INTERACTIVE CONTROL ENGINEERING
Principles and Applications

Volume 10 Control in Upstream and Midstream Oil and Gas

CHAPTER 1

Introduction: Pressure, Inventory and Limits From Well to Market

Pages 1–23

Edited by John D. Hedengren

How to cite this chapter. "Introduction: Pressure, Inventory and Limits From Well to Market," 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. 1, pp. 1–23.

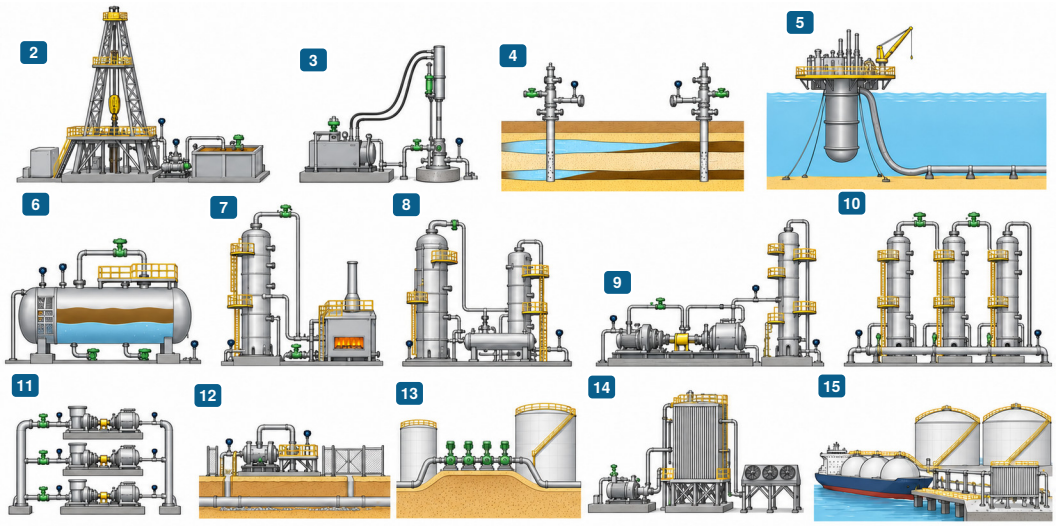
The page numbers in this file are those of the printed volume.

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

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

CHAPTER 1

Introduction: Pressure, Inventory and Limits From Well to Market



The chapters 2 drilling rig 3 rod-pumped well 4 smart well and flood 5 subsea flowline 6 offshore separator 7 glycol dehydrator 8 amine sweetener 9 turboexpander plant 10 NGL train 11 compressor station 12 gas pipeline 13 products pipeline 14 LNG liquefaction 15 LNG terminal

Notice that the rear row drills, lifts and floods the wells and brings their fluid up a riser, that the middle row separates it, treats the gas and splits its liquids, and that the front row moves gas and refined products by pipeline and gas by ship as LNG. Illustration, not to scale.

IN PLAIN TERMS **Moving oil and gas from the rock to the market**

The main idea. Oil and gas travel from a well through pipes, separators, cleaning plants and long pipelines before they are sold. At most steps a valve, a pump or a machine sets a pressure, and the flow then follows from the rock, the pipe or the buyer.

The challenge. Liquid arrives in surges, and a gas compressor gets into trouble when its flow falls too low. The limit that matters often lies deep in a well or inside a pipe where no instrument can reach, or it is known only minutes late.

Why it matters. Good control keeps people safe and pipes whole, and it decides how much oil and gas reaches the market.

THREE WAYS TO READ THIS CHAPTER

Explore	7 min	This page, Section 1.2 with its timeline and two History boxes, Figure 1.2, the Did you know? boxes in Sections 1.4.2, 1.4.4 and 1.4.5, and Key points in Section 1.8. No equations.
Study	40 min	The whole chapter in order, working each number of the example W-1, then Exercises 1 to 10.
Practice	10 min	This page, Sections 1.4.8 and 1.4.9 with Table 1.2, Section 1.5 with Figures 1.7 and 1.10, and Section 1.8.

1.1 The Question This Volume Answers

Oil and gas reach a market through a chain of plants that were never designed as one. A well is drilled and produced, its fluid crosses a flowline to a separator, the gas is treated, compressed and sent down a pipeline or liquefied, and the liquids leave by their own lines. Each plant in the chain has

a few actuators, many limits and fewer instruments than its limits need [1]. This volume answers one question: which control techniques recur along that chain, and where does what can be measured, rather than what can be actuated, set the limit of what control achieves?

Six control problems define the industry, and each chapter meets at least two of them.

- *Pressure is the handle and the flow follows.* A choke, a pump or a compressor sets a pressure, and the rate follows from a reservoir, a line or a market that the actuator cannot change.
- *Flow arrives in slugs and ramps.* Multiphase lines deliver liquid in surges, and gas demand rises every morning. Separators, surge drums, pipelines and tanks store the difference.
- *Machines work near an instability.* A centrifugal compressor surges when its flow falls too far, and a recycle valve that keeps it out of surge compresses gas for nothing.
- *Products are judged late.* Water, hydrogen sulfide, methane in a liquid product and ethane in propane are read by analyzers whose reports arrive minutes after the gas or liquid they describe.
- *The binding limit sits where no instrument is.* A hydrate margin at the foot of a riser, a fluid level at a pump nine hundred meters down and a temperature approach inside a welded exchanger are computed, not measured.
- *Many limits share one actuator.* A choke protects drawdown, sand, water and pipe pressure at once, and a pump station protects suction and discharge. Selectors and overrides decide which limit owns the actuator.

Fourteen chapters answer the question on fourteen plants, each with an app that the reader runs in a browser. Chapters 2 to 5 are upstream: a well being drilled, a rod-pumped well, a waterflood with a smart producer and a deepwater flowline. Chapters 6 to 10 follow the fluid through separation, gas treating and the recovery and fractionation of natural gas liquids. Chapters 11 to 15 are midstream: a compressor station, a gas transmission line, a refined products line and the two ends of the liquefied natural gas (LNG) chain.

A reader who has worked through the volume can do the following on a plant that is new to the reader.

- Find the operating point of a well and its choke, and say how much of a choke move reaches the rate.
- Judge from a few numbers whether a riser will slug, and size the level control that absorbs a slug without passing it downstream.
- Compute the line pack of a gas line and say why a liquid line needs overrides that a gas line does not.
- Read a compressor map, place the surge and control lines, and find the throughput below which the recycle opens.
- Explain the floor that the lean solvent sets on an absorber and the cold that an expander makes and a valve does not.
- Build an override on a shared actuator so that it takes over at the limit, and say what each limit is known through.

This chapter develops each of these on one small example and names the chapter that uses it. PID forms and tuning are taken as known from Volume 2, selectors and overrides in general from Volume 4 and predictive control from Volume 9.

1.2 A Short History

The milestones of Figure 1.1 come from three places: the gas plant and its lines, the pumped well and the drill floor. In 1930 a patent described a small group of amines that unite with acidic gases at a low temperature, give them up at a higher one and so circulate in a closed cycle [2]. In 1934 the crystalline plugs found in flowlines above the freezing point of water were shown to be hydrates. An equation for the conditions that form them followed, with inhibitors such as methanol and glycol [3]. Chapter 7 dries gas so that hydrates cannot form, and Chapter 5 doses methanol against them.

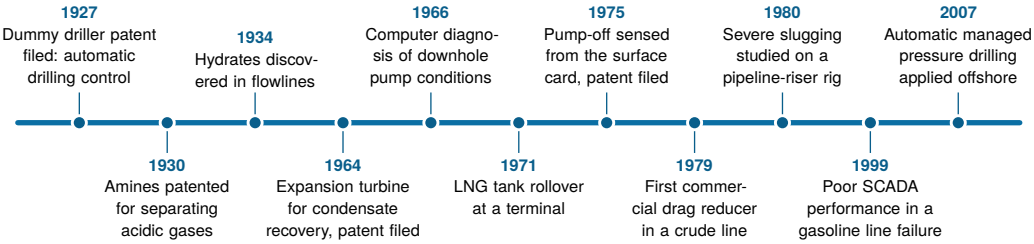


Figure 1.1 Milestones of upstream and midstream oil and gas and of their control. The dates are from [4] (1927), [2] (1930), [3] (1934), [5] (1964), [6] (1966), [7] (1971), [8] (1975), [9] (1979), [10] (1980), [11] (1999) and [12] (2007). A patent is dated by its filing where the card says so. Notice that the drill floor holds the first and the last milestone of the eighty years. Events are evenly spaced, not to a time scale.

About half of the gas transmission network of the United States was installed in the 1950s and 1960s, after the demand for natural gas more than doubled following the Second World War [13]. Gas plants learned to make their cold with a turbine. A patent filed in 1964 expanded cooled feed gas through a turbine that tolerates condensate and used the turbine's work to recompress the residue gas [5], the arrangement of Chapter 9. In 1966 a computer method turned the load and position measured at the top of a rod string into the card at the pump [6]. A patent filed in 1975 sensed that a well had pumped off from a fall in the power input, which the area of the surface card measures, so that the pump could be stopped [8]. Both rest on the two measurements of Chapter 3. The drill floor was automated in steps of its own.

> ⧖ HISTORY **1927 to 2007: Automating the Drill Floor**
>
> A patent filed in 1927 set out to keep a set pressure on a rotary drill hung partly from a cable. A reversible motor paid the cable out or picked it up, and too much pull made it pay out until the strain on the cable fell [4]. A patent filed in 1952 let the tension in the drilling line set how fast the drum unspooled. When the hydraulic pressure that read the tension fell, the bit carried too much weight, the drum slowed and the bit drilled off [14]. In 2007 a choke at the surface held the downhole pressure of a well drilled offshore under automatic control, because narrow drilling margins ask for highly accurate pressure control [12]. Chapter 2 closes all three loops, on weight, on rotation and on pressure.

Pipelines and terminals taught their lessons through events. In 1971 a cargo loaded beneath colder liquefied gas in a terminal tank formed two layers. The layers later mixed suddenly, and about 2000 tonnes of vapor left through the safety valves over a few hours [7]. That rollover is a case of Chapter 15. A polymer that reduces drag first raised the flow of a crude oil line commercially in 1979 [9]. In 1980 a test rig showed that a riser slugs at low gas and liquid rates and that choking removes the slugging [10]. The separator of Chapter 6 takes the slugs as they come, because its wells deliver an imposed stream.

> ⧖ HISTORY **1999: A Control Room That Could Not See**
>
> A pipeline is run from a control room. A control server at a control center gathers field data for the operator, while remote terminal units or programmable controllers at the field sites run the local loops [11]. In June 1999 a 16-inch gasoline line leaked 237 000 gallons, which ignited 1.5 hours later and caused three deaths. During the event the supervisory system performed poorly, which kept the pipeline controllers from seeing an abnormal operation develop. The investigators recommended an offline system on which to test changes to the supervisory database [11]. Chapter 12 and Chapter 13 run below such a room.

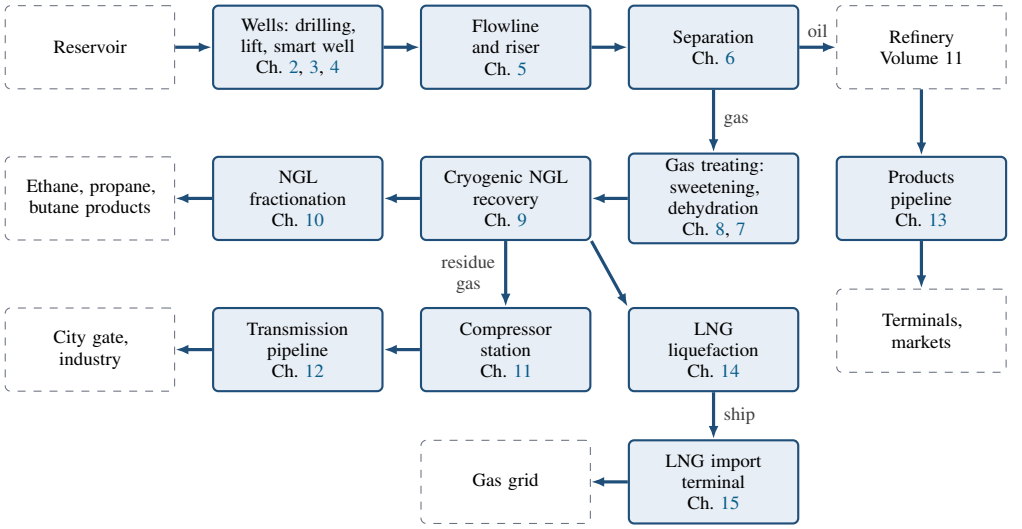


Figure 1.2 The Chain From Reservoir to Market. Filled Blocks Are the Plants of Chapters 2 to 15, Dashed Blocks Lie Outside the Volume. A Liquefaction Plant Usually Treats Its Own Feed Gas.

1.3 Notation and Conventions

Pressures are in bar. A chapter writes barg or bara where gauge and absolute differ, and this chapter uses absolute pressure throughout. Liquid rates are in m^3/h at line conditions, gas in standard cubic meters (Sm^3) or in kg/s, and temperatures in °C. The reservoir pressure is p_R, the flowing bottomhole pressure p_{wf}, the wellhead pressure p_{wh}, the productivity index J and the choke opening x in percent. The liquid holdup, the share of a pipe's volume that liquid fills, is H_L. A separator level is ℓ, because the framework's h is its integration step. A compressor's flow coefficient is ϕ, its head coefficient ψ, its speed n as a fraction of design and its surge margin SM. An absorber's absorption factor is $\mathscr{A}$. The framework's a is the controller action, so a wave speed here is c.

The chapters keep their own symbols, and several collide. Chapter 6 writes level heights as h, and Chapter 5 writes the holdup as H_i. Chapter 11 uses H and W for the shape of its compressor map. Chapter 13 writes the wave speed a, and Chapter 12 writes the isothermal wave speed c, as here. Chapter 3 writes the controller action α and Chapter 7 writes it σ. Several chapters give water content in pounds per million standard cubic feet, the unit of the trade, and this chapter uses mg/Sm^3.

1.4 From Reservoir to Market on One Well

1.4.1 The Path From Reservoir to Market

Figure 1.2 draws the chain and marks where each chapter's plant sits in it. Drilling makes the well, and the well brings oil, gas and water to the surface by its own pressure or with a pump. A flowline carries the mixture to a separator, which splits it into three streams. Oil goes to a refinery, whose control is the subject of Volume 11, and the refined products leave by pipeline. Gas is sweetened and dried, its heavier hydrocarbons are recovered at low temperature and fractionated into products, and the remaining residue gas is compressed into a transmission line or liquefied for a ship. At the other end of the voyage a terminal returns the LNG to gas.

The chapters take the plants one at a time, and the techniques they share are developed below on one running example. The example is a single oil well, W-1, on a small platform, followed from the reservoir to its export lines. The well produces through a choke into a riser and a separator, and its gas is compressed into a short export pipeline. Every number of the example is declared here or computed from these declarations by the script kept with this chapter's record.

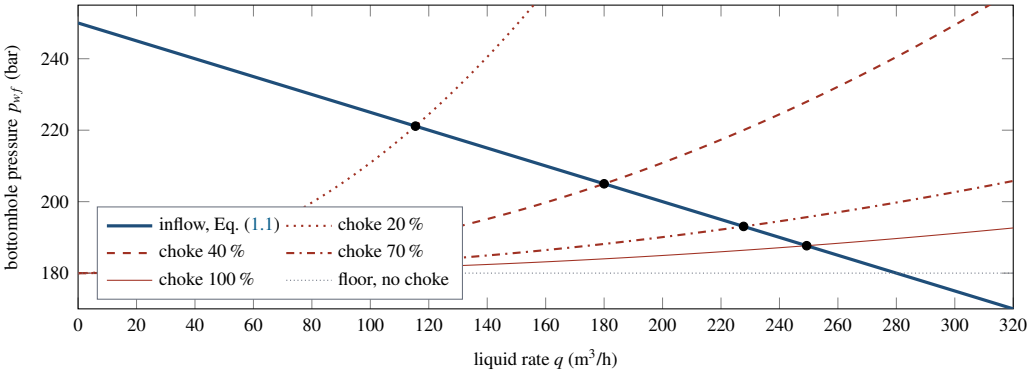


Figure 1.3 Operating Points of W-1. The Inflow Line (Thick Solid) Meets the Outflow Curve of Each Choke Opening (Dotted 20, Dashed 40, Dash-Dotted 70, Thin Solid 100 %) at the Dots. No Opening Takes the Bottomhole Pressure Below the Densely Dotted Floor.

1.4.2 Wellbore, Choke and Flowline Pressure

A reservoir delivers liquid in proportion to the pressure difference that draws it into the well. With the straight-line productivity index, the simplest inflow model, the rate is Equation (1.1).

$$q = J\,(p_R - p_{wf}) \tag{1.1}$$

where $p_R - p_{wf}$ is the drawdown. A well whose oil releases gas below its bubble point delivers less than the straight line at large drawdowns [15], and the example ignores this. W-1 has $p_R = 250\,\text{bar}$ and $J = 4\,\text{m}^3/\text{h}$ per bar.

The fluid then climbs the tubing, crosses the choke and flows to the separator. Each step costs pressure, as shown in Equation (1.2).

$$p_{wf} = p_{wh} + \Delta p_{col}, \qquad q = c\,x\sqrt{p_{wh} - p_{dn}}, \qquad p_{dn} = p_{sep} + \Delta p_{riser} + \Delta p_{fr} \tag{1.2}$$

where Δp_{col} is the weight of the column in the tubing, the choke is a valve with a linear characteristic and coefficient c, p_{dn} is the pressure just downstream of it, and the flowline adds the separator's pressure, the weight of the riser's contents and friction. W-1 has a tubing column of 120 bar, 2000 m of fluid at 0.06 bar/m, a separator at 10 bar and a riser that weighs 50 bar while it flows. Tubing and flowline friction are neglected, so $p_{dn} = 60\,\text{bar}$, and $c = 0.9$ in m^3/h per percent of opening per square root of a bar.

The operating point is where the reservoir delivers what the choke passes. At $x = 40\,\%$ the two meet at $q = 180\,\text{m}^3/\text{h}$, with $p_{wf} = 205\,\text{bar}$, $p_{wh} = 85\,\text{bar}$, 25 bar across the choke and a drawdown of 45 bar. Figure 1.3 draws the inflow line and the outflow curves for four openings. The curves have a floor. With no pressure lost at the choke the bottomhole pressure cannot fall below $\Delta p_{col} + p_{dn} = 180\,\text{bar}$, and the well cannot produce more than $J(p_R - 180) = 280\,\text{m}^3/\text{h}$. The choke wide open still takes 7.7 bar and gives $249.3\,\text{m}^3/\text{h}$.

A choke move therefore reaches the rate only in part. Opening the choke raises the flow through it at once, by $c\sqrt{p_{wh} - p_{dn}} = 4.5\,\text{m}^3/\text{h}$ per percent at the design point. The extra flow drains the wellbore, the bottomhole pressure falls, and the inflow takes back part of the increase. Linearizing Equations (1.1) and (1.2) gives the share that remains at steady state, as shown in Equation (1.3).

$$\frac{\Delta q_\infty}{\Delta q_0} = g = \frac{J}{J + \partial q/\partial p_{wh}}, \qquad \frac{\partial q}{\partial p_{wh}} = \frac{c\,x}{2\sqrt{p_{wh} - p_{dn}}} \tag{1.3}$$

At $x = 40\,\%$ the choke passes $3.6\,\text{m}^3/\text{h}$ more for each bar of wellhead pressure, so $g = 4/7.6 = 0.526$, and the steady gain is $2.37\,\text{m}^3/\text{h}$ per percent. The response is a lead-lag, a jump of 4.5 that relaxes to

2.37 as the wellbore drains. With a wellbore that stores 1.25 m^3 per bar, the relaxation has a time constant of $60 \times 1.25/7.6 = 9.87$ min. Wider open, the choke takes a smaller share of the pressure and the share g falls, so a rate loop tuned at one opening is slower at a wider one.

The choke sets a pressure, and the rate follows from a reservoir pressure the choke cannot raise. A rate loop holding 180 m^3/h opens the choke as the reservoir depletes: 51.6 % at 240 bar, 63.3 % at 235 bar and 89.4 % at 230 bar. At 229 bar it is wide open and the target is lost. Chapter 4 meets the same lead-lag when it reads the wellbore from the choke, and the same loss of travel when its injection plant trips. Chapter 5 balances a well, a choke and a line in the same way. Chapter 2 writes the bottomhole pressure of a well being drilled as the same column plus friction plus a choke's back pressure, held between the pore and fracture pressures, and there closing the choke raises the pressure, a negative process gain. In Chapter 3 the column above the pump turns the pump's intake pressure into a fluid level.

> **+ DID YOU KNOW? Most Oil Wells Need a Pump**
>
> More than seventy percent of the oil wells now producing are lifted artificially [16]. In 1999 sucker rod pumps, the pumps of Chapter 3, were installed in about 90 percent of all oil wells in the United States [17].

1.4.3 Multiphase Flow and Slugging

The riser's weight depends on how much of it is liquid. Gas moves faster than liquid up a pipe, so the liquid holdup H_L exceeds the liquid's share of the volume flow, and the mixture density is Equation (1.4).

$$\rho_m = H_L \rho_L + (1 - H_L)\, \rho_g, \qquad \Delta p_{riser} = \rho_m\, g\, z_r \tag{1.4}$$

Holdup correlations for inclined pipes give H_L from the flow rates [18]. On W-1, oil at 800 kg/m^3 and gas at 60 kg/m^3 give the flowing riser's 0.05 bar/m at a holdup of 0.608 over the riser's 1000 m. A riser full of liquid weighs 78.5 bar instead of 50 bar, and the well behind it feels the 28.5 bar at once. At the same choke opening the rate falls from 180 to 120.9 m^3/h, 32.8 % less. The coupling is a positive feedback. Less flow means slower gas, more holdup, a heavier riser and less flow again.

Severe slugging is the cycle this feedback drives on a flowline that runs down to the foot of a riser [19]. Liquid collects at the riser base and seals it. The gas behind the seal is trapped in the flowline and compresses while the liquid column grows. When the gas pressure exceeds the column's weight, the gas pushes the liquid out in one piece and the cycle starts again. Whether the seal grows faster than the gas can follow decides whether a slug forms. For an isothermal ideal gas trapped in the flowline's gas volume V_g, the two rates compare as in Equation (1.5).

$$\Pi = \frac{\mathrm{d}p_g/\mathrm{d}t}{\mathrm{d}p_{col}/\mathrm{d}t} = \frac{p_s\, Q_{g,s}/V_g}{\rho_L\, g\, q_L/A_r} \tag{1.5}$$

where $p_s\, Q_{g,s}$ is the gas flow at standard pressure, q_L the liquid flow and A_r the riser's bore. A ratio below one means the column outruns the gas and the riser slugs, a criterion of the kind published for flowline and riser systems [20]. W-1's 0.2 m riser fills in 628 s at 180 m^3/h, so its column grows by 0.125 bar/s. Its 18 000 Sm^3/h of gas in 150 m^3 of flowline gas volume raises the trapped pressure by only 0.034 bar/s, so $\Pi = 0.270$ and the riser slugs. It would need 3.70 times the gas, 66 600 Sm^3/h, to keep up. Gas injected at the riser base, a smaller gas volume upstream and feedback that throttles the topside choke on the riser base pressure are the remedies. The last has been tested on flowlines in service [21].

The flowline also delays whatever enters it. Liquid moves as a column, so a change made at the well reaches a point downstream once all the liquid between them has been pushed past. W-1's 5 km flowline of 0.25 m bore holds 147.3 m^3 of liquid at a holdup of 0.6, which takes 49.1 min to displace at 180 m^3/h and 98.2 min at half the rate. The delay is a volume divided by a flow, so it doubles when the flow halves. Chapter 5 doses methanol at the tree and judges it at the riser base hours later,

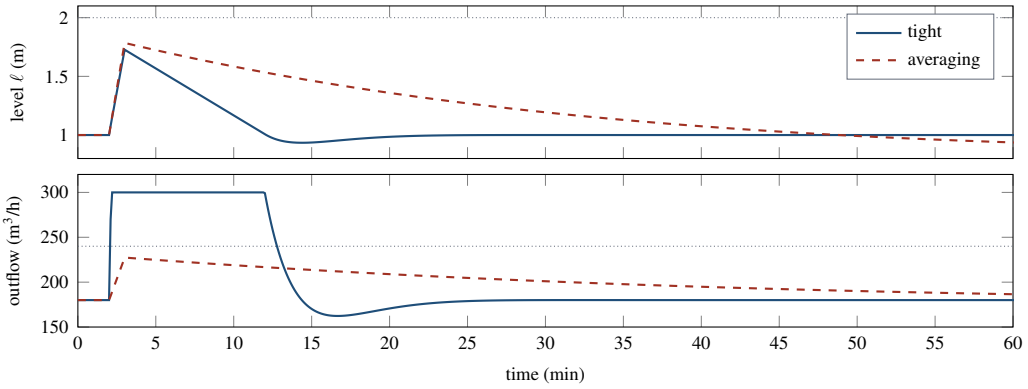


Figure 1.4 A 20 m^3 Slug Into W-1's Separator. Tight Tuning (Solid) and Averaging Tuning (Dashed) Store Nearly the Same Volume. The Dotted Lines Are the High-High Level Trip and the 240 m^3/h Limit of the Equipment Downstream.

through exactly this delay, and its constraint, the hydrate margin, is smallest where the riser's weight keeps the pressure high [22]. Chapter 6 receives the riser's slugs, and Chapter 13 carries its batches through the same displacement.

1.4.4 Separator Inventory and Averaging Level Control

The separator's liquid is an inventory between an inflow that the wells set and an outflow that a valve sets. W-1's separator has a liquid surface of 25 m^2, a normal level of 1.0 m, a high-high trip at 2.0 m, and an oil valve that passes 300 m^3/h wide open at the separator pressure. Its level is an integrator, as shown in Equation (1.6).

$$A_s \frac{\mathrm{d}\ell}{\mathrm{d}t} = q_{in} - q_{out}, \qquad q_{out} = q_{max}\,\frac{u}{100}, \qquad k = \frac{q_{max}}{100\,A_s} \tag{1.6}$$

so one percent of valve moves the level at $k = 0.12\,\mathrm{m/h}$. A PI controller on an integrator gives the characteristic polynomial $\tau_c\tau_I s^2 + \tau_I s + 1$ with $\tau_c = 1/(K_c k)$ and a damping ratio $\zeta = \frac{1}{2}\sqrt{\tau_I/\tau_c}$. Two tunings with the same damping, $\zeta = 0.894$, are compared. The tight one has $K_c = 400\,\%/\mathrm{m}$ and $\tau_I = 4\,\mathrm{min}$, so $\tau_c = 1.25\,\mathrm{min}$. The averaging one has $K_c = 20\,\%/\mathrm{m}$ and $\tau_I = 80\,\mathrm{min}$, so $\tau_c = 25\,\mathrm{min}$, twenty times slower.

A slug of 20 m^3 arrives at 2 min over one minute, an inflow of 1380 m^3/h against a valve that passes at most 300 m^3/h. Whatever the tuning, at least 18 m^3, 0.72 m of level, must be stored during the arrival. Figure 1.4 shows the two runs. The tight loop peaks at 1.727 m, with its valve wide open from 2.15 min, and returns within 5 cm of the setpoint by 11.5 min. It hands the slug to the second stage at the full 300 m^3/h, above a downstream limit of 240 m^3/h for 10.7 min. The averaging loop peaks only 5.8 cm higher, at 1.785 m, never passes more than 227.3 m^3/h and returns within 5 cm by 42.6 min. The largest rate of change of the outflow over six seconds falls from 900 to 47.9 m^3/h per minute. The vessel stores the slug under either tuning. The tuning decides only how fast the stored liquid is passed on.

DID YOU KNOW? Three Barrels of Water for Each Barrel of Oil

The water produced with oil rose from less than 30 million barrels a day in 1990 to about 100 million in 2015 [23]. Wells give up an estimated three barrels of water for every barrel of oil [24], and a separator like that of Chapter 6 must part them.

Averaging control was posed as an optimization: pass an inflow change to the outflow as smoothly as the level band allows [25]. For a step Δq in the inflow, the smoothest response ramps the outflow at a constant rate r. The level then gains the area of a triangle, $\Delta q^2/(2r)$ in volume, which must fit in the band $\Delta\ell$, as shown in Equation (1.7).

$$r \geq \frac{\Delta q^2}{2 A_s \Delta \ell} \tag{1.7}$$

A second well of $60\,\mathrm{m^3/h}$ brought on line, with 0.5 m of band, can be passed on at $2.4\,\mathrm{m^3/h}$ per minute, over 25 min. No tuning does better, and a tighter loop uses less of the band and ramps faster. Chapter 6 measures exactly this trade on a riser slug, with the extra cost that a higher level brings foam nearer the mist extractor. Chapter 10 detunes its surge drum's level loop so that the drum, not the three columns behind it, takes a rate change. Chapter 15 has an integrating inventory too, the vapor space of its tanks, whose outlet capacity is set by the send-out.

1.4.5 Pipeline Inventory: Line Pack and Pressure Waves

A pipeline is also an inventory, and whether it is a large one depends on what it carries. Gas is stored in proportion to its pressure. For an isothermal gas with wave speed $c = \sqrt{ZRT}$ the mass in a line of bore area A_p and length L is Equation (1.8).

$$M = \frac{A_p L \bar{p}}{c^2}, \qquad \frac{\mathrm{d}\bar{p}}{\mathrm{d}t} = \frac{c^2}{A_p L}\left(w_{in} - w_{out}\right) \tag{1.8}$$

where $\bar{p}$ is the mean pressure and w a mass flow. W-1's gas leaves through 50 km of 0.4 m line with $c = 350\,\mathrm{m/s}$, the speed Chapter 12 uses. The line holds 5.13 t of gas per bar, and 307.7 t at a mean of 60 bar. A deficit of 5 kg/s, a tenth of the line's 50 kg/s, lowers the mean pressure by 3.5 bar an hour, so a 2 bar margin lasts 34 min. The line rides through a change in demand on its own pack while the compressor catches up. In steady flow the square of the pressure falls linearly along the line, $p_{in}^2 - p_{out}^2 = K w^2$, and the pressure at the far end follows a change at the near end as a slow diffusion rather than a delay [26].

DID YOU KNOW? Three Million Miles of Gas Pipe

About 3 million miles of mainline and other pipelines link the gas fields and storage of the United States with consumers. In 2022 they delivered about 29.2 trillion cubic feet of natural gas to about 78.3 million consumers [13].

A liquid line stores almost nothing. Oil at $840\,\mathrm{kg/m^3}$ with a pressure wave speed of 1130 m/s in its pipe has an effective bulk modulus $\rho c^2 = 1.07\,\mathrm{GPa}$, so W-1's 20 km oil line of 0.3 m bore, $1414\,\mathrm{m^3}$, takes up only $0.132\,\mathrm{m^3}$ per bar. A mismatch of $18\,\mathrm{m^3/h}$, a tenth of the oil, moves the pressure by 2.3 bar a minute. A sudden stop is worse. A column moving at velocity v that is stopped raises the pressure by the water hammer rise of Equation (1.9) [27].

$$\Delta p = \rho\, c\, \Delta v \tag{1.9}$$

which is 19.0 bar for a 2 m/s stop, arriving at the other end of the 20 km line 17.7 s later. The two lines need different control. The gas line's pack gives the controller time and makes a forecast worth acting

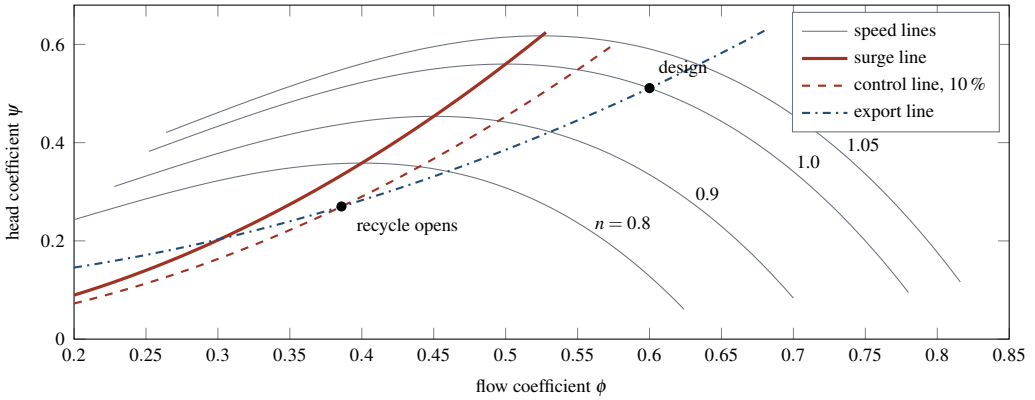


Figure 1.5 W-1's Export Compressor. Thin Speed Lines Scale by the Fan Laws. The Thick Surge Line Passes Through Their Peaks, the Dashed Control Line Lies 10 % to Its Right and the Dash-Dotted Curve Is the Head the Export Line Asks For.

on. The liquid line gives no time, so pumps in series must be held by fast limit loops. Chapter 12 plans its pack before the evening peak, and Chapter 13 joins suction and discharge limits at low selectors because one flow passes every station. Chapter 11 keeps a line pack behind its station, and the tank vapor space of Chapter 15 is an inventory of the same kind.

1.4.6 Compressor Maps and Surge

A centrifugal compressor's head falls as its flow rises, and the map draws that characteristic at each speed. The fan laws scale it, as shown in Equation (1.10).

$$\phi \propto n, \qquad \psi \propto n^2, \qquad P \propto n^3, \qquad \psi(\phi, n) = n^2\,\Psi(\phi/n) \tag{1.10}$$

where Ψ is the characteristic at design speed and P the power. W-1's export compressor uses the cubic characteristic of Chapter 11 [28], which peaks at $\phi/n = 0.5$. Left of the peak the head rises with flow, the compression system loses its steady operating point, and the flow cycles, reversing through the machine [29]. That is surge. Because the peak scales with the fan laws, the surge line is the parabola $\psi = 2.24\,\phi^2$ through the peaks, and the surge margin at the present speed is Equation (1.11).

$$SM = 100\,\frac{\phi - \phi_s(n)}{\phi_s(n)}, \qquad \phi_s(n) = 0.5\,n \tag{1.11}$$

The anti-surge controller holds a control line at 10 % margin, $\phi = 0.55\,n$, by opening a recycle valve that returns discharge gas to suction.

What the compressor must deliver is set by the system it feeds. W-1's export line asks for a head $\psi = 0.1 + 1.142\,\phi^2$: a static part, the pressure ratio the line holds even at low flow, and a part that grows with friction. The design point is $\phi = 0.6$ at full speed, with $\psi = 0.511$ and a margin of 20 %. Figure 1.5 draws the map. As the export flow falls, the speed falls along the system curve and the margin shrinks, 18.3 % at 90 % of design flow and 12.6 % at 70 %. The system curve meets the control line at $\phi = 0.386$ and $n = 0.70$, 64.3 % of design flow. Below it the machine stays on the control line and the recycle carries the difference. At 50 % of design flow the machine runs at $n = 0.608$ and recycles 11.5 % of the forward flow, power spent for nothing. The static head is what pushes the machine toward surge. A system of pure friction would follow a fan-law parabola parallel to the surge line and keep its margin at any speed.

Chapter 11 runs three such machines on one header, where every machine sees the same head and the sharing law decides which nears its surge line first. Chapter 12 places a compressor map at the

Table 1.1 Water Leaving W-1's Two-Stage Contactor (mg/Sm^3) From 1000 mg/Sm^3 at the Inlet, by Absorption Factor and Lean Floor

Absorption factor $\mathscr{A}$	1	1.5	2	3	4	8	∞
Share of the excess removed	0.667	0.789	0.857	0.923	0.952	0.986	1
Floor $y^* = 50$	366.7	250.0	185.7	123.1	95.2	63.0	50.0
Floor $y^* = 25$	350.0	230.3	164.3	100.0	71.4	38.4	25.0

inlet of its line, Chapter 15 stages two boil-off compressors that recycle what they do not send, and the liquefaction train of Chapter 14 is limited by the power of its refrigerant compressor's driver.

1.4.7 Absorption and Cryogenic Gas Processing

Gas leaving the separator carries water and, in a sour field, hydrogen sulfide and carbon dioxide. Both are removed by absorption in a solvent that is regenerated and recycled: glycol for water and an amine for the acid gases [30]. In a counter-current absorber of N ideal stages, the gas leaving the top cannot be cleaner than the gas in equilibrium with the lean solvent entering it, $y^* = mx_{lean}$, where m is the equilibrium ratio. The share of the removable excess that the column removes depends on the absorption factor, the solvent rate relative to the gas rate and the equilibrium, as shown in Equation (1.12).

$$\frac{y_{in} - y_{out}}{y_{in} - y^*} = \frac{\mathscr{A}^{N+1} - \mathscr{A}}{\mathscr{A}^{N+1} - 1}, \qquad \mathscr{A} = \frac{L}{mG} \tag{1.12}$$

which is the Kremser relation, with L and G the molar flows of solvent and gas.

W-1's gas enters a two-stage contactor with 1000 mg/Sm^3 of water, against an example specification of 110 mg/Sm^3. Table 1.1 lists the outlet for two lean purities. With a floor of $y^* = 50\,\text{mg/Sm}^3$, doubling the absorption factor from 2 to 4 lowers the outlet from 185.7 to 95.2 mg/Sm^3, and no circulation reaches below 50. The specification needs $\mathscr{A} = 3.38$. Purer solvent lowers the floor to 25 and the specification then needs $\mathscr{A} = 2.77$, 18 % less circulation. Circulation is the fast handle with a floor, and purity is the slow handle that moves the floor. Purity is made in the regenerator and reaches the contactor only as the solvent inventory turns over, four hours for a 6 m^3 surge drum at 1.5 m^3/h.

Chapter 7 finds the floor on its contactor, with stage efficiencies in place of ideal stages, and a ceiling on its reboiler that caps purity. Chapter 8 has two solutes on one solvent. Hydrogen sulfide approaches equilibrium quickly while carbon dioxide is absorbed at a rate set by contact time, so the solvent rate also sets the selectivity.

Ethane, propane and heavier hydrocarbons are recovered by cooling the gas until they condense. The cold comes from expansion. An expander that does work on a shaft follows, at best, the isentropic path, which for an ideal gas with heat capacity ratio γ ends at the temperature of Equation (1.13).

$$T_{2s} = T_1 \left(\frac{p_2}{p_1} \right)^{(\gamma-1)/\gamma}, \qquad T_2 = T_1 - \eta \left(T_1 - T_{2s} \right) \tag{1.13}$$

with η the isentropic efficiency. W-1's gas at −30 °C and 60 bar expanded to 25 bar with $\gamma = 1.3$ reaches −74.5 °C ideally and −67.8 °C at $\eta = 0.85$. A valve that throttles the same gas does no work and keeps its enthalpy. An ideal gas would not cool at all, and a real gas cools only by its Joule-Thomson effect. On the real-gas model of Chapter 9, expanding the same separator vapor to 23 bar, the expander reaches −70.9 °C where the valve reaches −49.2 °C. That chapter's plant trades recovery against recompression power through the column pressure and switches between recovering and rejecting ethane as prices move, and Chapter 10 splits the recovered liquid into products. Chapter 14 cools the whole stream to liquid, where the closest approach between the warming refrigerant and the cooling gas inside the exchanger decides the power the plant burns.

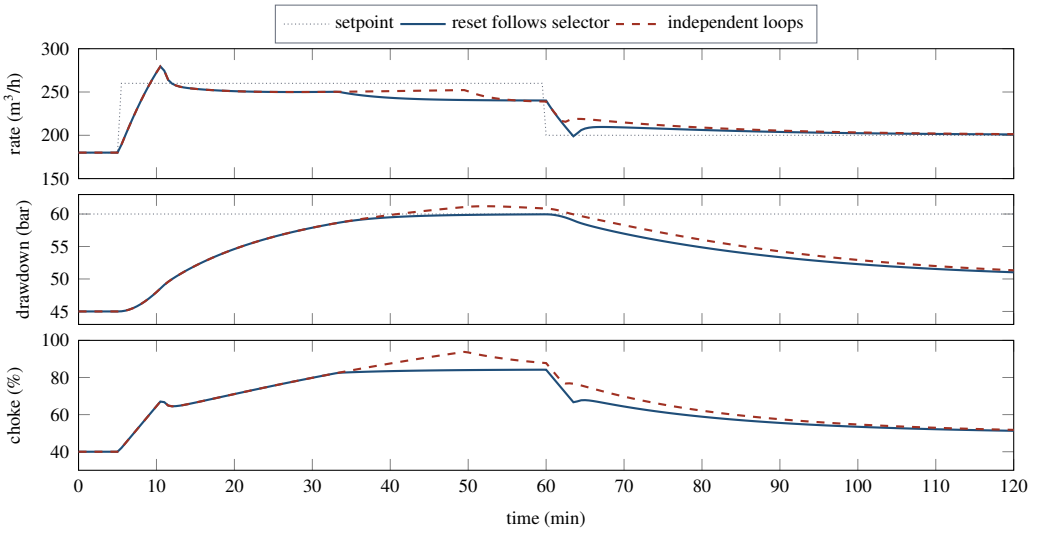


Figure 1.6 A Rate Loop and a Drawdown Override on W-1's Choke. With Both Reset Terms Following the Selected Output (Solid) the Override Holds the 60 bar Limit (Dotted). With Independent Loops (Dashed) the Waiting Override Has Wound Up and the Limit Is Broken.

1.4.8 Override and Constraint Control as Practiced

One actuator often has several jobs. W-1's choke holds the rate, but the rate must not draw the well down by more than 60 bar, a limit set by sand or water coning, so no more than 240 m^3/h may be produced. A rate loop (FIC) and a drawdown limit loop (PDIC) both compute a choke opening, and a low selector applies the smaller, as shown in Equation (1.14).

$$x = \min(u_F, u_D), \qquad u_j = K_{c,j}\, e_j + F_j, \qquad F_{j,k} = F_{j,k-1} + \left(1 - \mathrm{e}^{-\Delta t/\tau_{I,j}}\right)\left(x_k - F_{j,k-1}\right) \tag{1.14}$$

written in the external-reset form that Chapter 13 uses. Each loop's reset term F_j follows the opening actually applied, not the loop's own output. The selected loop is then exactly a PI controller. The other rests above the applied opening by its own proportional term, so it takes over the moment its error says it should [31]. The framework's back-calculation of Section 1.6.2 does the same when the applied output is the selected one.

The example's loops run every 0.1 min: FIC with $K_c = 0.5$ % per m^3/h and $\tau_I = 6$ min on a meter with a 2 min lag, PDIC with $K_c = 3$ %/bar and $\tau_I = 5$ min, and a choke that travels 5 %/min. The rate setpoint is raised from 180 to 260 m^3/h at 5 min, which the limit forbids, and lowered to 200 m^3/h at 60 min. Figure 1.6 compares two versions. In the first, both reset terms follow the applied opening. The rate overshoots to 282 m^3/h as the choke opens ahead of the draining wellbore, the override holds the choke from 33.3 min, and the drawdown peaks at 59.96 bar. In the second, each loop protects only itself from its own output limit, the arrangement of two independent controllers wired to a selector. The override's integral sits near 100 % while it waits, so it takes the choke only at 49.6 min. The drawdown passes 60 bar at 41.0 min, peaks at 61.2 bar and stays above the limit for 21.8 min. The same two loops with the same tunings protect the well in one wiring and not in the other.

Overrides recur in nearly every chapter, in three shapes. A low selector lets the most restrictive limit close a shared actuator: four loops on the choke in Chapter 4, rate and pressure on the choke in Chapter 5, suction and discharge limits on each pump in Chapter 13, the driver's power on the feed in Chapter 14, the recondenser on the compressor load in Chapter 15 and a high-limit override on the reboiler in Chapter 7. A high selector lets the most demanding protection open an actuator: the methanol dose in Chapter 5, the loading floor on the amine circulation in Chapter 8 and every

anti-surge valve in Chapter 11. A supervisor moves setpoints toward the active limit and backs away from it: the hill climber that backs off the most violated of six constraints in Chapter 2, the grid search above a protection layer that can only cut the pump in Chapter 3, and the rule-based supervisor of Chapter 6. Many chapters also offer a predictive controller, described in Section 1.6.3, as an option measured against the selectors.

1.4.9 Where Measurement Is the Limit

An override can protect only a limit it can see. In this industry the quantity that binds is often not the quantity a transmitter reads, and three cases recur. A product quality is measured late by an analyzer, so a loop on it acts on the past. A constraint sits where no instrument can be placed, so it is inferred from what can be measured, and the loop holds the inference exactly while the true value drifts [32]. An inventory is computed from a few transmitters, so its rate of change, which a leak detector needs, is uncertain during exactly the transients when it matters [33].

The example has one of each. The rate meter reads late, which is why the rate overshoots in Figure 1.6. The drawdown is measured only because the example puts a gauge at the bottom of the well. Without it the drawdown would be computed from the wellhead pressure and a column weight that changes with the holdup, and an error of a few bar in that weight would move the override's point of action by as much. The line pack of the export line is inferred from a few pressures. A leak detector compares inflow, outflow and the change in pack over a window, as shown in Equation (1.15).

$$I = \frac{1}{T_w}\int_{t-T_w}^{t} (w_{in} - w_{out})\,\mathrm{d}t' - \frac{\hat{M}(t) - \hat{M}(t-T_w)}{T_w} \tag{1.15}$$

An error of only 0.1 bar in the change of the estimated mean pressure over a 30 min window is 0.285 kg/s of apparent imbalance, 0.57 % of the line's flow, before any meter error is counted. A smaller leak cannot be told from the estimate's own error.

Table 1.2 lists, for each chapter, the limit that binds, how the controller knows it and the structures the chapter uses. Chapter 16 adds the measured results and asks which techniques recur from well to terminal and where measurement is the limit.

1.5 Instruments and Final Elements in This Volume

Every loop in this volume begins at a sensor and ends at a final element, the device that changes a flow of material or of energy. Figure 1.7 lists what the controllers of each chapter read and what they move. A flow is metered in thirteen chapters and a pressure read at the surface in ten, and eleven chapters act on a quantity that is inferred or estimated rather than measured. A valve or a choke is moved in eleven chapters, a pump in seven and a compressor or its driver in five.

Each device sets a limit that no tuning removes. Most pressures are read by a diaphragm that bends under the process pressure, with strain gauges bonded to it whose resistance changes as it deforms [34], as drawn in Figure 1.8. Two chapters read the pressure at the bottom of a well. Without such a gauge it is computed from the surface, as Section 1.4.9 explains for W-1. Three chapters read a multiphase meter, which gives the oil, water and gas of a well stream without separating it. A test separator is reliable and accurate but too slow for monitoring in real time, and the water and gas that a meter must read grow as a reservoir matures [35].

The analyzers bring the longest delays. A chromatograph separates the components of a sample in time, as shown in Figure 1.9, so cycles of minutes are not uncommon and some dead time is inevitable [34]. The moisture analyzer of Chapter 7 reports 10 min after its sample left the gas, and the product chromatograph of Chapter 9 samples every 10 min and reports 6 min later.

The final element is most often a choke or a control valve, and on lines and injection wells a variable-speed drive, which sets the speed of a motor by changing the frequency of the power that feeds it [34]. A drive turns the compressor of Chapter 12 and the four pumps of Chapter 13.

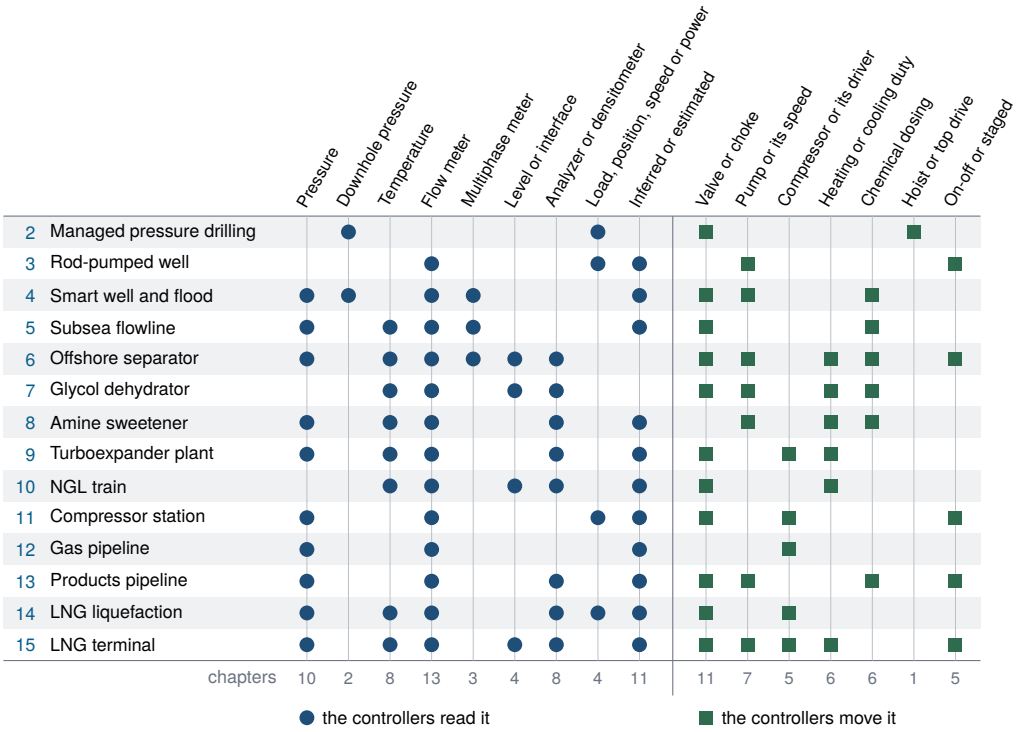


Figure 1.7 Sensors and final elements by chapter. A round mark means that the chapter's controllers or operator read the device, a square mark that they move it. A mark is set only where the chapter's variables table, its instrument text or its section on technology in practice names the device or its reading, so a blank says that the model leaves the device out and not that a real plant lacks it. Notice that eleven of the fourteen chapters act on an inferred or estimated quantity.

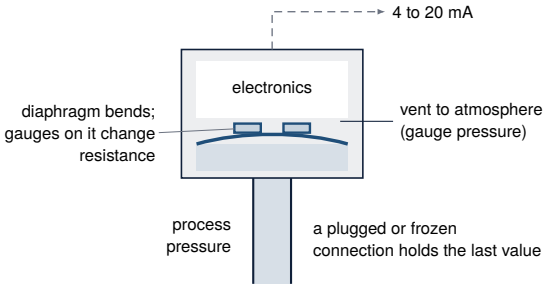


Figure 1.8 Principle of a gauge pressure transmitter. Notice that the reading comes from the bending of one diaphragm, so a metal element that has been overpressured can lose accuracy without any change in the process [34].

Table 1.2 For Each Chapter, the Limit That Binds, How the Controller Knows It and the Structures Used

Ch.	Binding limit	How it is known	Structures
2	Vibration, cuttings, torque, pore and fracture pressure	Six constraint ratios, two of them vibration indicators over 20 s windows	PI loops, supervisor backing off the worst limit, choke MPC
3	Pump fillage, fluid level, rod load	Inferred from the surface load cell through the downhole card	Timer, PI on the inferred level, grid search over a protection layer
4	Drawdown, sand, water, the flood fronts	Drawdown from a model, a delayed multiphase meter, fronts never measured	Four PI loops at a low selector, injection override, planner
5	Hydrate margin at the weakest point	Profile estimator from a few pressures and temperatures	Rate loop with pressure override, methanol feedforward with trim, sequence
6	Levels between trips, oil and water quality	Level transmitters, analyzers behind sample delays	Averaging level, inlet feedforward, rule-based supervisor, MPC
7	Water in the gas, reboiler ceiling	Moisture analyzer reporting late	Ratio feedforward, analyzer trim, high-limit override, economic MPC
8	Hydrogen sulfide in the gas, rich loading	Sampled analyzer, loading inferred from an acid-gas meter	Ratio, loading floor at a high select, analyzer trim, MPC
9	Methane in the liquid product, flooding, power	Chromatograph after a delay, tray temperature as a proxy	Two PI loops, analyzer trim, mode sequence, supervisor
10	Ethane in propane, shared steam	Chromatograph one column downstream, inferential estimate	Local, analyzer and cascade loops, averaging drum level, MPC
11	Surge margin of each machine	Computed from measured flow and head on the commissioned map	Pressure loop, sharing law, anti-surge overrides, sequencer
12	Delivery pressure floor, line pack	Pack inferred from three transmitters, a demand forecast	Cascade, forecast feedforward, pack window, leak detector, MPC
13	Suction and discharge limits, a full hill, batch cuts	Station pressures, densitometer at the terminal, volume balance	Low selectors in external-reset form, back pressure loop, MPC
14	Temperature approach in the exchanger, driver power	Computed from flows, pressures and a sampled refrigerant analysis	Cascade, inferential approach loop, low selector on feed, MPC
15	Tank pressure below the flare, recondenser capacity	Tank pressure, capacity from subcooling, boil-off observer	PI with recondenser override, staging, nomination preview, MPC

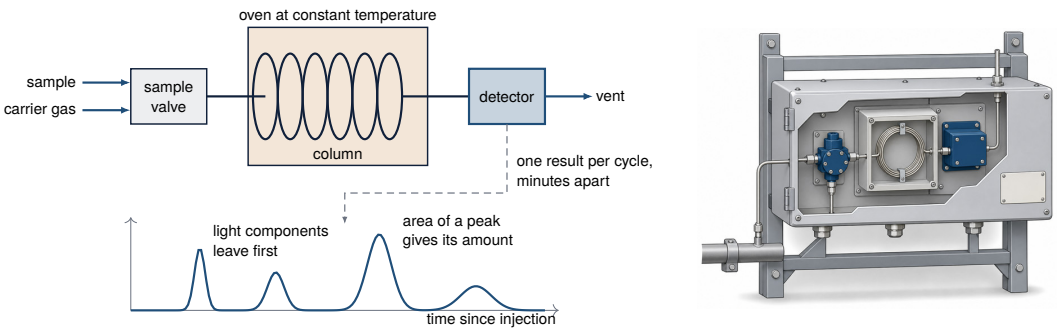


Figure 1.9 Principle of a process gas chromatograph, with a view of the device as installed. Notice that one sample gives one result at the end of the cycle, so a loop on this instrument acts on a value that is always old.

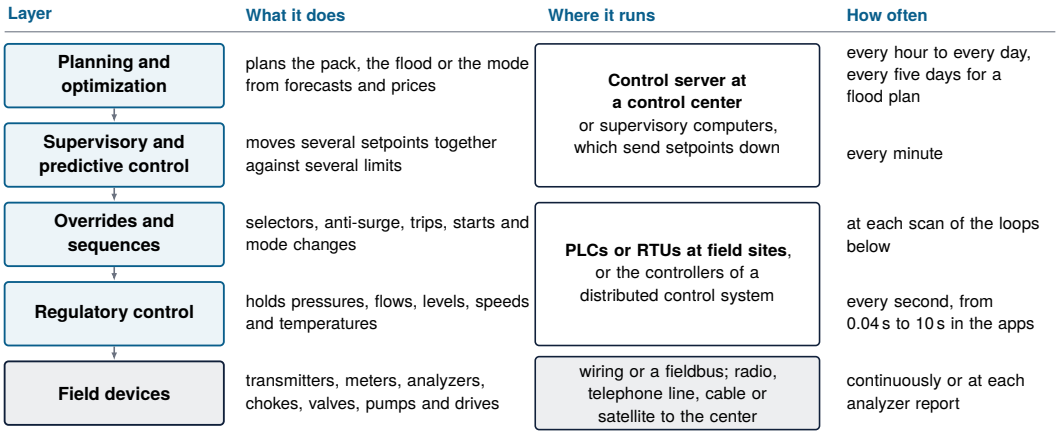


Figure 1.10 The control layers of this volume, the platform that hosts each and how often each acts. The periods of the regulatory, predictive and optimization layers are those of [36, 37], the platforms those of [11], and the ranges in the apps those of the chapters. Notice that each layer sends setpoints to the layer below and acts less often than it.

Table 1.3 What Every App Provides. Standalone Apps Are Described in Their Own Chapters.

Element	What it does
Run, Pause, Reset	Start, hold and restart the run. Reset keeps the reader's settings.
Speed	Simulated time per second of real time. It changes the pace and never the result.
Manual and Auto	Hand each actuator to the reader or to the controller, without a bump.
Scenario list	Numbered cases, each with a briefing, a completion rule and a reload button.
Control panel	Setpoints, manual outputs, tuning, options and disturbances, at a Beginner or an Advanced level.
Trends	SP and PV above the MV for each loop, with limits, event marks and a cursor.
Scorecard	IAE, overshoot, settling time, MV travel and time at a limit for each loop.
Learn card	Predict, Try It, two hints and Explain, opened by the reader in that order.
Show Model	Equations, live variable values, editable parameters and assumptions.
Export CSV	Every logged sample of the run as a text file with a header row.
Instructions	A separate page with background, an operating guide and the model's limitations.

The platform sets the last limit. A pipeline is run through a supervisory system. A control server at a control center reaches the field sites by radio, telephone line, cable or satellite, and remote terminal units or programmable controllers at those sites control the local process. Cycle times run from milliseconds to minutes [11]. Figure 1.10 places the layers on that hardware. Regulatory loops act about every second, predictive control every minute and optimization every hour to every day [36, 37]. The apps run their loops every 0.04 s in Chapter 3, every 1 s in Chapter 13, whose predictive controller acts every 2 min, and every 10 s in Chapter 12.

1.6 The Simulator Framework

Every chapter of this book has an app that runs in a web browser, collected with the others at https://apmonitor.com/apps/. About 240 apps exist across the series. All but a few are built on one shared framework, so the engine, the PID algorithm, the trends and the data export behave the same way in every chapter. The exceptions are standalone studios, such as the cascade studio and the two Temperature Control Lab apps, which follow the same ideas in their own code. This section describes the framework once and states where the standalone apps differ. Table 1.3 lists what a reader meets.

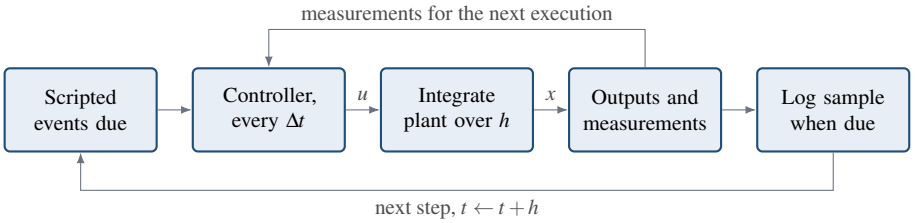


Figure 1.11 Work of the Engine in One Integration Step

1.6.1 The Simulation Engine

Each app carries a plant model written as ordinary differential equations $\frac{\mathrm{d}x}{\mathrm{d}t} = f(x,u,d,p)$ in the states x, the inputs u, the disturbances d and the parameters p. The engine integrates the model with the classical fourth-order Runge–Kutta method on a fixed step h set for the fastest dynamics of the app. Some apps supply their own step, such as explicit Euler for a discontinuous right-hand side or an exact update for a model driven by events.

Three clocks run at separate rates. The plant advances every step h. The controller executes every control period Δt, equal to h unless the app declares a slower controller, and holds its outputs between executions. The log takes one sample at a fixed interval of simulated time for the trends, the scorecard and the exported file. The order of work in one step is shown in Figure 1.11. The controller also executes once when a scenario is loaded, so every run starts from an initialized controller.

Drawing the screen is kept apart from these clocks. A speed of 5 means five units of the app's time base per second of real time. Speed, frame rate and pauses change how fast a run is shown and never what it computes, so a phone and a workstation produce the same trajectory and the same file.

Noise comes from a deterministic pseudorandom generator. Its seed is declared by the app, or by a scenario that needs its own noise record, and it is restored whenever a scenario is loaded or Reset is pressed. The framework gives the reader no seed setting, so *default seed* in this book means the seed the app ships with. A run repeats bit for bit when the same scenario is loaded and the same settings are made before Run. A change made by hand during a run lands on a different step each time, which is one reason numeric answers carry a tolerance. The standalone simulators draw noise from the browser's own generator, so a noisy run there repeats in its statistics and not sample for sample.

Run and Pause start and hold the clock. Reset restarts the run at $t = 0$, paused, and keeps the reader's settings: every setpoint, tuning value, option, model parameter, disturbance slider and mode the reader changed is put back before the plant and the controller are initialized. Scripted events replay, and a one-shot disturbance button pressed during the run is cleared. The reload button beside the scenario list restores the scenario's own settings instead.

1.6.2 The Velocity-Form PID Algorithm

One PID implementation serves the PID loops of every framework app. It is the velocity (incremental) form: each execution adds the change in the proportional and derivative terms to a running sum, and the integral is a separate term, as shown in Equations (1.16) to (1.19) for execution k and control period Δt.

$$e_k = a\,(\mathrm{SP}_k - \mathrm{PV}_k), \qquad P_k = K_c\,e_k, \qquad D_k = -a\,K_c\,\tau_D\,\frac{\mathrm{PV}^f_k - \mathrm{PV}^f_{k-1}}{\Delta t} \tag{1.16}$$

$$S_k = S_{k-1} + (P_k - P_{k-1}) + (D_k - D_{k-1}) \tag{1.17}$$

$$v_k = u_b + S_k + I_{k-1}, \qquad u_k = \mathrm{sat}(v_k) \tag{1.18}$$

$$I_k = I_{k-1} + \left[\frac{K_c}{\tau_I}\,e_k + \frac{u_k - v_k}{T_t}\right]\Delta t \tag{1.19}$$

where $a = \pm 1$ is the controller action, PV^f is the measurement after an optional first-order filter, u_b is the output bias, S is the running sum of proportional and derivative increments, v is the output before limiting, u is the output sent to the actuator, I is the integral term and T_t is the tracking time constant, equal to τ_I unless a loop sets another value. The function sat clamps v_k to the actuator range and then applies a rate limit if the loop declares one.

The limit acts on the emitted output u_k only. The sum S_k is never clamped, so it equals $P_k + D_k$ and the output before limiting is the position form $u_b + K_c e_k + D_k + I_{k-1}$. Anti-windup touches the integral alone. While the output is limited, the last term of Equation (1.19) drives I toward the value that reproduces the output applied, which is back-calculation [38]. With no integral action, Equation (1.19) is skipped. Clamping the running sum would instead discard each increment that crosses a limit and keep each one that returns, so symmetric noise would push the output off the limit even in a proportional-only loop.

The derivative acts on the measurement, so a setpoint step gives a proportional kick and no derivative kick. A change of K_c acts on the whole error at the next execution, as in the position form. The action $a = +1$ suits a process whose PV rises with the MV, a reverse-acting controller in the ISA convention, and $a = -1$ gives the direct-acting one.

In Manual the controller keeps executing and sets $I_k = u_{\mathrm{man}} - u_b - S_k$, so the output it would send equals the manual value u_{man}. The first automatic output differs from it by one period's change in P and D. In Auto the engine copies each controller output to the manual station, so a switch to Manual is also bumpless. The standalone studios with PID loops use velocity-form code of their own, with the limit on the output and anti-windup on the integral alone, in which a gain change applies only to later increments.

1.6.3 The Predictive Control and Estimator Options

More than 100 apps offer model predictive control (MPC) as an option beside the structure they open with, which stays the default. Volume 9 develops the theory. In most of them the controller plans with a small linear model in discrete time as shown in Equation (1.20).

$$x_{k+1} = A\,x_k + B\,u_k + E\,w_k + c, \qquad y_k = C\,x_k \tag{1.20}$$

where w is a measured or scheduled disturbance and c is a constant left by linearization. The model comes from a step test at the design point, from the app's balances linearized there or again at every decision, or from the model the reader fits. It never reads the true states or parameters of the run, so a parameter the reader changes in Show Model is model mismatch, as on a real plant.

At each decision the controller minimizes the objective of Equation (1.21) over a horizon of N steps.

$$J = \sum_{k=1}^{N} (y_k - r_k)^{\top} Q\,(y_k - r_k) + \sum_{j} \Delta u_j^{\top} R\,\Delta u_j + \sum_{i} \rho_i\, s_i^2 \tag{1.21}$$

where r is the setpoint, Q weights the tracking error, R is the move suppression on each input change Δu_j and ρ_i is the price of a slack $s_i \geq 0$ that lets an output band be broken when it cannot be held. Actuator limits are hard constraints. The quadratic program is solved in the browser by a dual active-set method [39]. The first move is applied, and the plan is made again at the next decision.

The plan starts from an estimate of the model state. About half of these apps use a steady-state Kalman filter on the model with an added disturbance state, which absorbs what the model did not predict and removes steady offset without an integral term [40]. Others correct the model output with a filtered bias to the same end. About a dozen apps also offer moving horizon estimation (MHE), which solves the estimation problem over a window of recent samples and can hold bounds that a Kalman filter cannot, such as a concentration that is never negative [41].

With the MPC selected, the tuning fields of the loops it replaces leave the Control panel. In most of these apps the Advanced level shows up to three settings, each a multiple of the shipped value: the prediction horizon, the move suppression and the band price. The trend draws the accepted plan

ahead of the present time, and a list of constraint prices shows which limits bind. A plan that is not optimal, not finite or outside its limits is rejected, and a fallback, in most apps the shipped controller, holds the actuator until a plan succeeds. The Model Fitting and Tuning Studio builds the same kind of controller from the reader's own data in its tab, MPC from the fit.

1.6.4 Scenarios, Progress and Data

An app teaches through scenarios, the numbered cases at the top of the Control panel, nine to thirteen in most apps. Loading one sets the model parameters, the controller settings, the disturbances, the mode, the initial state and any scripted events. The information button beside the list opens a briefing in four parts: what the case is, why it matters, what to do and one detail for the practicing engineer. The Learn card holds the lesson back, with Predict and Try It before the run, hints on request and Explain afterward.

Each scenario has a completion rule that the app judges from the run on the simulated clock: a run length, the scripted events passed, a setting made by the reader, a state held for a stated time or a target met. Choosing a case earns nothing. A completed case is marked in the list, and a badge row shows the count and a level, Explorer after one case, Practitioner at half and Completed at all. Progress is kept in the reader's browser with no account. A certificate of progress or of completion can be issued from it, as a record of self-paced work and not an assessed credential.

One case, usually the last, is the graded challenge. It states a target, such as an error band to hold through a sequence of upsets. Changing the model parameters or the disturbances voids the attempt in most apps, and a record made with Reveal suggested tuning says so.

Show Model opens the equations, live values of the variables, the editable parameters and the assumptions. The trends plot SP and PV above the MV for each loop, with actuator limits and marks at setpoint moves, disturbances, mode changes and scripted events. The scorecard reports for each loop the integral of the absolute error (IAE), overshoot, settling time, MV travel and share of time at a limit, with the app's own measures, and it does not combine them into one score. IAE and travel accumulate from Reset, so two runs are compared at the same clock reading.

Export CSV saves the log as a text file with a header row of column names that carry their units and one row per logged sample. In most apps the time column is in seconds, minutes or hours according to the run length. The log holds the most recent 60 000 samples unless the app sets another size. Apps with the MPC option begin the file with comment lines that start with #, which a script must skip.

An app exercise in this book names the scenario by its number and name as the app shows them, every setting that differs from the scenario's defaults, the run length in simulated time, the seed (the default seed unless stated), what to read and where (a trend, a scorecard tile or a column of the exported file) and a tolerance on a numeric answer. Settings are made before Run unless the exercise gives a time.

1.7 Working With AI Agents While Studying

A reader of this book is likely to have an AI assistant within reach, and the apps make its work easy to test. This section gives the editor's view of how to use one while working through a chapter. In short, the assistant may do the algebra and the scripting, and the reader keeps the checking.

What to delegate. Tasks with a clear input and a result that can be tested are good candidates. An assistant can derive the linearization of a chapter's balance at an operating point and the transfer function that follows from it. It can write a short script that reads the exported CSV file and fits a first-order-plus-dead-time model to a step test. It can propose a first tuning from that fit with a named rule, convert units, explain a term in a model or outline a report. The assistant knows nothing about the app until it is told. Give it the equations and parameter values from Show Model, the column names of the exported file and the scenario that produced the data, as described in Section 1.6.4. Ask it to state its assumptions and the units of every result.

What to check, and how. An assistant returns a fluent answer whether or not the answer is right, so each result needs a test that it could fail. Four tests cover most of the errors a chapter will produce.

- *Units.* Every term of an equation carries the same units, and a process gain is in PV units per MV unit. A gain computed with the valve as a fraction is 100 times the gain of a loop whose valve is in percent.
- *Signs.* Decide from the physics which way the PV moves when the MV rises. Compare that with the sign of the fitted gain and with the controller action of Section 1.6.2.
- *Steady state.* Set the derivatives to zero. The model must return the operating point that Show Model reports, and a linearization must give zero deviation there.
- *A run on the app.* Load the named scenario at the default seed, enter the proposed numbers and run. The trend and the scorecard decide. A tuning that reads well on paper and holds the valve at a limit for half the run has failed.

A fitted model gets one more test: compare it with data it was not fitted to, such as a second step in the opposite direction. When a check fails, report the failure to the assistant with the evidence and ask for the cause, then check the correction the same way.

How the exercises train the habit. One exercise in each chapter is marked *Review the agent*. It presents a tuning, a model or a claim as an AI assistant might propose it, with a flaw, and asks the reader to find the flaw in about five minutes. The flaws are ordinary ones: a unit, a sign, a steady state that does not close, a result measured on the wrong scenario or at a different run length, or a limit that was ignored. Each is found with one of the tests above. The purpose is not distrust. An engineer who can say why an answer is right can hand an assistant much larger tasks and still stand behind the result. That larger picture, in which agents do most of the routine work and the engineer supervises and checks, is the subject of Chapter 16.

1.8 Chapter Roadmap

Table 1.4 lists the chapters with the lesson each adds to this one. The level says how much of Section 1.4 a chapter assumes. The time is an estimate for reading the chapter once and running each case of its app once, at about eight minutes a case.

Chapters 6, 7, 12 and 15 form a first course: an inventory, an absorber, a line and a tank, each with one main loop and one limit. They rest on Sections 1.4.4 to 1.4.7. Chapters 4, 11 and 13 add the selectors of Section 1.4.8 and are best read together. Chapters 3, 8 and 10 turn on what is inferred or measured late, Section 1.4.9. Chapters 2, 5, 9 and 14 combine several of these and are best left until last. A reader who follows the fluid can read the volume in order, from the well to the terminal.

KEY POINTS

- From well to market a choke, a pump or a compressor sets a pressure, and the flow follows from a reservoir, a line or a market.
- A separator stores a slug under any tuning. The tuning decides only how fast the liquid is passed on.
- A gas line rides through a change in demand on its own pack. A liquid line stores almost nothing and needs fast limit loops.
- Two loops at a selector protect a limit only when the waiting loop's reset follows the output actually applied.
- An override can protect only a limit it can see, and the quantity that binds is often inferred or measured late.

Where else this applies. The same layers and limits return in every industry volume of the series. Selectors and overrides are developed in Volume 4, predictive control in Volume 9, and the refinery that takes the oil in Volume 11.

Table 1.4 Roadmap of the Volume

Ch.	Plant	The lesson	Level	Time
2	Managed pressure drilling rig	The best drilling rate lies on a constraint boundary that moves as the bit wears.	advanced	3 h
3	Rod-pumped oil well	A loop on an inferred level holds the inference, not the level.	intermediate	2.5 h
4	Smart producer and waterflood	Four limits share a choke, and the most oil now is not the most oil recovered.	intermediate	2 h
5	Deepwater flowline and riser	A minimum that moves along the line, dosed through a transport delay.	advanced	2.5 h
6	Offshore three-phase separator	No tuning keeps a slug out. Tuning decides how fast it leaves.	introductory	2.5 h
7	Glycol dehydrator	A fast handle with a floor, a slow handle with a ceiling.	introductory	2 h
8	Amine sweetener	A loading limit the product analyzer cannot see.	introductory	2 h
9	Turboexpander NGL plant	The order of two setpoint moves decides the specification.	advanced	2 h
10	NGL fractionation train	A specification made one column upstream of where it is measured.	intermediate	2.5 h
11	Compressor station	Equal surge margin shares the load where equal speed does not.	intermediate	2 h
12	Gas transmission line	The line is a tank, and a forecast is worth more than retuning.	introductory	2.5 h
13	Refined products line	One flow through four pumps, held by low selectors.	intermediate	2 h
14	LNG liquefaction train	Pairing on a valve with two effects, and a limit computed from a late analysis.	advanced	3 h
15	LNG import terminal	An integrating pressure whose outlet the market sets.	introductory	2 h

References

[1] B. Foss, "Process control in conventional oil and gas fields—challenges and opportunities," *Control Engineering Practice*, vol. 20, no. 10, pp. 1058–1064, 2012.

[2] R. R. Bottoms, "Process for separating acidic gases," U.S. Patent 1 783 901, Dec. 2, 1930, filed October 7, 1930.

[3] E. D. Sloan, "Hydrocarbon hydrate flow assurance history as a guide to a conceptual model," *Molecules*, vol. 26, no. 15, p. 4476, 2021.

[4] W. A. Loomis, "Dummy driller," U.S. Patent 1 779 480, Oct. 28, 1930, filed December 14, 1927.

[5] R. W. Bucklin, "Method and equipment for treating hydrocarbon gases for pressure reduction and condensate recovery," U.S. Patent 3,292,380, filed April 28, 1964, granted December 20, 1966, 1966. [Online]. Available: https://patents.google.com/patent/US3292380A/en

[6] S. G. Gibbs and A. B. Neely, "Computer diagnosis of down-hole conditions in sucker rod pumping wells," *Journal of Petroleum Technology*, vol. 18, no. 1, pp. 91–98, 1966.

[7] P. Arjomandnia, "Simulation of LNG rollover in storage tanks," Ph.D. dissertation, Curtin University, Perth, Australia, 2015, open access, hdl.handle.net/20.500.11937/1544.

[8] S. G. Gibbs, "Method for determining the pump-off of a well," U.S. Patent 3 951 209, Apr. 20, 1976, filed June 9, 1975.

[9] E. D. Burger, W. R. Munk, and H. A. Wahl, "Flow increase in the trans alaska pipeline through use of a polymeric drag-reducing additive," *Journal of Petroleum Technology*, vol. 34, no. 2, pp. 377–386, 1982.

[10] Z. Schmidt, J. P. Brill, and H. D. Beggs, "Experimental study of severe slugging in a two-phase-flow pipeline - riser pipe system," *Society of Petroleum Engineers Journal*, vol. 20, no. 5, pp. 407–414, 1980.

[11] 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.

[12] J.-M. Godhavn, "Control requirements for automatic managed pressure drilling system," *SPE Drilling & Completion*, vol. 25, no. 3, pp. 336–345, 2010.

[13] 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

[14] J. G. Abraham, “Drilling control,” U.S. Patent 2 759 702, Aug. 21, 1956, filed March 25, 1952.

[15] J. V. Vogel, “Inflow performance relationships for solution-gas drive wells,” *Journal of Petroleum Technology*, vol. 20, no. 1, pp. 83–92, 1968.

[16] S. Fakher, A. Khlaifat, M. E. Hossain, and H. Nameer, “A comprehensive review of sucker rod pumps’ components, diagnostics, mathematical models, and common failures and mitigations,” *Journal of Petroleum Exploration and Production Technology*, vol. 11, no. 10, pp. 3815–3839, 2021.

[17] R. P. Cutler and A. J. Mansure, “Fluid dynamics in sucker rod pumps,” Sandia National Laboratories, Albuquerque, NM, Report SAND99-0093C (Southwest Petroleum Short Course, Lubbock, TX), OSTI 3261, 1999.

[18] D. H. Beggs and J. P. Brill, “A study of two-phase flow in inclined pipes,” *Journal of Petroleum Technology*, vol. 25, no. 5, pp. 607–617, 1973.

[19] Y. Taitel, “Stability of severe slugging,” *International Journal of Multiphase Flow*, vol. 12, no. 2, pp. 203–217, 1986.

[20] B. F. M. Pots, I. G. Bromilow, and M. J. W. F. Konijn, “Severe slug flow in offshore flowline/riser systems,” *SPE Production Engineering*, vol. 2, no. 4, pp. 319–324, 1987.

[21] J.-M. Godhavn, M. P. Fard, and P. H. Fuchs, “New slug control strategies, tuning rules and experimental results,” *Journal of Process Control*, vol. 15, no. 5, pp. 547–557, 2005.

[22] E. D. Sloan, “Fundamental principles and applications of natural gas hydrates,” *Nature*, vol. 426, no. 6964, pp. 353–359, 2003.

[23] H. Gul Zaman, L. Baloo, R. Pendyala, P. Singa, S. Ilyas, and S. Kutty, “Produced water treatment with conventional adsorbents and MOF as an alternative: A review,” *Materials*, vol. 14, no. 24, p. 7607, 2021.

[24] D. Asante-Sackey, S. Rathilal, E. K. Tetteh, and E. K. Armah, “Membrane bioreactors for produced water treatment: A mini-review,” *Membranes*, vol. 12, no. 3, p. 275, 2022.

[25] K. A. McDonald, T. J. McAvoy, and A. Tits, “Optimal averaging level control,” *AIChE Journal*, vol. 32, no. 1, pp. 75–86, 1986.

[26] A. R. D. Thorley and C. H. 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.

[27] M. S. Ghidaoui, M. Zhao, D. A. McInnis, and D. H. Axworthy, “A review of water hammer theory and practice,” *Applied Mechanics Reviews*, vol. 58, no. 1, pp. 49–76, 2005.

[28] F. K. Moore and E. M. Greitzer, “A theory of post-stall transients in axial compression systems: Part I—development of equations,” *Journal of Engineering for Gas Turbines and Power*, vol. 108, no. 1, pp. 68–76, 1986.

[29] E. M. Greitzer, “Surge and rotating stall in axial flow compressors—part I: Theoretical compression system model,” *Journal of Engineering for Power*, vol. 98, no. 2, pp. 190–198, 1976.

[30] A. J. Kidnay, W. R. Parrish, and D. G. McCartney, *Fundamentals of Natural Gas Processing*, 3rd ed. CRC Press, 2019.

[31] A. H. Glattfelder and W. Schaufelberger, *Control Systems with Input and Output Constraints*, ser. Advanced Textbooks in Control and Signal Processing. London: Springer, 2003.

[32] B. Joseph and C. B. Brosilow, “Inferential control of processes: Part I. steady state analysis and design,” *AIChE Journal*, vol. 24, no. 3, pp. 485–492, 1978.

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

[34] 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.

[35] L. S. Hansen, S. Pedersen, and P. Durdevic, “Multi-phase flow metering in offshore oil and gas transportation pipelines: Trends and perspectives,” *Sensors*, vol. 19, no. 9, p. 2184, 2019.

[36] 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.

[37] S. Skogestad, “Plantwide control: The search for the self-optimizing control structure,” *Journal of Process Control*, vol. 10, no. 5, pp. 487–507, 2000.

[38] K. J. Åström and L. Rundqwist, "Integrator windup and how to avoid it," in *Proc. 1989 American Control Conference*, 1989, pp. 1693–1698.

[39] 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.

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

[41] C. V. Rao, J. B. Rawlings, and D. Q. Mayne, "Constrained state estimation for nonlinear discrete-time systems: Stability and moving horizon approximations," *IEEE Transactions on Automatic Control*, vol. 48, no. 2, pp. 246–258, 2003.

1.9 Exercises

1.1 Theory and calculation · 5 min

W-1's reservoir has fallen to 240 bar. All else is as in Section 1.4.2, with the choke held at 40 %.

(a) Find the rate. *Checkpoint:* 160.5 m^3/h.

(b) Find the drawdown and the pressure drop across the choke. *Checkpoint:* 40.1 bar and 19.9 bar.

1.2 Concept check · 2 min

W-1's rate loop holds 180 m^3/h on the choke while the reservoir pressure falls from 250 bar. Which statement is right?

(a) the choke stays at 40 %, because the rate it passes depends on the choke alone

(b) the choke opens further as the reservoir falls, and at 229 bar it is wide open and the rate starts to fall

(c) the rate can be held at any reservoir pressure above 180 bar, the floor under the bottomhole pressure

(d) the rate is held until the drawdown reaches its 60 bar limit, which the falling reservoir reaches first

1.3 Theory and calculation · 4 min

W-1's choke is opened to 70 %, where the well produces 227.7 m^3/h with 13.07 bar across the choke.

(a) Find the share g of an immediate choke response that remains at steady state, from Equation (1.3). *Checkpoint:* 0.315.

(b) Find the immediate and the steady gain of rate to choke opening. *Checkpoint:* 3.25 and 1.02 m^3/h per percent.

1.4 Theory and calculation · 4 min

W-1's sealed riser fills at 0.125 bar/s. Use Equation (1.5) with a standard pressure of 1.013 bar.

(a) The gas rate is doubled to 36 000 Sm^3/h by gas lift with the flowline gas volume still 150 m^3. Find Π. *Checkpoint:* 0.54.

(b) The gas volume upstream of the riser is also cut to 100 m^3. Find Π and say whether the riser still slugs. *Checkpoint:* 0.81, it still slugs.

1.5 Theory and calculation · 5 min

W-1's separator has a liquid surface of 25 m^2 and an oil valve that passes 300 m^3/h wide open, so $k = 0.12\,\mathrm{m/h}$ per percent.

(a) A PI level loop has $K_c = 100\,\%/\mathrm{m}$ and $\tau_I = 20\,\mathrm{min}$. Find τ_c and ζ. *Checkpoint:* 5 min and 1.0.

(b) A third well adds a step of 90 m^3/h to the inflow and the level may use 0.6 m of band. Find the slowest outflow ramp that keeps the level in the band, and how long the ramp lasts. *Checkpoint:* 4.5 m^3/h per minute for 20 min.

1.6 Concept check · 2 min

Which statement matches the two runs of Figure 1.4?

(a) the tight loop keeps most of the slug out of the vessel, because it opens the valve at once

(b) the averaging loop stores about 18 m^3 more than the tight loop, because it opens the valve slowly

(c) the averaging loop takes the level over the high-high trip, which is the price of a smooth outflow

(d) both store nearly the same volume, and the tuning decides only how fast the stored liquid leaves

1.7 Theory and calculation · 4 min

A gas transmission line is 80 km long with a bore of 0.5 m and an isothermal wave speed of 350 m/s.

(a) Find the gas it stores per bar of mean pressure. *Checkpoint:* 12.82 t per bar.

(b) Demand exceeds supply by 8 kg/s for two hours. Find the fall in mean pressure. *Checkpoint:* 4.49 bar.

1.8 Theory and calculation · 4 min

A two-stage glycol contactor sees 800 mg/Sm^3 of water at the inlet, and the lean glycol sets a floor of 40 mg/Sm^3.

(a) Find the outlet water at an absorption factor of 2, from Equation (1.12). *Checkpoint:* 148.6 mg/Sm^3.

(b) Find the outlet at an absorption factor of 4, and the least water any circulation can reach. *Checkpoint:* 76.2 mg/Sm^3, and 40 mg/Sm^3.

1.9 Concept check · 2 min

The export line of W-1 is replaced by one whose head is pure friction, $\psi = k\phi^2$ with no static part. What happens to the compressor's surge margin as the export flow falls?

(a) it stays the same at every speed, because the operating point follows a fan-law parabola

(b) it shrinks faster than before, because friction adds head at low flow

(c) it shrinks because the surge line moves to the right with the friction

(d) it grows as the speed falls, because the head falls with the cube of speed

1.10 Theory and calculation · 3 min

Natural gas at −20 °C and 70 bar is expanded to 30 bar. Treat it as an ideal gas with $\gamma = 1.3$.

(a) Find the outlet temperature of an isentropic expander. *Checkpoint:* −65.0 °C.

(b) Find the outlet temperature at an isentropic efficiency of 0.8. *Checkpoint:* −56.0 °C.