaves the plant at 82.5 m^3/h against a 500 m^3/h setpoint.
- The correct pairing has a relative gain of 1.000 05 and still deviates 0.937 points in recovery, because the gain matrix is triangular rather than diagonal.
- One static feedforward from the pump command to the valve reduces the peak to 0.160 points.

Where else this applies. Chapter 10 has another triangular plant, and Chapter 3 a plant coupled both ways, with a decoupler of both terms.

References

[1] M. Elimelech and W. A. Phillip, “The future of seawater desalination: Energy, technology, and the environment,” *Science*, vol. 333, no. 6043, pp. 712–717, 2011.

[2] D. E. Seborg, T. F. Edgar, D. A. Mellichamp, and F. J. Doyle, III, *Process Dynamics and Control*, 4th ed. Hoboken, NJ: John Wiley & Sons, 2016.

[3] B. W. Bequette, *Process Control: Modeling, Design, and Simulation*. Upper Saddle River, NJ: Prentice Hall, 2003.

[4] J. G. Wijmans and R. W. Baker, “The solution-diffusion model: a review,” *Journal of Membrane Science*, vol. 107, no. 1–2, pp. 1–21, 1995.

[5] K. J. Åström and R. M. Murray, *Feedback Systems: An Introduction for Scientists and Engineers*, 2nd ed. Princeton, NJ: Princeton University Press, 2021.

[6] N. Voutchkov, “Energy use for membrane seawater desalination: current status and trends,” *Desalination*, vol. 431, pp. 2–14, 2018.

[7] D. Giammar, S. Jiang, P. Xu, R. Breckenridge, T. Edirisooriya, W. Jiang, L. Lin, J. Macknick, N. Rao, D. Sedlak, J. Stokes-Draut, and X. Xu, “National alliance for water innovation (NAWI) municipal sector technology roadmap 2021,” U.S. Department of Energy, Report DOE/GO-102021-5565, 2021.

[8] B. W. Bequette, "Process control practice and education: Past, present and future," *Computers & Chemical Engineering*, vol. 128, pp. 538–556, 2019.

[9] J. A. Rossiter, B. Pasik-Duncan, S. Dormido, L. Vlacic, B. Jones, and R. Murray, "A survey of good practice in control education," *European Journal of Engineering Education*, vol. 43, no. 6, pp. 801–823, 2018.

[10] G. Migliorini and E. Luzzo, "Seawater reverse osmosis plant using the pressure exchanger for energy recovery: a calculation model," *Desalination*, vol. 165, pp. 289–298, 2004.

[11] I. C. Watson, O. J. Morin, Jr., and L. Henthorne, "Desalting handbook for planners," U.S. Department of the Interior, Bureau of Reclamation, Denver, CO, Desalination and Water Purification Research and Development Program Report 72, 2003. [Online]. Available: https://www.usbr.gov/research/dwpr/reportpdfs/report072.pdf

[12] I. M. Alatiqi, A. H. Ghabris, and S. Ebrahim, "System identification and control of reverse osmosis desalination," *Desalination*, vol. 75, pp. 119–140, 1989.

[13] S. Sobana and R. C. Panda, "Review on modelling and control of desalination system using reverse osmosis," *Reviews in Environmental Science and Bio/Technology*, vol. 10, no. 2, pp. 139–150, 2011.

[14] A. R. Bartman, P. D. Christofides, and Y. Cohen, "Nonlinear model-based control of an experimental reverse-osmosis water desalination system," *Industrial & Engineering Chemistry Research*, vol. 48, no. 13, pp. 6126–6136, 2009.

[15] A. Zhu, P. D. Christofides, and Y. Cohen, "Effect of thermodynamic restriction on energy cost optimization of RO membrane water desalination," *Industrial & Engineering Chemistry Research*, vol. 48, no. 13, pp. 6010–6021, 2008.

[16] A. Ruiz-García and I. Nuez, "Long-term performance decline in a brackish water reverse osmosis desalination plant. predictive model for the water permeability coefficient," *Desalination*, vol. 397, pp. 101–107, 2016.

[17] A. Antony, J. H. Low, S. Gray, A. E. Childress, P. Le-Clech, and G. Leslie, "Scale formation and control in high pressure membrane water treatment systems: A review," *Journal of Membrane Science*, vol. 383, no. 1–2, pp. 1–16, 2011.

[18] E. Bristol, "On a new measure of interaction for multivariable process control," *IEEE Transactions on Automatic Control*, vol. 11, no. 1, pp. 133–134, 1966.

[19] K. J. Åström and T. Hägglund, *Advanced PID Control*. Research Triangle Park, NC: ISA – The Instrumentation, Systems, and Automation Society, 2006.

[20] S. Skogestad, "Simple analytic rules for model reduction and PID controller tuning," *Journal of Process Control*, vol. 13, no. 4, pp. 291–309, 2003.

[21] 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.

[22] A. Abbas, "Model predictive control of a reverse osmosis desalination unit," *Desalination*, vol. 194, no. 1–3, pp. 268–280, 2006.

[23] A. R. Bartman, A. Zhu, P. D. Christofides, and Y. Cohen, "Minimizing energy consumption in reverse osmosis membrane desalination using optimization-based control," *Journal of Process Control*, vol. 20, no. 10, pp. 1261–1269, 2010.

[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] T. R. Kuphaldt, *Lessons in Industrial Instrumentation*, version 2.33 ed. Open textbook, Creative Commons Attribution 4.0 license, ibiblio.org/kuphaldt/socratic/sinst, 2024.

[26] K. Stouffer, M. Pease, C. Tang, T. Zimmerman, V. Pillitteri, S. Lightman, A. Hahn, S. Saravia, A. Sherule, and M. Thompson, "Guide to operational technology (OT) security," National Institute of Standards and Technology, Gaithersburg, MD, Special Publication 800-82r3, 2023.

[27] D. Sedlak, M. Mauter, J. Macknick, J. Stokes-Draut, P. Fiske, D. Agarwal, T. Borch, R. Breckenridge, T. Cath, S. Chellam, A. Childress, D. Dionysiou, D. Giammar, E. Hoek, S. Jiang, L. Katz, J. Kim, R. Kostecki, J. McCutcheon, Y. Polsky, Z. Stoll, and P. Xu, "National alliance for water innovation (NAWI) master technology roadmap," U.S. Department of Energy, Report DOE/GO-102021-5617, 2021.

2.11 Exercises

2.1 Concept check · 1 min

Which balance does Equation (2.8) express?

(a) The water and salt fluxes that cross the membrane wall
(b) The pressure balance on the high-pressure feed node
(c) The growth of fouling with the permeate flux
(d) The salt inventory of the feed and concentrate channel

2.2 Concept check · 2 min

In Table 2.4 the peak recovery deviation on the scenario 6 production step falls from 0.937 points with no decoupler to 0.160 at $D = 0.95$, and rises again to 0.864 at $D = 1.6$. Which explanation is right?

(a) At 1.6 the valve moves 1.75 times the cancelling move, so the error returns with the opposite sign
(b) At 1.6 the decoupler drives the valve onto its stop, so the recovery loop loses its actuator
(c) At 1.6 the decoupler closes a second feedback path through the recovery, which then oscillates
(d) At 1.6 the conductivity cell delay lets the recovery drift before the decoupler acts

2.3 Concept check · 2 min

A plant engineer wants to use the app before a site study on a seawater train. Which question can the app answer within the limitations stated in Section 2.8?

(a) How much boron the permeate carries, and whether a second pass is needed
(b) How far the recovery moves when production is stepped on the pump
(c) How the cartridge filter pressure drop rises before a fouling event
(d) How the flux of the lead element compares with that of the tail element

2.4 Theory and calculation · 7 min

Use the gain matrix of Equation (2.13), with $K_{11} = 17.147$, $K_{12} = -0.001$, $K_{21} = 0.7029$ and $K_{22} = -0.7687$.

(a) Compute the relative gain λ_{11} of the direct pairing, production on the pump and recovery on the valve. *Checkpoint:* 1.00005.

(b) Compute the Niederlinski index of the swapped pairing, production on the valve and recovery on the pump, and say what it predicts for loops with integral action that are each stable alone. *Checkpoint:* About -1.9×10^4. Unstable for any tuning.

(c) Find the steady gain from the valve to the permeate flow when the recovery loop on the pump is closed perfectly. *Checkpoint:* $+18.75$ m^3/h/%.

2.5 Theory and calculation · 6 min

Use the design-point gains of Equation (2.13).

(a) Find the decoupler gain D that cancels the steady effect of the pump on the recovery, and the steady gain from the pump to the recovery that the shipped value $D = 0.95$ leaves. *Checkpoint:* 0.914 and -0.027 %/%.

(b) In scenario 3 the recovery loop is open with the valve at 50 %, and production steps from 450 to 510 m^3/h. Estimate the pump move and the change in recovery. *Checkpoint:* 3.50 % of speed and 2.46 points.

(c) In scenario 4 both loops are closed and production steps from 450 to 520 m^3/h. Estimate the steady valve move the recovery loop must make, and compare it with the valve travel reported in Section 2.6. *Checkpoint:* 3.73 %, against 50.00 to 53.44 %.

2.6 Theory and calculation · 8 min

At the design point the feed is 35.0 g/L at 25 °C, the recovery is 45.00 %, the flux is 13.636 LMH and the permeate is 250 mg/L. Use Table 2.1 and Equations (2.1) to (2.4).

(a) Find the concentrate salinity from the steady salt balance of Equation (2.8), and the log-mean bulk salinity. *Checkpoint:* 63.43 g/L and 47.82 g/L.

(b) Find the polarization factor, the wall salinity, the wall osmotic pressure and the net driving pressure. *Checkpoint:* 1.100, 52.60 g/L, 40.55 bar and 19.70 bar.

(c) Confirm the permeate salinity and the rejection. Then raise B by 30 %, as the oxidant of scenario 10 does, at the same flux and wall salinity, and find the new permeate salinity. *Checkpoint:* 250 mg/L, 99.286 %, then about 325 mg/L.

2.7 App challenge · 8 min

Use the default seed and the default speed. Read the recovery from its trend or from the `rec_pct` column of the exported file.

(a) Load scenario 4, Both loops, and a relative gain of one, and run to 60 min. Production steps from 450 to 520 m^3/h at 45 min. Read the highest recovery after the step and the concentrate valve position (`valve_pct`) at 60 min. *Checkpoint:* 45.94 % and 53.43 %.

(b) Load scenario 6, A one-way decoupler, and run to 25 min. The same step comes at 10 min. Read the lowest and the highest recovery after the step. *Checkpoint:* 44.84 and 45.00 %.

(c) Set the Pump-to-valve decoupling gain to 1.6, press Reset and run to 25 min again. Read the lowest recovery after the step. *Checkpoint:* 44.14 %.

Accept recoveries within 0.01 point and a valve position within 0.05 %. Compare each result with Tables 2.3 and 2.4.

2.8 App challenge · 6 min

Load scenario 5, The pairing you would have guessed, and run to 120 min at the default seed and a speed of 20. Both loops go to Auto at 5 min and the production setpoint steps from 450 to 500 m^3/h at 10 min.

(a) Read the highest recovery and the time at which it occurs. *Checkpoint:* 76.3 % at about 13.5 min.

(b) Read the pump speed, the permeate flow and the permeate salinity at 120 min from the exported file (`pump_pct`, `flow_m3h`, `tds_mgL`). *Checkpoint:* 68.39 %, 82.5 m^3/h and 1241 mg/L.

Accept the recovery within 0.1 point, the pump within 0.05 %, the flow within 0.5 m^3/h and the salinity within 2 mg/L. Explain the end state with the swapped pairing's relative gain.

2.9 App challenge · 8 min

Load scenario 11, Constraints, not a pairing: MPC against the decoupled loops. Use the default seed and a speed of 20, and run each case for 300 min. From each exported file read the highest feed pressure (`feed_bar`), the minutes above the rating (`metric_overpressure_min`) and the delivered water (`metric_water_m3`) in the last row, and compute the specific energy over the run as `metric_energy_kwh` divided by `metric_water_m3`.

(a) Run the scenario as shipped, with Multivariable control on MPC. *Checkpoint:* 69.81 bar, 0 min, 2415 m^3, 2.757 kWh/m^3.

(b) Set Multivariable control to Loops, press Reset and run again. *Checkpoint:* 74.00 bar, 108.6 min, 2415 m^3, 2.675 kWh/m^3.

Accept pressures within 0.02 bar, minutes within 0.2 and specific energies within 0.002 kWh/m^3. Say what holds the rating, and what it costs.

2.10 Application area · 4 min

At the design point the train makes 450 m^3/h of permeate from 1000 m^3/h of feed at 62.00 bar, with an array drop of 2.50 bar. Take the other constants from Table 2.1.

(a) Find the three terms of Equation (2.9) and the specific energy. *Checkpoint:* 907.0, 114.4 and 86.8 kW, and 2.463 kWh/m^3.

(b) Find the specific energy without the pressure exchanger ($\eta_{PX} = 0$), and the electricity the exchanger saves in a year of continuous operation at this duty, and its value at the app's default electricity price of $0.085 per kWh. *Checkpoint:* 4.907 kWh/m^3, about 9.6 GW h and $0.82 million.

2.11 Application area · 4 min

In the graded challenge the feed salinity steps to 36.5 g/L at 130 min and the scaling tendency h rises from 1.00 to 1.10 at 230 min. The saturation index is C_b/C_{sat}, the ratio in Equation (2.11), with C_{sat} as given in Section 2.3.

(a) Find C_{sat} after the hardness step at an antiscalant dose of 5 mg/L and at 2 mg/L. *Checkpoint:* 70.24 g/L and 65.60 g/L.

(b) Estimate the concentrate salinity at 46 % recovery, neglecting the salt in the permeate, and the saturation index at each dose. Compare with Table 2.7. *Checkpoint:* 67.6 g/L, 0.962 and 1.030.

(c) Find the highest recovery that keeps the index below one at 5 mg/L after the hardness step. *Checkpoint:* About 48.0 %.

2.12 Review the agent · 5 min

An AI assistant is asked to simplify the controls of the train and replies as follows. "The relative gain of the direct pairing is 1.00005, so the production and recovery loops do not interact. The static decoupler adds nothing and can be removed. To save valve wear, slow the recovery loop to an integral time of 8 min. A production step will then move the pump and leave the recovery alone." Find the flaw in the proposal and say how it could be checked.