INTERACTIVE CONTROL ENGINEERING
Principles and Applications

Volume 10 Control in Upstream and Midstream Oil and Gas

CHAPTER 14

An LNG Liquefaction Simulator for Loop Pairing and Inferential Constraint Control

Pages 321–352

Edited by John D. Hedengren

How to cite this chapter. "An LNG Liquefaction Simulator for Loop Pairing and Inferential Constraint Control," 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. 14, pp. 321–352.

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

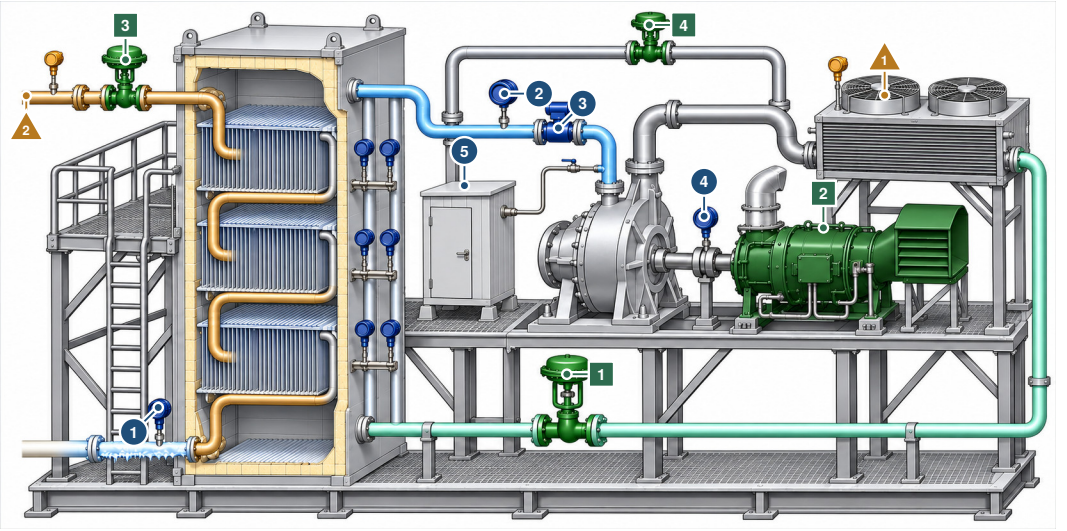
The app for this chapter
https://apmonitor.com/apps/lng-liquefaction/

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

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

CHAPTER 14

An LNG Liquefaction Simulator for Loop Pairing and Inferential Constraint Control



The controller moves 1 Joule-Thomson valve, which sets the circulation 2 compressor driver, which sets the speed 3 feed gas valve, cut back when the driver runs short 4 anti-surge recycle valve

The controller reads 1 LNG outlet temperature 2 suction pressure 3 refrigerant circulation 4 shaft power 5 refrigerant composition, every twelve minutes

What disturbs it ▲1 warm air at the air cooler, which also derates the driver ▲2 a richer feed, or a cut in it ▲ light ends lost from the refrigerant

Notice the throttling valve on the cold return, whose two effects the chapter separates, and that the gauges drawn on the cold box wall are not in the chapter, whose approaches inside the box are computed from the refrigerant analysis in the cabinet rather than measured. Illustration, not to scale.

IN PLAIN TERMS Cooling Gas Until It Turns to Liquid

The main idea. A liquefaction train chills natural gas until it becomes a liquid that ships can carry. A very cold fluid boils in a block of metal passages and draws heat out of the gas.

The challenge. One valve sets both how much cold fluid flows and how warm it boils, and the two pull in opposite directions. The closest gap between gas and fluid lies where no instrument reaches, so the plant computes it from a sample up to twelve minutes old.

Why it matters. The machine that drives the cooling is the largest on the site, so the energy it spends on each tonne matters most.

THREE WAYS TO READ THIS CHAPTER

Explore	12 min	This page, the History boxes in Section 14.2, the Did you know? boxes in Sections 14.1 and 14.6.5, the Try it boxes in Section 14.6, and Key points in Section 14.10. No equations.
Study	70 min	The whole chapter in order, stopping at the three pause points, then Exercises 1 to 12.
Practice	25 min	This page, Section 14.6 with Figure 14.6, Section 14.7 and Section 14.8 with its Open challenge box.

The control question The Joule-Thomson valve of a mixed refrigerant train sets the circulation and the suction pressure together, and the two reach the LNG temperature with opposite signs, while the closest approach inside the exchanger is computed from an analysis twelve minutes old. Which loop structure makes the valve a usable handle, and what does holding an inferred limit cost when its input goes stale?

Learning objectives

- Separate the two effects of the Joule-Thomson valve by experiment and explain why no PI tuning repairs a gain that changes sign.
- Compute the refrigerant boiling temperature, a zone duty and the bundle lag from the model's equations.
- Classify circulation as a capacity variable and suction pressure as the temperature trim, and the blend as a max-min problem.
- Explain how a stale analysis drives the computed cold-end approach the wrong way, and read the gap from the exported file.
- Compare the cascade with the predictive controller by time, depth and integral of the driver excess.

Prerequisites Sections 1.4.7 (cryogenic processing), 1.4.6 (surge margin), 1.4.8 (low selector), 1.4.9 (inferred limits), 1.6.2 and 1.6.4, with Section 1.6.3 only for the MPC. Chapter 11 for anti-surge recycle.

The app https://apmonitor.com/apps/lng-liquefaction/

See also Volume 6, Chapter 2 for a valve whose gain changes sign near design, and Volume 6, Chapter 8 for a plantwide structure on a plant with sampled analyzers. Volume 5, Chapter 3 and Volume 5, Chapter 10 for quality held through a late analysis. Volume 4, Chapter 2 for the surge cycle under a protective override, and Volume 4, Chapter 3 for limiters selected on a margin that is never measured.

14.1 Introduction

Every loop on a liquefaction train holds a quantity that an instrument reports. Suction pressure comes from a transmitter, refrigerant circulation from a flowmeter, and the product temperature from a thermocouple in the outlet nozzle. The quantity that decides how much power the plant burns per tonne is none of those. It is the minimum internal temperature approach, the closest the warming refrigerant and the cooling gas come to each other anywhere inside the exchanger, and in the coil-wound bundle modeled here three of the four places that minimum can sit are inside a welded aluminum block with no instrument in them. What the control room works with instead is an arithmetic result computed from measured flows, measured pressures and a refrigerant composition that arrives from a gas chromatograph once every twelve minutes. Holding a variable that is calculated rather than measured is inferential control, a structure posed and analyzed for chemical processes half a century ago [1] and now the ordinary arrangement wherever the economically interesting property is slow or impossible to measure [2].

The class is wider than cryogenics. A distillation column controls a tray temperature because the composition analyzer is slow, a fired heater controls a tube skin because the coke layer is invisible, and a polymerization reactor controls a jacket temperature because the melt index arrives from a laboratory once a shift. Volume 5, Chapter 3 controls a column through a slow analyzer, and Volume 5, Chapter 10 steers a dryer by its exhaust air while the moisture arrives from a laboratory. Each of those plants runs a loop on a proxy, and each inherits the same failure mode. The loop holds its setpoint exactly while the quantity it stands for walks away, because the model that connects the two has an input that is stale. Which controlled variable to hold, and how much the choice costs when the plant moves, is the self-optimizing control question that the plantwide literature poses in general terms [3] and that the LNG literature has posed for this particular plant [4, 5].

A second difficulty on this train is a pairing one, and it is the reason a transfer function is a poor substitute for the plant. The Joule-Thomson valve does two things at once. It sets how much refrigerant circulates, which is more cooling, and, because whatever it passes has to be taken away again by the compressor, it sets the pressure at which that refrigerant boils, which is warmer refrigerant. At steady state, opening the valve from 54.9 to 62 % puts 4.6 kg/s more refrigerant through the bundle

and raises the suction pressure from 2.90 to 3.17 bar, and the two effects together buy 4.1 °C of product temperature before the gain turns over and becomes negative. A first-order-plus-dead-time model fitted to this loop never shows that sign change, because the sign change is the thing the fit averages away. A poorly behaved loop is easily taken for a tuning problem, and here it is not one. No proportional-integral setting repairs a gain of 3.4 °C per percent in one direction, 0.27 °C per percent in the other and negative beyond 60 % valve. The repair is structural, and interaction measures have been the standard way of choosing among such structures since the relative gain array was proposed [6].

Liquefaction is where an LNG project spends its energy. The refrigerant compressor driver is the largest machine on the site, and specific power is the number a commercial team argues about, which is why optimization of the cycle and of its operation has its own substantial literature [7, 8] and why the fuel is assessed on the efficiency of the chain that delivers it [9]. Published treatments of LNG control select controlled variables and evaluate structures on steady-state or full process-simulator models that cannot be opened, run or perturbed in an hour. This chapter works with a browser-hosted single mixed refrigerant train whose every equation, constant and measured number is given in the text, so that its sixteen states, its three exchanger zones, its compressor map and its five loops can be rebuilt, and uses it to pull apart two effects that one valve delivers together. The pairing is measured: the open-loop product temperature turns over at 60 % valve travel at fixed compressor speed and becomes linear at 9.2 °C per bar once a suction pressure loop is closed, and refrigerant circulation is a capacity variable whose gain changes by a factor of 16 across its working range. The model is an effectiveness-NTU three-zone cold box with bundle metal inertia, a sampled composition analyzer, an inferential approach estimator and a low selector on throughput, written out as implemented rather than as intended. An optional model predictive controller is compared with the cascade on five measured cases, losses included. The simulator reports both the true minimum approach and the control room estimate of it, so the difference between the variable a loop holds and the variable that limits the plant can be watched on the same screen. The cycle, the effectiveness-NTU relation, the cascade and the predictive controller are all established methods.

> ✦ **DID YOU KNOW?** **Part of the Gas Pays for the Cold**
>
> A federal estimate puts the gas used for liquefaction at about 7 to 15 % of the gas delivered to United States LNG export plants, mostly to run the equipment on site. All but one of those plants drive their refrigerant compressors with gas turbines [10]. This chapter's train spends 0.321 kWh of shaft work per kilogram.

Section 14.2 traces the cycle from the classical cascade to the mixed refrigerant and the published control work. Section 14.3 states the train model, its constants and its assumptions. Section 14.4 measures the two effects of the valve, the pressure handle, the circulation as capacity and the blend as a max-min problem. Section 14.5 describes the five loops, the low selector, the instruments and the predictive controller. Section 14.6 measures the scenarios, from two handles by hand to the graded shift and the predictive controller under a driver limit. Section 14.7 sets the hardware of a real train beside the model, Section 14.8 lists what the model leaves out, Section 14.9 weighs the economics and the methods that fit the same equations, Section 14.10 sums up, and the exercises in Section 14.11 put the scenarios to work. The LNG leaving this train is returned to gas at the terminal of Chapter 15.

14.2 Background

Liquefying natural gas is an old problem with a young constraint. How the mixed refrigerant, a blend that boils over a range of temperature, replaced the classical cascade of pure fluids and its three machine trains with one is the subject of the first History box below [8, 11]. What the substitution bought in capital cost it charged for in operability, because the shape of the refrigerant warming curve is set by its composition, and composition is not an actuator that anybody moves in a hurry.

HISTORY 1934 to 1970: From a Staircase of Pure Fluids to a Blend

A patent application of 1934 proposed a refrigerant made of a mixture of gases and vapors, compressed by one machine and condensed in stages, whose composition could be adjusted while the system was still in operation [12]. Natural gas liquefiers first took another road. A 1966 application for a liquefier of gas to be shipped across water described the practice of its day: separate cycles of essentially pure propane, ethylene and methane, each over its own band of temperature, whose curve is a staircase that in its example needed nine exchangers. It proposed one fluid of methane, ethane, propane, butanes and pentanes, with perhaps a little nitrogen, whose composition could be changed to keep the plant near its optimum [13]. An application of 1970 named the gain: the curves of single refrigerants do not closely match the continuous cooling curve of the gas, and horsepower is wasted [14].

The constraint acquired a name and a method in a different discipline. Pinch analysis, developed for heat exchanger networks in the late 1970s and formalized as the pinch design method, established the minimum approach temperature as the single quantity that trades exchanger area against utility consumption, and established that the network can only be understood by plotting both composite curves and reading the gap between them [15]. Liquefaction is that argument compressed into one exchanger over 185 degrees of temperature. Refrigerant synthesis methods that choose a blend by matching composite curves over the whole temperature range followed, and they treat the minimum approach as a constraint on the optimization rather than as a setpoint [16]. The design approach on this train is 5 °C, and the reason it is a constraint and not a target is that the composition which maximizes the smallest approach is an interior maximum. Going lighter opens the cold end and closes the warm end, and going heavier does the reverse.

The published control work asks which variables a liquefaction plant should hold constant when the ambient temperature and the feed move, and answers it with steady-state loss evaluations rather than with a tuning rule, from the first study of optimal operation of a simple mixed refrigerant cycle, told in the second History box below [4], to controlled-variable selection for a full LNG process on a dynamic first-principles model of the TEALARC plant [5] and a control structure designed from a dynamic model with the same question in view [17]. Three results are common to that work and are the ones this application is arranged around. Suction pressure, and not refrigerant flow, is the handle that sets product temperature. The compressor driver is an active constraint for much of the year, so throughput is a manipulated variable rather than a schedule. And the approach temperature has to be inferred, so the estimator and its inputs are part of the control design rather than an instrument detail.

HISTORY 2006: Operating the Plant, Not Only Designing It

A 2006 study of the simplest single mixed refrigerant process began by noting that, for all the work that goes into the design of LNG processes, their operation received surprisingly little attention. It counted nine manipulated inputs and three constraints active at the optimum, one of them the LNG leaving at −155 °C or colder. In design the minimum approach is specified, but in operation it is free to move. With the design areas held, the optimum closed it from 1.2 to 0.537 °C and needed 2.60 % less compressor power. The study then asked which variables, held at constant setpoints, keep the plant close to that optimum as disturbances arrive [4].

The ambient constraint deserves its own sentence, because it is the one a control room meets daily. Warm air raises the temperature at which the refrigerant condenses, which raises the discharge pressure and therefore the work per kilogram, and it derates the gas turbine driving the compressor at the same time. The two effects move production in the same direction and are the subject of a dedicated body of mitigation work in the LNG industry [18]. On the train modeled here an ambient change from 18 to 36 °C raises the shaft power required from 23.1 to 26.9 MW while lowering what the driver can deliver from 26.0 to 22.7 MW, and nothing on the product temperature trend says so until the plant runs out of power.

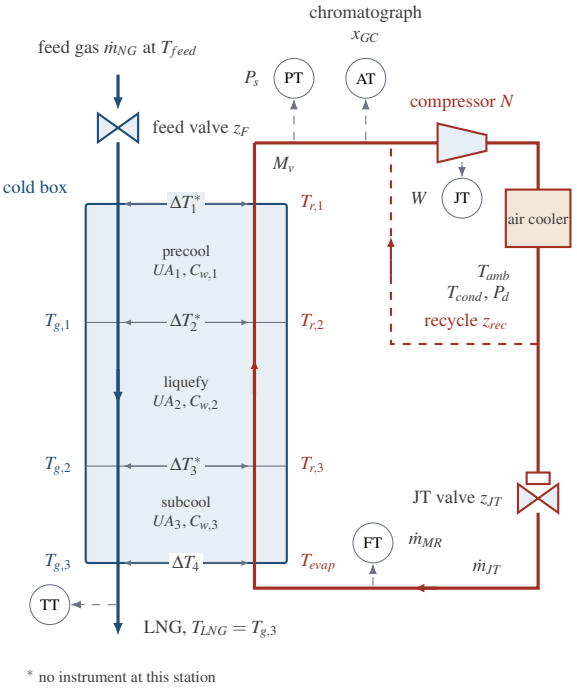


Figure 14.1 Single mixed refrigerant train with the cold box drawn as one block containing three counter-current zones. Feed gas runs down the left from T_{feed} to T_{LNG} and refrigerant runs up the right from T_{evap} to $T_{r,1}$, and the four stations where the two streams can come closest are marked ΔT_1 through ΔT_4. Only the cold end ΔT_4 has a temperature measurement on both sides of it.

14.3 Liquefaction Train Model

The application simulates a single mixed refrigerant train of the kind used at mid scale, taking 20.0 kg/s of natural gas at 30 °C and 55 bar and delivering LNG at −155.0 °C, which is 72.0 t/h and about 0.58 million tonnes a year once availability is allowed for. The refrigerant is compressed to 46.1 bar, condensed against ambient air, throttled through the Joule-Thomson valve to 2.90 bar and boiled back through the same exchanger it is cooling. The arrangement, the four places the two streams can pinch and the symbol for every quantity used below appear in Figure 14.1. The single-refrigerant cycle is used in preference to a propane-precooled arrangement because it carries the whole argument with one loop instead of three.

14.3.1 Variable Classification

Table 14.1 classifies every signal the loops and the scorecard interact with. Two entries need a word of explanation. The refrigerant light fraction x is a state and a setpoint at the same time, because the makeup system moves the inventory toward an operator target at a bounded rate of 0.010/min rather than setting it. The true minimum approach ΔT_{min} is an output that no instrument produces and that the implemented controllers never see, and it appears on the screen only because the simulator can compute what the plant cannot measure.

14.3.2 Three Counter-Current Zones

The cold box is one block carrying three lumped zones in series, precool, liquefaction and subcool, with the gas passing 1 to 3 and the refrigerant 3 to 1. Each zone is solved as a counter-current effectiveness-NTU exchanger rather than on a mean temperature difference. Heat-capacity rates

Table 14.1 Every Signal the Loops and the Scorecard Touch. The Nominal Column Is the Design Line-Up That Opens Each Scenario.

Symbol	Role	Nominal	Range
T_{LNG}	CV, measured	−155.0 °C	−168 to −146 SP
ΔT_4	CV, inferred	7.00 °C	2–16 SP
P_s	CV, measured	2.90 bar	1.9–5.2 bar
$\dot{m}_{MR}$	CV, measured	33.0 kg/s	14–48 kg/s
W	CV, measured	23.10 MW	16–28 MW SP
SM	CV, computed	44.9 %	6–30 % SP
z_{JT}	MV	54.9 %	0–100 %
N	MV	92.0 %	55–105 %
z_F	MV	100 %	40–110 %
z_{rec}	MV	0 %	0–100 %
T_{amb}	DV, measured	18.0 °C	0–45 °C
f_{rich}	DV, measured	1.00×	0.85–1.25×
f_{avail}	DV, measured	1.00×	0.40–1.05×
f_b	DV, unmeasured	1.00×	0.70–1.05×
f_{ac}	DV, unmeasured	1.00×	0.70–1.10×
$T_{g,1..3}$	states, °C	−35.0, −119.9, −155.0	−200 to +120
$T_{r,1..3}$	states, °C	+25.0, −40.0, −124.9	−200 to +120
$Q_{1..3}$	states, kW	3379, 9515, 2384	−4 to 40 MW
P_s	state	2.901 bar	0.6–14 bar
T_{cond}	state	30.16 °C	—
N_a, $z_{\cdot,a}$	actuator states	as MV	as MV
x	state and target	0.420	0.30–0.55
T_{evap}	output	−162.0 °C	—
ΔT_{min}	output, unmeasured	5.00 °C	—
$\Delta\hat{T}_{min}$	output, inferred	5.00 °C	—
w	output	0.321 kWh/kg	—

follow from the flows and the per-zone heat capacities, as shown in Equation (14.1).

$$C_{g,k} = \dot{m}_{NG}\, c_{g,k}, \qquad C_{r,k} = \dot{m}_{MR}\, c_{r,k} \tag{14.1}$$

where $\dot{m}_{NG}$ and $\dot{m}_{MR}$ are the gas and refrigerant mass flows in kg/s and $c_{g,k}$ and $c_{r,k}$ are the zone heat capacities in kJ/kg/°C, so every rate is in kW/°C. The smaller of the two rates is written $C_{\min,k}$. Feed richness raises the gas heat capacities by 20, 60 and 15 % of the excess $f_{rich} - 1$ in zones 1, 2 and 3, which puts most of the extra latent load in the middle zone. The refrigerant heat capacities move with composition, as shown in Equation (14.2).

$$c_{r,k} = \max\left(c_{r,k}^{\min},\; c_{m,k}\left[1 + s_k (x - x_0)\right]\right) \tag{14.2}$$

where x is the light fraction of the blend, x_0 is its design value, $c_{m,k}$ is the design heat capacity of zone k and s_k is the sensitivity of that zone to the blend, with $s_1 = -1.40$, $s_2 = -0.15$ and $s_3 = +1.60$ per unit of light fraction. The floors $c_{r,k}^{\min}$ are 0.35, 0.60 and 0.35 kJ/kg/°C. Those three signs are the physical content of the mixed refrigerant. Heavy components boil at the warm end and light ones at the cold end, so making the blend lighter moves duty from the top of the bundle to the bottom.

Zone effectiveness follows from the number of transfer units and the heat-capacity-rate ratio, as shown in Equation (14.3).

$$\varepsilon_k = \frac{1 - \mathrm{e}^{-\mathrm{NTU}_k (1 - C_{r,k}^{*})}}{1 - C_{r,k}^{*}\, \mathrm{e}^{-\mathrm{NTU}_k (1 - C_{r,k}^{*})}}, \qquad \mathrm{NTU}_k = \frac{U A_k\, f_b}{C_{\min,k}} \tag{14.3}$$

where $C^{*}_{r,k}$ is the ratio of the smaller heat-capacity rate to the larger, UA_k is the design zone conductance in kW/°C and f_b is the cold box condition multiplier. The balanced branch $\varepsilon = \mathrm{NTU}/(1+\mathrm{NTU})$ is used whenever $C^{*}_{r,k} > 0.999$, and on this plant that branch is the normal case rather than an edge case. At the design point $C_{g,1} = C_{r,1} = 52.0\,\mathrm{kW/^{\circ}C}$ and $C_{g,2} = C_{r,2} = 112.0\,\mathrm{kW/^{\circ}C}$, because matching the two rates in every zone is precisely what blending a mixed refrigerant is for. The duty a zone would deliver at its present inlets is then Equation (14.4).

$$Q_{ss,k} = \varepsilon_k C_{\min,k}\left(T_{g,in,k} - T_{r,in,k}\right) \tag{14.4}$$

with the gas inlets $T_{g,in} = (T_{feed}, T_{g,1}, T_{g,2})$ and the refrigerant inlets $T_{r,in} = (T_{r,2}, T_{r,3}, T_{evap})$, the two lists running in opposite directions because the block is counter-current. An effectiveness formulation cannot produce a crossed pair of streams at steady state. As the approach closes the effectiveness delivers less duty and the zone simply stops transferring, which is what the plant does, whereas a mean-driving-force cell will return a refrigerant colder than the gas at one end and warmer at the other and report a negative pinch that nobody can act on. In a transient the cold-end approach can still fall below zero for a while, because the boiling temperature of (14.9) follows the suction pressure and the blend at once while the gas holdup lags. It reaches $-0.52\,^{\circ}\mathrm{C}$ after the feed cut of scenario 10 and $-5.2\,^{\circ}\mathrm{C}$ at the instant of the relief event in the graded challenge run with its suggested settings.

14.3.3 Bundle Metal, Stream Holdup and the Duty Limiter

Five hundred tonnes of aluminum sit between the duty the conditions call for and the duty being delivered. The metal is carried as a first-order lag on each zone duty rather than as a temperature node, as shown in Equation (14.5).

$$\tau_k \frac{\mathrm{d}Q_k}{\mathrm{d}t} = Q_{ss,k} - Q_k, \qquad \tau_k = \frac{C_{w,k}}{60\,UA_k\,f_b} \tag{14.5}$$

where $C_{w,k}$ is the bundle heat capacity of zone k in kJ/°C and the factor 60 converts the seconds implied by kJ/°C over kW/°C into the minutes the clock runs in. At the design conductances the three lags are 4.44, 4.55 and 6.65 min, which is why nothing about a cold box answers quickly.

A lagged duty must never exceed what the two streams can physically exchange at the inlets they have at this instant, so a ceiling is computed at every step, as shown in Equation (14.6).

$$Q_{\lim,k} = C_{\min,k}\left(T_{g,in,k} - T_{r,in,k}\right) \tag{14.6}$$

The duty that reaches the temperature balances, $Q_{eff,k}$, is the state Q_k clipped between zero and $Q_{\lim,k}$ when $Q_{\lim,k}$ is positive and between $Q_{\lim,k}$ and zero when it is negative. Without (14.6) a sudden loss of feed would leave stale duty in the metal that cools the gas straight past the refrigerant and out the other side, and every approach in the application would go negative on a transient that a real bundle rides out. The two stream temperatures then follow from lumped holdups, as shown in Equations (14.7) and (14.8).

$$C_{NG}\frac{\mathrm{d}T_{g,k}}{\mathrm{d}t} = 60\left[C_{g,k}\left(T_{g,in,k} - T_{g,k}\right) - Q_{eff,k}\right] \tag{14.7}$$

$$C_{MR}\frac{\mathrm{d}T_{r,k}}{\mathrm{d}t} = 60\left[C_{r,k}\left(T_{r,in,k} - T_{r,k}\right) + Q_{eff,k}\right] \tag{14.8}$$

where C_{NG} and C_{MR} are the gas-side and refrigerant-side holdups per zone in kJ/°C. All six stream temperatures are clipped to -200 to $120\,^{\circ}\mathrm{C}$ and the three duties to -4 to 40 MW after every step.

14.3.4 Where the Refrigerant Boils

Two quantities set the cold end of the cycle and nothing else does, as shown in Equation (14.9).

$$T_{evap} = T_{evap,0} + k_P\left(P_s - P_{s,0}\right) - k_x\left(x - x_0\right) \tag{14.9}$$

where P_s is the suction pressure in bar, x is the light fraction, $k_P = 9.5\,^{\circ}\mathrm{C}$ per bar and $k_x = 120\,^{\circ}\mathrm{C}$ per unit of light fraction. Eight points of blend is therefore 9.6 °C of boiling temperature, which is more than the entire cold-end margin. Neither the Joule-Thomson valve nor the compressor speed appears in (14.9). They reach the boiling temperature only through P_s, which is the fact the pairing argument rests on.

14.3.5 Suction Inventory, Compressor and Driver

The suction pressure is a vapor inventory rather than a resistance, so its rate of change is set by the flow imbalance, as shown in Equation (14.10).

$$\frac{\mathrm{d}P_s}{\mathrm{d}t} = \frac{60\,P_s}{M_v}\left(\dot{m}_{JT} + \dot{m}_{rec} - \dot{m}_{comp}\right) \tag{14.10}$$

where M_v is the suction system vapor holdup in kg and the three flows are in kg/s. The inventory regulates itself, because a higher pressure feeds the compressor of (14.12) more vapor and leaves the Joule-Thomson valve a smaller pressure difference, so a speed change moves the pressure to a new value rather than ramping it. After a 1 % speed step with every loop open the pressure settles 0.043 bar lower, with 63 % of the change in 0.65 min. The two valves are square-root devices on the pressure difference across them, as shown in Equation (14.11).

$$\begin{aligned} \dot{m}_{JT} &= k_{JT}\,\frac{z_{JT}}{100}\sqrt{\max(0, P_d - P_s)} \\ \dot{m}_{rec} &= k_{rec}\,\frac{z_{rec}}{100}\sqrt{\max(0, P_d - P_s)} \end{aligned} \tag{14.11}$$

The compressor is a single-line map in which flow is proportional to speed and to suction density and is reduced by whatever head is being asked of it, as shown in Equation (14.12).

$$\dot{m}_{comp} = \dot{m}_{c0}\,n\,\frac{P_s}{P_{s,0}}\left[1 - a_M\left(\frac{P_d/P_s}{n^2\,P_{d,0}/P_{s,0}} - 1\right)\right] \tag{14.12}$$

where $n = N/N_0$ is the fractional speed, with the result floored at 0.5 kg/s and the pressure ratio floored at 1.05. The flow at the surge line for the same speed and suction condition is $\dot{m}_{surge} = \phi_s \dot{m}_{c0} n P_s/P_{s,0}$, and the surge margin reported to the operator and to the anti-surge loop is $\mathrm{SM} = \dot{m}_{comp}/\dot{m}_{surge} - 1$, the same fractional margin as Equation (1.11) of Section 1.4.6, with the surge flow in place of a map. Shaft power is polytropic head times the whole flow through the machine, recycle included, as shown in Equation (14.13).

$$W = \frac{\dot{m}_{comp}\,c_{p,g}\,(T_{r,1} + 273.15)}{\eta_c}\left[\left(\frac{P_d}{P_s}\right)^{\kappa} - 1\right] \tag{14.13}$$

where $\kappa = (\gamma - 1)/\gamma$ and the suction temperature is the refrigerant return temperature $T_{r,1}$. Recycled gas appears in $\dot{m}_{comp}$ and does no cooling, which is why anti-surge protection has a price. The air cooler and the driver close the loop on the weather, as shown in Equations (14.14) and (14.15).

$$\tau_{cond}\,\frac{\mathrm{d}T_{cond}}{\mathrm{d}t} = T_{amb} + \Delta T_{app} f_{ac}\left(\frac{\dot{m}_{comp}}{\dot{m}_{c0}}\right)^{0.8} - T_{cond} \tag{14.14}$$

$$\begin{aligned} P_d &= P_{d,0}\,\mathrm{e}^{k_d (T_{cond} - 30)} \\ W_{max} &= W_{driver}\left[1 - \beta\,(T_{amb} - 18)\right] \end{aligned} \tag{14.15}$$

where $\beta = 0.007/^{\circ}\mathrm{C}$ is the gas turbine derate. Both P_d and W_{max} move against the plant as the air warms.

14.3.6 Actuators, Makeup and Feed

Four actuator positions are states, each a first-order lag on its command with time constants of 0.15, 0.30, 0.08 and 0.20 min for the Joule-Thomson valve, the compressor speed, the recycle valve and the feed valve. The refrigerant inventory moves toward the makeup target at a bounded rate, as shown in Equation (14.16).

$$\frac{\mathrm{d}x}{\mathrm{d}t} = \mathrm{clamp}\left(\frac{x_{SP} - x}{\tau_{make}},\ -\dot{x}_{max},\ +\dot{x}_{max}\right) \tag{14.16}$$

with $\tau_{make} = 25\,\mathrm{min}$ and $\dot{x}_{max} = 0.010/\mathrm{min}$, so a blend correction of five points takes about half an hour. Feed is $\dot{m}_{NG} = \dot{m}_{NG,0}\,(z_F/100)\,f_{avail}$.

14.3.7 The Four Approaches and What the Control Room Sees

The four candidate pinch stations are the warm end, the two zone boundaries and the cold end, and the minimum internal temperature approach is the smallest of them, as shown in Equation (14.17).

$$\begin{aligned}\Delta T_{min} = \min\big(&T_{feed} - T_{r,1},\ \ T_{g,1} - T_{r,2},\\ &T_{g,2} - T_{r,3},\ \ T_{g,3} - T_{evap}\big)\end{aligned} \tag{14.17}$$

At the design point the four are 5.01, 5.00, 5.00 and 7.00 °C, so the bundle is balanced rather than pinched at any one station and the application names whichever is smallest at the time. Only the fourth term of (14.17) has a temperature measurement at both ends of it, and even that one requires T_{evap}, which (14.9) obtains from the composition. The chromatograph is a zero-order hold of period $t_{GC} = 12\,\mathrm{min}$, refreshed at the first controller scan at or after each multiple of t_{GC}, so the estimate the control room computes carries an age, as shown in Equation (14.18).

$$\hat{T}_{evap} = T_{evap}\big(\hat{P}_s, x_{GC}\big), \qquad \Delta\hat{T}_4 = \hat{T}_{LNG} - \hat{T}_{evap} \tag{14.18}$$

where $\hat{P}_s$ and $\hat{T}_{LNG}$ are the filtered transmitter readings and x_{GC} is the last analysis. That is the cold-end approach the approach controller holds. The estimated minimum reported beside the true one evaluates the same cold-end term with the plant's own suction pressure and product temperature, so its error is the composition error alone. The two internal boundary estimates are shifted by $0.55\,k_x\,(x_{GC} - x)$ to represent an estimator that runs the same zone model on a stale composition, the warm-end estimate is taken as observed, and the estimated minimum $\Delta\hat{T}_{min}$ is the smallest of the four estimates. The gap between ΔT_{min} and $\Delta\hat{T}_{min}$ is the whole subject of the application, and it is zero only while the analysis is current.

14.3.8 Parameters, Initialization and Numerics

Tables 14.2 and 14.3 list every constant. Seven are hidden and never displayed, and those seven set the inertia and the speed of the suction inventory that the scenarios explore. The rest are visible and adjustable from the parameter panel.

No scenario begins on a transient the operator did not cause. The initializer solves the suction pressure at which the Joule-Thomson valve passes exactly what the compressor takes away, by bisection on the residual $\dot{m}_{JT} + \dot{m}_{rec} - \dot{m}_{comp}$ over 0.6 to 14 bar, which is monotone because raising the pressure feeds the machine and starves the valve. The three duties then follow from (14.4) by direct solution rather than iteration, because marching the gas down and the refrigerant up makes the three duties linear in each other. At $z_{JT} = 54.9\,\%$ and $N = 92.0\,\%$ the initializer returns 2.901 bar, 32.99 kg/s of refrigerant, −154.95 °C of LNG and 23.09 MW of shaft power, and the design operating point specified directly as 2.90 bar and 33.0 kg/s returns −155.00 °C and 23.10 MW. The application's verification suite, which the shared test harness of the apps runs against the shipped code, checks that second point, the absence of crossed streams at steady state over a grid of 75 operating conditions, and eleven other properties.

Table 14.2 Feed, Cold Box and Refrigerant Parameters

Symbol	Meaning	Value	Unit	Visibility
$\dot{m}_{NG,0}$	feed at 100 % valve	20.0	kg/s	adjustable
T_{feed}	feed temperature	30	°C	adjustable
$c_{g,1}$	gas c_p, precool	2.60	kJ/kg/°C	adjustable
$c_{g,2}$	gas c_p, liquefy	5.60	kJ/kg/°C	adjustable
$c_{g,3}$	gas c_p, subcool	3.40	kJ/kg/°C	adjustable
UA_1	precool conductance	676	kW/°C	adjustable
UA_2	liquefy conductance	1904	kW/°C	adjustable
UA_3	subcool conductance	401	kW/°C	adjustable
$C_{w,1}$	precool bundle metal	180000	kJ/°C	hidden
$C_{w,2}$	liquefy bundle metal	520000	kJ/°C	hidden
$C_{w,3}$	subcool bundle metal	160000	kJ/°C	hidden
C_{NG}	gas holdup per zone	6000	kJ/°C	hidden
C_{MR}	refrigerant holdup	8000	kJ/°C	hidden
$c_{m,1}$	refrigerant c_p, warm	1.576	kJ/kg/°C	adjustable
$c_{m,2}$	refrigerant c_p, mid	3.394	kJ/kg/°C	adjustable
$c_{m,3}$	refrigerant c_p, cold	1.949	kJ/kg/°C	adjustable
s_1	warm zone blend slope	−1.40	per unit x	adjustable
s_2	mid zone blend slope	−0.15	per unit x	adjustable
s_3	cold zone blend slope	+1.60	per unit x	adjustable
x_0	design light fraction	0.420	—	adjustable
$T_{evap,0}$	boiling T at design	−162.0	°C	adjustable
k_P	boiling T per bar	9.5	°C/bar	adjustable
k_x	boiling T per unit x	120	°C	adjustable
τ_{make}	makeup response	25	min	adjustable
$\dot{x}_{max}$	fastest makeup rate	0.010	1/min	adjustable

The engine of Section 1.6.1 advances the state vector on a fixed 0.01 min step (0.6 s) under classical fourth-order Runge-Kutta, the six controllers execute on a 0.05 min (3 s) scan, and the historian keeps ten samples for every simulated minute. The clock runs in minutes, and the speed setting counts simulated minutes per second of real time. The noise seed is 20260913, and no scenario switches the noise on, so every result in this chapter is deterministic. A speed of 20, which runs of several hours need, alters only the rate at which simulated time is consumed and leaves the trajectory untouched.

14.3.9 Assumptions

The model is deliberately detailed on the pinch and thin on thermodynamics, and each simplification fails somewhere.

Three lumped zones stand in for a continuous bundle. The pinch can only appear at a zone boundary or at an end, and on a real coil-wound exchanger the tightest point is usually in the middle of a zone. The zone heat capacities are fixed by the design temperature ranges, so a long way off design the capacity assigned to each zone stops being physical.

The refrigerant is two numbers. A light fraction and the heat capacity it implies per zone replace four components and a bubble curve. The signs and the orders of magnitude in (14.2) are right, but a real blend has a phase envelope that a linear coefficient cannot represent, and a real makeup system adds one component at a time rather than moving a single fraction. Equation (14.9) is linear in pressure and in composition, which is a fair fit over the range the sliders allow and not beyond it.

The gas is a sensible-heat stream. The latent heat of condensation is carried by the elevated $c_{g,2}$ of the middle zone, so the application can show how much extra duty a rich feed brings but not what it does to the shape of the cooling curve. There is no process-side pressure drop.

The bundle metal is a lag on duty and not a temperature node. That carries the thermal inertia, which is what matters for control, and loses the internal metal profile, which is what matters for cool-down and for thermal stress.

Table 14.3 Machinery, Valve, Instrument and Limit Parameters

Symbol	Meaning	Value	Unit	Visibility
$\dot{m}_{c0}$	flow at rated speed	39.9	kg/s	adjustable
N_0	design speed	100	%	hidden
$P_{s,0}$	design suction	2.90	bar	adjustable
a_M	flow lost per excess PR	0.55	—	adjustable
ϕ_s	reduced flow at surge	0.62	—	adjustable
$c_{p,g}$	vapor heat capacity	1.90	kJ/kg/°C	adjustable
κ	$(\gamma-1)/\gamma$	0.230	—	adjustable
η_c	compressor efficiency	0.720	—	adjustable
W_{driver}	driver rating at 18 °C	26000	kW	adjustable
β	driver derate	0.0070	1/°C	adjustable
M_v	suction vapor holdup	1000	kg	hidden
$P_{d,0}$	discharge at 30 °C	46.0	bar	adjustable
k_d	discharge rise	0.018	1/°C	adjustable
ΔT_{app}	air cooler approach	14.15	°C	adjustable
τ_{cond}	air cooler lag	1.5	min	adjustable
k_{JT}	JT valve capacity	9.14	kg/s per $\sqrt{\text{bar}}$	adjustable
k_{rec}	recycle valve capacity	3.6	kg/s per $\sqrt{\text{bar}}$	adjustable
τ_{JT}	JT valve travel	0.15	min	adjustable
τ_N	speed response	0.30	min	adjustable
τ_{rec}	recycle valve travel	0.08	min	adjustable
τ_F	feed valve travel	0.20	min	adjustable
T_{spec}	tank bubble point	−150.0	°C	adjustable
ΔT^{lo}_{min}	pinch warning	1.5	°C	adjustable
SM_{alarm}	surge alarm	0.10	—	adjustable
τ_{Tm}	cryogenic TT response	0.35	min	adjustable
τ_{Pm}	pressure transmitter	0.08	min	adjustable
τ_{Wm}	power transducer	0.15	min	adjustable
t_{GC}	chromatograph cycle	12	min	adjustable
p_{LNG}	value of on-spec LNG	0.44	$/kg	adjustable
p_{fuel}	cost of shaft power	0.030	$/kWh	adjustable

The machinery is a single-point map. Equation (14.12) is one line with a surge line and no choke, there is no compressor curve to read, and the driver is a rating with a linear ambient derate rather than a thermodynamic model. The surge margin is reported and defended, and the reversal itself is not simulated.

The chromatograph is a perfect sample and hold. Real analyzers have a sample line with its own lag, a calibration that drifts, and a habit of failing in ways that leave the last good value on the screen. Failure of the analyzer is the most common cause of the inferential problem this application demonstrates, and it is outside the model.

Three disturbances are stylized. A relief event that removes five points of light ends at once is faster than any real release, but it separates the composition upset from the analyzer delay cleanly, which a slow drift would not. An ambient change of eighteen degrees arrives as a step, which the air cooler follows within about five minutes, rather than over the hours a real afternoon takes, because the constraint crossing is the subject and not the ramp rate. A feed cut of 35 % in a single step would trip most plants, and it is used because the over-refrigerated bundle it produces is the clearest illustration of specific power moving while every controlled variable stays on setpoint.

Cool-down from ambient is absent. Every scenario starts with a cold plant. The sixteen to thirty hours a real train takes to get there is the operation that most limits how often anybody is willing to shut one down.

The arithmetic is too reliable. Freezing the step and the scan gives the same trajectory on every

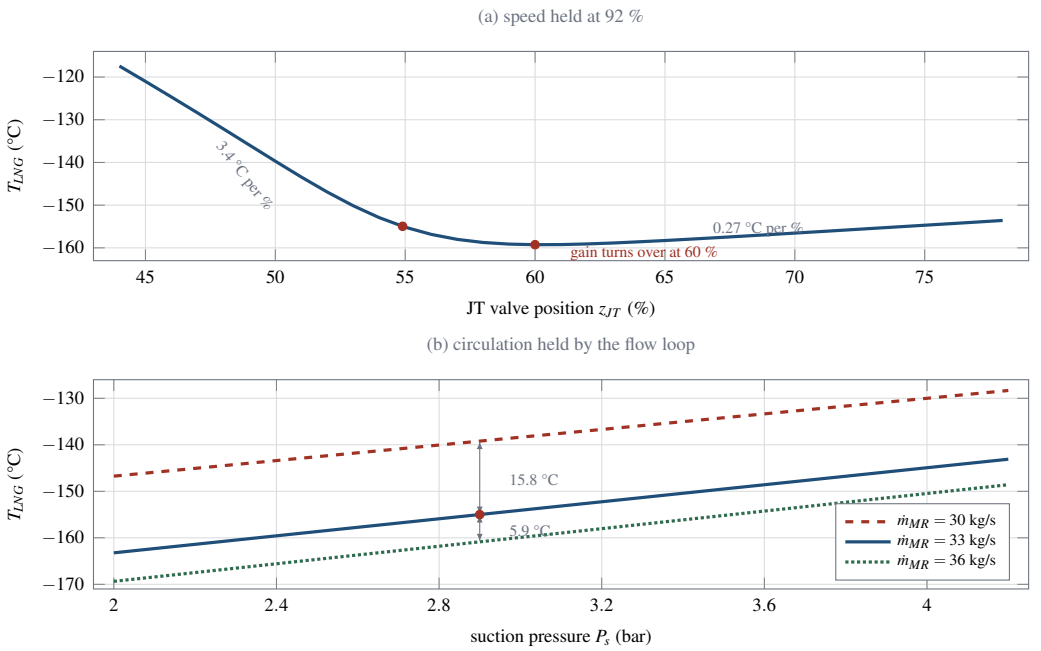


Figure 14.2 Steady-state characteristics computed from the implemented model. Panel (a) sweeps the Joule-Thomson valve with the compressor speed fixed, so that both of the valve's effects act, and the product temperature reaches a minimum of −159.26 °C at 60 % travel. Panel (b) sweeps the suction pressure with the circulation held, so that only one effect acts, and the three curves are straight and nearly parallel: 30 kg/s dashed, 33 kg/s solid and 36 kg/s dotted. The dots mark the design point and, in (a), the turning point.

machine, and a real control system delivers scan jitter, coarse measurement resolution and an execution order that moves.

14.4 Pairing Analysis

14.4.1 Two Effects on One Valve

The Joule-Thomson valve reaches the product temperature by two paths that oppose each other. More valve opening is more refrigerant through the bundle, which is more cooling and arrives over the four to seven minutes of the metal lags of (14.5). More valve opening also delivers more vapor to a compressor that cannot swallow it, so by (14.10) the suction pressure rises, and by (14.9) the refrigerant boils warmer. That second path is 63 % complete in 0.7 min and 95 % complete in 3.2 min after a step to 62 %. Figure 14.2(a) is the resulting open-loop characteristic with the compressor held at 92.0 % speed.

Closing the valve from 54.9 to 44 % costs 37.5 °C of product temperature at steady state, and in scenario 2 the product is 31.1 °C warmer than at 54.9 % an hour after the step and still moving, because the plant is short of refrigerant and the whole bundle has to warm up. Opening it by the same eleven points buys 3.0 °C, which the scenario delivers within eight minutes by putting the product onto the boiling temperature, and past 60 % the product gets warmer again, reaching −157.95 °C at 66 % against −159.08 °C at 62 %. The average gain over those two excursions is 3.4 °C per percent in one direction against 0.27 °C per percent in the other, a factor of 12, with a sign change in between. A relative gain array computed on this pair flags the interaction, because both candidate manipulated variables reach both candidate controlled variables [6], although a linear interaction measure cannot

show the change of sign. No proportional-integral setting repairs it, and the standard remedy is not a better tuning but a structural one, which is to give one of the two effects to another loop. The seawater train of Volume 6, Chapter 2 has a valve whose gain reverses sign a few percent either side of design for a similar reason, two terms of opposite sign on one handle.

14.4.2 What the Pressure Loop Buys

Closing a suction pressure loop on the compressor takes the pressure path away from the valve. What remains is Figure 14.2(b), and it is a straight line. Over 2.6 to 3.2 bar at 33 kg/s the product temperature moves 9.15 °C per bar, and the same sweep at 30 and 36 kg/s gives 8.37 and 9.44 °C per bar, so the gain of the outer loop is within 6 % of constant across the whole working range of circulation. The shaft power moves the other way over the same sweep, from 24.34 to 22.00 MW, so colder product is bought with power and the trade is visible on the same screen.

The vertical spacing of the three curves in Figure 14.2(b) carries the second result. At 2.90 bar the step from 30 to 33 kg/s is worth 15.8 °C while the step from 33 to 36 kg/s is worth 5.9 °C and the next three kilograms are worth 1.0 °C. Below about 31 kg/s there is not enough refrigerant to carry the duty at all, and above about 35 kg/s the cold end is already sitting on the boiling temperature and the surplus is pumped around the loop for nothing at about 1 MW of shaft power for every 3 kg/s. The ratio of the low-flow gain to the high-flow gain is 16.1. Circulation is therefore a capacity variable that decides how much cooling the plant has, and suction pressure is the trim variable that decides what temperature that cooling is delivered at, which is the pairing the published control studies of this plant arrive at from steady-state loss arguments [4, 5].

🕮 **Pause.** Try Exercise 14.3 before reading on.

14.4.3 The Blend Is a Max-Min Problem

Composition is the third handle and it is not a control handle. A heavy blend at $x = 0.360$ raises the boiling temperature to −154.80 °C, so at 2.90 bar no circulation can bring the product to its −155 °C setpoint. At 33 kg/s the product settles at −146.04 °C with the smallest approach at 3.22 °C on the subcool boundary, and 39 kg/s buys −154.66 °C only by closing the cold end to 0.14 °C. A light blend at $x = 0.500$ lowers the boiling temperature to −171.60 °C and gives away the warm end instead, where the approach closes to 1.22 °C because a light mixture has almost nothing left to boil at −40 °C. The design blend is neither. Searching x over 0.32 to 0.54 at 2.90 bar and 33 kg/s, the application returns $x = 0.419$ with a smallest approach of 4.97 °C, against 5.00 °C at the nominal $x_0 = 0.420$. The objective has an interior maximum, the makeup system needs half an hour to move a few points, and the measurement that says where the blend is arrives every twelve minutes, so a proportional-integral controller pointed at composition would not know which way to go. The simulator offers the search as a manual helper instead.

14.5 Implemented Control Structure

14.5.1 Five Loops and a Selector

Five loops are implemented, and the control structure switch selects how many of them are in service. Figure 14.3 shows the arrangement in its full form. The two inner loops hold pressure and flow on their own setpoints. The cascade adds the product temperature above the pressure loop and, when the approach mode is enabled, the inferred cold-end approach above the flow loop. The throughput switch puts the feed valve behind a low selector with the driver power controller. A fourth setting of the structure switch, never the default, hands the suction pressure and circulation setpoints and the makeup target to the model predictive controller of Section 14.5.5.

14.5.2 The Discrete Algorithm

All five loops and the anti-surge loop run the velocity-form PID of Section 1.6.2, Equations (1.16) to (1.19), executed every $\Delta t_c = 0.05$ min. Anti-windup is back-calculation with tracking time constant

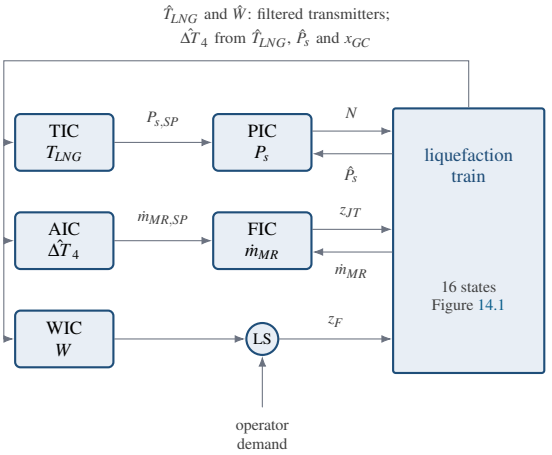


Figure 14.3 Control structure in its full form. The temperature controller drives the suction pressure setpoint, the inferred approach controller drives the circulation setpoint, and a low selector gives the feed valve to whichever of the operator and the power controller asks for less.

Table 14.4 Default Tuning of the Five Adjustable Loops. An Action of +1 Lowers the Output When the Measurement Rises.

CV	MV	K_c	τ_I (min)	a	Output limits
T_{LNG}	$P_{s,SP}$	0.075	14	+1	1.9–5.2 bar
P_s	N	14	1.6	−1	55–105 %
$\dot{m}_{MR}$	z_{JT}	2.0	0.4	+1	0–100 %
$\Delta\hat{T}_4$	$\dot{m}_{MR,SP}$	0.18	60	−1	14–48 kg/s
W	z_F	0.004	30	+1	40–110 %

$T_t = \tau_I$ on every loop [19]. Under a held error the integral comes to rest at the limit itself, so the demand before clipping exceeds the limit only by the proportional term and the output leaves the limit as soon as the error reverses. Away from the limits the recursion is algebraically identical to the positional law with the same gains, and on a limit it is not. Derivative action is zero on all six loops, so the derivative filter never runs, and no loop carries a rate limit, so the clipped output reaches the actuator model directly. In MANUAL the loops keep executing and track the operator's actuator position exactly, and in AUTO the operator's stations are set to the controller outputs after every execution, so transfer in either direction is bumpless with no separate logic. Output bias is zero on every loop and the integral is seeded at reset with the starting actuator value, so a loop that starts on its setpoint starts without a bump.

The action flag is the ISA sign convention. A loop with $a = +1$ is reverse acting, so a rising measurement lowers the output, and a loop with $a = -1$ is direct acting. Table 14.4 lists the five adjustable loops with their pairings, defaults and limits. The anti-surge loop is fixed at $K_c = 220$ and $\tau_I = 0.25$ min on the recycle valve, is reverse acting on surge margin, and is not offered for tuning because it is protection rather than regulation.

Three of the five tunings deserve a comment. The pressure loop carries 14 % of speed per bar with 1.6 min of reset, an order of magnitude more gain and an order of magnitude less reset than the temperature loop above it, because (14.10) is a fast, self-regulating inventory, 63 % complete in 0.65 min after a speed step, against the 4.44 to 6.65 min metal lags the temperature loop works through. The temperature loop carries 0.075 bar per degree with 14 min of reset, which is slow

because everything between the pressure and the thermocouple is aluminum. The approach loop carries 60 min of reset, because its measurement is refreshed by an analyzer every twelve minutes and a loop that reacts inside its own measurement interval is reacting to the hold rather than to the plant.

14.5.3 The Low Selector on Throughput

The feed valve belongs to the operator until the driver runs out of power. The power controller computes a demand against the smaller of the operator power setpoint and the derated driver rating W_{max} of (14.15), and the selector passes the lower of that demand and the operator's own, as shown in Equation (14.19).

$$z_F = \min(z_{F,op},\, u_W)\,, \qquad W_{lim} = \min(W_{SP},\, W_{max}) \tag{14.19}$$

The loser is told what was selected and tracks it through the same MANUAL expression, so the losing controller does not wind up while it is out of service and hands over without a bump when the selection changes, as the reset that follows the selector does in Section 1.4.8. The operator's own demand is never overwritten by the selected value, so the selector hands the valve back as soon as the driver can afford it. MANUAL holds the valve at the selected z_F (88.0 against 100 % at 40 min of scenario 9) and AUTO restores the demand, because the selector acts in AUTO only. That path is slow by construction. The feed valve does not reach the compressor directly. It changes the heat load, which changes the product temperature, which the temperature controller answers with suction pressure, which is what the machine feels, so the whole cold box is inside the power loop and the 30 min reset is a consequence rather than a choice.

14.5.4 Instruments and the Analyzer

The product thermocouple, the suction pressure transmitter and the power transducer are first-order filters of 0.35, 0.08 and 0.15 min, and the refrigerant flow measurement carries a fixed 0.08 min filter. Optional colored noise can be added to the temperature and pressure signals with a standard deviation of up to 2 °C and 0.15 bar. The composition analyzer is different in kind. It is a sample, a column and an integration, and it arrives as a step every 12 min carrying the plant that produced it. Between analyses the estimator of (14.18) runs on a number that is up to twelve minutes old, which is the mechanism the seventh scenario shows.

14.5.5 Model Predictive Control

A fourth setting of the structure switch, never the default and selected by no shipped scenario except scenario 12, replaces the temperature and approach controllers with one model predictive controller (MPC), an instance of the option described in general in Section 1.6.3. Every $T_s = 0.5$ min, it solves a quadratic program for the three setpoints $u = (P_{s,SP},\, \dot{m}_{MR,SP},\, x_{SP})$ and keeps only its first move, held for one decision period, the receding-horizon arrangement of industrial predictive control [20], also proposed in adaptive form for natural gas liquefaction [21]. The pressure and flow loops carry out the first two setpoints, the makeup law of (14.16) carries out the third, and the low selector and the anti-surge loop stay above the controller on the run's own settings.

The controller's model is the cold box of Section 14.3 written again with the declared values of Tables 14.2 and 14.3 and the design disturbances, $T_{amb} = 18\,^{\circ}\mathrm{C}$ and $f_{rich} = f_{avail} = f_b = f_{ac} = 1$, never with the run's, so a plant the reader changes is model mismatch. A tilde marks the controller's copy of a quantity. The inner loops enter as their identified responses to the setpoints, and the gas heat-capacity rates carry an estimated load factor $\tilde{L}$, as shown in Equation (14.20).

$$\begin{aligned} \tilde{P}_s &= \alpha\, \tilde{p}_1 + (1-\alpha)\, \tilde{p}_2, \qquad \tau_i \frac{\mathrm{d}\tilde{p}_i}{\mathrm{d}t} = P_{s,SP} - \tilde{p}_i \\ \tau_m \frac{\mathrm{d}\tilde{m}}{\mathrm{d}t} &= \dot{m}_{MR,SP} - \tilde{m}, \qquad \tilde{C}_{g,k} = \tilde{L}\, \dot{m}_{NG,0}\, \frac{\tilde{z}_F}{100}\, c_{g,k} \end{aligned} \tag{14.20}$$

where $\tilde{p}_1$ and $\tilde{p}_2$ are parallel lags, $\tilde{m}$ is the circulation and $\tilde{z}_F$ follows the feed command through its 0.20 min lag. The constants $\alpha = 0.413$, $\tau_1 = 0.70$ min, $\tau_2 = 3.90$ min and $\tau_m = 0.307$ min are fits to setpoint steps of either sign in scenario 3, with residuals of 1.2 and 2.1 % RMS of a unit step. The model carries 15 states, and its shaft power and compressor speed are (14.13) and (14.12) plus biases $\tilde{d}_W$ and $\tilde{d}_N$. Driven open loop on the declared plant through a 150 min sequence of steps, it predicts the LNG temperature within 0.036 °C RMS and every approach within 0.054 °C RMS. One linearization at the design point would put the LNG at −144.95 °C for 30 kg/s of circulation where the model settles at −139.21 °C, so every decision linearizes the model anew around the current estimate and samples it over T_s with an exact zero-order hold.

Each decision minimizes Equation (14.21) over $N = 120$ steps, a horizon of 60 min, with the moves held in nine blocks lasting 1, 1, 2, 2, 4, 6, 9, 15 and 80 steps, a program of 27 moves and 16 slacks.

$$
\begin{aligned}
\min_{u,\varepsilon} \sum_{k=1}^{N} \Big[& \big(\tilde{T}_{LNG,k} - T^*\big)^2 + q_A \big(\Delta\tilde{T}_{4,k} - \Delta T_4^*\big)^2 \\
& + r \sum_{j=1}^{3} g_j^2 \big(u_{j,k} - u_j^*\big)^2 \Big] + \lambda \sum_{b=1}^{9} \sum_{j=1}^{3} g_j^2 \Delta u_{j,b}^2 \\
& + \sum_{\ell=1}^{16} \rho_\ell \big(\varepsilon_\ell^2 + \varepsilon_\ell\big) \\
\text{s.t. } & u_j^{lo} \le u_{j,b} \le u_j^{hi}, \quad |\Delta u_{j,1}| \le \delta_j, \quad \varepsilon_\ell \ge 0 \\
& \tilde{y}_{i,k} \le \bar{y}_i + \phi_{i,k} + \varepsilon_\ell \quad \text{(LNG, power, speed)} \\
& \tilde{y}_{i,k} \ge \bar{y}_i - \phi_{i,k} - \varepsilon_\ell \quad \text{(approaches, surge margin)}
\end{aligned}
\tag{14.21}
$$

where k counts the steps and b the blocks, $\Delta u_{j,b}$ is the move of setpoint j into block b, the first measured from the setpoint in force, T^*, ΔT_4^* and u^* are the targets of (14.22), q_A is 1 when the circulation is the controller's and 0 when it is the operator's, $r = 0.01$, $\lambda = 10$, and $g = (9.2,\ 2.0,\ 145)$ °C per bar, per kg/s and per unit of light fraction prices each setpoint in the degrees of LNG it is worth near design. The bounds are the clamps the cascade and the makeup apply, 1.9 to 5.2 bar, 14 to 48 kg/s and 0.30 to 0.55, and the first-move cap is $\delta = (0.05\,\text{bar},\ 0.5\,\text{kg/s},\ 0.02)$, both hard. Every output limit $\bar{y}_i$ is soft. The LNG ceiling is $T_{spec} - 1$ °C, the three internal approaches are held at $\Delta T_{min}^{lo} + 1$ °C, which is 2.5 °C as shipped, the cold-end guard is 0.5 °C, the power limit is the selector's W_{lim} of (14.19), or 0.5 MW under it when the feed is the operator's, the speed limit is 104 % and the surge floor is 2 % above the anti-surge setpoint. The slack prices ρ_ℓ are 10^4 for the internal approaches, 10^3 for the LNG ceiling and the power and 10^2 for the rest, so the approach floor outranks the driver tenfold. The power and surge limits are soft because the shared solver has no hard output rows, which leaves the selector and the anti-surge loop as the hard protections. That solver also carries one slack per soft output, so each limit enters twice, over the first 10 min and over the rest, with a funnel $\phi_{i,k} = v_i\, \mathrm{e}^{-kT_s/\tau_i}$ that returns an output already past its limit by v_i, with τ_i of 2 to 5 min and 25 min for the internal approaches in the second window. No slack per step is claimed.

A steady-state Kalman filter, designed at the design point for the 3 s scan, corrects 14 states and the random-walk disturbances $\tilde{L}$, $\tilde{d}_W$ and $\tilde{d}_N$ at every scan from the five signals the shipped loops read, which are $\hat{T}_{LNG}$, $\hat{P}_s$, the circulation, the shaft power and the speed command. Such disturbance states keep a predictive controller free of steady offset under model error [22]. The blend is not estimated. It is the last analysis x_{GC} carried forward by (14.16), which is what the control room knows. The approach limit is therefore a limit on a prediction, because the warm end and the two zone boundaries of (14.17) have no instrument, and in scenario 7 the smallest approach the controller predicts misses the true minimum by up to 5.37 °C before the vented light ends reach an analysis. The cold end, with a measurement on both sides, is tracked to its setpoint instead and guarded only against a cross.

The targets are the solution of Equation (14.22) on the model's steady map, found by up to eight steps of sequential quadratic programming inside a trust region of 0.3 bar, 3 kg/s and 0.04.

$$\begin{aligned}\min_{u^*}\ & \left(\tilde{T}_{ss}-T_{LNG,SP}\right)^2+q_A\left(\Delta\tilde{T}_{4,ss}-\Delta T_{SP}\right)^2\\ &+2500\left(x^*-x_0\right)^2+c_W\,\tilde{W}_{ss}+10^{-3}\sum_{j=1}^{3}g_j^2\left(u_j^*-u_j\right)^2\\ &+\sum_i\rho_i\left(\varepsilon_i^2+\varepsilon_i\right)+10^4\left(\varepsilon_T^2+\varepsilon_T\right)\end{aligned}\tag{14.22}$$

where ss marks the model's steady value at u^*, u_j is the setpoint in force, ΔT_{SP} is the cold-end setpoint, ε_i carry the limits of (14.21) other than the cold-end guard at the same prices, and ε_T forbids an LNG target colder than the setpoint. While the selector has the feed, c_W rises from 0 to 10 per MW as the steady power climbs from 1.5 to 0.5 MW under W_{lim}, because a megawatt saved is feed the selector can give back. A squared degree of LNG error then costs as much as a tenth of a megawatt, and the target runs the LNG warm whenever power or a predicted approach is worth more.

A plan is discarded and marked unless the solver reports an optimum with every value finite and every move inside its bounds, and the temperature and approach controllers, which track the applied setpoints at every scan, then hold the setpoints until the next valid plan. Forced rejections hand over within 0.0024 bar of suction setpoint with the noise off, and every plan in the measured runs was accepted. The priming call at $t = 0$ already solves a plan. In Chrome a plan takes 3.43 ms on average in scenario 12, 0.68 % of a core at one simulated minute per second and 13.7 % at the top speed of 20. The plant on the run's settings, the instruments and the controller's copy on the declared values are thus three separate models, and the 8800 measurement packets of scenario 12 replayed over four other hidden plants give the same 800 plans bit for bit.

14.6 Results

The thirteen scenarios of Table 14.5 run as a laboratory sequence. Each one loads a starting mode, a control structure, five tunings, a set of disturbance positions and whatever timed events belong to it, and carries the briefing, completion rule and data export of Section 1.6.4. One hundred twenty-five columns, twenty-seven of them the predictive controller's, leave the historian at ten samples per simulated minute, among them the true minimum approach beside the estimated one, the light fraction beside the last analysis, the age of that analysis, all four approaches and the running specific power. The first two open in MANUAL with deliberately weak tunings loaded, so that a reader who switches to AUTO and closes the loops sees what a badly tuned plant looks like before the tuned settings arrive in scenario 3. No scenario switches the measurement noise on, and the results below are the same whether a case is loaded from the selector, reloaded or restarted.

14.6.1 Manual Operation and the Two Handles

Scenario 1 makes the pairing argument with two open-loop steps and no controller, each made here from the design line-up five minutes after Run, as the scenario asks (a station set before Run starts the plant at its steady state, so it would show no response). Opening the Joule-Thomson valve from 54.9 to 62 % raises circulation from 32.99 to 37.63 kg/s and the suction pressure from 2.901 to 3.167 bar within about five minutes, so the refrigerant that boiled at −161.99 °C now boils at −159.46 °C. Twenty minutes after the move the product is at −159.46 °C, on the boiling temperature itself, and the cold-end approach has closed from 7.04 °C to zero, which raises the temperature-cross alarm 24.3 min after the move. The whole 4.6 kg/s of extra refrigerant bought 4.51 °C. The model's steady state at 62 % is −159.08 °C with 0.38 °C left at the cold end, and the run creeps toward it only over hours (−159.17 °C 240 min after the move). Raising the compressor speed from 92.0 to 96.0 % instead moves the suction to 2.735 bar and the boiling temperature to −163.57 °C, and twenty minutes later the product is at −156.60 °C (−156.73 °C at steady state), with the four approaches

Table 14.5 The Thirteen Scenarios and What Each One Is Arranged to Show

#	Scenario	What it exercises
1	Two handles by hand	the valve's two effects, separated by experiment
2	The wrong pairing	direction-dependent gain and a sign change
3	Pressure and flow on setpoints	the linear handle, and a fast self-regulating inner loop
4	The cascade	temperature onto pressure, and what a degree costs
5	Circulation is capacity	gain varying by a factor of 16
6	The blend	an interior maximum in the smallest approach
7	Nitrogen to the flare	a stale analysis and an inferential error
8	A hot afternoon	two constraints crossing with no loop misbehaving
9	Throughput on the driver	a low selector with the plant inside the loop
10	Upstream cuts the feed	specific power as the diagnostic
11	Turndown into surge	protective override and what recycle costs
12	Three handles under one driver limit	one predictive plan for three setpoints, priced against the driver
13	A shift on the train (challenge)	380 minutes with four upsets and a heavy makeup target

Table 14.6 Measured Closed-Loop Responses, With the Model's Steady State in Parentheses. Times Are Simulated Minutes From the Step.

Step	Measured response
$P_{s,SP}$ −0.30 bar at $t = 0$, scenario 3	9.3 °C per bar, 63 % of the move in 8.3 min and 95 % in 18.0 min
P_s itself, same step	63 % of the move in 2.2 min
$T_{LNG,SP}$ −155 → −158 °C, scenario 4	after 70 min P_s is 0.32 bar lower and shaft power 1.31 MW higher, about 440 kW per degree, with the product at −157.83 °C
$\dot{m}_{MR,SP}$ 33 → 30 kg/s, scenario 5	product 11.4 °C warmer at 55 min and still moving (15.8 °C, to −139.2 °C)
$\dot{m}_{MR,SP}$ 30 → 36 kg/s	product 17.0 °C colder at 60 min, 5.6 °C below the start (21.7 °C, of which the last 3 kg/s contributes 5.9 °C)
$\dot{m}_{MR,SP}$ 36 → 39 kg/s, model only	(product cools 1.0 °C for 1 MW of shaft power)
Blend 0.420 → 0.360, scenario 6	at 105 min T_{evap} is −154.91 °C and the product −146.16 °C (−154.80 °C and −146.04 °C)
Blend 0.360 → 0.500	warm-end approach down to 0.63 °C at 88 min while the cold end keeps 5.4 °C or more (1.22 °C at 0.500)

essentially unchanged at 5.18, 5.09, 4.96 and 6.97 °C. One handle moved the capacity and spoiled the profile, and the other moved the temperature and left the profile alone.

🕮 **Pause.** Try Exercise 14.1 before reading on.

14.6.2 Closed-Loop Steps

Table 14.6 collects the measured responses of the tuned loops. Every entry is a change read at a stated time in the scenario or a time to a stated fraction of the final change, with the model's steady state beside it where the scenario moves on before the plant settles, so none of them is an accumulated integral. The application also reports integral of absolute error, integral of squared error, overshoot, settling time, actuator travel and saturated fraction on its scorecard, and the two integral measures keep accruing after a transient ends. An IAE figure from this plant is only interpretable with the observation window printed beside it, so the comparisons below are made with peak deviations, settling times, actuator travel and power instead.

The inner loop is faster than the outer one by the margin a cascade needs. Suction pressure reaches 63 % of a setpoint move in 2.2 min against 8.3 min for the product temperature, a ratio of 3.7, and everything between the two is aluminum.

Table 14.7 Measured Disturbance Results. Each Entry Compares the Condition Before the Upset With the Stated Later Instant.

Quantity	Before	After	Scenario
Ambient 18 to 36 °C, feed fixed, read at 170 min			
Discharge pressure (bar)	46.1	63.8	8
Shaft power (MW)	23.1	26.9	8
Driver rating (MW)	26.0	22.7	8
Same afternoon with the low selector, read at 200 min			
Feed valve (%)	100	88.9	9
Production (t/h)	72.0	64.0	9
Shaft power (MW)	23.1	22.7	9
Upstream removes 35 % of the feed, read 40 min later			
Product temperature (°C)	−155.0	−156.2	10
Shaft power (MW)	23.1	17.0	10
Production (t/h)	72.0	46.8	10
Specific power (kWh/kg)	0.321	0.362	10
Turndown to 20 kg/s, read at 60 and 120 min			
Surge margin (%)	45	6.5	11
Recycle valve (%)	0	10.6	11
Recycle shaft power (MW)	0	1.8	11
Product temperature (°C)	−155.0	−80.5	11

14.6.3 What an Inferential Measurement Costs

Scenario 7 is the central result of the chapter. The cascade is in service with the approach controller on circulation, the plant is settled at the design point, and at $t = 15$ min a relief event vents five points of light ends from the refrigerant inventory at once. Nine minutes later, just before the next analysis lands, the refrigerant is boiling at −160.3 °C instead of −162.0 °C, the product has walked warm to −151.9 °C, and the cold-end approach the control room computes has gone *up* to 12.7 °C, because it is still subtracting the boiling temperature the last analysis implies. The estimated minimum approach reads 6.5 °C while the true minimum is 4.2 °C.

🕮 **Pause.** Try Exercise 14.2 before reading on.

Both controllers then do the right thing for the wrong reason. The approach controller sees a wide approach and adds refrigerant, which is the correct action, and the temperature controller pulls the suction pressure down to chase the product back. The analysis at 24 min closes the gap between the estimated and true minimum to under 0.5 °C, but the makeup keeps moving the blend between analyses and the gap reopens to 0.9 °C by 36 min and stays above 0.5 °C until 48 min. By the hour mark the analysis reads 0.407 against a true 0.412 and the plant is back at −155.6 °C, on 23.7 MW instead of 23.1 MW for the same product. Figure 14.4 shows the screen at 26 min, two minutes after the first analysis that follows the event. The application's verification suite reproduces the transient directly and requires the gap between the estimated and true minimum approach to exceed 1.5 °C five minutes after the event and to have fallen by more than half by 45 min. Section 1.4.9 places this inferred constraint among the other measurement limits of the volume.

14.6.4 Constraints, Throughput and Turndown

Table 14.7 collects the disturbance scenarios. Each row reports the settled condition before the upset and the condition at the stated later instant, so the figures are comparisons rather than accumulations.

The eighth scenario is the one with no misbehaving loop in it. Warm air raises the condensing temperature, so the discharge pressure of (14.14) climbs from 46.1 to 63.8 bar and every kilogram costs a sixth more to compress, while (14.15) takes the driver from 26.0 to 22.7 MW. The two

numbers cross going in opposite directions, the temperature controller holds its setpoint by spending power the plant no longer has, and nothing on the temperature trend says so. Scenario 9 hands the feed valve to the selector of (14.19) and the plant settles at 88.9 % feed, 64.0 t/h instead of 72.0 t/h, with the shaft on the derated driver and the product on specification through the afternoon. Getting there is slow by construction. The shaft stays above the derated driver for 76 min after the ambient step, by as much as 1.31 MW, before the power controller brings it under. When a cold front arrives at $t = 200$ min the selector returns the valve on its own and the feed is back at 100 % within a minute, faster than the temperature cascade can follow, so the product runs warm to −146.6 °C at 203.7 min and spends 4.4 min off specification. By 400 min the shaft is at 21.8 MW of an available 27.8 MW.

> **TRY IT IN THE APP** **8 · A hot afternoon**
>
> Load the case and press Run. Touch nothing. Watch the Shaft power trend with its derated driver line. When the air warms from 18 to 36 °C, the driver falls from 26.0 to 22.7 MW while the shaft climbs from 23.1 to 26.9 MW. The LNG outlet temperature trend shows nothing wrong.

Scenario 10 is the case where the trends look perfect and the plant is worse off. Removing 35 % of the feed leaves the refrigerant loop circulating the same 33 kg/s into two thirds of the heat load, so the product runs down onto the boiling temperature and the cold-end approach collapses to nothing. The temperature controller answers by raising the suction pressure, which is exactly right, and forty minutes later the plant is at −156.2 °C on 17.0 MW, drifting slowly toward −155.1 °C and 16.5 MW by the time the feed returns at 140 min. Dividing rather than reading gives the diagnosis. That 17.0 MW is buying 46.8 t/h instead of 72.0 t/h, so specific power has gone from 0.321 to 0.362 kWh/kg, 13 % more energy per tonne with every loop on setpoint, and it is still 10 % above design at 140 min.

> **TRY IT IN THE APP** **10 · Upstream cuts the feed**
>
> Load the case and press Run. Upstream takes 35 % of the feed away. Watch the Power and Rate chips above the plant. Forty minutes after the cut the shaft is down from 23.1 to 17.0 MW, and production from 72.0 to 46.8 t/h, so each kilogram costs 0.362 against 0.321 kWh. Every loop stays on its setpoint.

Scenario 11 separates a machine limit from a process limit. Cutting circulation to 20 kg/s with the anti-surge system bypassed takes the surge margin from 45 to 6.5 %, inside the alarm at 10 % and inside the region where a real machine risks reversing. Restoring the recycle loop at 60 min buys the margin back at 10.6 % valve opening and 1.8 MW of shaft power, all of it spent compressing refrigerant that is throttled straight back to the suction. The product is the other half of the answer. At 20 kg/s the outlet is at −80.5 °C by 120 min and at 16 kg/s it is at −58 °C, so nothing on specification is produced for the entire scenario, which is the real answer to deep turndown on a single-refrigerant train.

The graded challenge runs 380 min from a plant whose makeup target is four points too heavy, whose circulation setpoint is two kilograms low and whose bundle has lost a tenth of its conductance, and applies an afternoon warming from 22 to 33 °C at $t = 30$, a rich feed at $t = 80$, a relief event at $t = 130$ and an evening cooling to 14 °C at $t = 180$ min. The stated target has four parts, and the application judges all four exactly as stated. The product must stay within ±2 °C of −155 °C and the true minimum approach above 1.5 °C over the last two hours, from 260 to 380 min, both read from the true values logged ten times a minute rather than from the thermocouple or the estimate. The shaft must not exceed the derated driver at any integration step of the whole run, and the specific power, shaft energy divided by on-specification product since $t = 0$, must finish below 0.36 kWh/kg. The approach is not judged before 260 min, which matters, because even with the suggested settings the cold end touches the boiling temperature twice in the first 80 min and the relief event crosses it by up to 5.2 °C for 2.5 min. The heavy makeup target and the relief event both act on the blend rather than on any loop, and the blend is the part no controller handles.

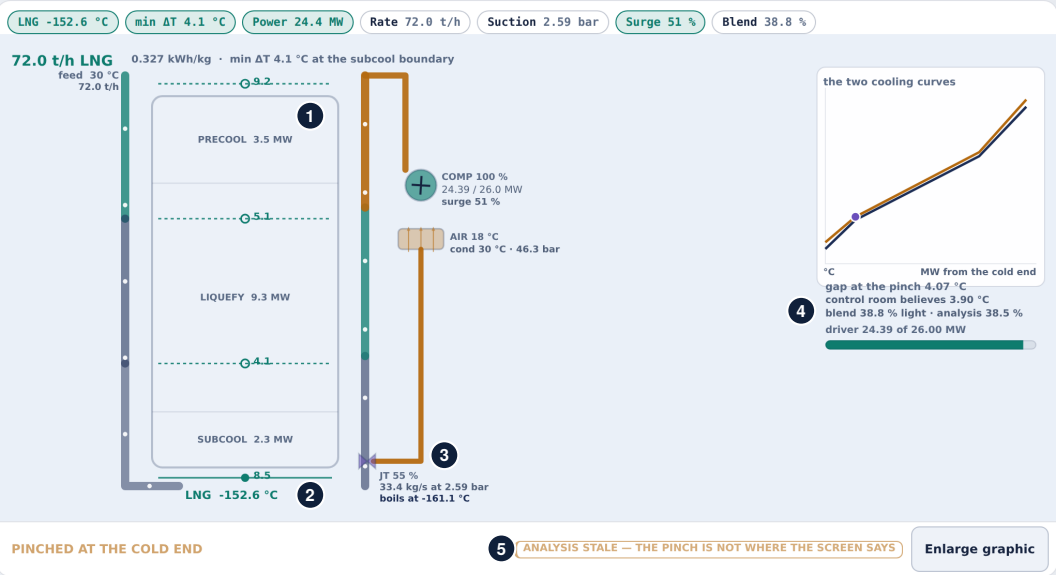


Figure 14.4 The stale-analysis transient of scenario 7 at $t = 26$ min, eleven minutes after the relief event took five points of light ends out of the refrigerant, with the cascade in service and the cold-end approach controller holding the circulation setpoint at 33.34 kg/s. A fresh analysis has since landed, so the readout beside the inset reads 38.8 % light against an analysis of 38.5 % and a true pinch of 4.07 °C against the 3.90 °C the control room computes for it. The product temperature has walked warm to −152.57 °C against a −155.00 °C setpoint and a −150.00 °C tank limit, and the temperature controller has pulled the suction pressure setpoint down to 2.563 bar after it. The cold box animation marks the four approach stations, three of them drawn hollow because no instrument exists there (1), and the one measured station at the cold end (2). The Joule-Thomson valve carries its opening, the circulation and the boiling temperature (3), the inset plots the two cooling curves against cumulative duty with the true and the computed pinch beside it (4), and the status line warns that the analysis is stale (5).

Each scenario is credited when the reader performs the actions its completion rule lists and the simulation reaches the stated state. The challenge is judged on the case as given, so changing a model parameter or a disturbance during the attempt voids it. With the shipped settings (pressure and flow loops only, weak tunings, makeup target 0.380) every part fails. The worst product error over the last two hours is 31.2 °C, the smallest approach 0.47 °C, the shaft spends 149.1 min above the derated driver and the specific power ends at 4.39 kWh/kg. With the suggested settings (the full structure of Figure 14.3 with the default tunings, makeup target 0.420 and a 20.5 MW power limit) every part passes, with a worst error of 0.49 °C, a smallest approach of 3.84 °C, no time over the driver (0.59 MW of margin at the closest) and 0.338 kWh/kg. Revealed before Run or carried through Reset, they give the same run.

14.6.5 Three Handles Under One Driver Limit

Scenario 12 is the afternoon of scenario 9, ambient 18 to 36 °C at 25 min and a cold front at 200 min with the selector in service, run first with the MPC and then, after Reset, with the cascade and its approach controller. Table 14.8 compares the structures on scenarios 12, 8, 10 and 7 and the challenge, the MPC on the settings of scenario 12 without retuning. In scenario 12 the model-information row gives the cascade the MPC's steady-state blend target, and the hindsight row is the best of 20 fixed

Table 14.8 Cascade Against the MPC, Default Seed, No Noise, Full Duration of Each Case

Structure	Over driver (min)	Peak over (MW)	Excess (MW min)	Under 1.5 °C (min)	Smallest ΔT_{min} (°C)	On-spec LNG (t)	Off spec (min)	LNG IAE (°C min)	w (kWh/kg)	Smallest SM (%)
Scenario 12, 400 min, hot afternoon and selector, 800 MPC plans										
MPC	18.6	2.78	20.4	41.5	−0.37	455.1	4.6	274.0	0.320	14.8
cascade	85.1	1.31	28.9	61.5	−0.16	450.5	4.4	275.3	0.329	18.6
cascade, MPC's blend target[a]	49.5	1.30	21.9	45.4	−0.14	455.1	3.7	321.5	0.323	16.8
hindsight cascade[b]	53.3	0.37	2.4	12.3	−0.50	457.5	1.5	140.6	0.312	16.2
Scenario 8, 260 min, feed on the operator, 520 plans per MPC row										
MPC	150.0	2.10	25.5	0.0	2.12	312.0	0.0	602.0	0.315	15.5
cascade	150.8	4.43	620.5	0.0	4.09	312.0	0.0	20.2	0.351	30.3
MPC, no power limit[c]	150.7	4.18	618.2	0.0	4.13	312.0	0.0	4.1	0.351	31.1
Scenario 10, 240 min, feed cut and return, 480 MPC plans										
MPC	0.0	0.00	0.0	36.4	−1.09	215.9	18.1	587.7	0.360	15.4
cascade	0.0	0.00	0.0	152.6	−0.59	230.4	6.0	337.4	0.337	25.9
cascade at 21.0 kg/s[d]	0.0	0.00	0.0	175.5	−0.59	118.7	99.1	1446.4	0.656	12.9
Scenario 7, 180 min, stale analysis, 360 MPC plans										
MPC	0.0	0.00	0.0	2.8	1.03	216.0	0.0	62.0	0.323	44.9
cascade	0.0	0.00	0.0	0.2	1.06	216.0	0.0	88.9	0.326	44.9
Scenario 13, 380 min, challenge, 760 plans per MPC row										
MPC, suggested, passes	0.0	0.00	0.0	77.3	−4.32	401.3	17.2	908.6	0.325	14.2
cascade, suggested, passes	0.0	0.00	0.0	177.2	−5.20	388.8	0.0	239.5	0.338	18.7
MPC, rest as shipped, fails	150.0	3.03	317.9	293.7	0.00	358.6	81.2	1854.9	0.408	23.6

Over driver counts every physics step with the true shaft power above the derated driver, and the approach columns read the true minimum. SM is the smallest surge margin. No MPC plan was rejected. [a]Makeup target set to the MPC's blend target as it moved, no plan. [b]Best by time under 1.5 °C of makeup 0.40 to 0.48, shipped or twice-as-fast temperature and approach loops, approach setpoint 5 or 7 °C, here 0.44, fast, 5 °C. [c]A copy with no power row or price. [d]Circulation forced to 21.0 kg/s over 100 to 139.5 min.

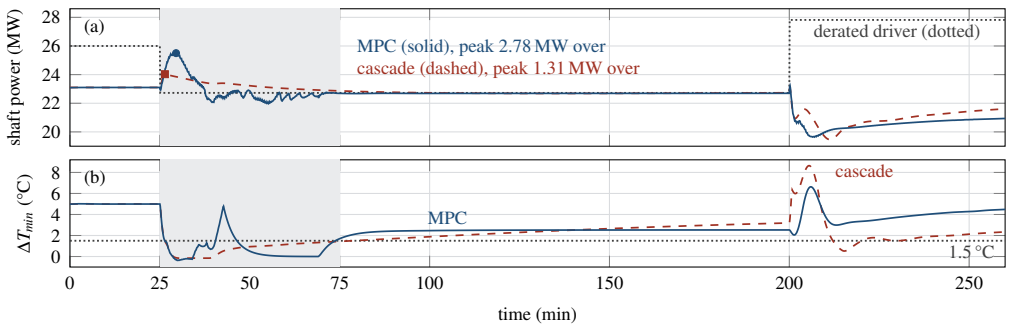


Figure 14.5 Scenario 12 with the MPC that the scenario loads and with the cascade selected after Reset, default seed, with the window from 25 to 75 min shaded. (a) The true shaft power against the derated driver. The MPC is over for 18.6 min with a peak of 2.78 MW and 20.4 MW min, and the cascade for 85.1 min with a peak of 1.31 MW and 28.9 MW min. (b) The true minimum approach against the 1.5 °C warning. In both panels the MPC is solid, the cascade dashed and the limit dotted, and the dot and the square mark the two peaks.

cascade settings chosen after the runs. Figure 14.6 sets the MPC and cascade rows of the five cases side by side.

Under the MPC the shaft is over the derated driver for 18.6 against 85.1 min, all of it before 73.2 min, but it goes further over, with a peak of 2.78 against 1.31 MW and 8.6 against 6.2 min more than 1 MW over, while the excess is 20.4 against 28.9 MW min, as shown in Figure 14.5(a). Most of

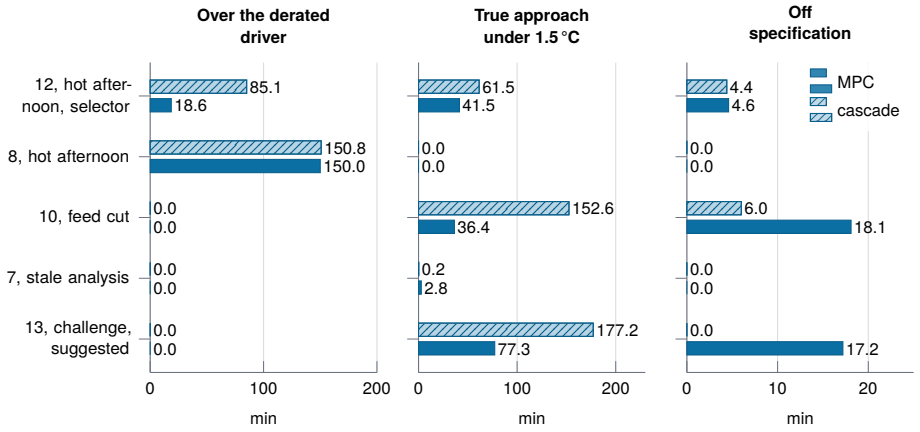


Figure 14.6 The MPC and the cascade of Table 14.8 on the five cases, labeled by scenario number, solid bars for the MPC and hatched bars for the cascade, with the table's value beside each bar. Notice that the MPC cuts the time over the driver markedly only in scenario 12, and that it spends more time off specification than the cascade in scenarios 10 and 13 while spending less time with the approach under 1.5 °C. The other rows of the table are not drawn.

the cascade's time is the selector holding the shaft just above its limit, more than 0.1 MW over for only 55.1 min. The deeper excursion is the price order at work. In the seven minutes after the ambient step the plan raises the circulation setpoint from 32.5 to 36.0 kg/s and the makeup target from 0.42 to 0.55 to hold its predicted internal approaches, accepting up to 1.84 MW of power slack, the selector cuts the feed to 80.2 % against 87.0 %, and the shaft reaches 25.50 MW against 24.03 MW. With the approach floor at 10^2 or the power row at 10^4 the peak falls to 2.11 or 2.09 MW, but neither variant ships, because repricing after measuring would be tuning the case.

The true minimum approach is under 1.5 °C for 41.5 against 61.5 min, crossing at the cold end to −0.37 °C against −0.16 °C near 30 min, and the plant makes 455.1 against 450.5 t of LNG at 0.320 against 0.329 kWh/kg. The LNG IAE is nearly equal, 274.0 against 275.3 °C min, but 223.9 against 121.9 in the afternoon and 48.8 against 120.8 after the cold front, and the MPC does not hold the setpoint. Its target sits more than 0.5 °C warm for 166.5 of the 400 min, as warm as −151.00 °C, because near the selector's limit a tenth of a megawatt or a shortfall of predicted approach is worth more to it than a degree of LNG.

The production gain is the model's blend and not the plan, because the cascade given only the MPC's blend target makes the same 455.1 t, with 49.5 min over the driver at a peak of 1.30 MW and 21.9 MW min. What the plan adds is time, 18.6 against 49.5 min over the driver and 41.5 against 45.4 min under 1.5 °C, bought with the deeper peak. The hindsight cascade beats the MPC on everything but minutes over the driver, with 457.5 t, 12.3 min under 1.5 °C and 53.3 min over at a peak of 0.37 MW and 2.4 MW min. The ranking holds at half the integration step, 17.8 min over at the same 2.78 MW peak and 20.4 MW min, and at three seeds with instrument noise, where the MPC peaks at 2.65 to 2.92 MW over the driver against 1.28 to 1.34 MW.

Scenario 8 is the loss the design accepts. With the feed on the operator only the power row protects the driver, and the MPC holds the LNG at its ceiling, −151.01 °C on average from 40 to 170 min, an IAE of 602.0 against 20.2. Both structures spend about 150 min over the driver, but the MPC is more than 0.1 MW over for only 23.9 of them, at a peak of 2.10 against 4.43 MW and 25.5 against 620.5 MW min, and handed no power limit it tracks the LNG with an IAE of 4.1, so the whole temperature loss is the price of the power row. In scenario 10 the MPC matches the circulation to the reduced feed, 21.0 kg/s against 29.2 kg/s just before the feed returns unannounced at 140 min, so the bundle is then short of refrigerant, the LNG warms to −119.3 °C against −139.3 °C, and the

MPC ends 18.1 against 6.0 min off specification and 14.5 t behind. The cascade forced to the same 21.0 kg/s does worse, 99.1 min off specification, so the loss is the circulation chosen by a controller that sees the feed only through its load estimate, and its surge margin falls to 15.4 against 25.9 %, above the 10 % alarm. In scenario 7 the stale prediction costs 2.8 against 0.2 min under 1.5 °C.

DID YOU KNOW? Less Refrigerant, Nearly Twice the Energy per Tonne

From the simulator (Table 14.8). In scenario 10 a cascade made to circulate only 21.0 kg/s through the feed cut ended at 0.656 kWh/kg, against 0.337 kWh/kg for the cascade left alone. Less refrigerant drew less power, and it made 118.7 t of on-specification LNG against 230.4 t.

With the challenge's suggested settings and the MPC selected the target is met, at a worst error of 1.69 °C against the cascade's 0.49 °C, but limits the rule does not score are broken on the way. The warm target keeps the product off specification for 17.2 min where the cascade has none, the true approach is under 1.5 °C for 77.3 against 177.2 min, and at the instant of the relief event at 130 min the cold end crosses to −4.32 °C under the MPC and −5.20 °C under the cascade. Half a minute later the MPC's model, on the analysis taken before the event, still put the cold end at 1.92 °C against a true −3.19 °C. With every other setting as shipped the MPC fails, at 7.05 °C with 150.0 min over the driver at a peak of 3.03 MW and 317.9 MW min.

14.7 Technology in Practice

Figure 14.7 draws the five loops of Section 14.5 and the anti-surge loop with the hardware of a liquefaction train in place of the schematic's symbols.

What Measures It

The loops read four instruments and one analyzer. A thermocouple in the outlet nozzle gives the product temperature, a transmitter the suction pressure, a flowmeter the circulation and a transducer the shaft power. The gas chromatograph on the suction line reports the blend, and the cold-end approach that the approach controller holds is computed from that analysis, the suction pressure and the product temperature. The other three stations of the bundle have no instrument at all.

TECHNOLOGY Sensor: The Thermocouple in the LNG Outlet

A thermocouple is a junction of two unlike metals whose voltage depends on its temperature. It is read against a second, reference junction, so it measures a difference and the instrument must compensate for the reference [23]. Thermocouples are more rugged than resistance thermometers and span more temperature, though usually less accurate. Copper against constantan, type T, covers −300 to 700 degrees Fahrenheit and resists both oxidizing and reducing gases well at low temperatures, even wet. The most common failure is to fail open, called burnout. An open input picks up noise that reads as a wildly varying temperature, so the instrument is set to drive high or low on burnout, whichever is safer [23]. The chapter gives the product thermocouple a 0.35 min filter.

What Moves It

Four final elements and one slow system carry the handles. The Joule-Thomson valve sets the circulation, the compressor speed sets the suction pressure, the feed valve sets the heat load under the low selector, and the recycle valve keeps the machine away from surge. The makeup system moves the blend, at no more than 0.010 of light fraction per minute. Behind the speed stands the driver, whose rating falls as the air warms, and that rating is the limit the selector defends.

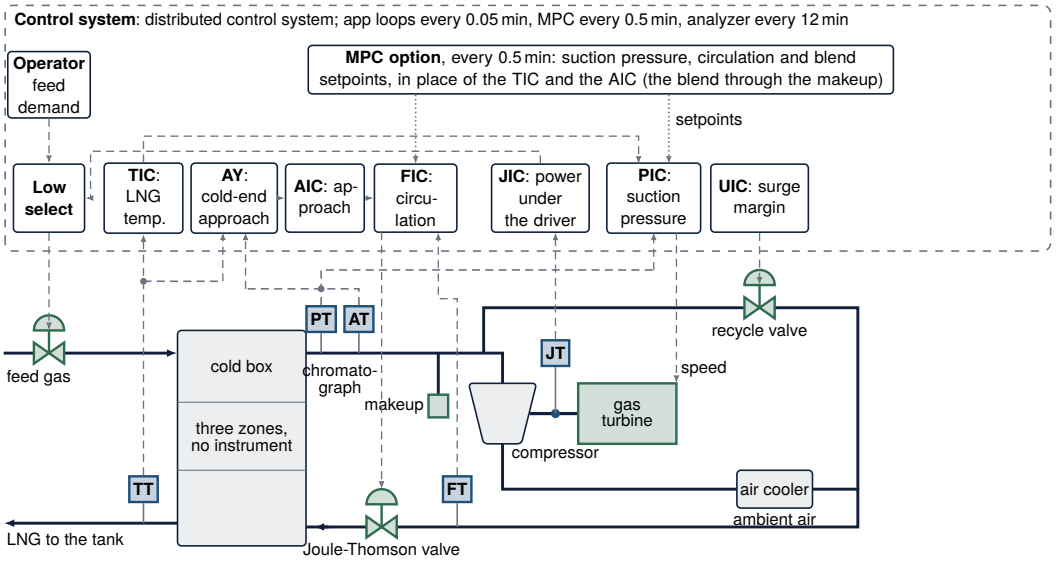


Figure 14.7 The loops of the train with their hardware. Lettered boxes are sensing devices with the chapter's tags: TT on the LNG outlet, PT and AT on the suction line, FT on the circulation and JT on the shaft. Green symbols are final elements: the Joule-Thomson valve, the gas turbine that sets the compressor speed, the feed valve, the recycle valve and the makeup of the blend. Dashed arrows are signals, dotted arrows the three setpoints that the MPC option writes in place of the temperature and approach controllers. Notice that the approach controller reads an analyzer that reports every 12 min, and that no instrument stands inside the cold box. Not to scale, and not a piping diagram.

> **TECHNOLOGY** **Final Element: The Gas Turbine on the Compressor Shaft**
>
> In oil and gas service a gas turbine often turns a compressor directly, with no generator between. Its own compressor squeezes the inlet air, fuel heats it, and the expansion turbine drives both that compressor and the load [24]. Warm air is less dense, so less of it passes, and it takes more work to compress, so power and efficiency both fall as the inlet warms. Ratings are quoted at sea level and 59 degrees Fahrenheit, and inlet-air cooling by refrigeration, a water spray or stored cold recovers part of the loss. At part load the turbine inlet temperature is lowered, which costs efficiency too [24]. The chapter keeps a speed lag of 0.30 min and a rating that falls with the air, from 26.0 to 22.7 MW between 18 and 36 °C.

Where the Controller Runs

A distributed control system controls production within one location, and its supervisory and lower-level loops run with cycle times from milliseconds to minutes [25]. Control of a process plant is usually layered by time scale, with supervisory and predictive control acting in minutes and regulatory control in seconds [3]. The app's loops execute every 0.05 min, its MPC every 0.5 min, and its analyzer reports every 12 min, so the chapter spans all three of those speeds. Its structure maps onto standard blocks: a temperature PI cascaded onto a pressure PI on speed, a flow PI on the valve under an approach PI that reads a calculation block, a power PI and the operator's demand into a low selector with tracking, and an anti-surge PI on the recycle valve.

> **IN PRACTICE What the Federal LNG Rules Ask of the Controls**
>
> United States rules for LNG facilities require an attended control center from which operations and warning devices are watched, and from which every remotely actuated and automatic shutdown control can be worked [26]. Operations are watched for gas, temperature, pressure, vacuum and flow alarms [27]. Each control system must be adjusted to operate within design limits, one out of service for 30 days or more is tested before it returns, and those normally in operation in a base load plant are inspected and tested every calendar year, at intervals of 15 months at most [28]. Cooldown must keep thermal stresses within design limits, and the cold piping is then checked for leaks [29]. The chapter starts every scenario cold.

What Is Modernizing It

Three developments touch this chapter. The first is the drive. Liquefaction uses a share of the feed itself, and the one United States export plant whose refrigerant compressors run on electric motors converts most of its feed gas into LNG [10]. New plants are increasingly driven by electricity under a power management system, and a modern drive can cut the power to the process almost instantaneously [30]. The second is the estimate. Soft sensors are used beside online analyzers, among other things to validate them, and a survey of industrial applications names model maintenance as their most critical issue [31]. The moving horizon estimator of Section 14.9 is the chapter's version of that idea. The third is the predictive layer: receding-horizon control is established in industry [20] and has been proposed in adaptive form for natural gas liquefaction [21], and scenario 12 shows that the trade it makes is set by its prices.

14.8 Limitations

Sizing an exchanger or a machine from this model is not a legitimate use of it, because the model is a control study and not a process simulator. The compositions are not a property package, the refrigerant is two numbers rather than four components with a phase envelope, and the boiling temperature relation of (14.9) is a linear fit that is credible only over the range the sliders allow. Three lumped zones cannot put the pinch anywhere except at a zone boundary or at an end, which is the simplification that most limits what the application can say about a real coil-wound bundle, where the tightest point usually sits inside a zone. The zone heat capacities are constants fixed by the design temperature ranges, so far from design the duty split between zones stops being physical. The compressor is a single line with a surge line and no choke, and the surge reversal itself is not simulated, so scenario 11 reports the approach to the limit and not what happens after it. Volume 4, Chapter 2 simulates the surge cycle itself, with a protective override on a recycle valve. The gas turbine is a rating with a linear derate rather than a thermodynamic model, which is adequate for showing two constraints crossing and inadequate for predicting the crossing point on a particular machine. Cool-down from ambient is absent, and the sixteen to thirty hours it takes is the operation that decides whether a train is shut down at all on low feed. The analyzer is a perfect sample and hold, so the failure modes that cause most real inferential errors, a plugged sample line or a calibration that has drifted, cannot be demonstrated. Perfect repeatability is itself a departure from practice, since one frozen step and one frozen scan hand back the same trajectory everywhere while a distributed control system never does. The features that survived the simplification, the two-effect valve, the fast suction pressure inventory, the metal inertia, the interior maximum in the blend, the sampled composition and the two constraints that move together with the weather, all survived because a control engineer has to design around each of them. Section 14.3.9 states where each individual simplification stops being true.

The predictive controller inherits the inferential problem rather than solving it. Its approach floor is on predictions from the last analysis, wrong by up to 5.37 °C in scenario 7, and the effect of the blend on the boiling temperature of (14.9) is taken from that analysis carried forward by the makeup law, with nothing inferring the blend between analyses. Its model holds the declared parameters, so with the bundle fouled to 0.85 its warm-end approach is off by up to 1.78 °C. It has no feed measurement

and no model of the selector, its power and surge limits are soft with one slack per window, and its prices rather than its setpoints decide what it gives up.

> **OPEN CHALLENGE** **Running On the Power That Is Left**
>
> New LNG plants are increasingly driven by electric motors under a site power management system. A 2022 steady-state study found that a mixed fluid cascade plant could keep running on about 30 % less power, 60 of 200 MW, instead of shutting down. Its model left the transient open: the LNG could warm toward saturation and the refrigerant compressors move toward surge while the plant moves to its new point. Progress, as the authors set it out, means tests on an industrial simulator and in plant experiments, and a study of the dynamics and of the need for more complex control structures [30]. The chapter's low selector meets a slower form of the limit in scenario 9.

14.9 Discussion

The economic argument on a liquefaction train is a ratio and not a rate. At the design point the plant turns 20.0 kg/s of feed into 72.0 t/h of LNG for 23.10 MW of shaft power, which is 0.321 kWh/kg, and that ratio is what a percent of efficiency is measured against. Scenario 10 moves it to 0.362 kWh/kg without a single loop leaving its setpoint, and scenario 7 moves it from 23.1 to 23.7 MW for the same product while the approach display reads better than the truth. Both of those are invisible on the controlled variables and both are visible in the ratio, which is the argument for putting specific power on the wall of a control room rather than production. The same reasoning is why the LNG operation literature evaluates candidate controlled variables by the steady-state loss they cause rather than by their closed-loop response [3–5].

Methods beyond the implemented controllers fit the same equations. Sixteen states with an active driver constraint and a bounded approach make a natural nonlinear model predictive control problem, because the controller can trade the approach against the driver limit directly instead of relying on a selector to notice the collision, and the successively linearized controller of Section 14.5.5 shows in scenario 12 that the trade it makes is set by its prices. Moving horizon estimation on the same equations, using the sampled composition as an infrequent measurement alongside the continuous temperatures, is the standard remedy for the error that scenario 7 puts on the screen, and it is the direction the soft-sensor literature takes for exactly this situation [1, 2]. An economic formulation would place the value of on-specification LNG against the cost of shaft power, both of which the application already carries as adjustable parameters, and would decide the feed rate and the approach setpoint together rather than separately.

14.10 Conclusion

This chapter studied a single mixed refrigerant liquefaction train as an interactive laboratory for inferential constraint control. The simulator integrates sixteen states covering three counter-current effectiveness-NTU zones, the bundle metal as a lag of 4.44, 4.55 and 6.65 min on the three zone duties, a self-regulating suction pressure inventory, four actuator positions and the refrigerant light fraction, and it reports the true minimum internal temperature approach beside the estimate a control room would compute from an analysis that is refreshed every 12 min. Swept open loop with the compressor speed held, the product temperature reaches a minimum of −159.26 °C at 60 % valve travel, giving 3.4 °C per percent in one direction, 0.27 °C per percent in the other, and a sign change in between. Closing a suction pressure loop removes one of the valve's two effects and leaves a handle worth 9.2 °C per bar that is straight across the working range, with the pressure reaching 63 % of a setpoint move in 2.2 min against 8.3 min for the product temperature. Circulation is a capacity variable whose gain falls by a factor of 16 between 30 and 39 kg/s, and the refrigerant blend is an interior maximum at $x = 0.419$ rather than a setpoint. A relief event that removes five points of light ends sends the computed cold-end approach up to 12.7 °C while the true minimum falls to 4.2 °C, and the plant recovers on 23.7 MW instead of 23.1 MW for the same product. An optional model

predictive controller spends 18.6 rather than 85.1 min over the derated driver in scenario 12 but goes further over it, 2.78 MW at its peak against 1.31 MW with 20.4 against 28.9 MW min, owes its 4.6 t of extra LNG to the blend its model chooses, and loses scenario 10 with 18.1 against 6.0 min off specification. A fuller model would add a continuous bundle discretization so the pinch can sit inside a zone, a component-wise makeup system, an analyzer failure mode, and a moving horizon estimator to be compared against the sample-and-hold inference. The liquefaction train puts a plant whose limiting variable is computed rather than measured within reach of anyone with a browser.

KEY POINTS

- One valve sets both how much refrigerant circulates and how warm it boils, and the two effects pull the LNG temperature opposite ways.
- A suction pressure loop takes one effect away and leaves a handle worth 9.2 °C per bar, straight across the working range.
- Circulation is a capacity variable, and the blend is an interior maximum rather than a setpoint.
- A relief event sends the computed cold-end approach up to 12.7 °C while the true minimum falls to 4.2 °C.

Where else this applies. A valve whose gain changes sign returns in Volume 6, Chapter 2, and quality held through a late analysis in Volume 5, Chapter 3.

References

[1] 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.

[2] P. Kadlec, B. Gabrys, and S. Strandt, “Data-driven soft sensors in the process industry,” *Computers & Chemical Engineering*, vol. 33, no. 4, pp. 795–814, 2009.

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

[4] J. B. Jensen and S. Skogestad, “Optimal operation of a simple LNG process,” *IFAC Proceedings Volumes*, vol. 39, no. 2, pp. 241–246, 2006.

[5] F. A. Michelsen, B. F. Lund, and I. J. Halvorsen, “Selection of optimal, controlled variables for the TEALARC LNG process,” *Industrial & Engineering Chemistry Research*, vol. 49, no. 18, pp. 8624–8632, 2010.

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

[7] B. Austbø, S. W. Løvseth, and T. Gundersen, “Annotated bibliography—use of optimization in LNG process design and operation,” *Computers & Chemical Engineering*, vol. 71, pp. 391–414, 2014.

[8] W. Lim, K. Choi, and I. Moon, “Current status and perspectives of liquefied natural gas (LNG) plant design,” *Industrial & Engineering Chemistry Research*, vol. 52, no. 9, pp. 3065–3088, 2013.

[9] S. Kumar, H.-T. Kwon, K.-H. Choi, W. Lim, J. H. Cho, K. Tak, and I. Moon, “LNG: An eco-friendly cryogenic fuel for sustainable development,” *Applied Energy*, vol. 88, no. 12, pp. 4264–4273, 2011.

[10] U.S. Energy Information Administration, “Liquefied natural gas,” Natural Gas Explained, last updated June 21, 2024, 2024, accessed November 18, 2026. [Online]. Available: https://www.eia.gov/energyexplained/natural-gas/liquefied-natural-gas.php

[11] G. Venkatarathnam, “Natural gas liquefaction processes,” in *Cryogenic Mixed Refrigerant Processes*, ser. International Cryogenics Monograph Series. New York, NY: Springer, 2008, pp. 149–220.

[12] W. J. Podbielniak, “Art of refrigeration,” U.S. Patent 2,041,725, filed July 14, 1934, granted May 26, 1936, 1936. [Online]. Available: https://patents.google.com/patent/US2041725A/en

[13] J. C. Perret, “Method and apparatus for the cooling and low temperature liquefaction of gaseous mixtures,” U.S. Patent 3,364,685, filed March 16, 1966, granted January 23, 1968, 1968. [Online]. Available: https://patents.google.com/patent/US3364685A/en

[14] L. S. Gaumer, Jr. and C. L. Newton, “Combined cascade and multicomponent refrigeration system and method,” U.S. Patent 3,763,658, filed January 12, 1970, granted October 9, 1973, 1973. [Online]. Available: https://patents.google.com/patent/US3763658A/en

[15] B. Linnhoff and E. Hindmarsh, "The pinch design method for heat exchanger networks," *Chemical Engineering Science*, vol. 38, no. 5, pp. 745–763, 1983.

[16] G. C. Lee, R. Smith, and X. X. Zhu, "Optimal synthesis of mixed-refrigerant systems for low-temperature processes," *Industrial & Engineering Chemistry Research*, vol. 41, no. 20, pp. 5016–5028, 2002.

[17] A. Singh and M. Hovd, "Dynamic modeling and control structure design for a liquefied natural gas process," in *2007 American Control Conference*, 2007, pp. 1347–1352.

[18] S. Huang, M. Chapeaux, and B. Martinez, "Strategies for mitigating impacts on LNG plant capacity by variations in ambient temperature," in *Proceedings of the 2nd Annual Gas Processing Symposium*. Elsevier, 2010, pp. 455–463.

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

[20] 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.

[21] W. Won and K. S. Lee, "Delta-operator-based adaptive model predictive control and online optimization of a natural gas liquefaction process," *Chemical Engineering Science*, vol. 162, pp. 21–32, 2017.

[22] K. R. Muske and T. A. Badgwell, "Disturbance modeling for offset-free linear model predictive control," *Journal of Process Control*, vol. 12, no. 5, pp. 617–632, 2002.

[23] 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.

[24] U.S. Environmental Protection Agency Combined Heat and Power Partnership, "Catalog of CHP technologies, section 3. technology characterization: Combustion turbines," U.S. Environmental Protection Agency, Washington, DC, Tech. Rep., 2017, march 2015, revised September 2017. [Online]. Available: https://www.epa.gov/sites/default/files/2015-07/documents/catalog_of_chp_technologies_section_3._technology_characterization_-_combustion_turbines.pdf

[25] 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.

[26] Pipeline and Hazardous Materials Safety Administration, "49 CFR 193.2441, liquefied natural gas facilities: Control center," Code of Federal Regulations, Title 49, Part 193, edition of October 1, 2025, govinfo.gov, 2025.

[27] ——, "49 CFR 193.2507, liquefied natural gas facilities: Monitoring operations," Code of Federal Regulations, Title 49, Part 193, edition of October 1, 2025, govinfo.gov, 2025.

[28] ——, "49 CFR 193.2619, liquefied natural gas facilities: Control systems," Code of Federal Regulations, Title 49, Part 193, edition of October 1, 2025, govinfo.gov, 2025.

[29] ——, "49 CFR 193.2505, liquefied natural gas facilities: Cooldown," Code of Federal Regulations, Title 49, Part 193, edition of October 1, 2025, govinfo.gov, 2025.

[30] I. Y. Xuan, C. Skourup, J. B. Jensen, T. Haugen, and N. F. Thornhill, "Flexible operation of a mixed fluid cascade LNG plant for electrical power management," *Energy*, vol. 250, p. 123794, 2022.

[31] M. Kano and K. Fujiwara, "Virtual sensing technology in process industries: Trends and challenges revealed by recent industrial applications," *Journal of Chemical Engineering of Japan*, vol. 46, no. 1, pp. 1–17, 2013.

14.11 Exercises

14.1 ○ Concept check · 2 min

Scenario 1 starts in Manual at the design line-up. Five minutes after Run the Joule-Thomson valve is opened from 54.9 to 62 % with the compressor speed left at 92.0 %. Twenty minutes later the circulation has risen by 4.6 kg/s, yet the product is only 4.51 °C colder and sits on the refrigerant boiling temperature. Why?

(a) The flow loop holds the circulation at its setpoint, so most of the extra opening is lost

(b) The extra refrigerant reaches the product through the bundle lag, so most of the gain is still to come

(c) The extra vapor raises the suction pressure, so the refrigerant boils warmer and the product rests on it

(d) The air cooler warms with the extra flow, so the discharge pressure falls and the valve passes less

14.2 Concept check · 2 min

In scenario 7 a relief event at 15 min vents five points of light ends. At 24 min, just before the next analysis lands, the cold-end approach the control room computes reads 12.7 °C, wider than its 7 °C setpoint, while the true minimum approach has fallen to 4.2 °C. What makes the computed approach move the wrong way?

(a) The product thermocouple lags by 0.35 min, so the computed approach trails the true one

(b) The approach controller has raised the circulation, and the extra refrigerant has opened the cold end

(c) The suction transmitter reads high after the event, so the estimate puts the boiling point too low

(d) The estimate subtracts a boiling temperature from the last analysis, which predates the vented light ends

14.3 Concept check · 2 min

At a suction pressure held at 2.90 bar, circulation steps of 3 kg/s from 30 to 33, 33 to 36 and 36 to 39 kg/s are worth 15.8, 5.9 and 1.0 °C of product temperature. Which role does this chapter give the circulation?

(a) A capacity setting that carries the duty, with the suction pressure as the temperature trim

(b) The temperature handle, with a gain schedule that spans the factor of 16 in its gain

(c) A ratio to the feed rate with no feedback, since the duty follows the feed

(d) The surge protection handle, held at the flow that gives the widest surge margin

14.4 Theory and calculation · 5 min

Equation (14.9) gives the refrigerant boiling temperature with $T_{evap,0} = -162.0\,°\mathrm{C}$, $k_P = 9.5\,°\mathrm{C}$ per bar, $k_x = 120\,°\mathrm{C}$ per unit of light fraction, $P_{s,0} = 2.90\,\mathrm{bar}$ and $x_0 = 0.420$.

(a) In scenario 1 the valve step raises the suction pressure from 2.901 to 3.167 bar at the design blend. Compute the boiling temperature before and after. *Checkpoint:* −161.99 °C and −159.46 °C.

(b) A relief event takes 0.05 of light fraction out of the blend at 2.90 bar. Compute the new boiling temperature. *Checkpoint:* −156.0 °C, 6.0 °C warmer.

(c) Find the suction pressure that brings the boiling temperature back to −162.0 °C at $x = 0.370$, and say whether the temperature controller's output limits of Table 14.4 allow it. *Checkpoint:* 2.27 bar, inside the 1.9 to 5.2 bar range.

14.5 Theory and calculation · 6 min

Zone 1 (precool) runs at the design point with $\dot{m}_{NG} = 20.0\,\mathrm{kg/s}$, $\dot{m}_{MR} = 33.0\,\mathrm{kg/s}$, $c_{g,1} = 2.60\,\mathrm{kJ/kg/°C}$, $c_{r,1} = c_{m,1} = 1.576\,\mathrm{kJ/kg/°C}$ at $x = x_0$, $UA_1 = 676\,\mathrm{kW/°C}$, $C_{w,1} = 180000\,\mathrm{kJ/°C}$, $T_{feed} = 30\,°\mathrm{C}$ and a refrigerant inlet $T_{r,2} = -40.0\,°\mathrm{C}$.

(a) Compute the two heat-capacity rates and their ratio $C^*_{r,1}$, and say which branch of Equation (14.3) applies. *Checkpoint:* 52.0 and 52.0 kW/°C, $C^* = 0.9998$, the balanced branch.

(b) Compute NTU_1, ε_1 and the duty of Equation (14.4), then the refrigerant leaving the zone and the warm-end approach. *Checkpoint:* 13.0, 0.929, 3380 kW, then 25.0 °C and 5.0 °C.

(c) Compute the metal lag τ_1 of Equation (14.5) at $f_b = 1$ and for the challenge's bundle at $f_b = 0.9$. *Checkpoint:* 4.44 min and 4.93 min.

14.6 Theory and calculation · 7 min

The ambient air steps from 18 to 36 °C, as in scenario 8. Use Equations (14.14), (14.15) and (14.13) with the constants of Table 14.3, $f_{ac} = 1$ and the compressor flow held at 33.0 kg/s.

(a) Compute the derated driver rating W_{max} at 36 °C. *Checkpoint:* 22.72 MW.

(b) Compute the steady condensing temperature and the discharge pressure at 18 and at 36 °C. *Checkpoint:* 30.16 °C and 46.13 bar, then 48.16 °C and 63.78 bar.

(c) With the suction pressure at 2.90 bar and the suction temperature unchanged, find the factor by which the shaft power rises, and the shaft power that follows from the design 23.10 MW. *Checkpoint:* 1.164, 26.9 MW against 22.7 MW available.

14.7 App challenge · 8 min

Load scenario 1, Two handles by hand (Manual), default seed, at the default speed of 2. Press Run, pause at any time from 5 to 7 min, set the JT valve to 62 %, and run on until 20 min after the move. Export the CSV file.

(a) In the row 20 min after the move, read the circulation `mr_kgs`, the suction pressure `psuc_bar`, the boiling temperature `tev_C`, the product `lng_C` and the cold-end approach `app_cold_C`. *Checkpoint:* 37.63 kg/s, 3.167 bar, −159.46 °C for both temperatures and 0.00 °C, each within 0.02.

(b) Press the reload button, run again, and in the same window set Compressor speed to 96.0 % instead, leaving the valve at 54.9 %. Read `psuc_bar` and `lng_C` 20 min after the move, and compute the product change per bar of suction change from the first row (2.901 bar, −154.95 °C). *Checkpoint:* 2.735 bar and −156.60 °C, about 10 °C per bar.

(c) Compare the warm-end approach `app_warm_C` 20 min after each move with its starting 4.99 °C. *Checkpoint:* 17.76 °C after the valve move and 5.18 °C after the speed move, within 0.1 °C.

14.8 App challenge · 8 min

Load scenario 7, Nitrogen to the flare, and a stale analysis, and change no setting. Run to 60 min at the default speed with the default seed and export the CSV file.

(a) In the row at 24.0 min, read `tev_C`, `lng_C`, the computed cold-end approach `appr_meas_C`, `mita_est_C`, `mita_C`, `x_light` and `x_gc`. *Checkpoint:* −160.33 °C, −151.89 °C, 12.67 °C, 6.52 °C, 4.22 °C, 0.385 and 0.420, temperatures within 0.02 °C.

(b) In the same row, compare $k_x(x_{GC} - x)$ with the difference between `appr_meas_C` and the true cold-end approach `app_cold_C`. *Checkpoint:* 4.19 °C against $12.67 - 8.44 = 4.23$ °C.

(c) Read `appr_meas_C` and `x_gc` in the row at 24.1 min, and the shaft power `power_kW` at 15 and 60 min. *Checkpoint:* 8.49 °C and 0.385, then 23.10 MW and 23.72 MW.

14.9 App challenge · 10 min

Load scenario 12, Three handles under one driver limit, which opens with Control structure on MPC and Throughput on the driver. Run 400 min at speed 20 with the default seed and export the CSV file. Press Reset, set Control structure to Cascade before Run, run the same 400 min and export again.

(a) For each file, read `metric_overpower_min` at 400 min. Then take the excess of `power_kW` over `driver_kW` in every row, find its largest value, and sum its positive values times the 0.1 min row spacing. *Checkpoint:* MPC 18.6 min, 2.78 MW, 20.4 MW min, and cascade 85.1 min, 1.31 MW, 28.9 MW min, sums within 0.1 MW min.

(b) Read `metric_pinch_min`, the minutes with the true minimum approach under 1.5 °C, and `metric_lng_t` at 400 min. *Checkpoint:* 41.5 min and 455.1 t under the MPC, 61.5 min and 450.5 t under the cascade.

(c) The three driver measures rank the structures differently. Say why, and which one a plant manager should be shown.

14.10 Application area · 4 min

Table 14.3 prices on-specification LNG at 0.44 $/kg and shaft power at 0.030 $/kWh.

(a) At the design specific power of 0.321 kWh/kg, what does the shaft power cost per kilogram of LNG, and what share of the LNG's value is that? *Checkpoint:* 0.0096 $/kg, 2.2 %.

(b) In scenario 9 the low selector holds the feed on the derated driver. At the design specific power, how much LNG per hour does one megawatt of shaft power saved let the selector give back, what is it worth, and what does the megawatt cost in fuel? *Checkpoint:* 3.12 t/h, about 1370 $/h, against 30 $/h.

14.11 Application area · 3 min

The Limitations section notes that the model's analyzer cannot fail. Suppose the sample line plugs and the chromatograph keeps reporting its last good value, 0.420, while the inventory drifts to 0.400 light over a shift with the suction pressure at 2.90 bar. The approach controller holds the computed cold-end approach at 7 °C.

(a) By how much, and in which direction, is the computed approach wrong, and where does the true cold-end approach settle? *Checkpoint:* 2.4 °C too wide, so the true approach settles near 4.6 °C.

(b) Name one check that would expose the fault.

14.12 Review the agent · 5 min

An AI assistant reviews scenario 1 and writes: *Opening the JT valve from 54.9 to 62 % made the LNG 4.5 °C colder, a process gain of −0.63 °C per percent with a time constant of a few minutes. Pair the LNG temperature directly with the JT valve under a PI loop tuned from that step, leave the compressor at fixed speed, and drop the suction pressure loop as unnecessary.* Find the flaw and say how you would check it in the app.