INTERACTIVE CONTROL ENGINEERING
Principles and Applications

Volume 10 Control in Upstream and Midstream Oil and Gas

CHAPTER 12

A Gas Transmission Line Simulator for Line Pack, Forecast Feedforward and Leak Detection

Pages 276–297

Edited by John D. Hedengren

How to cite this chapter. "A Gas Transmission Line Simulator for Line Pack, Forecast Feedforward and Leak Detection," in *Control in Upstream and Midstream Oil and Gas*, John D. Hedengren, Ed., vol. 10 of *Interactive Control Engineering: Principles and Applications*. Advanced Process Solutions, LLC, 2026, ch. 12, pp. 276–297.

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

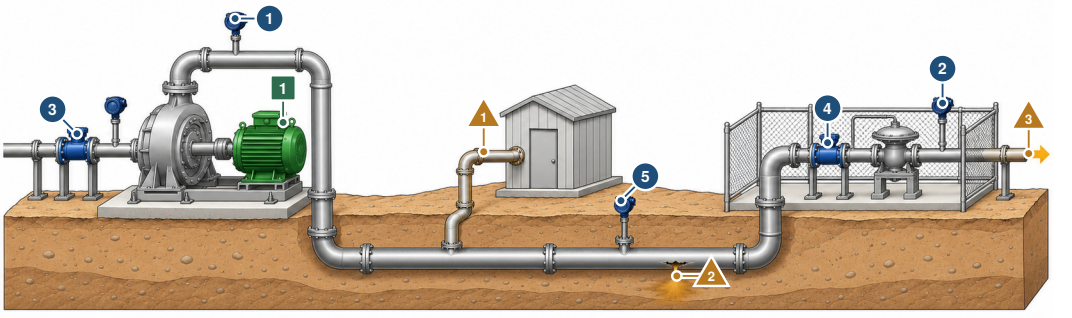
The app for this chapter
https://apmonitor.com/apps/gas-pipeline/

The volume online: contents, other chapters and citation records
https://apmonitor.com/apps/books/v10-oil-gas/#gas-pipeline

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

CHAPTER 12

A Gas Transmission Line Simulator for Line Pack, Forecast Feedforward and Leak Detection



The controller moves 1 compressor speed, through the motor's drive

The controller reads 1 pressure where the gas leaves the station 2 pressure at the city gate
3 gas metered in, for the leak detector 4 gas metered out, for the leak detector
5 pressure at the middle of the line, for the line pack

What disturbs it 1 gas taken by a side branch 2 a leak 3 what the city takes

Notice that the controller moves the compressor on the left but must hold the pressure at the city gate on the right, with a long buried pipe of stored gas between them that only the midpoint transmitter looks into. Illustration, not to scale.

IN PLAIN TERMS A Pipe That Is Also a Tank

The main idea. A long gas pipeline also stores gas, and the higher its pressure, the more it holds. A machine called a compressor pushes gas in at one end, and a city takes it out at the other.

The challenge. A push at the start takes most of an hour to reach the city, so waiting for the city's pressure to fall is too late. No meter weighs the stored gas. So a leak is found by comparing gas in with gas out, and a meter that reads a little low looks like a leak.

Why it matters. The city must get its gas above a promised pressure, without wasted power or false alarms.

THREE WAYS TO READ THIS CHAPTER

Explore	12 min	This page, the History box in Section 12.2, the Did you know? boxes in Sections 12.3 and 12.5, the Try it boxes in Section 12.5, and Key points in Section 12.8. No equations.
Study	45 min	The whole chapter in order, stopping at the three pause points, then Exercises 1 to 12.
Practice	20 min	This page, Section 12.5 with Figure 12.4, Section 12.6 and Section 12.7 with its Open challenge box.

The control question A transmission line stores its own gas and takes most of an hour to pass a station change to the city gate, which must stay above a contract floor. What must the station's control see, and how small a leak can a mass balance on three pressure transmitters find?

Learning objectives

- Compute the steady pressure profile, the line pack and the three time scales of a gas transmission line from its friction law.
- Explain why a station loop on its discharge leaves the delivery short, and why a cascade cannot avoid the dip.
- Use a demand forecast as feedforward and trade the pressure margin it saves against compression cost and the lead time.
- Predict the detection time and the smallest detectable leak of a windowed mass balance, and explain its false alarms.
- Judge what a predictive controller that carries the floor as a constraint buys over a cascade with a higher setpoint.

Prerequisites Sections 1.4.5 (line pack), 1.4.9 (inferred inventories and leak detection), 1.6.2 (the framework's PI) and 1.6.4 (the apps). Section 1.6.3 is needed only for the predictive comparison. Chapter 11 describes the station at the head of such a line.

The app https://apmonitor.com/apps/gas-pipeline/

See also Volume 3, Chapter 2 for loop structure matched to where a disturbance enters, Volume 3, Chapter 6 for feedforward hours ahead of the feedback, Volume 6, Chapter 11 for an inventory in transit ordered from a forecast, and Volume 5, Chapter 11 for a quantity inferred from the measurements around it.

12.1 Introduction

Gas pipelines are usually drawn as a pipe with a pressure drop, which is correct only at steady state. Between steady states a transmission line is a long, pressurized storage vessel. Every bar of pressure along its length holds gas, the inventory rises and falls through the day, and that stored gas carries the city through the morning ramp while the compressor catches up [1]. The transient behavior of such lines has been modeled for decades, from implicit and explicit schemes for networks to reviews of compressible flow in long pipelines [2–4].

The control problem follows from the storage. The delivery pressure at the city gate is the quantity the contract names, but the compressor acts at the other end of a line whose response is a slow diffusion of pressure rather than a transport delay or a first-order lag. A station that holds its own discharge pressure keeps its own number and loses the customer's. A delivery loop above it has to act through tens of minutes of line response, and demand for gas is largely predictable, which makes the forecast worth more than any retuning [5]. Scheduling of compressors against forecast demand, electricity price and pressure limits is the central optimization problem of gas transmission and has a long literature of its own [6–8].

The misconception the simulator targets is that a pipeline is a pipe. The broader class of problem is a distributed inventory controlled from its boundaries and measured at a few points: water mains, oil lines, district heating networks and irrigation canals all share it. Volume 9, Chapter 5 releases water ahead of demand along a canal, Volume 9, Chapter 6 sends heat ahead of a morning peak, and Chapter 13 runs a liquid line with a leak detector of its own. Three consequences carry across the class. A far-end quantity needs its own loop above the local one. A forecastable load belongs in feedforward. An inventory that is not measured has to be inferred, and the inference is wrong during exactly the transients in which a leak detector most needs it [9, 10].

This chapter studies a transmission line simulator in which the station loop, the delivery loop, the forecast, the pack window and the leak detector can each be changed while a winter day runs, and the result read as time under the contract floor, compression cost and alarms. The flow model, cascade control, forecast feedforward and computational pipeline monitoring are all established methods. The

simulator joins them in a method-of-lines model with a compressor map and a leak detector that infers the line pack from three transmitters, as a simple real system does, and its scenarios run from a manual speed step through a full day to the detection limits of a mass balance. The application also offers a model predictive controller (MPC) that sets the discharge setpoint in place of the delivery loop, a method with its own literature on gas networks [6, 11], and its measured comparison with the cascade is reported here.

Section 12.2 traces transmission operations from telephone dispatch to transient simulation and leak detection. Section 12.3 states the twenty-cell line model, the compressor, the demand and the leak. Section 12.4 describes the cascade, the forecast feedforward, the pack window, the mass-balance leak detector and the predictive controller. Section 12.5 measures the scenarios, from the line as a tank to the floor carried as a constraint. Section 12.6 sets the hardware of a real line beside the model, Section 12.7 lists what the model leaves out, Section 12.8 sums up, and the exercises in Section 12.9 put the scenarios to work.

12.2 Gas Transmission Operations

Long-distance gas transmission grew with welded large-diameter pipe and centrifugal compressors in the middle of the twentieth century. Early lines were run by station operators who held discharge pressures by hand, with dispatchers setting those pressures by telephone from a few readings at the delivery end. The flow was steady enough, and the pipe forgiving enough, that the transient behavior of the line could be left to experience.

Numerical transient simulation changed that, and the optimization of compressor operation grew up beside it [3, 6], as the History box tells.

> ⧖ HISTORY **1971 to 2015: Optimizing a Line That Is Never Steady**
>
> Computing how a gas line moves between steady states came first. A technique of 1984 calculated the transient flow in high-pressure transmission systems, with and without compressors, from a partial differential equation and a numerical scheme for it [2]. Optimizing that motion proved harder. Hierarchical control of transient flow in gas networks was proposed in 1971, and an algorithm of that kind from 1998, tested on part of the British transmission network over 24 hours in steps of 2 hours, could not solve the day exactly in reasonable time [7]. For a long time the operators could plan a day as if consumption stayed constant, because cities and industry drew gas predictably on day-ahead contracts, and early studies optimized steady flows. By 2015 that assumption no longer held, and the transient problem was posed as optimal control of the whole network [8].

Leak detection developed alongside. The mass-balance method compares metered inflow and outflow with the change in computed inventory, and model-based methods for pipelines were formalized in the 1980s [9]. Its limits are the meter accuracy and the inventory estimate, and a survey of gas leak detection places it among hardware-based and acoustic alternatives with different sensitivities and detection times [10]. Section 1.4.9 places the computed inventory among the other inferred and late measurements of this volume. Line pack itself has more recently been treated as a quantity of its own, with dynamic depletion models for networks under stress [1].

12.3 Process Model

A single compressor station takes gas at 40 bar and pushes it through 100 km of 0.9 m pipe to a city gate that must never fall below 45 bar. At a steady 210 kg/s with the delivery at 48 bar, the station discharges at 60.35 bar, the compressor runs at 73.2 % speed and 17.3 MW, the line holds 2825 t of gas, and the gas itself takes about 3.7 h to cross. The arrangement is shown in Figure 12.1.

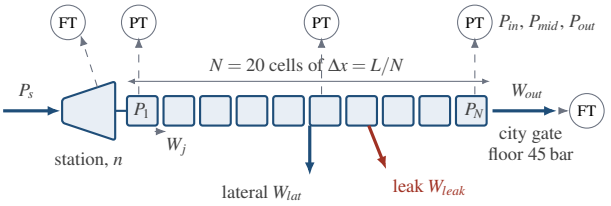


Figure 12.1 Transmission line model. A compressor at speed n takes gas at suction pressure P_s and discharges into the first of N pressure cells of length Δx, joined by faces carrying mass flows W_j. A lateral at 50 km takes W_{lat}, a leak at 60 km loses W_{leak}, and the city gate delivers W_{out} as long as P_N allows. Three pressure transmitters PT at the inlet, midpoint and outlet and two flow meters FT feed the leak detector.

Table 12.1 Variables at the Design Operating Point

Symbol	Role	Nominal	Range
P_{out}	CV, delivery pressure	48 bar	SP 45–60
P_{in}	CV (inner), discharge	60.35 bar	$\leq$ MAOP 72
$P_{SP,dis}$	MV of outer loop	60.35 bar	20–70
n	MV, compressor speed	73.2 %	50–100
$P_1 \ldots P_{20}$	states, cell pressures	—	≥ 1 bar
$W_1 \ldots W_{19}$	states, face flows	210 kg/s	—
W_{dem}	DV, city demand	140–285 kg/s	daily profile
W_{lat}	DV, lateral takeoff	0 kg/s	0–60
W_{leak}	DV, leak at 60 km	0 %	0–10
b_{in}, b_{out}	DV, meter biases	0 %	± 3

> ✦ **DID YOU KNOW? Three Million Miles of Pipe**
>
> In 2022 about 3 million miles of mainline and other pipelines in the United States delivered about 29.2 trillion cubic feet of natural gas to about 78.3 million consumers. About half of the mainline transmission network was laid in the 1950s and 1960s, after demand more than doubled following World War II [12].

The variables are classified in Table 12.1. Forty states are integrated: twenty cell pressures, nineteen internal face flows and the compressor speed.

12.3.1 Mass and Momentum on Twenty Cells

Isothermal one-dimensional flow is discretized by the method of lines. Each cell stores gas in proportion to its pressure, and each internal face carries a flow driven by the pressure difference and opposed by friction, Equations (12.1) and (12.2).

$$\frac{A\,\Delta x}{c^2}\frac{\mathrm{d}P_i}{\mathrm{d}t} = W_i - W_{i+1} - W_{lat}\,\delta_{i,i_{lat}} - W_{leak}\,\delta_{i,i_{leak}} \tag{12.1}$$

$$\frac{\Delta x}{A}\frac{\mathrm{d}W_j}{\mathrm{d}t} = P_{j-1} - P_j - \frac{f\,c^2\,\Delta x\,W_j|W_j|}{2DA^2\,\bar{P}_j} \tag{12.2}$$

where A is the pipe area, $c = \sqrt{ZRT} = 350\,\mathrm{m/s}$ the isothermal wave speed, f the Darcy friction factor, D the diameter and $\bar{P}_j$ the mean of the two neighboring cell pressures. The inlet face carries the compressor flow and the outlet face the delivered flow. In steady flow the square of the pressure falls linearly along the line, as shown in Equation (12.3).

$$P_{in}^2 - P_{out}^2 = K\,W^2, \qquad K = \frac{f\,L\,c^2}{DA^2} \tag{12.3}$$

Table 12.2 Line and Station Parameters, Marking Those the Interface Hides

Symbol	Meaning	Value	Class
L, D	length, diameter	100 km, 0.90 m	visible
f	Darcy friction factor	0.0095	visible
c	isothermal wave speed	350 m/s	visible
MAOP	pressure limit	72 bar	visible
—	contract floor	45 bar	visible
—	curtailment onset	38 bar	visible
P_s	suction pressure	40 bar	visible
W_r, r_{max}	rated flow, ratio	300 kg/s, 1.90	visible
α	map curvature	0.40	hidden
τ_N	speed response	20 s	visible
—	speed ramp limit	1.0 %/s	visible
η	compressor efficiency	0.80	visible
x_{lat}, x_{leak}	lateral, leak location	50, 60 % of L	visible

Between the first and last cell centers, 95 km apart, the discretized line gives $(N-1)KW^2/N$, which sets the design discharge pressure. The line pack is $M = \sum_i A\,\Delta x\,P_i/c^2$, the cell-by-cell form of Equation (1.8). Three time scales follow from Equations (12.1) and (12.2). The acoustic wave crosses in under five minutes, friction damps it within about half a minute of travel, and what reaches the outlet is a diffusion of pressure with a time scale of order $L^2 f|W|/(2DA\bar{P})$, about 53 min at design with the mean line pressure of 54.4 bar, and shorter at low flow. The gas velocity is a few meters per second, so the gas crosses in about 3.7 h.

Pause. Try Exercise 12.2 before reading on.

12.3.2 Compressor, Delivery and Leak

The compressor map gives the pressure ratio $r = P_1/P_s$ as a function of speed fraction $s = n/100$ and flow, and inverting it gives the flow the station delivers against the ratio the line imposes, Equation (12.4).

$$W_{in} = W_r \sqrt{\frac{s^2 - (r-1)(1-\alpha)/(r_{max}-1)}{\alpha}} \tag{12.4}$$

with rated flow $W_r = 300\,\mathrm{kg/s}$, ratio $r_{max} = 1.9$ at full speed and rated flow, and curvature $\alpha = 0.4$. The flow falls as the ratio rises, which is the only self-regulation the line has. Surge is flagged below $0.5\,W_r s$. The shaft power is $W_{in} c_p T_s (r^{0.3/1.3} - 1)/\eta$ with $c_p = 2.3\,\mathrm{kJ/(kg\,K)}$, $T_s = 288\,\mathrm{K}$ and $\eta = 0.80$, and the compression cost integrates it against a daily electricity price of 35 to 140 \$/MWh. The speed follows its command as a 20 s lag limited to 1 %/s.

The city takes its demand as long as the pressure allows. Below 38 bar the delivery is curtailed linearly to zero at 33 bar. The demand follows a winter weekday profile with a morning peak near 270 kg/s at 08:00 and an evening peak near 285 kg/s at 19:00, scaled by a demand factor and multiplied by a slow random wander of 2 % with a 600 s time constant. The leak is a fixed opening, $W_{leak} = (\ell/100)\,(210\,\mathrm{kg/s})\sqrt{P/60\,\mathrm{bar}}$, with ℓ in percent of design flow.

Pressure transmitters carry 0.02 bar of colored noise and the flow meters 0.3 % with their biases. The inlet pressure is flagged above the 72 bar maximum allowable operating pressure (MAOP) and the delivery below the 45 bar floor.

Table 12.2 gives the parameter values, and every scenario opens on the steady profile of Equation (12.3) for its initial demand.

12.3.3 Assumptions

Flow is isothermal with constant wave speed and friction factor. Temperature and compressibility changes shift the line pack and the profile by a few percent on a real line [3].

Twenty cells of 5 km reproduce the diffusion behavior and smear a wave front. The model is not suitable for surge or rupture transients, which need a much finer grid and the full momentum equation [4].

One compressor with a quadratic map and fixed suction pressure represents the station. Parallel units, recycle, driver limits and fuel gas consumption are omitted. Chapter 11 places parallel machines with recycle valves on shared headers ahead of such a line.

Demand follows a known profile with a slow wander, so the forecast is nearly perfect. Real forecast errors are larger and weather-driven [5].

The city network behind the gate is not modeled, and delivery is curtailed by a linear rule below 38 bar.

12.4 Control and Monitoring

The controllers run every 10 s and the physics advances by fourth-order Runge-Kutta steps of 2 s, on the engine of Section 1.6.1. Both loops are reverse acting in the ISA sense, because both process gains are positive, so a measurement below its setpoint raises the output ($a = +1$). Both use the velocity-form PI of Section 1.6.2, with no derivative term in either loop and the forecast of Section 12.4.2 as PIC-D's output bias u_b in Equation (1.18). Its running sum is never clipped, and at a limit only the integral is corrected, by back-calculation with a tracking time $T_t = \tau_I$, so each scan of Δt_c removes $\Delta t_c / \tau_I$ of the clipped excess, 1/240 for the delivery loop and 1/12 for the station loop. The speed lag and ramp limit act after the station loop's output and are not seen by its integral, but the loop's own 50 to 100 % range is the tighter limit.

12.4.1 The Cascade

The inner loop PIC-S holds the station discharge pressure at its setpoint by adjusting the compressor speed, with $K_c = 8$ % of speed per bar, $\tau_I = 120$ s and a speed range of 50 to 100 %. The outer loop PIC-D holds the delivery pressure by moving that setpoint, with $K_c = 1.0$ bar per bar, $\tau_I = 2400$ s and an output range from 20 bar to 2 bar below MAOP. In manual the operator sets the speed and both loops track, each setting its integral so that the output before limiting equals the value applied, the operator's speed for PIC-S and the measured discharge pressure for PIC-D. In automatic the manual speed is set to the station loop's output after every execution, so transfer is bumpless in both directions. A manual start holds the steady speed unless the reader sets another.

12.4.2 Forecast Feedforward and the Pack Window

The feedforward is the steady inlet pressure that would deliver the demand forecast a lead time T_{lead} ahead at the delivery setpoint, the static inverse of Equation (12.3), and it enters as the output bias u_b of PIC-D, Equation (12.5).

$$P_{ff}^2 = P_{SP}(t + T_{lead})^2 + K\left[W_{dem}(t + T_{lead}) + W_{lat}\right]^2 \tag{12.5}$$

The bias uses the full-length K, so at design flow it sits 0.58 bar above the 60.35 bar the twenty-cell line needs, and the integral carries the difference. The default lead is 30 min, a little over half the line's own response time, and switching the feedforward moves the integral by the change in bias so the setpoint does not bump. The pack window raises the delivery setpoint by a boost ΔP_{pack} between two hours of the day, storing gas in the line ahead of a peak.

12.4.3 The Mass-Balance Leak Detector

Every simulated minute the detector computes the line pack from the three transmitters by assuming the square of the pressure varies linearly between them, which is exact only in steady flow, Equation (12.6), where $P^2(x)$ is the piecewise linear interpolant through P_{in}^2, P_{mid}^2 and P_{out}^2 evaluated at the cell centers x_i.

$$\hat{M} = \frac{A\,\Delta x}{c^2} \sum_{i=1}^{N} \sqrt{P^2(x_i)} \tag{12.6}$$

The imbalance over a window T_w is the metered inflow less the metered outflow and the lateral, less the change in computed pack, expressed in percent of 210 kg/s, Equation (12.7), the form of Equation (1.15) with the lateral metered.

$$I = \frac{100}{210\,T_w}\Big[\int_{t-T_w}^{t} (\hat{W}_{in} - \hat{W}_{out} - W_{lat})\,dt \\ -\hat{M}(t) + \hat{M}(t-T_w)\Big] \tag{12.7}$$

An alarm is raised when $|I|$ exceeds the threshold for the persistence time, and cleared when it falls below 70 % of the threshold. The defaults are a 30 min window, a 2 % threshold and a 10 min persistence. Until a full window of samples exists the balance uses the shorter record it has, so the first half hour of a run is the noisiest. A meter bias enters I one for one, and the error in $\hat{M}$ during a ramp enters as a false imbalance.

12.4.4 Model Predictive Control

The delivery control has a second setting, never the default and pinned only by scenario 10, in which a model predictive controller (MPC) writes the discharge setpoint in place of PIC-D. Every $T_s = 10$ min it plans the setpoint over $N_p = 72$ steps (12 h), and only the first move of each plan reaches the line, the receding-horizon scheme of industrial predictive control [13]. The move is ramped over its period, the applied setpoint on scan $j = 0,\ldots,59$ being $u_a + (u_1 - u_a)(j+1)/60$, with u_a the setpoint when the plan was made and u_1 its first value, and the plan models the ramp exactly. PIC-S still turns the setpoint into speed. The setpoint is clamped at every scan to PIC-D's range, 20 bar to MAOP less its 2 bar margin as the run sets them, so MAOP stays a protection above the plan, and PIC-D tracks the applied setpoint with no preload. Section 1.6.3 describes the MPC option of the series in general. The controller here is written for this line, with a nonlinear model and its own estimator.

The controller's model is the line's own mass and friction balance reduced to five nodes and written from the declared parameter defaults of Table 12.2, never from the run's values, as shown in Equation (12.8).

$$\begin{aligned} \tau_S \frac{\mathrm{d}p_0}{\mathrm{d}t} &= u + d_{in} - p_0 \\ C_j \frac{\mathrm{d}p_j}{\mathrm{d}t} &= Q_j - Q_{j+1} - \delta_{j,2}\,(W_{lat} + d_{mid}), \quad j = 1,\ldots,4 \\ Q_j &= \frac{\Delta_j}{\sqrt{k_j\left(|\Delta_j| + k_j W_\varepsilon^2\right)}}, \qquad \Delta_j = p_{j-1}^2 - p_j^2 \\ Q_5 &= \hat{W}_{dem} + d_{out} + d_{pass} \end{aligned} \tag{12.8}$$

where p_0 to p_4 are pressures in bar at the centers of cells 1, 6, 11, 16 and 20 (2.5, 27.5, 52.5, 77.5 and 97.5 km), so that p_0, p_2 and p_4 sit at the three transmitters, u is the discharge setpoint, $\tau_S = 66$ s is the closed-loop lag of the station loop identified from setpoint steps on the plant at design flow, Q_j is the flow through segment j, $k_j = Kn_j/N$ with n_j equal to 5, 5, 5 and 4 cells and $K = 0.031\,95$ bar^2 per (kg/s)^2 from Equation (12.3), $W_\varepsilon = 5$ kg/s keeps the flow law smooth through zero flow and changes a 100 kg/s flow by 0.13 %, C_j is the cell capacity $A\,\Delta x/c^2 = 2597$ kg/bar times the node's share of 3, 5, 5, 4.5 and 2.5 cells, $\delta_{j,2}$ places the lateral at the midpoint node, $\hat{W}_{dem}$ is the city's nomination, and d_{in}, d_{mid}, d_{out} and d_{pass} are the disturbances of the estimator below. In steady flow the model's squared pressures fall linearly as in Equation (12.3), and at design it needs 60.36 bar at the inlet for 48 bar delivered, against 60.35 bar on the twenty-cell line.

Stepped by 1 bar up and down with the delivery loop frozen, the plant and the model have steady gains of 1.2572 and 1.2573 at 210 kg/s and 1.4236 and 1.4233 at 280 kg/s, 63 % times of 54 and 55 min and of 68 and 70 min, and RMS differences of 0.66 and 0.68 % of the response, 6.9 % at 120 kg/s, where the plant's falling step is not linear. One identified lag would not serve, because the

gain rises 1.3 times from 120 to 280 kg/s (case 4). Run open loop from its hourly estimate, the model predicts the delivery 1 to 12 h ahead within 0.02 to 0.06 bar RMS with the wander and the noise removed and within 0.14 to 0.51 bar RMS as shipped.

The estimator is a steady-state Kalman filter on the three pressure transmitters, and no flow meter enters it. The model is augmented with a station bias d_{in} and unexplained offtakes d_{mid} at the midpoint node and d_{out} at the city gate, all three random walks, and with a passing offtake d_{pass} at the city gate that the plan lets decay with $\tau_W = 1$ h, so that a swing of the wander is not held for the whole horizon. As many integrating disturbances as transmitters absorb a persistent mismatch at each measured point, which makes the tracking offset-free [14]. Each scan predicts with the nonlinear model and corrects, as shown in Equation (12.9).

$$\begin{bmatrix} \hat{\mathbf{p}}_k \\ \hat{\mathbf{d}}_k \end{bmatrix} = \begin{bmatrix} \mathbf{F}\left(\hat{\mathbf{p}}_{k-1}, u_{k-1}, \mathbf{w}_{k-1}, \hat{\mathbf{d}}_{k-1}\right) \\ \mathbf{E}\,\hat{\mathbf{d}}_{k-1} \end{bmatrix} + \mathbf{K}_f\left(\mathbf{y}_k - \mathbf{H}\mathbf{F}_k\right) \tag{12.9}$$

where $\hat{\mathbf{p}} = (p_0, \ldots, p_4)$ and $\hat{\mathbf{d}} = (d_{in}, d_{mid}, d_{out}, d_{pass})$ are the estimates, $\mathbf{F}$ integrates Equation (12.8) over one scan Δt_c with the setpoint and the nominations $\mathbf{w} = (\hat{W}_{dem}, W_{lat})$ that held over it, $\mathbf{F}_k$ is that prediction, $\mathbf{H}$ selects p_0, p_2 and p_4, $\mathbf{y}_k = (P_{in}, P_{mid}, P_{out})$ is the transmitter reading, $\mathbf{E} = \mathrm{diag}(1, 1, 1, e^{-\Delta t_c/\tau_W})$, and $\mathbf{K}_f$ is the 9×3 steady-state gain of the model linearized at design flow and sampled at Δt_c, with per-scan variances of 10^{-6} bar^2 on each node, 10^{-4} bar^2 on d_{in}, 10^{-3} (kg/s)^2 on d_{mid} and d_{out} and 0.8 (kg/s)^2 on d_{pass}, and 4×10^{-4} bar^2 on each transmitter, the square of its 0.02 bar noise. The gain, from 110,318 Riccati iterations, is stored, and the filter runs under either structure so that the MPC, once selected, starts from a current estimate.

The preview $\mathbf{w}$ is the dispatcher's nominations, the day's demand profile on the demand scale (or the flat demand) without its wander and the lateral, from the function that gives the feedforward of Equation (12.5) its forecast, read at the middle of each step, and with the preview off the present nomination is held. The realized demand, the wander, a leak and the meters are never read, the challenge's lateral at 10:00 is absent from the 09:50 plan, and a change to a nomination reaches the preview when it is made, as it reaches the feedforward.

A decision is the quadratic program of Equation (12.10), with the setpoint held in twelve blocks of $1, 1, 1, 1, 2, 2, 3, 4, 6, 9, 15, 27$ steps.

$$\begin{aligned} \min_{\mathbf{u}, s} \; & \sum_{k=1}^{N_p} \left(\hat{p}_{4,k} - r_k\right)^2 + \lambda \sum_{b=1}^{12} \left(u_b - u_{b-1}\right)^2 \\ & + w_E \sum_{k=1}^{N_p} e_k\, u_{b(k)} + \rho\, s^2 + \sigma s \\ \text{s.t. } & 20\,\text{bar} \le u_b \le P_{MAOP} - 2\,\text{bar}, \quad s \ge 0 \\ & \hat{p}_{4,k} \ge P_{fl} + m_k - s \\ & m_k = 0.3 + 0.5\left(1 - e^{-kT_s/\tau_m}\right) \end{aligned} \tag{12.10}$$

where $\mathbf{u} = (u_1, \ldots, u_{12})$ are the block values of the setpoint, $\hat{p}_{4,k}$ is the predicted delivery pressure at the end of step k, u_0 is the setpoint applied now, $u_{b(k)}$ is the block that holds step k, r_k is the dispatcher's delivery setpoint at that time with the pack window included, raised to $P_{fl} + m_k$ where it lies below so that no target pulls against the floor, $\lambda = 8$ weighs a squared bar of setpoint move against a squared bar of delivery error at one step, e_k is the compression energy in MWh that one more bar of discharge pressure costs over step k at the nominated flow, from the power law of Section 12.3 with the declared suction pressure and efficiency, $w_E = 0.05$ bar^2 per MWh makes that energy a tie-breaker, $\rho = \sigma = 1000$ price the slack s in bar, $P_{MAOP} = 72$ bar, $P_{fl} = 45$ bar is the contract floor, and m_k is a margin in bar with $\tau_m = 1$ h, worth 0.38 at the first step, 0.62 at 1 h, 0.78 at 3 h and 0.80 at 12 h. The margin grows with the look-ahead because the forecast's error, the wander, grows with

it. Where the delivery is already below $P_{fl} + 0.3$ bar the floor row is lowered by the shortfall times $e^{-kT_s/45\,\mathrm{min}}$, a funnel back to the band. The bounds on u are hard and the floor is soft, with one slack over the whole horizon because the shared solver carries one per output, not one per step.

The prediction is the nonlinear model run from the estimate under a nominal plan, with the passing offtake decaying, plus the departure from that plan through the model linearized at the estimate and sampled exactly over a ramped step. Two passes are made, around the last plan moved on by one decision and then around the first pass's result. Over scenario 10's 144 plans the program's prediction under the accepted plan departs from the nonlinear model's by a median of 0.0095 bar and at most 0.105 bar after the priming plan, which departs by 1.43 bar, and three to six passes move that plan's first move by under 0.001 bar and leave the day unchanged.

The margin's limit of 0.8 bar, λ and τ_W were fixed once, before the comparison, on case 7's day over seeds 1 to 10, which the comparison does not use, as the least mean cost ($40,443) of 60 combinations (limits of 0.6 to 1.5 bar, λ of 2 to 16, τ_W of 10 to 60 min) with no minute under the floor and no false alarm on any seed.

A plan is kept only if the solver reports an optimal status, checked even on a successful return, the arrays have the problem's shapes, the first value equals the first move, every block lies inside the clamp before rounding, and every value returned and every prediction published is finite. A specification test injects three such faults and rejects each. Otherwise PIC-D, which has tracked the setpoint, takes it over and the historian records a fallback, until a valid plan resumes on a ramp from where PIC-D left the setpoint. With the noise off, a rejection forced at 17:10 of scenario 10, the return, and switches to and from the cascade and from manual move the setpoint 0.009 to 0.014 bar, about one scan of either structure's own motion. No measured run rejected a plan, and the first plan is made at $t = 0$. In Chrome a plan takes 1.3 ms, under 0.8 % of a core at the top speed of 3600 simulated seconds per second.

The plant on the run's parameters, the noisy transmitters and the five-node model on the declared defaults are kept apart, so a reader's change to the pipe is mismatch the estimator carries, and replaying the same transmitter readings over another hidden plant, leak, meter bias and noise gives the same plans bit for bit.

12.5 Results

The scenarios are listed in Table 12.3 and the measured outcomes in Table 12.4. Section 1.6.4 describes how a scenario, its briefing and its completion rule work, and the exercises in Section 12.9 are built on these cases. Every figure quoted was measured on the application's own engine at its 2 s step with the default seed, with each setting changed before the run starts unless stated otherwise. Results that depend on a small margin were repeated on five other seeds of the demand wander and the instrument noise, and the spread is given beside them. Cases 1 to 4 use a flat demand with the wander removed so that the step is the only disturbance.

12.5.1 The Line as a Tank

A speed step in manual from the steady 73.16 % to 74.2 % reaches the delivery pressure only after about 12 min (a rise of 0.01 bar), and an hour later the delivery is 0.30 bar up and still rising at 0.35 bar/h while the inlet is 0.47 bar up. Most of the extra gas has gone into line pack. Holding the discharge constant with the station loop alone, a 30 kg/s demand increase takes the city gate under the floor 47 min later and leaves it near 43.5 bar, with the station on its setpoint at 60.35 bar within 0.02 bar. The delivery loop above it removes the offset, with a dip of about 1.8 bar to 46.2 bar 28 min after the step and a return to within 0.2 bar of setpoint after 124 min. Doubling K_c and halving τ_I raises the bottom of the dip only to 46.47 bar and overshoots to 48.59 bar, because the plant between the two loops is a diffusion that tuning cannot speed up.

🕮 **Pause.** Try Exercise 12.1 before reading on.

Table 12.3 Scenario Sequence

#	Scenario	Teaching point
1	Step the compressor	the line as a tank
2	The station loop alone	the wrong quantity held
3	The delivery loop	cascade
4	Night flow	flow-dependent dynamics
5	A winter day, feedback	margin carried by feedback
6	The same day, forecast	feedforward
7	A peak the line cannot carry	storage against a constraint
8	A meter that reads low	false alarm
9	A real leak	detection limits
10	The floor as a constraint	MPC against cascade
11	Challenge	graded day

Table 12.4 Measured Results

Case and condition	Result
1, speed to 74.2 %	outlet moves after 12 min, +0.30 bar at 1 h
2, station loop, +30 kg/s	delivery settles near 43.5 bar
3, cascade, same step	dip 1.8 bar to 46.2, recovery 2 h
4, 63% time at 120 / 210 / 280 kg/s[a]	34 / 54 / 68 min
4, gain at 120 / 210 / 280 kg/s[a]	1.09 / 1.25 / 1.42
4, cascade to 90%, night / design	83 / 58 min
5, feedback, 48 bar	low 44.3 bar, 52 min under floor, $37,300
6, forecast 30 min, 48 bar	low 47.0 bar
6, forecast 30 min, 46 bar	low 45.0 bar, $34,400
6, forecast 90 min, 46 bar	low 42.2 bar
7, demand +8 %, forecast, 46 bar	55 min under floor, discharge at limit
7, with 4 bar pack 12:00–21:00	low 45.0 bar, +$1,400, 1 false alarm
8, meter −1.5 %, threshold 1%	false alarm after 34 min
8, same, threshold 2%	no alarm in 24 h
9, leak 3% at 3 h, threshold 2 / 1%	found in 34 / 23 min
9, leak 1.5% from start, threshold 2 / 1%	not in 3 h / 32 min
9, leak 1% from start	not in 3 h at either

[a] Model test with the delivery loop off and the discharge setpoint stepped 1 bar directly, which the interface does not offer.

The line changes with flow. With the station holding a fixed discharge, the delivery answers a 1 bar change with a 63 % time of 34, 54 and 68 min at 120, 210 and 280 kg/s, and with gains of 1.09, 1.25 and 1.42, measured by stepping the discharge setpoint directly in the model. The cascade is nevertheless slower at night, 83 min to 90 % of a 1 bar setpoint rise against 57 min rising and 58 min falling at design, because the lower gain matters more than the faster response.

12.5.2 A Winter Day

Over a winter day at a 48 bar setpoint, feedback alone lets the delivery fall to about 44.3 bar on the morning ramp and spend 52 min under the floor (40 to 62 min on the other seeds), for about $37,300 of compression. With the 30 min forecast the same setpoint holds the delivery above 47.0 bar (46.7 to 47.2 bar on the other seeds), as shown in Figure 12.2, and the setpoint can be lowered to 46 bar with the lowest delivery at 45.01 bar, for about $34,400, 8 % less. That margin is 0.01 bar, and on two of the five other seeds the same day spends 25 and 38 min under the floor. The forecast reduces the margin feedback must carry, and the margin is paid for in compression all day. At 46 bar a lead of 15 min gives a lowest delivery of 45.15 bar, and 45 min already falls to 44.41 bar. A lead of 90 min is too long, because the setpoint then moves ahead of the line's own response and the delivery sags to about 42.2 bar.



Figure 12.2 The transmission line running scenario 6, a winter day with the demand forecast as feedforward at a 48 bar delivery setpoint, at 23:59, cropped to the plant view. Callout 1 marks the pressure along the line, with MAOP and the contract floor dashed and the three leak detector transmitters as dots. Callout 2 is the compressor with its speed, and the pipe runs from it to the city gate at callout 3 with tracers moving at the gas velocity. Callout 4 sets the true line pack against the pack computed from three transmitters, and callout 5 the leak detector imbalance against its threshold, beside the clock with the electricity price. The run's scorecard, cut from this view, read a lowest delivery of 47.05 bar, no time below the floor, a highest inlet of 70.03 bar, no leak alarms and $36,412 of compression.

> **TRY IT IN THE APP** **5 · A winter day, feedback only**
>
> Load the case and press Run at a high speed. Touch nothing. Watch the Delivery pressure trend against its floor line through the morning and evening ramps. On the morning ramp the delivery falls to about 44.3 bar, and over the day it spends 52 min below the floor. Read the Time below the floor and Compression cost tiles at the end of the day.

12.5.3 A Peak the Line Cannot Carry

With demand raised 8 %, the evening peak of 308 kg/s equals the line's steady capacity with 70 bar in and 45 bar out, and the demand wander takes it past that. With the forecast and a 46 bar setpoint, the discharge setpoint sits at its 70 bar limit for 76 min and the delivery still spends 55 min under the floor, with a low of 44.27 bar. No tuning changes that outcome, because the constraint and not the controller sets it. A 4 bar pack from 12:00 to 21:00 stores gas in the line during the afternoon and raises the evening low to 45.58 bar for about $1,400 more compression, and the pressure ramp it causes sets off one false leak alarm at 12:40. The lowest delivery of the packed day, 45.00 bar, then comes at 21:55, after the window closes and the setpoint falls back 4 bar. The capacity-limit comparison is within the spread of the wander: on the five other seeds the unpacked day spends 0 to 60 min under the floor and the packed day 8 to 32 min, on the morning ramp or after the window closes.

Pause. Try Exercise 12.3 before reading on.

Table 12.5 Challenge Outcomes Over One Day

Configuration	Low (bar)	Under floor	False	Cost ($)
as loaded	44.33	52 min	7	42,758
threshold 2.5%	44.33	52 min	0	42,758
forecast, 48 bar, 2.5%	46.05	0	0	41,917
forecast, 47 bar, 2.5%	45.62	0	0	40,934

12.5.4 What the Mass Balance Can Find

The leak detector is as good as its meters and its inventory estimate. An outlet meter reading 1.5 % low with a 1 % threshold produces an alarm 34 min later with no leak at all (30 to 34 min on the other seeds), and a 2 % threshold does not alarm in a day. A 3 % leak at 60 km is found 34 min after it starts at a 2 % threshold and 23 min after at 1 % (31 to 37 and 18 to 23 min on the other seeds). The leak sizes are set at 60 bar, so at the 52.6 bar of the leak site the 3, 1.5 and 1 % leaks lose 2.81, 1.41 and 0.94 % of design flow. Because the scripted leak at 3 h replaces whatever the slider holds, the smaller leaks are started at $t = 0$. A 1.5 % leak is not found in the three hours before the scripted leak at 2 %, and at 1 % the alarm comes 32 min after it starts (11 to 47 min on the other seeds). A 1 % leak is not found in three hours at either. Each of these times includes the 10 min persistence, and by its alarm the 3 % leak has released 12.1 t of gas at the 2 % threshold and 8.2 t at 1 %. The threshold that avoids the false alarm from the biased meter is the threshold that misses the 1.5 % leak, which is the trade the method is known for [9, 10].

> **TRY IT IN THE APP** **9 · A real leak, and how long the balance takes**
>
> Load the case and press Run. At three hours a leak of 3 % of design flow opens at 60 km. Watch the Leak detector imbalance trend climb toward its threshold, then read the Leak found after tile: about 34 min. Press Reset, set Leak alarm threshold to 1 % before Run, and run again. The alarm now comes after about 23 min.

12.5.5 The Graded Challenge

The challenge runs a winter day with a gas turbine plant starting on the lateral at 10:00 and the outlet meter reading 1 % low, starting with feedback only at 48 bar and a 1 % leak threshold. It is met with the delivery never under the floor, the inlet never over MAOP, no false alarms and a compression cost of $41,500 or less, on the case as given: changing a model parameter or a disturbance, or firing an extra disturbance event, voids the attempt. Table 12.5 reports the measured outcomes. As loaded, the day fails on three counts, with 52 min below the floor, seven false alarms and $42,758, $1,258 over budget.

A higher threshold alone removes the seven false alarms and nothing else. The forecast at 48 bar removes the floor violations but costs $41,917, over the target, and only the combination of the forecast with the setpoint lowered to 47 bar meets every condition, at $40,934, $566 under budget, with a lowest delivery of 45.62 bar and a highest inlet of 70.08 bar. At a 2 % threshold the same day raises two false alarms, at 07:48 on the morning ramp and at 10:40 after the lateral starts. On the five other seeds the suggested configuration passes four times and fails once, on one false alarm at 10:40 at 2.5 %.

12.5.6 The Floor as a Constraint

Scenario 10 runs case 7's day (demand 8 % above its profile, a 46 bar setpoint, the forecast on) under the MPC. Table 12.6 compares it with the cascade on the same forecast and with variants that locate the difference, and sets the two structures side by side on cases 5 and 6 and the challenge, at the default seed and on seeds 11 to 30, each a full day with the same transmitters, noise and nominations for both, and with the MPC on scenario 10's settings throughout.

Table 12.6 Model Predictive Control Against the Cascade Over Full Days, at the Default Seed and on Seeds 11 to 30

	Default seed							Seeds 11 to 30			
Structure	Under (min)	Lowest (bar)	Cost ($)	False alarms	IAE (bar h)	Travel (%)	Plans	Under, min mean (worst)	Days under	Cost ($)	False alarms
Case 10, case 7's day, demand 108 %, 46 bar											
cascade + forecast (case 7)	54.6	44.27	40,213	0	7.65	306	—	20.2 (75.6)	16	39,938	0
cascade + forecast at 46.4 bar	21.9	44.44	40,593	0	7.93	303	—	1.5 (25.8)	3	40,337	0
cascade + forecast at 47.2 bar	9.6	44.76	41,347	0	8.61	297	—	—	—	—	—
+ pack 4 bar 12:00 to 21:00 (case 7)	0.0	45.00	41,633	1	14.71	417	—	22.4 (53.4)	17	41,441	19
+ pack 2 bar 16:00 to 19:00, hindsight[a]	0.0	45.11	40,294	0	8.24	338	—	31.8 (98.5)	18	40,043	0
cascade, feedback only	333.6	41.57	41,315	0	31.79	306	—	343.8 (365.1)	20	41,033	0
MPC (scenario 10)	10.4	44.73	40,623	0	9.54	217	144	1.4 (25.2)	2	40,381	0
MPC, preview off	301.6	42.79	41,690	0	30.64	247	144	269.6 (326.1)	20	41,449	1
MPC, 1 h horizon	24.7	44.40	41,372	0	16.90	232	144	5.1 (29.6)	5	41,102	0
MPC without the floor row[b]	96.5	44.30	40,103	0	10.58	187	144	63.7 (162.4)	19	39,820	0
Case 5, 48 bar, no forecast											
cascade, feedback only	51.6	44.33	37,330	0	26.56	298	—	54.7 (86.3)	20	37,001	0
MPC, preview off	48.0	44.09	37,417	0	31.06	213	144	61.7 (112.4)	20	37,138	0
Case 6, 48 bar											
cascade + forecast	0.0	47.05	36,418	0	5.78	294	—	0.0 (0.0)	0	36,093	0
MPC	0.0	46.64	36,245	0	9.17	178	144	0.0 (0.0)	0	35,979	0
Case 6, 46 bar											
cascade + forecast	0.0	45.01	34,396	0	5.89	299	—	9.5 (32.4)	11	34,069	0
MPC	0.0	45.26	34,762	0	8.64	218	144	0.4 (7.3)	1	34,431	2
Challenge (case 11), as loaded											
cascade, feedback only	51.6	44.33	42,758	7	26.82	285	—	54.7 (86.3)	20	42,524	128
MPC	0.0	46.13	41,734	7	10.69	180	144	0.0 (0.0)	0	41,507	133
Challenge (case 11), suggested settings[c]											
cascade + forecast	0.0	45.62	40,934	0	7.85	324	—	0.0 (0.0)	0	40,659	7
MPC	0.0	45.62	40,782	0	10.01	184	144	0.0 (0.0)	0	40,517	1

Under is the time with the true delivery below the 45 bar floor, counted at every 2 s step, and false alarms are leak alarms with no leak. IAE is the integrated absolute delivery error and travel the total travel of the compressor speed, as the application's loop metrics report them. Each MPC row makes 144 plans a day and rejects none, at a mean of 1.3 ms a plan in Chrome. [a]The cheapest of 288 pack windows on this day, chosen knowing the outcome. [b]A copy of the controller with the floor constraint removed, tracking the 46 bar setpoint with the same forecast. [c]Forecast with a 30 min lead, 47 bar and a 2.5 % threshold, with the MPC selected in its row. The challenge is met in both suggested rows and in neither as-loaded row, and on seeds 11 to 30 with the suggested settings on 13 days under the cascade and 19 under the MPC.

On the shipped day the MPC spends 10.4 min under the floor against 54.6 min for the cascade, for $40,623 against $40,213, and on the further days it goes under on 2 against 16. The floor in its problem raises the setpoint sooner, as shown in Figure 12.3. At 17:40 its setpoint is at the 70 bar clamp and the cascade's at 68.1 bar, with 3029 t against 2983 t in the line, but the extra gas is modest, 30 t at 17:00 and 43 t at 18:00. The wander decides the evening, taking the demand to 327.6 kg/s at 19:06 against a nominated 307.8 kg/s, and both deliveries are lowest at 19:08, 44.73 bar under the MPC and 44.27 bar under the cascade. With the wander removed the two days stay at or above 45.71 and 45.32 bar.

The ablations credit the gain to the floor carried as a constraint, not to packing as a plan. Without the floor row the same controller tracks 46 bar and spends 96.5 min under the floor, more than the cascade, and without its preview it spends 301.6 min, against 333.6 min for feedback alone (48.0 against 51.6 min on case 5's day). A 1 h horizon gives 24.7 min for $749 more, and 3 and 6 h give 13.8 and 9.9 min. A fixed margin buys most of the same. Raised 0.4 bar, the cascade spends 21.9 min

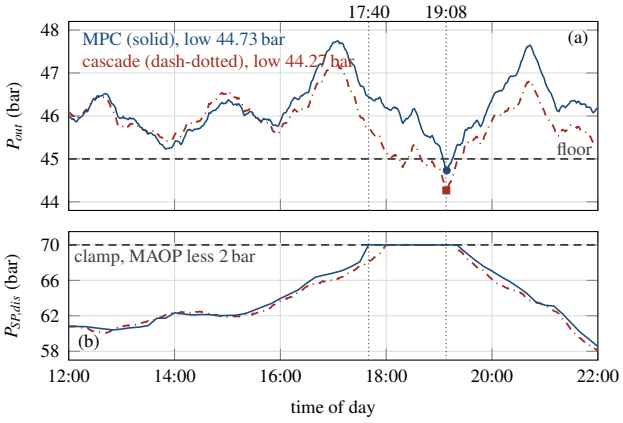


Figure 12.3 Scenario 10's afternoon and evening under the MPC (solid) and under the cascade with the same forecast (dash-dotted), at the default seed. (a) True delivery pressure against the 45 bar floor (dashed), with each day's lowest value at 19:08 marked, a circle for the MPC and a square for the cascade. (b) The discharge setpoint applied, against its 70 bar clamp (dashed), which the MPC reaches at 17:40 and the cascade at 18:00. Dotted verticals mark 17:40 and 19:08.

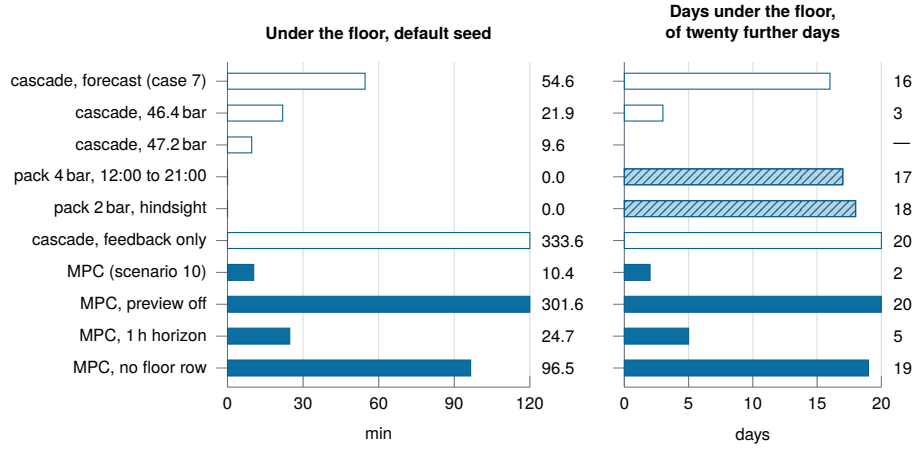


Figure 12.4 Case 10's day under each structure of Table 12.6: minutes under the 45 bar floor at the default seed (left) and days under it among the twenty further days of seeds 11 to 30 (right). Open bars are the cascade, hatched bars the two pack schedules and solid bars the MPC, and every row but feedback only carries the forecast. Two bars of the left panel are cut at the axis end, the column beside each panel gives the table's value, and the row at 47.2 bar was run at the default seed only. Notice that the pack schedules keep the floor on the shipped day and fail on most further days, while the MPC and the cascade raised 0.4 bar go under on 2 and 3.

under the floor on the shipped day, but on the further days 1.5 min on average against 1.4 and 3 days against 2, for \$40,337 against \$40,381. Only on the shipped day, a hard one, is the plan clearly ahead, because the cascade needs 47.2 bar and \$724 more to reach 9.6 min. Figure 12.4 sets the shipped day beside the further days for each structure.

Storage scheduled by hand does better on the day it is placed for. Case 7's 4 bar pack from 12:00 to 21:00 keeps the floor for \$41,633, and a 2 bar pack from 16:00 to 19:00, the cheapest of 288 windows in hindsight, for \$40,294, \$329 less than the MPC. On the further days the two leave the

delivery under the floor on 17 and 18 days, because their lost minutes fall outside their windows, 166 min in sum on the morning ramp for both, 282 min after 21:00 for the 4 bar pack and 434 min between 16:00 and 21:00 for the 2 bar pack, where the MPC loses 27 min. The 4 bar pack also raises 19 false alarms on those days.

The MPC loses where the forecast alone suffices. On case 6's day at 48 bar it tracks the setpoint less closely than the cascade with the forecast, an IAE of 9.17 against 5.78 bar h, with 178 against 294 % of speed travel. The move price is the cause, since with $\lambda = 2$ the same plan reaches 5.28 bar h for 210 % of travel and $275 more. At 46 bar both keep the floor on the shipped day and the MPC costs $366 more. On the further days it goes under on 1 against 11, but on one of them it raises two false leak alarms, at 08:23 and 19:04, where the cascade raises none. Its setpoint had swept 4.4 and 2.7 bar in the 40 minutes before each, and the pack computed from three transmitters trailed the true pack by 12 to 14 t, the mechanism of case 7's false alarm.

The MPC reads no meter and leaves the leak detector as it was. Through case 9's 3 % leak of 5.91 kg/s its persistent unexplained offtake lies between 4.00 and 6.17 kg/s from 03:30 on, and the balance finds the leak 34.0 min after it starts under either structure (23.0 min at 1 %). With case 8's meter 1.5 % low its offtake estimate is −0.19 kg/s at 06:00, and the false alarm comes 34.0 min after the fault under both.

As loaded, the challenge fails under the MPC as it does under the cascade. The floor holds, against 51.6 min under it for the cascade, but the 1 % threshold on a 1 % meter error raises seven false alarms under both, and the MPC's compression, $41,734, is $234 over budget. With the suggested settings both meet it at the default seed, the MPC at $40,782 against $40,934, and on the further days the MPC meets it on 19 and the cascade on 13, the cascade failing on false alarms at 10:40 after the lateral starts and the MPC once, at 08:43. Of the limits the rule leaves unscored, the inlet passes the 70 bar discharge clamp by at most 0.01 bar for 3.1 min under the MPC and 0.08 bar for 44.0 min under the cascade. MAOP stays a protection above the plan. No run at the default seed took the applied setpoint over the clamp, and the highest inlet of any measured day, 70.13 bar on scenario 10 under the MPC (70.10 bar under the cascade), is the station loop's overshoot of its 70 bar setpoint, 1.87 bar under MAOP.

> ✦ DID YOU KNOW? **More Than a Hundred False Alarms**
>
> From the simulator (Table 12.6). Run as loaded, with a 1 % leak threshold on an outlet meter reading 1 % low, the challenge day raises 128 false leak alarms over the twenty further days under the cascade and 133 under the MPC. With the suggested settings the counts fall to 7 and 1.

12.6 Technology in Practice

Figure 12.5 draws the station loop, the delivery loop and the leak detector of Section 12.4 with the hardware of a transmission line in place of the schematic's symbols.

What Measures It

The station loop reads the pressure transmitter at the discharge and the delivery loop the one at the city gate, so the contract's quantity is measured where no controller sits. The leak detector adds the midpoint transmitter and both flow meters. A gas meter at a city gate can be a transit-time ultrasonic meter: sound pulses sent with and against the flow arrive at different times, the difference gives the velocity, several paths average the flow profile, and the speed of sound the meter infers can be checked against the gas composition to catch a faulty sensor [15]. The app's case 8 names this meter, losing a path and reading low.

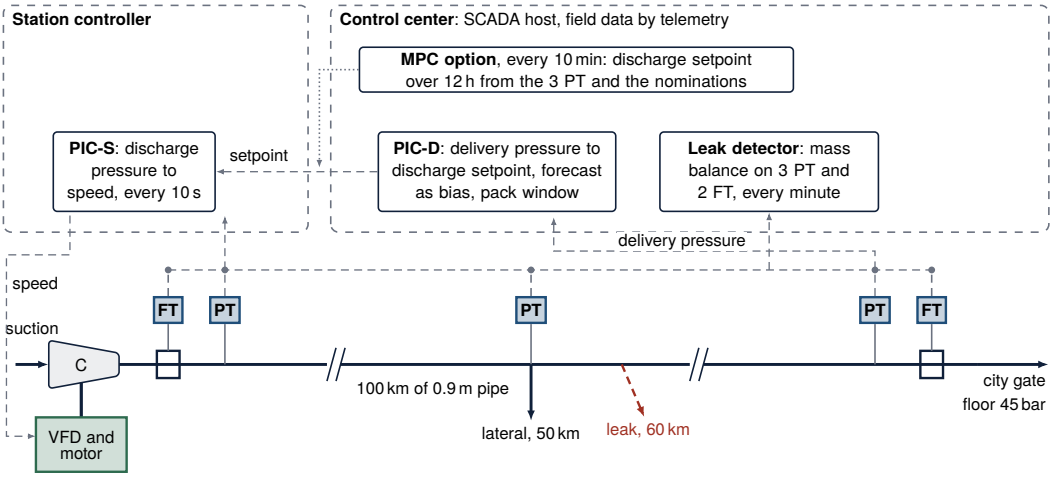


Figure 12.5 The loops of the line with their hardware. Lettered boxes are sensing devices with the chapter's tags: a PT at the station, at the midpoint and at the city gate, and an FT at each end. The green box is the final element, the variable-frequency drive and motor that turn the compressor. Dashed arrows are signals, the dotted arrow the discharge setpoint that the MPC option writes in place of PIC-D. Notice that the delivery loop reads one transmitter 100 km from the machine it moves, through telemetry, while the leak detector reads all five instruments and moves nothing. Not to scale, and not a piping diagram.

> **TECHNOLOGY** **Sensor: The Pressure Transmitter at the City Gate**
>
> An electronic pressure transmitter turns pressure into an electrical signal through a strain gauge, a capacitance cell or a resonating element. In the capacitance cell a taut diaphragm between two fixed plates flexes with the pressure, and the transmitter sends the reading on a 4 to 20 mA loop or as a digital signal [15]. The live zero of the loop shows a broken wire as no current. A transmitter that fails drives its current below 3.6 mA or above 21 mA, so its trend jumps off scale, and the control system can put the loop in manual [15]. The model's three transmitters carry 0.02 bar of colored noise, and the delivery loop reads the one 100 km from the station.

What Moves It

One handle moves the whole line, the compressor's speed. The delivery loop writes only the discharge setpoint, clamped 2 bar below MAOP, and that clamp stands for a legal limit. No segment of a United States gas line may run above its maximum allowable operating pressure, the lowest of several limits that include the design pressure of its weakest element and its test pressure after construction divided by a factor of 1.1 to 1.5 [16].

TECHNOLOGY Final Element: The Compressor's Variable-Frequency Drive

Most pipeline compressors burn gas taken from the line, but electric drives have been growing, mainly for environmental reasons [17]. An induction motor turns at a speed set by the frequency of its supply. A variable-frequency drive rectifies the line power to direct current, smooths it and inverts it back to alternating current at the frequency the speed calls for. Such drives limit the rate of acceleration and deceleration, and they shut themselves down on overcurrent or overvoltage [15]. The model's 20 s speed lag and 1 %/s ramp limit stand for such a drive, and the station loop's output is its speed command.

Where the Controller Runs

A transmission line is run through supervisory control and data acquisition, which lets the control room direct flows and hold pressures [17]. A control server at a control center gathers the field data, remote terminal units or programmable logic controllers at the field sites control each local process, and loops run with cycle times from milliseconds to minutes [18]. The app's loops execute every 10 s, the detector every simulated minute and the MPC every 10 min. The chapter's structure maps onto standard blocks: a PI in the station's controller from discharge pressure to speed, a PI in the control center from delivery pressure to that setpoint with the forecast on its bias input and the pack window on its setpoint, and the mass balance as a pipeline monitoring application beside them.

IN PRACTICE What the Control Room Rules Ask

In the United States an operator whose controllers run a transmission line with a compressor station through supervisory control and data acquisition must follow written control room procedures. Its alarm plan identifies, at least once each calendar month, the safety points that generated false alarms, had alarms inhibited, were taken off scan or held forced values. Alarm set points and descriptions are verified each calendar year, within 15 months, and the controllers' workload is reviewed so that they have time to react to alarms [19]. Case 8's false alarm from a meter reading low is the kind of point the monthly review lists.

What Is Modernizing It

The first of three developments is transient planning: a 2015 formulation posed a day of compressor operation as optimal control of the network's transient flow, minimizing compression cost or the deliveries left unmet, with gains in security and efficiency over methods that assume steady flow [8]. The MPC option of Section 12.4.4 is one line's version of that plan. The second is leak detection beyond the mass balance. Optical fiber laid with a pipeline picks up the temperature around it, and the travel time of the light locates a leak, with a spatial resolution of the order of meters on long-distance systems [20]. The third is security. A penetration test of a natural gas utility's business network strayed into a part connected to its control system, locked that system up, and left the utility unable to send gas through its pipelines for four hours [18].

12.7 Limitations

The flow is isothermal and the friction factor fixed, so the line pack carries errors of a few percent against a model with an energy balance and a real gas equation of state [3].

The grid is coarse. Wave fronts are smeared over several cells and the model says nothing about surge or rupture.

The forecast is nearly perfect because the demand is a known profile with a small wander. With realistic forecast errors, part of the benefit shown here returns to the feedback loop [5].

The station is one compressor with no fuel use and no recycle, and the electricity price is a fixed daily shape.

The leak detector is the simplest computational pipeline monitor: a windowed balance with a steady-state inventory estimate. Transient-model detectors and acoustic methods have different limits [10].

The predictive controller's model leaves out what decides the floor on the peak day. Its preview carries no wander, which on scenario 10's day takes the evening demand 19.8 kg/s above the nominated peak and both structures under the floor, while without the wander neither goes under. A real forecast's own error, larger and driven by the weather [5], would fall on the same margin. The model also omits the inertia of the gas, the curvature of the profile between nodes and the compressor map, its last block spans the final 4.5 h of a plan, the solver carries one slack over the horizon rather than one per step, and it has no model of the leak detector, whose pack estimate from three transmitters its ramps can outrun.

> ⚠ **OPEN CHALLENGE** **Gas-Fired Power Plants on the Line**
>
> The chapter's day follows a known profile, and its challenge adds one power plant on the lateral. Real lines increasingly feed gas-fired generators, which change their output quickly and are dispatched to balance wind and solar power. Their intermittent dispatch creates rapidly changing withdrawals and stresses that ad hoc methods find harder and harder to keep within design limits, and gas interrupted to generators can force electric load shedding [8]. Optimization of transient flow remains one of the most challenging research areas, still in a developing phase [7]. Progress would be a line that plans its pack from the power system's schedule and still holds every delivery floor when a plant starts unannounced.

12.8 Conclusion

A transmission line is a storage vessel with pressure transmitters at three points, and the simulator of this chapter makes that consequential. A 1 % speed step takes 12 min to reach the city gate and an hour later is still arriving. A station loop that holds its own discharge leaves the delivery at 43.5 bar after a demand step, and a delivery loop above it recovers with a 1.8 bar dip over two hours. Over a winter day, feedback alone at 48 bar spends 52 min under the 45 bar floor, while a 30 min forecast lets the setpoint fall to 46 bar with no violation on the default seed and 8 % lower compression cost. When the peak exceeds the line's capacity at its pressure limit, only storage helps, and a 4 bar afternoon pack keeps the floor for about $1,400, with a margin the demand wander can erase. The mass balance finds a 3 % leak in about 34 min and cannot tell a 1.5 % leak from a 1.5 % meter error. The optional predictive controller, holding the floor as a soft constraint over a twelve-hour plan, spends 10.4 min under it on scenario 10's day against 54.6 min for the cascade with the same forecast and goes under on 2 of twenty further days against 16, although a setpoint 0.4 bar higher matches it on those days and a pack chosen in hindsight keeps the floor on the shipped day for less. The model is isothermal, coarse and forecast-friendly. A fuller model would add an energy balance, several compressor units, forecast error and a transient-model leak detector. As it stands, the line is a plant on which to study the control of an inventory that cannot be measured.

KEY POINTS

- A gas line stores its own gas, so a speed step at the station takes 12 min to reach the city gate.
- A station loop on its discharge leaves the delivery short after a demand step, and a delivery loop above it recovers only after a 1.8 bar dip.
- A 30 min forecast lets the setpoint fall from 48 to 46 bar at 8 % lower compression cost.
- The mass balance cannot tell a 1.5 % leak from a 1.5 % meter error.
- The predictive controller goes under the floor on 2 of twenty further days against 16, but a setpoint 0.4 bar higher matches it on those days.

Where else this applies. Water released ahead of demand along Volume 9, Chapter 5, heat sent ahead of a morning peak in Volume 9, Chapter 6, and a liquid line with a leak detector of its own in Chapter 13 share the same distributed inventory.

References

[1] S. Chevalier and D. Wu, "Dynamic linepack depletion models for natural gas pipeline networks," *Applied Mathematical Modelling*, vol. 94, pp. 169–186, 2021.

[2] A. Osiadacz, "Simulation of transient gas flows in networks," *International Journal for Numerical Methods in Fluids*, vol. 4, no. 1, pp. 13–24, 1984.

[3] A. Thorley and C. Tiley, "Unsteady and transient flow of compressible fluids in pipelines—a review of theoretical and some experimental studies," *International Journal of Heat and Fluid Flow*, vol. 8, no. 1, pp. 3–15, 1987.

[4] T. Kiuchi, "An implicit method for transient gas flows in pipe networks," *International Journal of Heat and Fluid Flow*, vol. 15, no. 5, pp. 378–383, 1994.

[5] B. Soldo, "Forecasting natural gas consumption," *Applied Energy*, vol. 92, pp. 26–37, 2012.

[6] D. Marqués and M. Morari, "On-line optimization of gas pipeline networks," *Automatica*, vol. 24, no. 4, pp. 455–469, 1988.

[7] R. Z. Ríos-Mercado and C. Borraz-Sánchez, "Optimization problems in natural gas transportation systems: A state-of-the-art review," *Applied Energy*, vol. 147, pp. 536–555, 2015.

[8] A. Zlotnik, M. Chertkov, and S. Backhaus, "Optimal control of transient flow in natural gas networks," in *2015 54th IEEE Conference on Decision and Control (CDC)*, 2015, pp. 4563–4570.

[9] L. Billmann and R. Isermann, "Leak detection methods for pipelines," *Automatica*, vol. 23, no. 3, pp. 381–385, 1987.

[10] P.-S. Murvay and I. Silea, "A survey on gas leak detection and localization techniques," *Journal of Loss Prevention in the Process Industries*, vol. 25, no. 6, pp. 966–973, 2012.

[11] A. Gopalakrishnan and L. T. Biegler, "Economic nonlinear model predictive control for periodic optimal operation of gas pipeline networks," *Computers & Chemical Engineering*, vol. 52, pp. 90–99, 2013.

[12] U.S. Energy Information Administration, "Natural gas pipelines," Energy Explained, last updated March 19, 2024, 2024, accessed November 18, 2026. [Online]. Available: https://www.eia.gov/energyexplained/natural-gas/natural-gas-pipelines.php

[13] 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.

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

[15] 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.

[16] Pipeline and Hazardous Materials Safety Administration, "49 CFR 192.619, maximum allowable operating pressure: Steel or plastic pipelines," Code of Federal Regulations, Title 49, Part 192, edition of October 1, 2025, govinfo.gov, 2025.

[17] U.S. Energy Information Administration, "Natural gas pipelines: Transportation process and flow," About U.S. Natural Gas Pipelines, archived analysis, based on data through 2007/2008 with selected updates, 2008, accessed November 18, 2026. [Online]. Available: https://www.eia.gov/naturalgas/archive/analysis_publications/ngpipeline/process.html

[18] 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.

[19] Pipeline and Hazardous Materials Safety Administration, "49 CFR 192.631, control room management," Code of Federal Regulations, Title 49, Part 192, edition of October 1, 2025, govinfo.gov, 2025.

[20] J. Li and M. Zhang, "Physics and applications of Raman distributed optical fiber sensing," *Light: Science & Applications*, vol. 11, p. 128, 2022.

12.9 Exercises

12.1 Concept check · 2 min

In scenario 2 the station loop holds the discharge at 60.35 bar while the city demand steps from 210 to 240 kg/s. The delivery pressure settles near 43.5 bar and stays there. Why does it not recover?

(a) With the inlet held, $P_{out}^2 = P_{in}^2 - KW^2$, so a larger flow leaves a lower delivery for good, and no loop measures it.

(b) The extra flow drains the line pack at a steady rate, so the delivery keeps sinking slowly toward the curtailment pressure.

(c) The compressor reaches its 100 % speed limit, so the station cannot supply the extra 30 kg/s.

(d) The city gate curtails its offtake below the floor, which pins the delivery at the curtailment pressure.

12.2 Concept check · 1 min

A pressure wave crosses the 100 km line in under five minutes, yet in scenario 1 a 1 % speed step first moves the delivery pressure after about 12 min, and an hour later it is still rising. Which statement explains this?

(a) Pressure travels with the gas, which takes about 3.7 h to cross, so the delivery moves only as the new gas arrives.

(b) Friction damps the wave within about half a minute of travel, and what reaches the outlet is a slow diffusion of pressure as the line pack fills.

(c) The speed follows its command through a 20 s lag and a 1 %/s ramp, which holds back the extra flow for several minutes.

(d) The colored noise of the delivery transmitter hides the first movement until the change grows larger than the noise.

12.3 Concept check · 2 min

In scenario 7 with a 4 bar pack from 12:00 to 21:00, the leak detector raises one alarm at 12:40, although nothing leaks and both meters read true. What produces the imbalance?

(a) The 0.3 % noise of the two flow meters adds up over the 30 min window until it crosses the 2 % threshold.

(b) The higher delivery setpoint raises the flow the city takes, and the outlet meter does not register the extra delivery.

(c) The pack computed from three transmitters assumes a steady profile, so it trails the true pack during the ramp, and gas stored in the line is counted as lost.

(d) The compressor nears surge as the window opens, and the recycled gas passes the inlet meter twice.

12.4 Theory and calculation · 7 min

Use the line of Table 12.2: $L = 100\,\text{km}$, $D = 0.90\,\text{m}$, $f = 0.0095$, $c = 350\,\text{m/s}$, $N = 20$ cells.

(a) Compute K of Equation (12.3) in bar^2 per $(\text{kg/s})^2$. *Checkpoint:* $K = 0.031\,95$.

(b) For 210 kg/s delivered at 48 bar, find the inlet pressure with the full-length K and with the twenty-cell line's $(N-1)K/N$. Which does the feedforward of Equation (12.5) use, and what carries the difference? *Checkpoint:* 60.93 bar and 60.35 bar.

(c) Take the mean pressure of a steady line as $\bar{P} = \frac{2}{3}(P_{in}^3 - P_{out}^3)/(P_{in}^2 - P_{out}^2)$. Find the line pack from Equation (1.8) and the time the gas takes to cross, M/W.

12.5 Theory and calculation · 7 min

Case 4 measures the delivery's answer to a 1 bar step of a held discharge: 63 % times of 34, 54 and 68 min and gains of 1.09, 1.25 and 1.42 at 120, 210 and 280 kg/s. Take the delivery at 48 bar, the twenty-cell $0.95K$ with $K = 0.031\,95$, and $A = 0.6362\,\mathrm{m}^2$.

(a) Differentiate Equation (12.3) at fixed W to find the steady gain $\mathrm{d}P_{out}/\mathrm{d}P_{in}$, and evaluate it at the three flows. *Checkpoint:* 1.09, 1.26 and 1.43.

(b) Evaluate the diffusion time $L^2 fW/(2DA\bar{P})$ of Section 12.3 at the three flows, with $\bar{P}$ from the formula of the previous exercise. *Checkpoint:* 33, 53 and 66 min.

(c) The line answers faster at night, yet the cascade tuned at design reaches 90 % of a 1 bar setpoint step in 83 min at night against 58 min at design. Explain.

12.6 Theory and calculation · 6 min

The leak detector of Section 12.4.3 uses a 30 min window, a 10 min persistence and the leak law of Section 12.3. Case 9 opens a 3 % leak at 60 km, where the line is at 52.6 bar.

(a) Find the leak flow in percent of design flow and in kg/s. *Checkpoint:* 2.81 %, 5.90 kg/s.

(b) Model the imbalance as rising linearly from zero to the leak over one window. When does it cross a 2 % and a 1 % threshold, and when would the alarm come? Compare with the measured 34 and 23 min. *Checkpoint:* 21.4 and 10.7 min, alarms at 31.4 and 20.7 min.

(c) How much gas has the leak released by the measured alarms? What is the smallest leak setting that a 2 % threshold can ever see here?

12.7 App challenge · 7 min

Load scenario 1, "Step the compressor, time the outlet (MANUAL)", default seed, at speed 60 (one simulated minute per second). Press Run. When the Clock tile reads between 00:05 and 00:15, set Compressor speed to 74.2 % on the station card and note the time. Run one more hour of simulated time, pause and export the CSV.

(a) After the step, find the first row in which `p_out_bar` is at least 0.01 bar above its value at the step, and the changes in `p_out_bar` and `p_in_bar` one hour after the step. *Checkpoint:* 12 to 13 min, then +0.30 and +0.47 bar.

(b) Read `pack_t` at the step and an hour later. Use Equation (1.8) to turn the gain into a rise of the mean pressure, and compare it with (a).

(c) Over this hour, does the delivery behave more like a first-order lag or like an integrator? What will finally stop the rise?

12.8 App challenge · 9 min

Load scenario 6, "The same day with a forecast", default seed. Set Delivery pressure setpoint to 46 bar before Run. Run the full day at speed 3600 (one simulated hour per second) and export the CSV. In the row at `time_s` = 86400 read `metric_min_pout_bar`, `metric_below_floor_s` and `metric_cost_usd`. The Lowest delivery pressure, Time below the floor and Compression cost tiles show the same values at 24:00, but the cost keeps counting after midnight. Repeat with Forecast lead set to 15, 45 and 90 min, each after a Reset and before Run.

(a) Tabulate the lowest delivery, the minutes below the floor and the cost for the leads of 15, 30, 45 and 90 min. *Checkpoint:* 30 min lead: 45.01 bar, 0, $34,396.

(b) Case 5 at 48 bar without the forecast costs $37,330. What does the 30 min forecast at 46 bar save? Is its margin one to operate on?

(c) Explain the trend of the lowest delivery and the cost with the lead.

12.9 App challenge · 9 min

Load scenario 10, "The floor as a constraint: MPC against the cascade with the same forecast", default seed. It opens with Delivery control on MPC. Run one day at speed 3600, export the CSV and read `metric_below_floor_s`, `metric_min_pout_bar` and `metric_cost_usd` in the row at `time_s` = 86400. Then Reset, set Delivery control to Cascade before Run, and repeat. Then Reset, set Delivery pressure setpoint to 46.4 bar before Run, and repeat.

(a) Tabulate the three runs. *Checkpoint:* MPC: 10.4 min, 44.73 bar, $40,623.

(b) In the first two files, read `discharge_sp_bar` at 17:40 (`time_s` = 63600) and `pack_t` at 17:00 and 18:00. What does the MPC do differently?

(c) Reload the scenario, set MPC preview to Off (an Advanced control on the delivery card) before Run and run the day again. Where does the MPC's advantage come from, and what would you tell a manager who asks whether this line needs it?

12.10 Application area · 4 min

In scenario 7 the demand is 8 % above its profile, so the evening peak is $1.08 \times 285\,\mathrm{kg/s}$.

(a) With the inlet at the 70 bar clamp and the delivery at the 45 bar floor, what steady flow can the twenty-cell line carry, from Equation (12.3) with $0.95K$? *Checkpoint:* 307.8 kg/s.

(b) Packing is the control room's own answer. From the chapter, what does case 7's 4 bar pack cost, what does it do to the leak detector, and how does it fare on other days?

12.11 Application area · 3 min

A control room's detector runs at a 1 % threshold and calls the field crew out every time a meter drifts. The supervisor proposes 2 %. Using cases 8 and 9, state what the change gains and loses, and name two ways to keep the sensitivity without the false alarms.

12.12 Review the agent · 5 min

An AI assistant writes about the challenge: "Set Delivery control to MPC and raise the Leak alarm threshold to 2.5 %. As loaded, the MPC keeps the delivery above the floor all day, so its only failure is the seven false alarms from the 1 % meter error, and the higher threshold removes them. Keep the 48 bar setpoint for margin." Find the flaw, and say how you would check it in the app.