INTERACTIVE CONTROL ENGINEERING
Principles and Applications

Volume 13 Control in Chemical Plants

CHAPTER 11

Real-Time Optimization of a Williams–Otto Reactor Whose Model Lacks a Reaction

Pages 255–281

Edited by John D. Hedengren

How to cite this chapter. "Real-Time Optimization of a Williams–Otto Reactor Whose Model Lacks a Reaction," in *Control in Chemical Plants*, John D. Hedengren, Ed., vol. 13 of *Interactive Control Engineering: Principles and Applications*. Advanced Process Solutions, LLC, 2026, ch. 11, pp. 255–281.

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

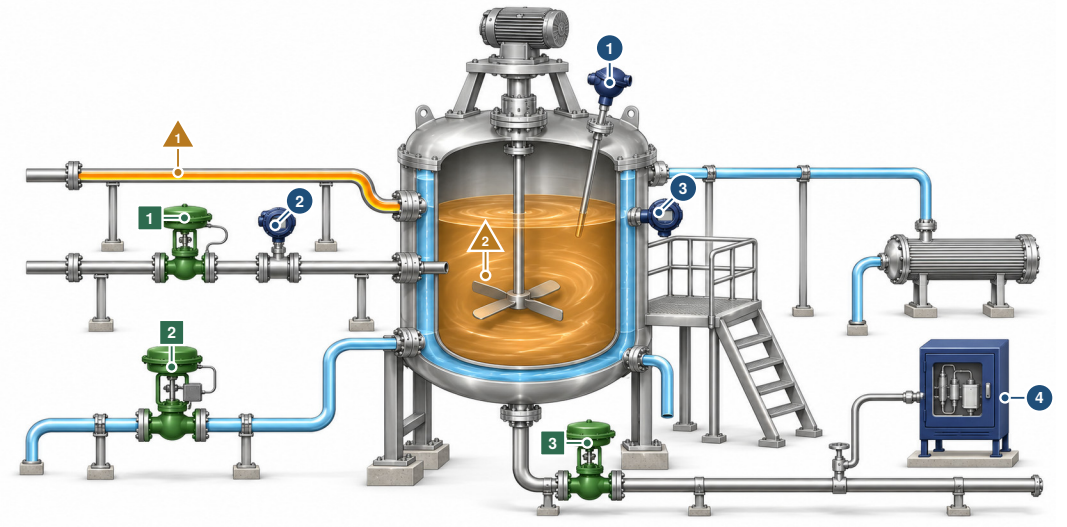
The app for this chapter
https://apmonitor.com/apps/williams-otto-rto/

The volume online: contents, other chapters and citation records
https://apmonitor.com/apps/books/v13-chemical-plants/#williams-otto-rto

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

CHAPTER 11

Real-Time Optimization of a Williams–Otto Reactor Whose Model Lacks a Reaction



The controller moves 1 B feed valve, whose flow setpoint the optimizer chooses
2 jacket water supply, which sets the reactor temperature the optimizer chooses 3 outlet valve, which holds the level

The controller reads 1 reactor temperature 2 B feed flow 3 level
4 analyzer, the optimizer's only view of what the reactor makes

What disturbs it 1 a change in the A feed, which a meter reports
2 a waste reaction that slows, which the optimizer is not told of

Notice the analyzer cabinet at the right: the optimizer above the three loops sees the product only through its reports, minutes old and noisy. Illustration, not to scale.

IN PLAIN TERMS An Optimizer That Stops Short

The main idea. A computer above the control loops of a chemical reactor picks the temperature and feed rate that should earn the most money. It decides with a simple model of the reactor, which it adjusts after each new set of readings.

The challenge. The model leaves out one of the three reactions in the tank, so it matches what the plant makes but not how it changes when a setting moves. Finding out takes small test moves, a wait for the plant to settle, and noisy readings.

Why it matters. A computer that stops in the wrong place loses money every hour while its own model says all is well.

THREE WAYS TO READ THIS CHAPTER

Explore	11 min	The first page, the history box in Section 11.2, the try-it and fact boxes in Section 11.6, and the key points
Study	60 min	The whole chapter with its pause points and the twelve exercises
Practice	20 min	The first page, Section 11.6, Section 11.7, and Section 11.8 with its open challenges

The control question A real-time optimizer moves the setpoints of a reactor's loops to the best steady state its model can find, and its model lacks one of the plant's three reactions. Refitted at every steady state, the model agrees with the plant where it stands but not on the slopes. Where does such an optimizer stop, and what does it cost to correct the slopes from plant data?

Learning objectives

- Explain why a two-step optimizer with a structurally wrong model stops where the fitted model's slopes vanish, not at the plant's optimum.
- Write the modifier-adaptation corrections and estimate a plant gradient from three steady states, with the noise such an estimate carries.
- Set the period, steady-state check, filter gain and back-off of an optimization layer, and judge the profit each costs or saves.
- Compare an economic MPC with modifier adaptation after a change the layer is told of and after one it can only infer.

Prerequisites Sections 1.4.7 (two-step optimization and modifier adaptation), 1.6.2 (the loops beneath the layer), 1.6.4 (scenarios, scorecard and data) and 1.6.3 (needed only for the economic MPC).

The app https://apmonitor.com/apps/williams-otto-rto/

See also Volume 9, Chapter 8 for an optimum that sits on a moving constraint and Volume 9, Chapter 10 for where a supervisory setpoint belongs. Volume 1, Chapter 11 for a model that fits a record without the right structure, and Volume 1, Chapter 12 for fitting models from plant data. Volume 7, Chapter 2 for a jacketed reactor whose heat release can outrun its cooling.



11.1 Introduction

Feed B to a stirred tank at 4.00 kg/s with 1.8275 kg/s of A, hold it at 75 °C, and the simulated Williams–Otto reactor earns 151.12 $/s, 84.64 % of the 178.540 $/s its best permitted point makes at 4.9747 kg/s and 84.322 °C, where the waste sits on its 8 % limit. Hand the two setpoints to a real-time optimizer whose model is refitted at every steady state, and after six hourly runs it settles at 4.4454 kg/s and 78.915 °C, earning 165.736 $/s, or 92.83 % of the best. The model fitted there predicts 180.92 $/s and a profit slope of zero in both directions, while the plant's own slopes there are –24.44 $/s per kg/s of B and +5.76 $/s per degree.

That run corrects a belief that most engineers carry into optimization, that a model good enough to reproduce the plant's measurements is good enough to optimize it. The model of this optimizer has two adjustable constants, it is refitted to five measured fractions at every run, and it agrees with the plant on the waste to within half a point at its stopping point. What it does not reproduce is how the profit and the waste change when the setpoints move, and an optimizer moves along those changes. The class is general. Every steady-state optimizer above the control system of a refinery unit, an ethylene plant or a utility runs on a model that is wrong somewhere, and a model updated only in its parameters or its output biases can converge to a point that is not the plant's optimum [1–3]. Correcting the model's gradients with plant gradients removes that error at convergence [3, 4], and estimating plant gradients from noisy steady states is where the remaining difficulty lies [3].

However, the benchmark on which these methods are usually compared is a static model, on which a steady state is available at a call, while a plant must be waited for, its steady states must be detected from noisy data [5], and its analyzer reports what it sampled minutes earlier. This chapter places the reactor-only Williams–Otto benchmark under regulatory loops in the app, with an optimization layer that reads only the measurements, and measures what the two-step method and modifier adaptation do there. The plant equations, the two-reaction model and both methods are published [4, 6, 7], and Section 1.4.7 introduces both methods on the running example of this volume. The findings are measurements. The two-step fixed point is reached from 21 starting points and is set by the model's structure. The perturbation rule a running modifier adaptation needs costs about three percent of profit on every third run. The settings first suggested for the graded challenge pass at only 21 of 31

noise seeds, while a lower filter gain with a waste back-off of 0.15 point passes at all 31. The plant is dynamic, its steady-state map is exactly the published one, checked against an independent solver to every printed digit, and a measurement boundary keeps the optimizer away from the simulator's truth. The application also offers an economic model predictive controller (MPC) as a fourth method of the layer, on the optimizer's own model with the reactor's holdup and the same modifiers. The two tie on profit, and the economic controller's one gain, a lower waste peak after a metered feed rise, comes from its dynamic model. Ten cases run from a hand search to a graded 100 h challenge.

Section 11.2 follows the benchmark and the methods from 1960 to economic MPC. Section 11.3 sets out the plant: its three reactions, the energy balance and jacket the app adds, the profit and the measurements. Section 11.4 gives the optimizer's two-reaction model and shows where a model refitted at every steady state stops. Section 11.5 describes the loops, the steady-state detector, two-step, modifier adaptation with its safeguards, and the economic MPC (Section 11.5.6). Section 11.6 measures the ten cases, the economic MPC against the other methods, the worst cases and the challenge. Section 11.7 names the instruments, valves and computers behind the loops and the layer. Section 11.8 states the limitations and Section 11.9 the conclusions. The simulator framework the app runs on is described in Section 1.6.

11.2 From a Benchmark Plant to Modifier Adaptation

Williams and Otto proposed their plant in 1960 because the computer-control schemes then offered to the chemical and petroleum industries could not be compared directly, and the joint studies of computer makers and process companies were proprietary [6]. Their answer was a generalized chemical processing model on which such schemes could be compared in the open. The part that survived into the optimization literature is the reactor alone, a stirred tank fed with A and B in which a desired product P is made through an intermediate C and lost to a waste G, with a byproduct E that sells for little. It has become a common test case of real-time optimization, used for gray-box neural models [8], for evaluations of optimization strategies and online parameter adaptation [9, 10], and for modifier adaptation with Bayesian and derivative-free optimizers [7], in the form whose constants this simulator takes from those sources.

> ⧖ HISTORY **A Computer Above the Loops (1959)**
>
> In March 1959 a digital computer went online on the catalytic polymerization unit of a refinery in Texas, among the first to control a process plant in closed loop [11, 12]. Like the optimization layer of this chapter, it did not drive the valves. The first process control computers were supervisory: they adjusted the setpoints of the controllers below them, so the plant kept working when the computer was off [11, 12]. Product quality from process stream analyzers was its feedback, and every hour it adjusted 19 parameters. Conversion of propylene to gasoline rose from 80 to 91 percent, and the computer paid for itself in less than three years. By 1965 the United States had 400 such computers [11]. The benchmark of this chapter was proposed a year after that first unit came online.

Industrial practice arrived at the layer the benchmark tests from a different direction. The argument that a steady-state optimizer must sit above multivariable control to capture a plant's profit was made in 1983 [13], on-line optimization was described as a two-phase approach in an application study of 1987 [14], and identifying a steady state on line from noisy data became a subject of its own [5]. Industrial practice of the layer was reviewed and assessed in 2011 [15].

The weakness of that two-step structure had been pointed out before the industrial argument was published. In 1979 Roberts showed that a model which is not an accurate representation of the process can still lead to the optimum if the optimization includes a term comparing the model's output derivatives with the process's, a modified two-step approach in his words, and demonstrated the algorithm on the feed rate and temperature setpoints of a simulated exothermic reactor [1]. In 1994 the model accuracy an economic optimizer needs was examined for the case of a bias update [2], and adaptation strategies for real-time optimization were reviewed in 2009 [16]. Modifier adaptation,

Table 11.1 Variables by Role, With Values in Case 3 at 8 h

Variable	Role	Value, range
ϕ	objective, profit (\$/s)	165.75, best 178.54
x_G	constraint, waste by mass	7.19 %, limit 8
T_{SP}, $F_{B,SP}$	MVs of the RTO, decisions u	78.93, 4.447 in a box of 70–100, 3–7
T	CV of TIC (°C)	78.91
F_B	CV of FIC (kg/s)	4.438
L	CV of LIC (%)	50.0, SP 40–60
T_{ws}, v_B, v_o	MVs of the loops	66.9 °C, 44.5, 41.8 %
F_A	DV, A feed (kg/s)	1.8275, 1.5–2.4
a_3	DV, waste reaction activity	100 %, 60–140
p_P, p_E, p_A, p_B	DV, prices (\$/kg)	published
σ_{an}	DV, analyzer noise	0.2 %, 0–2
m_j	states, six component masses	$W = 2105$ kg
T, T_j, T_{ws}	states, temperatures	78.9, 73.7, 66.8 °C
v_B, v_o	states, valve positions	—
θ_1, θ_2	RTO model constants	−2.519, −1.797
ε, λ_ϕ, λ_G	modifiers (MA only)	zero under two-step

published the same year, adds to the model's cost and constraints the differences between the plant's and the model's values and gradients, so that at convergence the modified problem satisfies the plant's conditions of optimality [4], and it needs no parameter estimation [3]. The gradients are the price. Dual formulations of 2010 and 2011 [17, 18], which a 2016 thesis states as constraints keeping each new point off the line through the last two [9], keep the operating points spread enough to estimate them, and a 2016 review gathers the variants, the estimators of gradients from noisy measurements and the applications that followed [3].

The wait for steady state has drawn two answers of its own. Modifier adaptation was reformulated in 2014 to use transient measurements, and on a simulated stirred tank it then converged in about one settling time of the plant where the static scheme needed more than five iterations [19]. Economic MPC takes the other route and puts the profit itself into a receding-horizon controller on a dynamic model, in place of a setpoint handed down from a steady-state layer. Its average performance is never worse than the best steady operation, an analysis of 2012 shows under its assumptions [20], and a tutorial review of 2014 collects its formulations [21].

11.3 The Reactor as Modeled

The simulated plant is a jacketed stirred tank fed with pure A at F_A and pure B at F_B, as shown in Figure 11.1. A level loop (LIC) holds the holdup W through the effluent valve v_o, a flow loop (FIC) holds the B feed through its valve v_B, and a temperature loop (TIC) holds the reactor temperature T by the temperature T_{ws} of the tempered water sent to the jacket. An analyzer (AT) on the effluent reports the six mass fractions x_j every 5 min, 5 min after it drew the sample. Above the loops the real-time optimization (RTO) layer sets the two setpoints T_{SP} and $F_{B,SP}$. Table 11.1 classifies the variables, with their values in case 3 at 8 h, where two-step has settled.

11.3.1 Reactions and Component Balances

The plant runs the three reactions of the benchmark, A + B → C, B + C → P + E and C + P → G, at rates in kilograms per second given by Equation (11.1).

$$\begin{aligned} r_1 &= k_1 x_A x_B W, \quad r_2 = k_2 x_B x_C W, \quad r_3 = a_3 k_3 x_C x_P W \\ k_i &= A_i \exp\big(-E_i/(T+273.15)\big) \end{aligned} \tag{11.1}$$

where $x_j = m_j/W$ is the mass fraction of component j, $W = \sum_j m_j$ the holdup in kilograms, T in °C, and a_3 the activity of the waste reaction that the drift disturbance lowers. The constants, $A_i = 1.6599 \times 10^6$, 7.2117×10^8 and 2.6745×10^{12} s^{-1} with $E_i = 6666.7$, 8333.3 and 11 111 K, are

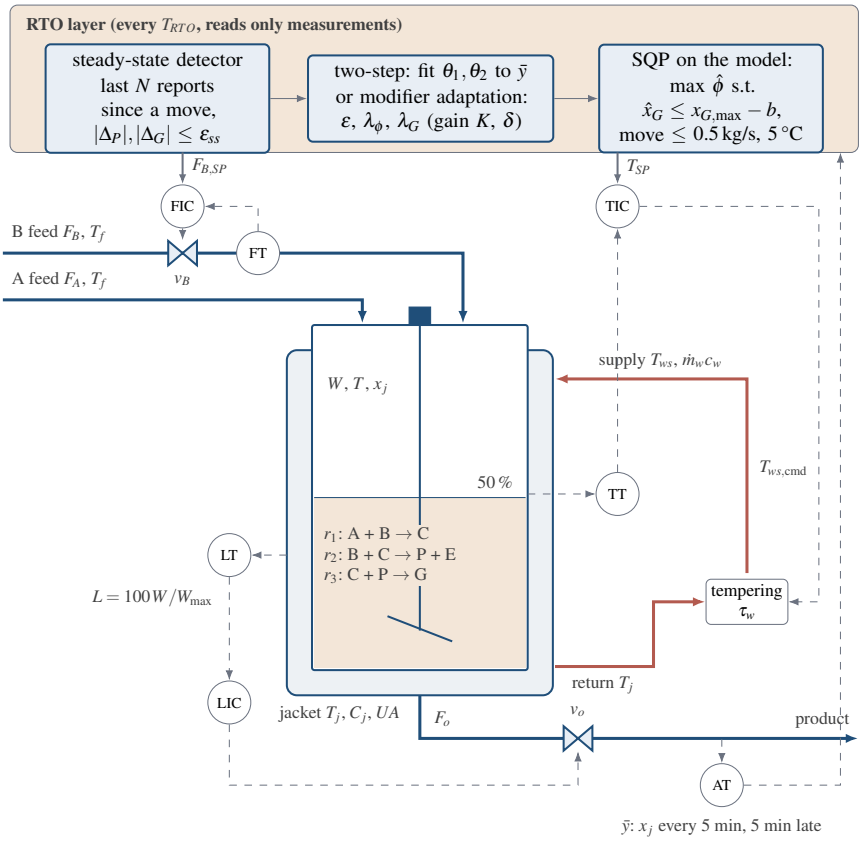


Figure 11.1 The Williams–Otto reactor under its three loops and the RTO layer, after the application's Reactor view of case 3 at 8 h with two-step settled. Feeds A and B at T_f enter the tank of holdup W at temperature T, where the three reactions run at r_1 to r_3. TIC moves the jacket supply T_{ws} through a tempering lag τ_w, the jacket of capacity C_j and conductance UA returns at T_j, FIC moves v_B and LIC moves v_o, and the analyzer reports the fractions x_j that the layer averages into $\bar{y}$. At that instant $T = 78.91\,^\circ$C (setpoint 78.93), $F_B = 4.438$ kg/s (setpoint 4.447), $T_{ws} = 66.8\,^\circ$C, $T_j = 73.7\,^\circ$C, $F_o = 6.260$ kg/s, the level is 50.0 %, the analyzer reads $x_G = 7.18$ %, and the plant earns 165.75 \$/s against the best 178.54 \$/s. Process model only, not to scale.

those of the reactor-only benchmark as the sources give them [8, 9]. The component masses follow Equation (11.2).

$$\frac{\mathrm{d}m_j}{\mathrm{d}t} = F_{j,in} + \sum_{i=1}^{3} \nu_{ij} r_i - F_o x_j$$

$$\nu = \begin{bmatrix} -1 & -1 & 2 & 0 & 0 & 0 \\ 0 & -1 & -2 & 2 & 1 & 0 \\ 0 & 0 & -1 & 0 & -\tfrac{1}{2} & \tfrac{3}{2} \end{bmatrix} \tag{11.2}$$

with columns in the order A, B, C, E, P, G, $F_{A,in} = F_A$, $F_{B,in} = F_B$ and the others zero. Each row of ν sums to zero, because the benchmark gives no molecular weights and the balances are written in mass, so mass is conserved reaction by reaction. Over 3 h with a 10 °C step and the A feed raised to 2.2 kg/s, the holdup differs from the integrated flows by 1.7×10^{-10} kg, and over 100 h by 2.9×10^{-8} kg.

11.3.2 Energy, Jacket and Valves

The benchmark is a steady-state problem with the temperature as a decision. The simulator turns the temperature into a controlled variable with an energy balance and a jacket, as shown in Equation (11.3).

$$\begin{aligned} W c_p \frac{\mathrm{d}T}{\mathrm{d}t} &= (F_A + F_B) c_p (T_f - T) \\ &\quad + \textstyle\sum_i (-\Delta H_i) r_i - UA(T - T_j) \\ C_j \frac{\mathrm{d}T_j}{\mathrm{d}t} &= \dot{m}_w c_w (T_{ws} - T_j) + UA(T - T_j) \\ \tau_w \frac{\mathrm{d}T_{ws}}{\mathrm{d}t} &= T_{ws,\mathrm{cmd}} - T_{ws}, \qquad \tau_v \frac{\mathrm{d}v}{\mathrm{d}t} = v_{\mathrm{cmd}} - v \end{aligned} \tag{11.3}$$

where the last equation holds for both valves, $F_B = F_{B,\max} v_B$ and $F_o = F_{o,\max} v_o$. Every constant in it is a teaching value, listed in Table 11.2. At the best point the reactions release 439.2 kW, heating the feeds from 80 °C absorbs 88.2 kW, and the jacket holds 79.93 °C with a supply of 74.08 °C, so the jacket cools across the operating box. The level loop returns the holdup to 2105.2 kg and the temperature loop puts T on its setpoint, so every steady state of the dynamic plant is the published one. The heats decide only how long the plant takes to get there. The balance has the form of the jacketed reactor of Volume 7, Chapter 2, where heat release can outrun the jacket and the reactor ignites. Here the jacket only sets the pace of each move.

11.3.3 Profit, Measurements and Numerics

The profit pays for P and E in the effluent and charges for both feeds, as shown in Equation (11.4).

$$\begin{aligned} \phi &= (p_P x_P + p_E x_E)(F_A + F_B) - p_A F_A - p_B F_B \\ x_G &\le x_{G,\max} \end{aligned} \tag{11.4}$$

in dollars per second, the benchmark's unit, with $p_P = 1143.38$, $p_E = 25.92$, $p_A = 76.23$ and $p_B = 114.34\,\$/\mathrm{kg}$ [9, 10]. The 8 % waste limit is a declared teaching value. The best permitted point earns 178.540 \$/s, about 643 000 \$/h. The analyzer draws a sample every 5 min and reports it one cycle later, each fraction multiplied by $1 + \sigma_{an} n$ with n a unit normal deviate, so a reading is always 5 to 10 min old. The thermowells have a lag of 0.3 min, the reactor reading 0.02 °C of noise, and the flow meters a lag of 0.1 min and 0.2 % relative noise. The noise comes from the framework's seeded generator (Section 1.6.1), and every execution draws the same count of deviates whatever the noise settings, so changing one noise level leaves the other channels' deviates as they were. The 13 states, the six masses, three temperatures, two valve positions and two flow ledgers that check the mass balance, advance by fourth-order Runge–Kutta steps of 6 s. The spectral radius of the Jacobian is 646 h^{-1} at 7 kg/s and 100 °C, a time constant of 5.6 s, and 279 to 281 h^{-1} at the best point and at three other corners of the box, so $\lambda \Delta t$ is at most 1.08 against the method's real-axis limit of 2.785. With the step halved, the temperature and waste 1.5 h after a 10 °C step differ by 0.0026 °C and 1.2×10^{-5}. Each case starts lined out at its setpoints, the steady state found by Newton's method on the six component balances.

11.3.4 Assumptions

The tank is perfectly mixed, the reactor-only benchmark carries no separator or recycle, and every balance is in mass. The steady state is assumed unique. The plant's Newton solve from one fixed start converged at all 5440 points of a grid over the box (B feed every 0.25 kg/s, temperature every 2 °C, the activity from 60 to 140 % and the A feed from 1.5 to 2.4 kg/s), which shows convergence but does not prove uniqueness, and a plant with several steady states would need a different initialization. The energy balance and the jacket are not part of the benchmark. They make each move take about half an hour, which is what an optimization layer above real loops has to wait for, and they cannot change where the plant settles. The waste limit, the prices' use as disturbances and the noise levels

Table 11.2 Constants. P: Published, Shown in the Model Dialog. A: Adjustable in the Parameters or Disturbance Panel. H: Hidden

Symbol	Meaning	Value	Mark
A_i, E_i	Arrhenius constants	see text	P
W_0	holdup at 50 % level	2105.2 kg	P
$W_{\max}$	holdup at 100 % level	4210.4 kg	H
F_A	A feed	1.8275 kg/s	P, A
$p_{P,E,A,B}$	prices ($/kg)	1143.38, 25.92, 76.23, 114.34	P, A
c_p	heat capacity of the mass	3.0 kJ/(kg K)	A
T_f	feed temperature	80 °C	A
$\Delta H_{1,2,3}$	heats of reaction	−150, −250, −100 kJ/kg	A
UA	jacket conductance	80 kW/K	A
C_j	jacket heat capacity	2000 kJ/K	A
$\dot{m}_w c_w$	jacket water flow capacity	60 kW/K	A
τ_w, τ_v	tempering, valve lags	2, 0.5 min	A
$F_{B,\max}$, $F_{o,\max}$	valve capacities	10, 15 kg/s	A
ψ_1, ψ_2, T_{ref}	model temperature terms	−17, −29, 383.15 K	P
θ^0	model constants at start	(−3, −4)	P
$x_{G,\max}$	waste limit	8 %	A
Δt	integration and control step	6 s	fixed

are teaching values. A real permit and its margin would come from the plant's operating license and its own data.

11.4 The Optimizer's Model and the Mismatch

11.4.1 Two Reactions Where the Plant Has Three

The RTO's model is the published two-reaction approximation, A + 2B → P + E and A + B + P → G, with no C, whose rate laws and reparametrized constants the application takes from the supporting code of a 2021 study [7], as shown in Equation (11.5).

$$\begin{aligned} \hat{r}_1 &= \hat{k}_1 x_A x_B^2 W, \qquad \hat{r}_2 = \hat{k}_2 x_A x_B x_P W \\ \hat{k}_i &= \exp(\theta_i)\exp\big(\psi_i\,(T_{ref}/(T+273.15)-1)\big) \end{aligned} \tag{11.5}$$

with mass coefficients A $-\hat{r}_1 - \hat{r}_2$, B $-2\hat{r}_1 - \hat{r}_2$, E $2\hat{r}_1$, P $\hat{r}_1 - \hat{r}_2$ and G $3\hat{r}_2$, each reaction conserving mass, and the same profit and limit. The model has five fractions, two constants to fit and fixed temperature dependences, so it can match the plant's compositions only approximately. Its predictions are written $\hat{x}_j(u,\theta)$, $\hat{\phi}(u,\theta)$ and $\hat{x}_G(u,\theta)$ with $u = (F_B, T)$. A model with the wrong structure can still fit a record well, as the fitted radiative term of Volume 1, Chapter 11 does in data generated without one.

11.4.2 The Published Point and the Plant Optimum

At 4.787 kg/s and 89.70 °C the plant gives the fractions 0.0875, 0.3896, 0.0153, 0.2906, 0.1095 and 0.1075 for A, B, C, E, P and G, within 0.00075 of the operating point a 2019 thesis tabulates [10], and earns 190.991 $/s. That table labels its fourth and fifth fractions as P and E, against its own state order and the balances, which make them E and P, and it lists the point as an initial condition and prints no profit. That the point is the unconstrained optimum is a computation. The application's SQP finds it at (4.7875 kg/s, 89.703 °C) with 10.75 % waste, and with the 8 % limit the optimum moves to (4.9747 kg/s, 84.322 °C), 178.540 $/s, with the waste on the limit. A solver written independently of the application (steady states by time integration polished by Newton, optima by Nelder–Mead and a boundary search on the limit) reproduces both points and both profits to every printed digit. The residence time $W/(F_A + F_B)$ is 5.16 min at the best point and 5.30 min at the published one.

11.4.3 Where a Refitted Model Stops

Section 1.4.7 shows on the running example S-1 why a refitted model can stop short of the plant's optimum. Two-step RTO from (4.0 kg/s, 75 °C) and $\theta^0 = (-3, -4)$, iterated on the static plant, converges to (4.4454 kg/s, 78.915 °C) with $\theta = (-2.5189, -1.7964)$. The plant earns 165.736 \$/s there, 92.83 % of its best, with 7.187 % waste, while the model, fitted there, predicts 7.645 % and 180.92 \$/s. The same point is reached from 20 other starts spread over the box, to within 4.9×10^{-8}, from $\theta^0 = (-1, -1)$, and by the independent solver from five starts. The model's optimum is interior, so the limit plays no part. The stop is where the fitted model's gradient vanishes, and the plant's gradient there is (−24.44 \$/s per kg/s, +5.763 \$/s per degree), which says that, locally, the plant would earn more with less B and a hotter reactor. Modifier adaptation with exact plant gradients and a filter gain of 0.5, from the same start, reads 84.641, 96.510, 98.648, 99.465 and 99.803 % of the best at its first iterations and 100.000 % from the tenth, on the waste limit, as the independent solver confirms.

Pause. Try Exercise 11.1 before reading on.

11.5 Loops and Optimization Layer as Implemented

11.5.1 Regulatory Loops

The three loops execute every 6 s on the velocity-form algorithm of Section 1.6.2, with no derivative action and with the tracking time of the back-calculation equal to each loop's integral time. TIC acts on $T_{SP} - T$ with $K_c = 2.5$ °C per °C and $\tau_I = 6$ min and commands T_{ws} between 20 and 140 °C. FIC uses 6 % of valve per kg/s and 0.5 min. Both are reverse acting, with action $a = +1$ in Equation (1.14). LIC uses 6 % per percent and 6 min and is direct acting, $a = -1$, so that a high level opens the effluent valve. From 75 °C, a 3 °C setpoint step overshoots by 0.47 °C and stays within 0.1 °C of the setpoint from 13.5 min after the step, and a 5 °C step overshoots by 0.78 °C and settles in 21.9 min. The composition follows with its 5 min residence and the analyzer reports 5 min late, so a move takes about half an hour to read. The header's Manual places all three loops on their stations and pauses the layer, and each transfer is bumpless, as Section 1.6.2 describes.

11.5.2 When the Layer Runs

The layer reads only the measurement packet: the analyzer reports, the meters, and the prices and limit it is configured with. It executes no sooner than the period T_{RTO} after its last setpoint move and, with the steady-state check on, only when the detector of Equation (11.6) passes, over the last N reports drawn since that move.

$$\begin{aligned} \Delta_j &= \hat{\beta}_j\,(t_N - t_1) \\ \text{steady} &\iff |\Delta_P| \le \varepsilon_{ss} \text{ and } |\Delta_G| \le \varepsilon_{ss} \end{aligned} \tag{11.6}$$

where $\hat{\beta}_j$ is the least-squares slope of fraction j against the sample times $t_1 \ldots t_N$, so Δ_j is how far the fitted line moves over the window. The defaults are $T_{RTO} = 1$ h, $N = 6$ and $\varepsilon_{ss} = 0.2$ percentage points. The steady state the run uses, $\bar{y}$, is the mean of the last N reports with the mean measured flows. Every move is bounded to 0.5 kg/s and 5 °C by a move box inside the decision box of 3 to 7 kg/s and 70 to 100 °C, for both methods.

11.5.3 Two-Step

Each two-step run fits the two constants by least squares on five fractions and then moves to the fitted model's best permitted point, as shown in Equation (11.7), the method of Equation (1.12) with the waste limit added.

$$\begin{aligned} \theta_k &= \arg\min_{\theta} \textstyle\sum_{j\in\{A,B,E,P,G\}} \left(\bar{y}_j - \hat{x}_j(u_k, \theta)\right)^2 \\ u_{k+1} &= \arg\max_{u} \hat{\phi}(u, \theta_k) \;\; \text{s.t. } \hat{x}_G(u, \theta_k) \le x_{G,\max} - b \end{aligned} \tag{11.7}$$

where b is the waste back-off, zero unless the learner sets it. The fit is Gauss–Newton from the last θ with a step capped at one unit and halved until the residual falls. The optimization is a sequential quadratic program (SQP) in two variables scaled so that 1 kg/s of B counts as 10 °C, with central-difference gradients, a central-difference Hessian of the Lagrangian shifted to be positive definite, an exact-penalty backtracking line search, and each quadratic subproblem handed to the shared solver of the MPC option (Section 1.6.3), an implementation of the dual algorithm of Goldfarb and Idnani [22]. The linearized waste row is elastic, a slack penalized 10^5 times the profit gradient's scale, which is the application's own device so that a linearization that cannot be met inside the move box yields the step that violates it least instead of no step. Over the 100 h challenge on the suggested settings the layer made 99 runs and 255 SQP solves, none of them failing, at 3.43 iterations each on average.

11.5.4 Modifier Adaptation

Modifier adaptation keeps the model as last fitted (fitting it once if it never was) and corrects it at every run k, as shown in Equation (11.8).

$$\begin{aligned} \phi_m(u) &= \hat{\phi}(u) + \lambda_\phi^\top (u - u_k) \\ G_m(u) &= \hat{x}_G(u) + \varepsilon_k + \lambda_G^\top (u - u_k) \\ \varepsilon_k &= \bar{y}_G - \hat{x}_G(u_k) \\ \lambda &\leftarrow (1 - K)\,\lambda + K\,(\nabla_p - \nabla\hat{}) \end{aligned} \tag{11.8}$$

and moves to the maximum of ϕ_m subject to $G_m \le x_{G,\max} - b$. Against Equation (1.13), the profit carries no value offset, which would not move the maximum, and the predicted waste carries one. The waste offset ε is taken whole at every run, because filtered at K it left the steady waste 0.10 to 0.13 points above the limit with no noise at all, against 0.02 to 0.04 unfiltered, as the application's development record notes. The gradient modifiers λ are filtered with gain K, 0.5 by default. The model's gradients $\nabla\hat{}$ come from central differences, and the plant's from the last three steady states read, as shown in Equation (11.9).

$$\begin{bmatrix} (u_{k-1} - u_k)^\top \\ (u_{k-2} - u_k)^\top \end{bmatrix} \nabla_p \phi = \begin{bmatrix} \phi_{k-1} - \phi_k \\ \phi_{k-2} - \phi_k \end{bmatrix} \tag{11.9}$$

with the same system for the waste. The estimate is used only when the inverse condition number of the scaled difference matrix is at least 0.1. On three exact plant points around (4.6 kg/s, 82 °C) it returns (−19.03, +4.081) against central differences of (−18.81, +4.117), and with $K = 1$ the corrected gradient equals the plant's, which confirms the sign conventions.

11.5.5 Three Safeguards and a Switching Rule

Three points on a line give no slope across it, and a running optimizer that moves straight toward an optimum lines its points up. Without excitation the three-point estimate of the application's first version collapsed to undefined modifiers and then to a corner of the box. The application therefore keeps the next point at least δ (default 1 °C, with 1 kg/s counted as 10 °C) off the line through the last two, solving the corrected problem once on each side of that line with $a^\top (u - u_k) \ge \delta$ as an extra row and taking the better side, unless the uncorrected move already lies that far off the line, and the first run after a fresh start probes 0.1 kg/s per degree of δ in B. The rule resembles the minimum-perturbation constraint of the dual approach of Rodger and Chachuat [18] as a 2016 thesis states it in its Equations 2.16 and 2.17 [9], without that approach's bound on the step length and with three-point differences rather than Broyden updates, so it is presented as the application's design and not as a reproduction of a published algorithm. The move box is the second safeguard. The third is a plant-change test of the application's own design: once three plant gradients have been used, a steady state whose profit change differs from the corrected model's prediction by more than 4 % of the profit, or whose waste differs by more than one point, clears the gradient history. The move limits, the test

and the elastic row were added after the challenge's drift drove the application's first version, as its development record notes, to (3.0 kg/s, 100 °C) with 16 points of excess waste on the default seed.

🕮 **Pause.** Try Exercise 11.2 before reading on.

Switching the method to modifier adaptation also clears its steady-state history, its gradient count, its last prediction and its modifiers. A usability review of the application found why. Two-step's last steady states sit within 0.002 kg/s and 0.03 °C of one another, so slopes estimated from them are noise, and before the rule a switch at 8 h in case 3 sent the plant to 55.8, 5.1, −25.6 and −7.8 percent of the best at the next four runs. With the rule the same switch reads 91.22, 90.60, 97.48 and 98.37 percent, and runs 14 to 24 average 98.87 %.

11.5.6 Economic Model Predictive Control

The economic MPC is the fourth and last method of the layer, and no case except one starts with it. Where the MPC option of Section 1.6.3 tracks setpoints with the objective of Equation (1.19), this controller prices the profit itself and chooses the setpoints. It writes the same two setpoints every $T_s = 5$ min, the analyzer cycle, counted as 50 executions of the 6 s loops, and plans them over 2 h, $N = 24$ steps grouped into nine move blocks, four single steps and then blocks of 2, 2, 3, 4 and 9 steps. The three loops run underneath as they do under the other methods. Its model is the optimizer's two-reaction model of Equation (11.5) made dynamic by the holdup, with the two setpoint loops as first-order lags, as shown in Equation (11.10).

$$\begin{aligned} \frac{\mathrm{d}\hat{x}_j}{\mathrm{d}t} &= \frac{1}{W_0}\left(F_{j,in} + \textstyle\sum_{i=1}^{2} \hat{\nu}_{ij}\hat{r}_i - (F_A + F_B)\,\hat{x}_j\right) \\ \frac{\mathrm{d}T}{\mathrm{d}t} &= \frac{T_{SP} - T}{\tau_T}, \qquad \frac{\mathrm{d}F_B}{\mathrm{d}t} = \frac{F_{B,SP} - F_B}{\tau_B} \end{aligned} \tag{11.10}$$

where j runs over A, B, E, P and G, $\hat{\nu}_{ij}$ are the mass coefficients listed after Equation (11.5), the rates use the fitted θ, and $W_0 = 2105.2$ kg is the declared holdup that the level loop holds, not the run's. The lags come from closed-loop step tests of ± 2 °C and ± 0.2 kg/s at the plant's best point with the layer off and the analyzer noise off, averaged over the two signs and fitted as first order plus dead time. The temperature loop gives a gain of 0.988, a lag of 1.06 min and a dead time of 1.33 min, with a residual of 0.045 of the step (rms) that is mostly the loop's overshoot, and the B feed loop a gain of 1.000, 0.71 min and 0.04 min, with 0.002. Each lag in the model is the sum of the fitted lag and dead time, $\tau_T = 2.39$ min and $\tau_B = 0.74$ min. The model's steady state is exactly the optimizer's, which a 40 h integration confirms to 7×10^{-15} in mass fraction. The model's five fractions come from an observer, not from the analyzer, which integrates them at every execution from the measured temperature, B feed and A feed (the last filtered over one decision period), starting from the model's steady state when the method takes over, and the analyzer reaches the controller only through the offset ε read at steady states.

Every decision begins by linearizing the model, by central differences, at its steady state for the plan's target u_t and discretizing it for a zero-order hold over T_s, with its affine remainder, as $z_{k+1} = A_d z_k + B_d u_k + c_d$ for $z = (\hat{x}_A, \hat{x}_B, \hat{x}_E, \hat{x}_P, \hat{x}_G, T, F_B)$ and $u = (F_{B,SP}, T_{SP})$. An analytic Jacobian with a separate scaling-and-squaring exponential reproduces A_d, B_d and c_d to 1.2×10^{-11}. At the target of case 9's decision at 13.0 h, (5.459 kg/s, 84.95 °C), the residence time is 5.18 min and the steady gains of the waste fraction are −0.02127 per kg/s of B and +0.004277 per degree, equal to the static model's. The plan maximizes the corrected profit of Equation (11.8) along the horizon, as

shown in Equation (11.11).

$$\begin{aligned}
\min_{u} \; & \sum_{k=1}^{N} \Big[-T_s\, g^{\top} y_k + \tfrac{T_s}{2} (s_k - u_t)^{\top} \tilde{H} (s_k - u_t) \Big] \\
& \quad + \textstyle\sum_{k=2}^{N} \left(\rho_1 \sigma_k + \rho_2 \sigma_k^2 \right) + \sum_k \Delta u_k^{\top} S \, \Delta u_k \\
\text{s.t. } & \hat{x}_{G,k} + \varepsilon + \lambda_G^{\top} (\zeta_k - u_r) \le x_{G,\max} - b + \sigma_k \\
& \sigma_k \ge 0, \qquad k = 2, \ldots, N \\
& u_{lo} \le u_k \le u_{hi}, \qquad |\Delta u_k| \le \Delta u_{\max} \\
& u_k = u_t \text{ from the first block that } \Delta u_{\max} \text{ allows}
\end{aligned} \tag{11.11}$$

where $y_k = (\hat{x}_P, \hat{x}_E, F_B, T)$ and $\zeta_k = (F_B, T)$ are predicted, s_k are the setpoints in force at step k, u_r are the setpoints of the steady state last read, and g is the gradient of the corrected profit at the target's steady state $\bar{x}$, with $\bar{F} = F_A + F_{B,t}$,

$$g = \left(p_P \bar{F},\; p_E \bar{F},\; p_P \bar{x}_P + p_E \bar{x}_E - p_B + \lambda_{\phi,1},\; \lambda_{\phi,2} \right). \tag{11.12}$$

The matrix $\tilde{H}$ is the model's steady-profit curvature at the target, by central differences and floored as the layer's SQP floors it (in scaled variables, each eigenvalue at least 1 % of the trace), and the modifiers add none because they are linear. At the decision at 13.0 h it is $[[28.377, -2.369], [-2.369, 0.773]]$ in \$/s per $(\mathrm{kg/s})^2$, per kg/s and degree, and per degree squared, and the floor did not act. The move suppression is a tenth of one period's curvature, $S = 0.1\, T_s\, \mathrm{diag}(\tilde{H})$. The rate limit $\Delta u_{\max}$ is 0.125 kg/s and 1.25 °C a decision, so eight moves cross the move box corner to corner and the target can always be reached inside the horizon. The box $[u_{lo}, u_{hi}]$ is modifier adaptation's move box about the last read point intersected with the decision box. The waste ceiling is the corrected waste of Equation (11.8) applied to the predicted transient and is soft, with one slack per step priced at $\rho_1 = 10^8$ and $\rho_2 = 10^{11}$ dollars per unit of mass fraction and per unit squared, so that one point of excess held for one step costs about $\$1.1 \times 10^7$, some 500 times in the linear term alone what a point of waste is worth near the limit (6.69 \$/s, or \$2006 a step). The first step is not judged, because the earliest the waste can answer a move is through the loop and the holdup, and the ceiling is judged from 10 min ahead. Judged only from 30 min, the plan ran hot inside the unjudged half hour and the waste peak after the feed rise of Section 11.6 was 9.05 % on three seeds against 8.63 %. A ridge of 10^{-6} keeps the program strictly convex, and the shared dual active-set solver [22] computes it, so each decision is one SQP for the target and one quadratic program for the path, with no nonlinear solver.

The target u_t is modifier adaptation's own decision. Under this method the layer keeps running the estimator of Equations (11.8) and (11.9) with its period, detector, filter, perturbation rule and change test, but hands its decision to the plan instead of moving. A probe or a perturbation is held at the requested point at once, because planned economically it earned nothing on the model, was put off at every decision, and left the plant at 74.8 to 75.5 % of the best for want of a gradient. A corrected optimum is reached along the plan. When the measured A feed has moved more than 1 %, or a price or the waste limit has changed, since the last steady state was read, the target is re-solved at every decision by the layer's SQP on the corrected model at the conditions then in force, inside the same move box. This is the only information the economic MPC acts on that modifier adaptation must wait an hour to use. A plan whose first move and target both lie within 0.02 kg/s and 0.2 °C of the setpoints holds them, and a move larger than that restarts the detector and the period, so that steady states are read at setpoints that have held.

The arrival rows of the last line of Equation (11.11) pin every block from the first one the rate limits allow to the target, and the setpoint that binds takes its first move as a full step. With the arrival left to the end of the horizon, or to a block counting down toward it, each new plan put the arrival off by one more decision. Without the arrival rows the plan earned 99.6 to 101.1 % of the best by riding

the transient, and held the true waste above 8.1 % at steady state for 5.8 to 7.4 h in 24 h (three seeds each, with no change, a feed rise or a price cut).

The modifiers are validated before the plan takes them. Modifier adaptation arms its change test only after three gradients, and in case 7 the plan's steady states fell so that one slope spanned the drift, giving $\lambda_{\phi,1} = +90\,\$/\mathrm{s}$ per kg/s and walking the first version of the plan to the corner of its move box. The rule adopted applies the same test, with the same thresholds of 4 % of the profit and one point of waste, to every steady state read under the plan from the second on, against the plan's own prediction of it. A new gradient estimate enters the plan's λ, through the filter of Equation (11.8), only when the two newest of the three steady states behind it passed. Otherwise the plan keeps the λ it last accepted and, while its modifiers differ from the layer's, makes modifier adaptation's decision itself with its own. The offset ε is taken at every read, as modifier adaptation takes it. The rule was fixed before it was run, and its results, good and bad, are in Section 11.6.5.

A plan is accepted only when the target is finite, the solver reports an optimal solution, every array has its shape, every planned setpoint is finite and inside its box, the first move is inside the rate limit, the last block sits on the target, every slack is finite and not negative, and the cost and the predicted waste are finite. Anything else is rejected and named, the setpoints hold where they are, and modifier adaptation moves them at its next steady state until a plan is accepted again, which then returns at the rate limit. A test injects each fault and finds each rejected without a step in any setpoint. The solver's plans were accepted in every run this chapter reports. In Chrome on the real frame loop of case 9, with another process loading the host, the decisions took 1.38 ms on average and 1.7 % of a core at the default speed, and 9.4 % at ten times that speed, close to the apps' budget of 10 % of a core for one controller.

11.5.7 Plant, Model and What the Optimizer May Read

Three models are kept apart. The plant of Section 11.3 is the first. The identified model is the two-reaction approximation with its fitted θ, and the corrected model adds ε and λ. The plant's best point, the profit map, the gap to the best and the challenge's yardstick are computed by the simulator from its own steady state, for display and grading, and no controller or optimizer reads them. The comma-separated export marks those columns with a `truth_` prefix. The economic MPC's model is the identified one with the declared holdup and the two identified loop lags, and it reads what the layer reads together with the measured temperature and B feed and the setpoints in force. One of the application's tests raises each of the 13 parameters, each disturbance setting other than the prices and the hidden masses of C and E by 20 % in place, without advancing the plant, in case 9 at 13 h, and the controller's inputs and decision stay the same bit for bit.

11.6 Ten Cases, Measured

The cases are listed in Table 11.3, and the Exercises section turns several of them into runs for the reader. Each number in this section was measured by stepping the engine directly, outside any browser, at the application's seed (the default seed of Section 1.6.1) unless six seeds are named, with settings entered as each case's instructions say, and every case runs identically whether it is opened first, after another case, or again through Reset. A percentage quoted for run k is the plant's true profit, as a share of the best, at the steady state run k read, which is where run $k-1$ left the plant.

11.6.1 A Hand Search and a Model That Points Away

In case 1 at 4.00 kg/s of B, raising the setpoint from 75 to 78.0 °C takes the true profit from 151.1 to about 168.1 \$/s, with the waste reaching 8 % at 78.002 °C. The profit would keep rising, to 182.51 \$/s at 84.70 °C, but with 11.25 % waste, so the limit and not the peak ends the search, as the lowest-cost point of Volume 9, Chapter 8 sits on a constraint. A move past the limit shows quickly. A setpoint of 79 °C entered at 0.1 h raises the waste alarm at 0.33 h, 0.23 s later at the default speed of 1 h per second. Case 2 starts at the plant's best point and lets the layer fit the model there, 35 min in, with a misfit of 0.408 points rms over five fractions. The fitted model predicts 8.35 % waste where the

Table 11.3 The Ten Cases

#	Start and events	Teaching point
1	(4.00, 75), RTO off	hand search, active limit
2	at the best point, RTO off	plant against model map
3	(4.00, 75), two-step from 0.1 h	where two-step stops
4	two-step point, MA	modifier adaptation
5	best point, MA, noise 0.2 to 1 % at 6 h	noise and the gain K
6	(4.79, 89.7), MA, limit 12 % then 8 % at 12 h	a limit turns active
7	best point, MA, a_3 to 80 % at 6 h	unannounced drift
8	(4.00, 75), MA every 0.25 h, check off	period and detector
9	best point, economic MPC, A feed 1.83 to 2.2 kg/s at 12 h, P price −10 % at 24 h	planning through the transient
10	Challenge: 100 h, drift at 30 h, P price −10 % at 60 h, two-step every 0.25 h as shipped	the whole problem

Table 11.4 Measured Outcomes at the Application's Seed

#	Setting	Outcome
1	78.0 °C at 4.00 kg/s	168.1 $/s, waste 8.00 %
2	model fitted at the best	model best (4.52, 80.46), 195.07 $/s
3	two-step, runs 6–12	165.65–165.99 $/s, 92.78–92.97 %
4	MA, runs 10–19	mean 98.96 %, dips 96.66–97.09 %
	$\delta = 0$, 0.5, 2 °C	mean 81.44, 99.33, 97.99 %
5	$K = 0.5$ as shipped, 24 h	Profit so far 97.18 %
	$K = 0.3$, 0.2 from 6.2 h	98.40, 98.68 %
6	reads after the cut	8.10, 7.97 % waste at 99.66, 99.58 %
7	MA against Off, 24 h	98.50 against 95.26 %
8	0.25 h, check off	94 runs, 87.78 %, waste +1.97
	check on	38 runs, 97.68 %, waste +0.21
	period 1 h	23 runs, 96.60 %, waste +0.04

plant makes 8.00 %, so it believes the plant is over its limit, and it places its best point at (4.52 kg/s, 80.46 °C), promising 195.07 $/s. Two-step switched on at 0.7 h goes there in one run, and the next steady state is 96.61 % of the best. The path ends at (4.45 kg/s, 78.9 °C), the model fitted at the right point having pointed away from it.

11.6.2 Two-Step Against Modifier Adaptation

In case 3 two-step switched on at 0.1 h makes its first run at 1 h, to (4.387 kg/s, 77.65 °C), and by the sixth it has stopped at (4.445 kg/s, 78.93 °C), the static fixed point. The steady states it reads from 6 to 12 h earn 165.65 to 165.99 $/s, 92.78 to 92.97 % of the best, while the fitted model promises 180.49 to 182.05 $/s at the same runs. At 8 h the model's best point (4.448 kg/s, 78.93 °C) sits on the setpoints (4.447 kg/s, 78.93 °C), as the application's Reactor view of that instant shows in Figure 11.3. Switched to modifier adaptation at that instant, the same run climbs toward the plant's best point and reads 99.46 % of it at 20 h, at (4.676 kg/s, 82.36 °C), as shown in Figure 11.2. The corrected model there carries $\varepsilon = -0.75$ points and a temperature slope of −1.24 $/s per degree in its uncorrected part, against a plant estimate of +4.36 $/s per degree.

> **TRY IT IN THE APP 3 · Two-Step RTO Settles Below the Best**
>
> Load the scenario, press Run and set RTO method to Two-step. Watch the numbered path on the profit map: it climbs for a few hourly runs and by the sixth it stops near 4.45 kg/s and 78.9 °C, short of the star. At 8 h Gap to best reads 7.16 %. Switch Map shows to RTO model, and the model's own best point sits where the path stopped.

Case 4 starts at the two-step point with modifier adaptation. The steady states runs 1 to 6 read are 91.21 to 98.20 % of the best while the gradient estimates build up. From run 7 on, two runs in three

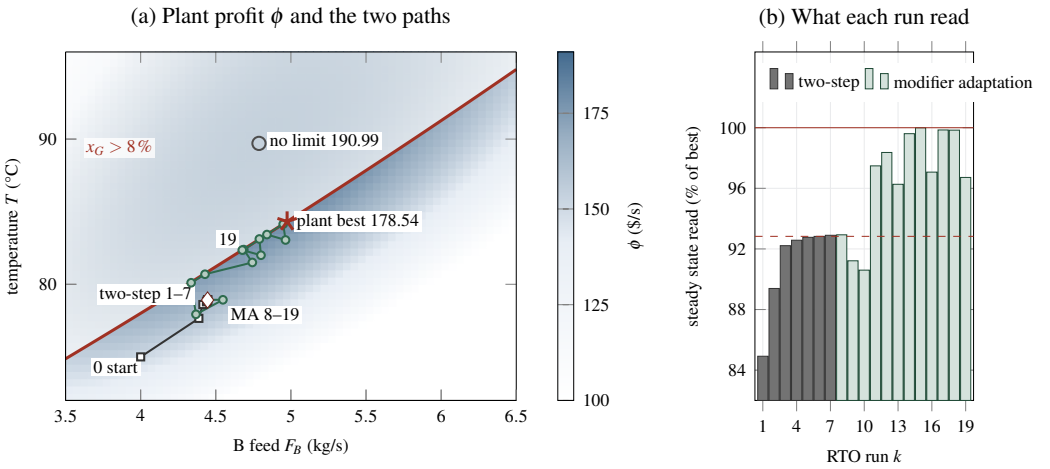


Figure 11.2 The lesson of case 3, computed from the application's own functions. (a) The plant's steady profit over B feed and temperature, with the 8 % waste limit (thick solid line) and the region beyond it faded, the plant's best point (star, 178.54 $/s), the two-step fixed point where the refitted model's best point sits (diamond), and the unconstrained optimum (open circle, 190.99 $/s). Two-step (squares) from the start point 0 stops at run 7, and modifier adaptation (filled circles) from 8 h climbs to run 19 at (4.676 kg/s, 82.36 °C), 99.46 % of the best at 20 h. (b) The steady state each run read, as a share of the best, two-step as dark bars and modifier adaptation as light bars, with the static two-step value of 92.83 % dashed and the best solid. Every third modifier-adaptation run reads lower because the run before it perturbed off the line of its predecessors.

read 98.61 to 100.09 %, with runs 13, 14, 16, 17 and 20 within 0.1 %, and every third, after a move the perturbation rule pushed off the line, reads 96.66 to 97.09 % (runs 9, 12, 15 and 18). Runs 10 to 19 average 98.96 %. Over runs 8 to 19 the model's own temperature slope is -2.56 to -1.08 $/s per degree while the plant estimate is $+3.13$ to $+5.41$, so the model would have cooled the reactor and the correction carried it hotter. The dips are the price of the perturbation and scale with it. With $\delta = 0$ the estimate collapses and runs 10 to 19 average 81.44 %, with $\delta = 0.5$ °C they average 99.33 %, and with 2 °C they average 97.99 % with dips to 93.35 %. The shipped 1 °C earns less here than 0.5 °C, and the application's development record keeps it for the margin it leaves the waste limit under noise in the challenge.

> **+ DID YOU KNOW? The Price of Probing, in Another Plant**
>
> The desalter of Volume 11, Chapter 2 pays the same price: its extremum-seeking controller perturbs the mixing pressure drop, and with its learning switched off the probing alone leaves more salt than a fixed setting. On the Williams–Otto reactor, a 2022 review found such model-free estimators reach the optimum without a model, but significantly more slowly than model-based ones [23].

11.6.3 Noise, a New Limit, Drift and the Period

In case 5 the analyzer noise rises from 0.2 to 1 % of the reading at 6 h. At 1 % one report of P scatters by about 0.0011 in mass fraction and the mean of six by 0.00044, about 3.4 $/s of profit, and a difference of two such means over 1 °C is about 4.85 $/s per degree, as large as the plant's own slope of 3.89 $/s per degree near its best point. As shipped ($K = 0.5$) the plant-change test reads the noise step as a plant change at 7 h, a false alarm (profit +4.0 against +13.6 $/s predicted), runs 14 to 16 read 91.7, 89.2 and 90.9 %, and Profit so far ends at 97.18 % at 24 h. With K lowered to 0.3 at 6.2 h it ends at 98.40 %, and with 0.2 at 98.68 %. The lesson depends on the noise drawn. Over the default

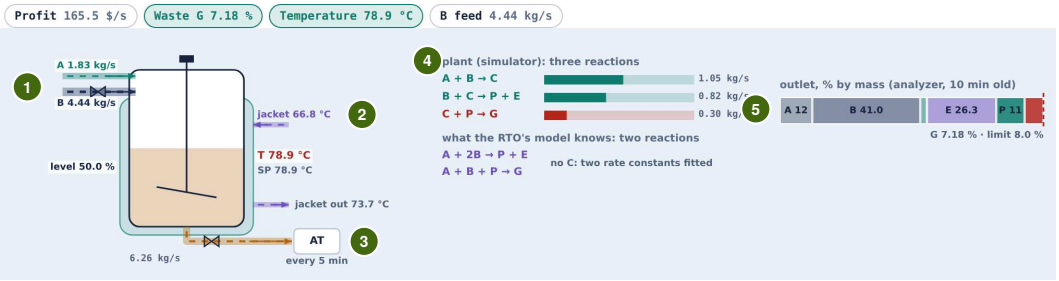


Figure 11.3 Case 3 at $t = 8$ h on the Reactor view, two-step settled after seven hourly runs, cropped to the plant view and its status chips. (1) The measured A and B feeds, 1.83 and 4.44 kg/s. (2) The jacket supply at 66.8 °C and its return at 73.7 °C. (3) The level at 50.0 % and the effluent at 6.26 kg/s to the analyzer. (4) The plant's three reaction rates (1.05, 0.82 and 0.30 kg/s) beside the model's two reactions with no C. (5) The outlet bar, the analyzer's report, 10 min old, with $x_G = 7.18$ % under the 8 % limit.

seed and seeds 1 to 5 the gain from $K = 0.3$ over 0.5 is $+1.22$, $+1.89$, -0.03, $+0.68$, $+0.47$ and $+2.05$ points, so one realization in six shows no benefit, and the case runs at its own seed so that a reader sees the first of these.

In case 6 the waste permit falls from 12 to 8 % at 12 h. The run at 12 h reads 11.53 % waste (11.55 % true), and the offset ε, measured at one steady state and taken whole, puts the corrected model's limit onto the plant's at once. The next two steady states hold 8.10 and 7.97 % waste at 99.66 and 99.58 % of the new best. In case 7 the waste reaction slows to 80 % at 6 h, unannounced, and the best possible rises from 178.54 to 242.88 $/s at (5.261 kg/s, 89.24 °C), 36.0 % more. The plant-change test fires at 7 h, a real change this time, and the runs at 9 and 10 h read 99.37 and 100.07 % of the new best. Profit so far at 24 h is 98.50 % against 95.26 % with the layer off, where the plant stays at 228.6 $/s with 6.96 % waste and its limit no longer binding.

Case 8 runs modifier adaptation every 15 min with the check off, so each run reads a plant still moving from the last. It makes 94 runs in 24 h, ends at 87.78 %, and leaves the waste up to 1.97 points over its limit at steady state. With the check on it makes 38 runs and earns 97.68 % with at most 0.21 point over, and hourly with the check off 23 runs and 96.60 % with at most 0.04. In case 3 the detector says steady only while the true fractions move less than 0.0073 h^{-1}, and under 1 % noise in case 5 while they move up to 0.0102 h^{-1}, so the detector alone admits a little transient that a longer period removes.

11.6.4 Economic MPC Against Modifier Adaptation

The comparison starts every method at the plant's best point, with the method chosen before the run, applies one change at 12 h and judges the true profit from 12 to 24 h as a share of the clairvoyant best. The changes are a rise of the A feed from 1.83 to 2.2 kg/s, which the A meter reports, a 10 % cut of the P price, which the layer is told at once, and no change, which measures each method's own running cost. Every row uses the default seed and five more, with the same analyzers and noise, and the economic MPC runs with the settings case 9 ships, which are the layer's defaults, retuned for no case. The results are listed in Table 11.5. At the hourly period the steady-state check makes no difference to either RTO method, because the detector has always passed by the time the period allows a run, so the rows with the check on and off are identical. Holding the setpoints earns more than the best permitted point only by running over the waste limit, for 11.6 h after the feed rise.

On profit the economic MPC ties modifier adaptation, and it does not earn more. After the feed rise it earns 98.98 % against 99.38 %, after the price cut 98.12 % against 97.81 %, and with no change 98.83 % against 98.88 %, each difference inside the spread of the six seeds. The premise that an RTO

Table 11.5 The Economic MPC Against the Layer's Other Methods Over Six Seeds. Profit Is the True Profit From 12 to 24 h as a Percentage of the Best (Mean, With the Range After the Feed Rise). Peak Is the Largest True Waste in That Window and h Over the Longest Time at Steady State Above 8.1 % in 0–24 h, Over the Six Seeds

	A feed rise (metered)			P price cut		No change	
Method	profit %	peak %	h over	profit %	h over	profit %	h over
Off, setpoints held	102.88	9.36	11.58	97.03	0.00	100.02	0.00
Two-step, 1 h, check on	92.15 (92.13–92.18)	8.14	0.57	90.19	0.00	92.84	0.00
Two-step, 1 h, check off	92.15 (92.13–92.18)	8.14	0.57	90.19	0.00	92.84	0.00
MA, 1 h, check on (shipped)	99.38 (98.77–99.79)	9.38	1.27	97.81	0.00	98.88	0.00
MA, 1 h, check off	99.38 (98.77–99.79)	9.38	1.27	97.81	0.00	98.88	0.00
MA, 15 min, check on	99.05 (98.25–99.47)	9.55	2.12	97.13	1.95	99.09	1.72
MA, 15 min, check off	93.45 (68.78–98.74)	10.98	16.21	39.09	17.74	78.23	17.74
Economic MPC	**98.98** (97.95–99.42)	**8.90**	**1.56**	**98.12**	**0.00**	**98.83**	**0.00**
Economic MPC without ε, λ	95.65 (95.59–95.72)	8.24	0.00	100.08	11.42	96.52	0.00
Economic MPC without the holdup	99.10 (98.22–99.50)	9.42	1.46	98.06	0.01	98.86	0.00
Economic MPC without the slope check	99.08 (98.84–99.37)	8.81	1.58	97.93	0.00	98.87	0.00

loses profit while it waits for steady state is not borne out on this plant, where a move settles in about half an hour and an hourly layer gives up little to the wait. Two counterfactual runs of the same controller attribute what it does. Without the modifiers, optimizing its fitted model alone, it earns 95.65 % after the feed rise and 96.52 % with no change, and after the price cut it earns 100.08 % only by holding the waste above the limit for 11.4 h, because without ε the model's waste is the one two-step trusts. The profit is therefore modifier adaptation's. Without the holdup, planning on the static map, it earns the same within the seed spread but lets the waste peak at 9.42 % after the feed rise, against 8.90 % with the holdup and 9.38 % for modifier adaptation. The lower peak comes from acting at once on the A meter through a model of how the outlet answers a move. After the feed rise the economic MPC still spends longer just over the limit at steady state, at most 1.56 h against 1.27 h, because both methods learn the new waste only from a steady state read after the change. Modifier adaptation every 15 min with the check on, the cheapest way to make the shipped method react sooner, earns 99.05 % after the feed rise with a waste peak of 9.55 % and up to 2.1 h over, and with the check off it acts as early as the economic MPC but reads transients as steady states, earning 39.09 to 93.45 % across the three changes with up to 17.7 h over. With the integration and control step halved, on three seeds, the peaks after the feed rise rank the same way, 8.83 % against 9.45 %. Figure 11.4 draws the profits and the waste peaks of Table 11.5.

🕮 **Pause.** Try Exercise 11.3 before reading on.

In case 9 itself, at the application's seed, the A feed rises at 12 h and the P price falls 10 % at 24 h. The economic MPC moves 5.1 min after the rise, at its next decision, and first lowers the temperature setpoint to 83.19 and then 82.64 °C while it raises the B feed, which holds the waste down while the extra A arrives, before it takes the setpoints to (5.459 kg/s, 84.98 °C) by 12.34 h, as shown in Figure 11.5. Modifier adaptation ran at 12.00 h, from a window the rise had not yet reached, and holds that move for an hour. The true waste peaks at 8.63 % under the economic MPC and 9.33 % under modifier adaptation, and at steady state the two spend 1.53 h and 1.23 h above 8.1 % in the 36 h run. Over 12 to 24 h modifier adaptation earns 99.45 % of the best and the economic MPC 99.22 %, over 24 to 36 h, after the price cut, 96.72 % and 97.25 %, and Profit so far at 30 h reads 98.46 % against 98.45 %. At 24 h it reads 98.66 % under the economic MPC and 98.83 % under modifier adaptation. The same case without the modifiers earns 95.62 % over 12 to 24 h, and without the holdup its waste peaks at 9.39 %.

A change that no instrument reports removes the economic MPC's advantage. Case 7, whose drift

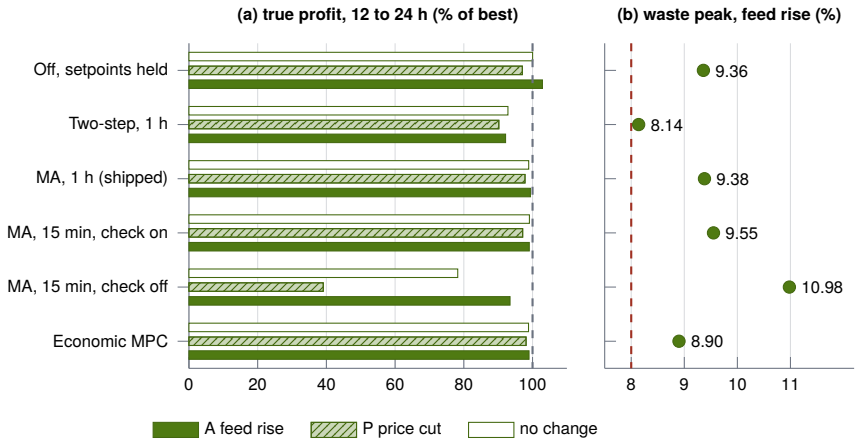


Figure 11.4 Table 11.5 drawn as a chart, mean of six seeds. (a) The true profit from 12 to 24 h as a share of the best after the metered A feed rise (dark bars), the P price cut (hatched) and with no change (white), with the best dashed. Two-step stands for its identical rows with the check on and off. (b) The largest true waste after the feed rise, with the 8 % limit dashed. Notice that the economic MPC's bars match those of hourly modifier adaptation within a point while its waste peak is lower. Two-step keeps the waste lowest but stops short on profit. Only holding the setpoints earns more than the best, and only by staying over the limit.

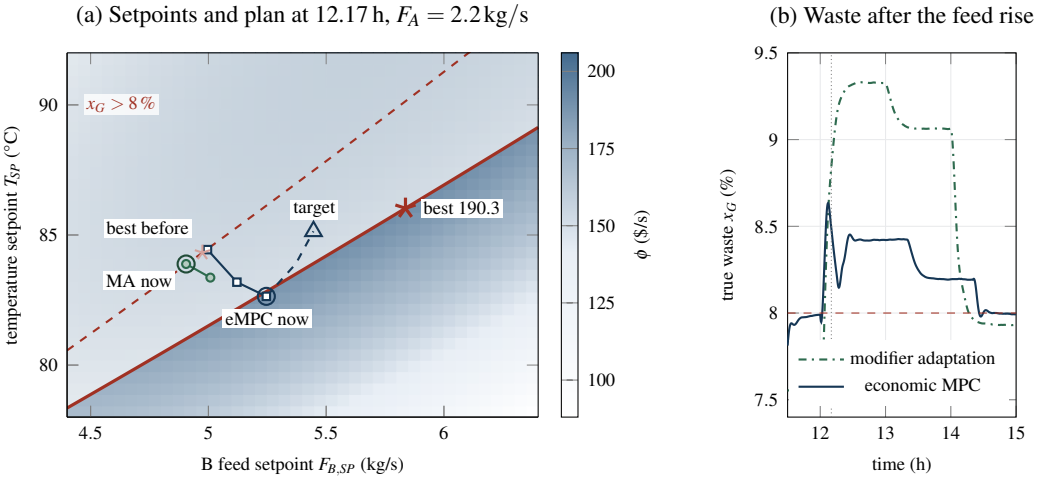


Figure 11.5 Case 9 at $t = 12.17\,\mathrm{h}$, ten minutes after the A feed rose to 2.2 kg/s, drawn from the engine's data for the Landscape view at the application's seed. (a) The plant's profit map and 8 % waste boundary (thick solid line) at the new feed, the boundary before the rise (dashed), and the plant's best point before (pale star, 178.54 $/s) and after (star, 190.33 $/s). The economic MPC (squares) has moved twice since 12 h, to (5.246 kg/s, 82.64 °C), and plans (dashed) to its target (triangle) at (5.446 kg/s, 85.14 °C) with a predicted corrected waste of 8.00 % at the end of its horizon. Modifier adaptation (filled circles) holds the point it chose at 12.00 h, (4.905 kg/s, 83.89 °C), until 13.00 h. (b) The true waste of both runs, peaking at 8.63 % under the economic MPC (solid) and 9.33 % under modifier adaptation (dash-dotted), with the limit dashed and the instant of (a) dotted.

at 6 h the layer is not told of, run to 24 h under both methods, each chosen before the run, gives a peak of the true waste after 6 h of 8.21 % under the economic MPC against 8.15 % under modifier

Table 11.6 Hours at Steady State With the True Waste Above 8.1 %, Mean (Max) Over Six Seeds, and Mean Profit of the Whole Run. eMPC, First: Without the Slope Check

Case	MA	eMPC, first	eMPC	profit %, MA, eMPC
1 and 3	0.00 (0.00)	0.22 (0.72)	0.00 (0.00)	97.41, 96.54
2	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)	98.68, 98.59
4	0.18 (0.58)	0.43 (1.44)	0.24 (1.44)	98.22, 98.04
5	1.39 (3.34)	1.32 (3.18)	0.84 (2.59)	96.02, 93.80
6	0.40 (1.06)	0.73 (1.47)	0.32 (0.67)	98.35, 97.70
7	1.92 (4.18)	1.59 (3.65)	0.00 (0.00)	98.15, 98.51
8	23.99 (32.50)	12.84 (28.44)	24.32 (30.69)	83.09, 81.20
10, shipped	5.71 (6.30)	4.76 (6.49)	4.45 (6.04)	98.40, 98.31
10, suggested	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)	97.43, 97.14

adaptation, and Profit so far of 97.72 % against 98.50 %. Both methods must wait for steady states to see what the model gets wrong, so acting at once gains nothing. Cases 9 and 7 together separate a disturbance the controller is told of from a plant change it can only infer.

> **TRY IT IN THE APP** **9 · Economic MPC Against Modifier Adaptation**
>
> Run the scenario to 30 h and watch the Waste G trend after the A feed rises at 12 h. Reset, set RTO method to Modifier adaptation and run again. The economic MPC moves at once and holds the true waste peak to 8.63 % against 9.33 %, yet Profit so far at 30 h reads 98.46 % against 98.45 %. The trend, logged every 0.05 h, peaks slightly lower.

11.6.5 Worst Cases and the Slope Check

Because the economic MPC shares the waste limit with every other method, each existing case was also run under it and under modifier adaptation with that case's own settings, events and duration (36 h, the challenge 100 h), over six seeds, and the time the true waste spends above 8.1 % at steady state is listed in Table 11.6. The first version of the plan took modifier adaptation's slopes whenever the layer formed them. At the application's seed it then spent 1.59 h over in case 7, where modifier adaptation spends none, and 26.84 h against 6.53 h in case 8. The cause in case 7 is the unarmed window of the change test. The plan's steady states at 5.25, 6.67 and 7.84 h fell across the drift at 6 h, the slope formed at 7.84 h gave $\lambda_{\phi,1} = +90.4\,\$/\mathrm{s}$ per kg/s, and the test that fired at the next read cleared the history but not that λ. Modifier adaptation has the same weakness. With the drift moved to 3, 4 or 4.5 h it spends 2.08, 2.06 and 2.05 h over, and over six seeds of case 7 it spends 1.92 h on average and up to 4.18 h, with the true waste reaching 11.12 % at one seed.

With the slope check of Section 11.5.6, case 7 spends no time over at any of the six seeds, case 8 spends 0.09 h at the application's seed, and over six seeds the economic MPC is at or under modifier adaptation's mean in seven of the nine rows of Table 11.6. The two exceptions are case 4, 0.24 against 0.18 h, and case 8, 24.32 against 23.99 h, a configuration that reads transients under every method. Cases 1 and 3 share their start, so with the method chosen before the run they are one row. The check has a price. Under the fivefold analyzer noise of case 5 more than half of the slopes fail it (10.2 of 18 on average), learning slows, and the run earns 93.80 % of the best against 96.02 %, and 89.15 % against 97.35 % at the application's seed. It costs a little in cases 1, 3 and 6 as well (96.54 against 97.41 %, 97.70 against 98.35 %), and gains in case 7 (98.51 against 98.15 %). On the challenge's suggested settings both methods keep the waste at steady state within the tolerance at every seed. The economic MPC meets both of the challenge's targets at the application's seed with 97.18 %, below modifier adaptation's 97.47 %, but earns as little as 96.88 % at another seed, short of the 97 % target.

11.6.6 The Challenge and Its Verdict

The challenge runs 100 h from (4.00 kg/s, 75 °C), with the waste reaction slowed to 80 % at 30 h and the P price cut 10 % at 60 h, so the best point moves from 178.5 \$/s at (4.97, 84.3) to 242.9 at

Table 11.7 The Challenge Over 100 h. Waste Is the Worst Excess Over 8 % at Steady State, in Points

Settings	Default seed	Other seeds
two-step, 0.25 h (shipped)	fails, 92.83 %	seeds 0–5: 92.82–92.85 %
two-step, 19 settings	fail, 86.10–92.89 %	—
MA, 1 h, K 0.2, b 0.15 (suggested)	passes, 97.47 %, +0.000	31/31 pass, 97.18–97.53 %, at $\Delta t/2$ 31/31, 97.10–97.53 %
MA, 1 h, K 0.3, b 0	passes, 98.19 %, +0.080	21/31 pass, at $\Delta t/2$ 16/31
MA, 1 h, K 0.5	passes, 98.24 %, +0.055	seeds 3, 4, 5 fail, +0.103 to +0.163
MA, 0.25 h, K 0.3	fails, 98.62 %, 4.97 h over	seeds 0–5 fail
RTO off	fails, 78.54 %	—

(5.26, 89.2) and then 152.4 at (4.77, 85.7). Its target is a cumulative true profit of at least 97 % of the clairvoyant best's, switched the instant a condition changes, and a true waste that never exceeds 8.1 % at steady state, which is the case's 8 % limit with a 0.1-point tolerance for an optimizer that holds the limit through a noisy analyzer. Steady state means the true fractions of P and G both move less than $0.002\,\mathrm{h}^{-1}$. A lined-out plant moves at most $0.001\,55\,\mathrm{h}^{-1}$ under its own noise, so it is always judged, and judgment resumes 27 min after a 3 °C move. Two verdict rules were added in a technical review of the application. The waste is judged against the case's 8 % and not against the learner's limit control, because with that control raised to 9 % the suggested settings had earned the certificate with the true waste up to 9.36 %. A run whose limit, disturbances or model constants leave the case's values beyond its two scripted events, or in which an event button is pressed, reads "not graded". With the limit at 9 % the run now reads not graded, and with a learner's back-off of 0.3 point, which is allowed, it is graded and fails at 96.31 %. The outcomes are listed in Table 11.7.

As shipped, two-step every 15 min earns 92.83 %. It cannot pass with any setting tried: over periods of 0.25 to 4 h, the check on and off, windows of 6 and 12 reports, a back-off and a tighter threshold, 19 settings in all, it earns 86.10 to 92.89 %, the lowest with a 12-report window every 15 min and the check off, which reads transients. Its fixed point belongs to its model. The application first suggested modifier adaptation hourly with $K = 0.3$, which passes at the default seed with 98.19 %. The technical review ran it at 31 noise seeds and found it passing at only 21, and at 16 of 31 with the step halved, the default seed itself failing at $\Delta t/2$ by 0.135 point. Every failure at Δt began between 5.39 and 8.61 h, after runs made before the gradient estimates were good, and left the waste at up to 8.28 % for 19 to 68 min. The suggestion became $K = 0.2$ with a 0.15-point waste back-off, a margin of the kind Chapter 10 keeps from a limit that its laboratory sees late. It was the only one of fifteen settings in the review's sweep to pass all 62 runs, at 97.18 to 97.53 % with at most 0.052 point over at Δt and 97.10 to 97.53 % with at most 0.059 at $\Delta t/2$, and it passes at $\Delta t/4$. At the default seed it earns 97.47 %, against 98.19 % for the old suggestion, so the margin on waste is paid for in profit. On the suggested settings the 100 h run takes about 0.2 s of computing in Node on the development host, and at 10 h per second the app plays it in ten seconds.

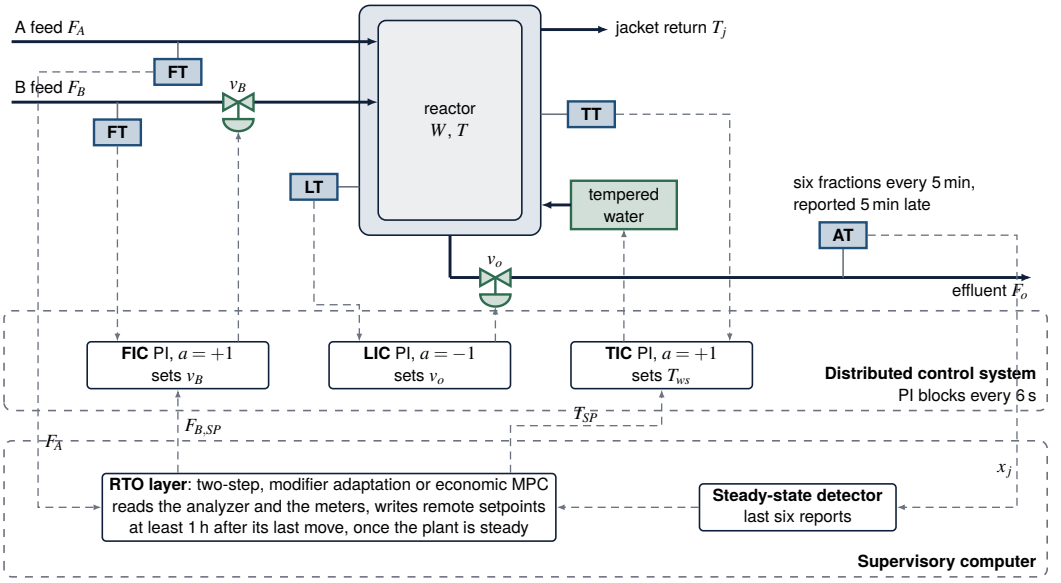


Figure 11.6 The loops and the optimization layer of the Williams–Otto reactor drawn with their hardware. Three PI blocks of the distributed control system hold the level, the B feed and the reactor temperature, the last through a tempered-water unit that sets the jacket supply. A supervisory computer reads the analyzer and the A feed meter and writes two remote setpoints, T_{SP} and $F_{B,SP}$. Notice that the layer reads no temperature or valve: it sees the plant through the analyzer, whose reports are minutes old, and it moves the plant only through the setpoints of loops that do the regulating.

DID YOU KNOW? The Most Profitable Setting Fails

From the simulator: in the 100 h challenge of Table 11.7 the highest profit, 98.62 % of the best, belongs to modifier adaptation run every 15 min with a filter gain of 0.3, and it fails, with the true waste over the tolerance for 4.97 h. With the layer off the plant earns 78.54 %.

11.7 Technology in Practice

The loops and the layer of Section 11.5 run on real instruments, valves and computers, drawn together in Figure 11.6.

What Measures It

The layer's eyes are the analyzer. In the model it reports six mass fractions every 5 min, each one a cycle after its sample was drawn, with a relative noise of 0.2 %. That timing is a process chromatograph's. Chromatographs separate the species in time, so cycle times of minutes are common and some dead time is inevitable [24], as Chapter 2 describes. A spectroscopic analyzer is the faster alternative for a liquid product.

TECHNOLOGY Sensor: A Near-Infrared Analyzer on the Product

A near-infrared analyzer reads the overtone and combination bands of bonds such as C–H and O–H, so nearly any molecule that carries them shows up in the spectrum. It is fast, one minute or less per sample, does not destroy the sample and suits in-line use, and optical fibers let one instrument watch several points of a plant up to 100 m away [25]. Its numbers come from a multivariate model calibrated on samples whose composition a reference method has determined, so it can never be better than that reference [25]. Its weak

point is the sample. When the line that feeds an analyzer plugs, the instrument measures old, stale sample, has no way of knowing it, and reports meaningless numbers [24].

The optimizer averages six reports before it trusts a steady state, and a move takes about half an hour to read, against lags of 0.3 and 0.1 min on the thermowells and flow meters.

What Moves It

The temperature loop does not move a valve in the model. Its output is the temperature of the tempered water sent to the jacket, which follows its command with a lag of 2 min. A higher supply warms the reactor, so the loop is reverse acting. The action follows the final element: a reactor whose loop throttles an air-to-open steam valve to its jacket needs reverse action, and a vessel whose loop opens a chilled glycol valve needs direct action [24]. The B feed and effluent valves follow their commands with a lag of 0.5 min.

Where the Controller Runs

The three PI blocks run in a distributed control system, every 6 s in the model. Its supervisory level sends set points to the field controllers and requests their data, at cycle times from milliseconds to minutes [26].

TECHNOLOGY Platform: A Supervisory Computer Writing Setpoints

An optimization layer runs on a supervisory computer. It sends setpoints over a network to the loop controllers and reads their measurements back, a two-layer system with the PID loops below and the computer and its process models above. Each loop then has three modes: manual, automatic with a local setpoint from the operator, and automatic with a remote setpoint from the computer [24]. In the classical hierarchy a unit's optimizer runs every hour, the predictive controller every minute and the basic loops every second [27]. Because the loops do the regulating, the plant keeps working when the computer stops [11]. On a trend the setpoints simply go flat while every loop holds them.

In this chapter the layer runs no sooner than 1 h after its last move and only once the detector has passed, and the economic MPC decides every 5 min. TIC and FIC take their setpoints from it, LIC keeps a setpoint of its own, and the header's Manual pauses the layer and places all three loops on their stations.

IN PRACTICE Squeeze the Variance, Then Shift the Setpoint

An output that must not cross its limit, even for a moment, is run with a back-off: its setpoint sits a margin inside the limit, and any back-off from an active constraint costs money, the constraint's shadow price times the margin [28]. The rule of thumb that follows is called squeeze and shift: improve control to squeeze the variance, then shift the setpoint closer to the limit [28]. A margin is added for imperfect control and noise, and a constraint whose shadow price is small can be given a large one [23]. The waste back-off of 0.15 point in this chapter's challenge is such a margin, and its price is measured in Section 11.6.

What Is Modernizing It

Three developments bear on this layer. The first, feedback optimizing control, turns the economic objective into controlled variables, needs no numerical optimization online, and avoids the wait for steady state by using measurements as they come [23]. A 2022 review compares such methods on the Williams–Otto reactor with the waste held at 0.08 by the reactor temperature, the limit of this chapter [23]. The second learns the mismatch instead of modeling it. The same review corrects the two-reaction model with a radial basis function network fitted to the gap between measured and predicted cost [23], and Gaussian processes serve as modifiers in schemes that steer their own probing [7]. The third is economic model predictive control, which this chapter tests. One review

advises it only for small problems or where the optimization and control layers cannot be separated in time scale, because for practical implementation and tuning the combined layer is rarely a good solution [28]. On this plant it tied modifier adaptation on profit.

11.8 Limitations

The plant is the reactor-only benchmark, perfectly mixed and in mass, with no separator or recycle, and its energy balance, jacket, heats of reaction and valves are teaching values that set the dynamics and not the steady states. The mismatch is the literature's two-reaction model, a single structural error of a known kind, while an industrial model is wrong in many places at once. Plant gradients come only from past steady states, so the layer has no dedicated plant tests, no Broyden or quadratic surrogate updates and no dual formulation with a step bound [3, 17]. The perturbation rule, the move limits, the plant-change test and the elastic waste row are the application's design and were tuned on this plant. Every result is a run at the application's seed, checked at 31 seeds for the challenge and at six for the noise case, whose lesson is seed dependent. The waste tolerance of 0.1 point and the steady-state definition of the verdict are the application's choices. The analyzer is unbiased and never fails, and the prices and limit reach the optimizer exactly, which an industrial layer cannot assume.

> ⚠ **OPEN CHALLENGE** **Why Plants Seldom Keep Their Optimizers**
>
> The business benefits of real-time optimization are not disputed, yet its use has stayed relatively low [7]. A review names four technical reasons: the cost of building the model, the wait for a steady state before each model update, numerical failures of large optimizations, and conflicts with the planning layer above [23]. An optimizer needs a team of skilled engineers to maintain and monitor it, and without one the application is often turned off [23]. Progress would be an optimizer that learns from routine data while the plant moves, flags its own model when it drifts, and stays switched on without a dedicated team.

The economic MPC's model leaves out the intermediate C and the third reaction, which is the shared mismatch and stays the lesson, the energy balance and jacket, which it replaces by a lag, the temperature loop's overshoot, the analyzer's delay, and any retuning of the loops after the step tests, which becomes mismatch in the lags. Its composition is an open-loop observer corrected only through ε at steady states, so it learns the drift, as modifier adaptation does, only from steady states. It is linearized once per decision at its target and solved as one quadratic program, not as a nonlinear program. The arrival rows, the slope check, the dead band and the waste prices are the application's design, chosen on this plant before the comparison. During a fallback the layer moves with its own modifiers, which may include a slope the plan had rejected, a path that no run reported here exercised.

> ⚠ **OPEN CHALLENGE** **Slopes From Noisy Data**
>
> Inferring the plant's gradients from noisy measurements is hard, and it decides how well and how reliably modifier adaptation works. Newer schemes learn the mismatch between model and plant with Gaussian processes and steer their probing with an acquisition function, but their theory and practical performance are yet to be analyzed in depth, and probing that favors the feasible side of the process constraints is left to future work [7]. The challenge of this chapter shows the stakes: a setting chosen at one seed passed at 21 of 31. Progress would be an estimator that probes only as hard as the noise demands and says how far its slopes can be trusted.

11.9 Conclusion

This chapter ran the Williams–Otto reactor under three loops and an optimization layer that reads only its analyzer and meters. Its steady states are the published ones to every digit an independent solver prints, 178.540 \$/s at (4.9747 kg/s, 84.322 °C) under the 8 % waste limit. A model with two reactions, refitted at every steady state, takes the plant to (4.4454 kg/s, 78.915 °C) and stops at 92.83 % of that from any start, where its own slopes vanish and the plant's are -24.4 and $+5.76$. Modifier adaptation

reaches 100.000 % with exact gradients and averages 98.96 % on the running plant, every third run about three points short because it must perturb.

A filter gain, the detector, the period and a plant-change test decide how much of that survives noise, drift and a short period, and the challenge's suggested settings pass at 31 of 31 seeds only with a waste back-off of 0.15 point, after a setting without it passed at 21. An economic MPC on the same model and modifiers ties modifier adaptation on profit, 98.98 % against 99.38 % after a metered feed rise over six seeds, and owes its lower waste peak, 8.90 % against 9.38 %, to its dynamic model.

The safeguards are the application's own and the mismatch is a single known one. A fuller study would add a dual formulation with a step bound, gradient estimation by dedicated plant tests, an analyzer that drifts or fails, and the separator and recycle of the original plant. The simulator shows that an optimizer's model must be right about slopes, not only about values.

KEY POINTS

- An optimizer whose model is wrong in structure can stop short of the plant's best, here at 92.83 % of it, where its model's slopes vanish.
- Modifier adaptation corrects the model with plant gradients and reaches 100.000 % with exact ones, 98.96 % on average on the running plant.
- Every third run is about three points short, the price of the perturbation the gradient estimate needs.
- The challenge's suggested settings pass at 31 of 31 seeds only with a waste back-off of 0.15 point.
- An economic MPC on the same model ties on profit and owes its lower waste peak to its dynamic model.

Where else this applies. An optimum on a moving constraint in Volume 9, Chapter 8, and a slope learned by probing the plant in Volume 11, Chapter 2.

References

[1] P. D. Roberts, "An algorithm for steady-state system optimization and parameter estimation," *International Journal of Systems Science*, vol. 10, no. 7, pp. 719–734, 1979.

[2] J. F. Forbes and T. E. Marlin, "Model accuracy for economic optimizing controllers: The bias update case," *Industrial & Engineering Chemistry Research*, vol. 33, no. 8, pp. 1919–1929, 1994.

[3] A. Marchetti, G. François, T. Faulwasser, and D. Bonvin, "Modifier adaptation for real-time optimization—methods and applications," *Processes*, vol. 4, no. 4, p. 55, 2016.

[4] A. Marchetti, B. Chachuat, and D. Bonvin, "Modifier-adaptation methodology for real-time optimization," *Industrial & Engineering Chemistry Research*, vol. 48, no. 13, pp. 6022–6033, 2009.

[5] S. Cao and R. R. Rhinehart, "An efficient method for on-line identification of steady state," *Journal of Process Control*, vol. 5, no. 6, pp. 363–374, 1995.

[6] T. J. Williams and R. E. Otto, "A generalized chemical processing model for the investigation of computer control," *Transactions of the American Institute of Electrical Engineers, Part I: Communication and Electronics*, vol. 79, no. 5, pp. 458–473, 1960.

[7] E. A. del Rio Chanona, P. Petsagkourakis, E. Bradford, J. E. Alves Graciano, and B. Chachuat, "Real-time optimization meets Bayesian optimization and derivative-free optimization: A tale of modifier adaptation," *Computers & Chemical Engineering*, vol. 147, p. 107249, 2021.

[8] F. A. Cubillos, G. Acuña, and E. L. Lima, "Real-time process optimization based on grey-box neural models," *Brazilian Journal of Chemical Engineering*, vol. 24, no. 3, pp. 433–443, 2007.

[9] J. E. A. Graciano, "Real time optimization in chemical process: evaluation of strategies, improvements and industrial application," Ph.D. dissertation, Escola Politécnica, Universidade de São Paulo, São Paulo, Brazil, 2016.

[10] J. O. A. Matias, "Real-time optimization with persistent parameter adaptation using online parameter estimation," Ph.D. dissertation, Escola Politécnica, Universidade de São Paulo, São Paulo, Brazil, 2019.

[11] M. Pitt, "Taking a look back at control: Part 2," *The Chemical Engineer*, Oct. 2023, published online 26 October 2023. [Online]. Available: https://www.thechemicalengineer.com/features/taking-a-look-back-at-control-part-2/

[12] Computer History Museum, "Computers in control," CHM Revolution, online exhibition, section Real-Time Computing, accessed 18 November 2026. [Online]. Available: https://www.computerhistory.org/revolution/real-time-computing/6/130

[13] C. R. Cutler and R. T. Perry, "Real time optimization with multivariable control is required to maximize profits," *Computers & Chemical Engineering*, vol. 7, no. 5, pp. 663–667, 1983.

[14] C. Y. Chen and B. Joseph, "On-line optimization using a two-phase approach: an application study," *Industrial & Engineering Chemistry Research*, vol. 26, no. 9, pp. 1924–1930, 1987.

[15] M. L. Darby, M. Nikolaou, J. Jones, and D. Nicholson, "RTO: An overview and assessment of current practice," *Journal of Process Control*, vol. 21, no. 6, pp. 874–884, 2011.

[16] B. Chachuat, B. Srinivasan, and D. Bonvin, "Adaptation strategies for real-time optimization," *Computers & Chemical Engineering*, vol. 33, no. 10, pp. 1557–1567, 2009.

[17] A. Marchetti, B. Chachuat, and D. Bonvin, "A dual modifier-adaptation approach for real-time optimization," *Journal of Process Control*, vol. 20, no. 9, pp. 1027–1037, 2010.

[18] E. A. Rodger and B. Chachuat, "Design methodology of modifier adaptation for on-line optimization of uncertain processes," *IFAC Proceedings Volumes*, vol. 44, no. 1, pp. 4113–4118, 2011.

[19] G. François and D. Bonvin, "Use of transient measurements for the optimization of steady-state performance via modifier adaptation," *Industrial & Engineering Chemistry Research*, vol. 53, no. 13, pp. 5148–5159, 2014.

[20] D. Angeli, R. Amrit, and J. B. Rawlings, "On average performance and stability of economic model predictive control," *IEEE Transactions on Automatic Control*, vol. 57, no. 7, pp. 1615–1626, 2012.

[21] M. Ellis, H. Durand, and P. D. Christofides, "A tutorial review of economic model predictive control methods," *Journal of Process Control*, vol. 24, no. 8, pp. 1156–1178, 2014.

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

[23] D. Krishnamoorthy and S. Skogestad, "Real-time optimization as a feedback control problem: A review," *Computers & Chemical Engineering*, vol. 161, p. 107723, 2022.

[24] 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.

[25] C. Pasquini, "Near infrared spectroscopy: fundamentals, practical aspects and analytical applications," *Journal of the Brazilian Chemical Society*, vol. 14, no. 2, pp. 198–219, 2003.

[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] 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.

[28] S. Skogestad, "Advanced control using decomposition and simple elements," *Annual Reviews in Control*, vol. 56, p. 100903, 2023.

11.10 Exercises

11.1 ○ Concept check · 2 min

In case 3 the two-step layer stops at (4.4454 kg/s, 78.915 °C), where the plant earns 92.83 % of its best with 7.19 % true waste under the 8 % limit, while the model refitted there predicts 180.92 $/s. What stops the layer at that point?

(a) The waste limit, which the fitted model believes is active there

(b) The move box, which caps every run at 0.5 kg/s and 5 °C

(c) Analyzer noise, which hides the remaining slope from the least-squares fit

(d) The fitted model's slopes vanish there, though the plant's slopes do not

11.2 ○ Concept check · 1 min

Why does the application's modifier adaptation keep each new operating point at least δ off the line through its last two points?

(a) Three points on one line give no estimate of the plant's slope across that line
(b) The waste offset ε can be measured only after a move of at least δ
(c) The steady-state detector passes only after a setpoint move of δ or more
(d) The SQP needs a step of at least δ to leave a stationary point

11.3 Concept check · 2 min

After the metered A feed rise of Table 11.5, the true waste peaks at 8.90 % under the economic MPC and at 9.38 % under modifier adaptation. The same controller without the reactor's holdup in its model peaks at 9.42 %. What gives the economic MPC its lower peak?

(a) Its modifiers, which carry the plant's slopes into the plan
(b) Acting at once on the A meter through a model of how the outlet answers a move
(c) Its slope check, which rejects gradients formed across the change
(d) Its soft waste ceiling, priced far above what a point of waste earns

11.4 Theory and calculation · 6 min

Near the plant's best point the outlet holds a P fraction of 0.1079 and $F_A + F_B = 6.80\,\mathrm{kg/s}$, with $p_P = 1143.38\,\$/\mathrm{kg}$. In case 5 the analyzer noise is 1 % of the reading (one standard deviation), and the layer averages $N = 6$ reports drawn independently.

(a) Compute the standard deviation of one report of x_P, of the mean of six, and of the profit that mean implies through the P term of Equation (11.4). *Checkpoint:* 0.00108, 0.00044 and 3.4 \$/s.

(b) A slope is the difference of two such means taken 1 °C apart. Compute its standard deviation in \$/s per degree and compare it with the plant's own slope of 3.89 \$/s per degree near its best point. *Checkpoint:* About 4.84 \$/s per degree, as large as the slope.

(c) The gradient modifiers follow $\lambda \leftarrow (1-K)\lambda + K\,(\nabla_p - \nabla\hat{})$. How many runs does it take before a step change in the true correction is 90 % taken up, for $K = 0.5$, 0.3 and 0.2? How many hours is that at the default period? *Checkpoint:* 4, 7 and 11 runs; 4, 7 and 11 h.

11.5 Theory and calculation · 6 min

At the best permitted point, 4.9747 kg/s of B and 84.322 °C with $F_A = 1.8275\,\mathrm{kg/s}$, the reactions release 439.2 kW. Use the energy balance of Equation (11.3) at steady state and the constants of Table 11.2.

(a) Compute the heat taken up in warming the feeds from T_f to T, and the residence time $W_0/(F_A + F_B)$. *Checkpoint:* 88.2 kW; 5.16 min.

(b) Compute the heat the jacket removes and the jacket temperature T_j. *Checkpoint:* 351.0 kW; 79.93 °C.

(c) Compute the tempered-water supply T_{ws} the jacket needs, and say why none of these constants moves the plant's best point. *Checkpoint:* 74.08 °C.

11.6 Theory and calculation · 6 min

The economic MPC of Section 11.5.6 decides every $T_s = 5\,\mathrm{min}$ over $N = 24$ steps grouped into move blocks of 1, 1, 1, 1, 2, 2, 3, 4 and 9 steps, with rate limits of 0.125 kg/s and 1.25 °C a decision inside a move box of 0.5 kg/s and 5 °C either side of the last steady state read.

(a) Check that the blocks cover the horizon, give its length, and count the moves each setpoint may make in one plan. *Checkpoint:* 24 steps, 2 h; 9 moves each.

(b) How many decisions does a setpoint need to cross the move box corner to corner, and how long is that in minutes? Can one plan reach such a target before its horizon ends? *Checkpoint:* 8 decisions, 40 min; yes.

(c) The ceiling's slack costs $\rho_1 = 10^8$ per unit of mass fraction and $\rho_2 = 10^{11}$ per unit squared, each step. Compute the cost of one point of excess waste held for one step, and compare its linear part with what a point of waste earns near the limit, 6.69 \$/s, over one step. *Checkpoint:* $\$1.1 \times 10^7$; about \$2000 a step, a ratio near 500.

11.7 App challenge · 8 min

Open scenario 1, *1 · Find the best temperature by hand*, at the default seed and the default speed of 1 h per second. Leave the B feed setpoint at 4.00 kg/s. A value typed before Run starts the case lined out at that setpoint.

(a) Type 78.0 into the *Temperature setpoint* before Run and run to 2 h. Read the *Best possible*, *Gap to best* and *Waste G, by mass* tiles, and compute the true profit. Give the gap to within 0.1 point and the waste to within 0.05 point. *Checkpoint:* 178.5 \$/s, 5.82 %, 7.99 %; about 168.1 \$/s.

(b) Reload the case, type 79.0 before Run and run to 2 h. Read the same tiles. Why does the higher profit not count? *Checkpoint:* 3.51 %, 8.45 %, waste tile in alarm.

(c) The completion rule asks for a temperature at which the true profit is at least 165 \$/s and the true waste at most 8.1 % for half an hour. Interpolate the two waste readings to find the highest setpoint on the control's 0.1 °C grid that the rule accepts. *Checkpoint:* 78.2 °C.

11.8 App challenge · 8 min

Run the two methods of Section 11.5 at the default seed and 1 h per second, each with its RTO method chosen before Run. Export the CSV where asked: its time column, `Time (h)`, is in hours, and its first lines start with #.

(a) Open scenario 3, *3 · Two-step RTO settles below the best*, set *RTO method* to *Two-step* and run to 8 h. Read *Gap to best* and *Profit so far* to within 0.15 point, and compare `FB_sp_kgps` and `T_sp_C` in the last row of the CSV with `rto_model_opt_FB_kgps` and `rto_model_opt_T_C`. *Checkpoint:* About 7.1 % and 91.3 %; the two points agree to the third decimal, near (4.447 kg/s, 78.93 °C).

(b) Explain from part (a) why more runs of two-step will not close the gap.

(c) Open scenario 4, *4 · Modifier adaptation climbs to the best*, set *RTO method* to *Modifier adaptation* and run to 20 h. Read *Profit so far* to within 0.05 point. On the RTO view, which of runs 7 to 19 read lowest, and why? *Checkpoint:* 97.66 %; runs 9, 12, 15 and 18.

11.9 App challenge · 9 min

Compare the economic MPC and modifier adaptation after a change the layer is told of and after one it is not. Use the default seed, 2 h per second, and read the peak of the true waste from the truth line of the *Waste G* trend or the CSV column `truth_xG_pct`. Give every reading to within 0.05 point.

(a) Open scenario 9, *9 · Economic MPC against modifier adaptation: a feed rise and a price cut*, as shipped and run to 24 h. Read *Profit so far*, pausing by 24.5 h, and the peak of the true waste after 12 h. Press Reset, set *RTO method* to *Modifier adaptation* and run to 24 h again. *Checkpoint:* 98.66 % and 8.59 %; 98.83 % and 9.33 %.

(b) In the two exported files, find the rows after 12 h where `FB_sp_kgps` or `T_sp_C` first change. Which of those moves could have used the new A feed? *Checkpoint:* Economic MPC from 12.10 h; modifier adaptation at 12.05 and 13.05 h, of which only the second.

(c) Open scenario 7, *7 · Plant drift: the RTO retracks*, set *RTO method* to *Economic MPC* and run to 24 h. Then reload the case, which runs modifier adaptation, and run to 24 h. Read the same two numbers, with the peak after 6 h, and explain the difference from part (a). *Checkpoint:* 97.72 % and 8.17 %; 98.50 % and 8.15 %.

11.10 Application area · 3 min

At the benchmark's prices the best permitted point earns 178.54 \$/s, and two-step stops at 165.74 \$/s, 92.83 % of it. Modifier adaptation averages 98.96 % of the best on the running plant of case 4.

(a) What is the two-step gap worth per hour and per day? *Checkpoint:* About \$46 100 an hour and \$1.11 million a day.

(b) What does modifier adaptation's perturbation cost per hour against the best, and what does it recover per hour against two-step? Name one reason from Section 11.8 why a plant's layer would recover less. *Checkpoint:* About \$6700 and \$39 400 an hour.

11.11 Application area · 4 min

Section 11.8 assumes an unbiased analyzer. Suppose the G channel reads 0.3 point high under modifier adaptation with no back-off. Near the best point the plant's best profit changes by about 10 \$/s for each point the waste limit moves.

(a) Where does the true waste settle, and what does the bias cost? *Checkpoint:* About 7.7 %; about 3 \$/s.

(b) If the channel read 0.3 point low instead, what would happen, and which setting of the layer and which practice protect the permit?

11.12 Review the agent · 5 min

An AI assistant reviews the challenge and writes: "Modifier adaptation hourly with $K = 0.3$ and no back-off passes at the default seed with 98.19 %, more than the suggested setting's 97.47 %. The 0.15-point back-off only costs profit. Ship $K = 0.3$ with $b = 0$." Find the flaw in the recommendation and say how you would check it.