INTERACTIVE CONTROL ENGINEERING
Principles and Applications

Volume 13 Control in Chemical Plants

CHAPTER 2

Ammonia Synthesis Loop Simulator for Recycle Amplification and Inventory Pairing

Pages 24–52

Edited by John D. Hedengren

How to cite this chapter. "Ammonia Synthesis Loop Simulator for Recycle Amplification and Inventory Pairing," in *Control in Chemical Plants*, John D. Hedengren, Ed., vol. 13 of *Interactive Control Engineering: Principles and Applications*. Advanced Process Solutions, LLC, 2026, ch. 2, pp. 24–52.

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

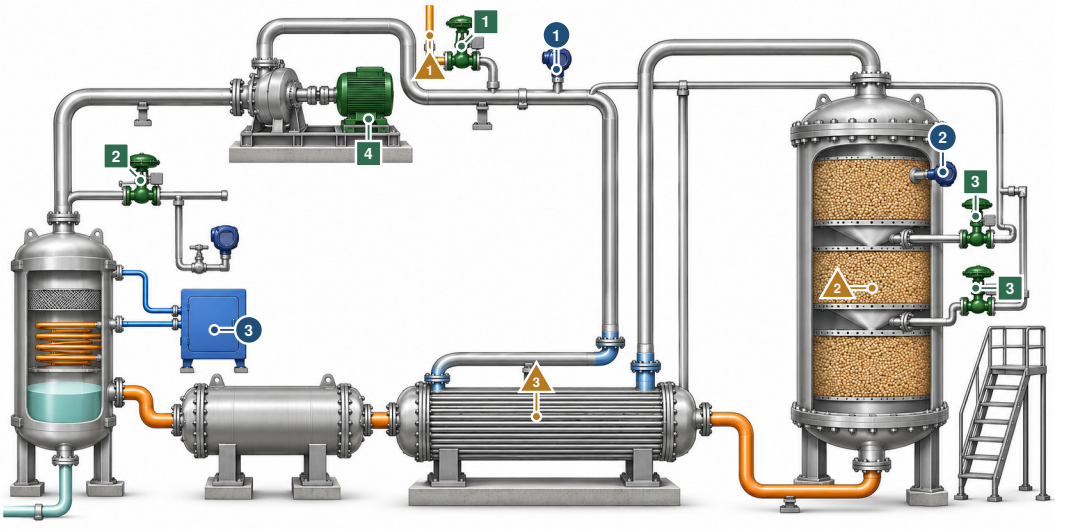
The app for this chapter
https://apmonitor.com/apps/ammonia-loop/

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

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

CHAPTER 2

Ammonia Synthesis Loop Simulator for Recycle Amplification and Inventory Pairing



The controller moves 1 makeup gas valve 2 purge valve 3 quench gas between the beds 4 circulator speed

The controller reads 1 loop pressure 2 hot spot in the beds

3 chromatograph: inerts and the hydrogen to nitrogen ratio, every quarter hour

What disturbs it 1 makeup gas that is dirtier, off ratio or scarcer 2 aging catalyst

3 an exchanger that loses effectiveness

Follow the gas around: it leaves the separator, passes the circulator, takes on makeup, is heated in the exchanger and reacts in the converter, so most of it goes around again and only the small purge lets the inerts out. Illustration, not to scale.

IN PLAIN TERMS **A Loop That Magnifies Small Mistakes**

The main idea. An ammonia plant turns only a little of its gas into ammonia on each pass through the reactor. So it sends the rest around again in a loop, and a small bleed lets out the gases that cannot react.

The challenge. A small error in the fresh gas grows much larger inside the loop, and it takes hours to arrive. The valve that tops up the loop can hold the pressure or the output steady, and the trouble then shows up in the other one.

Why it matters. No alarm sounds while the plant quietly makes less ammonia or uses up its safety margin.

THREE WAYS TO READ THIS CHAPTER

Explore	12 min	the first page; the history and fact boxes of Sections 2.1 and 2.2; the try-it boxes of Sections 2.4 and 2.6; the fact box of Section 2.6; the key points
Study	60 min	the whole chapter with its pause points and the exercises
Practice	20 min	the first page; Section 2.6 with its result chart; Section 2.7; Section 2.8 with its open challenge

The control question An ammonia converter reacts about four percent of the gas on each pass, so the plant is a recycle with a converter inside it. How far and how slowly does the loop amplify a small error in the makeup gas, and which variable should the makeup valve hold: the loop pressure or the production rate?

Learning objectives

- Compute the steady inert level and the composition time constant of the loop from its holdup, its makeup and its purge.
- Explain why an exothermic equilibrium converter has a best temperature and, because it heats its own feed, a second steady state in which it is out.
- Measure the amplification of a front-end ratio error and explain why it is larger for a hydrogen-rich error than for a nitrogen-rich one.
- Predict where a disturbance appears under each pairing of the makeup valve, and read constraint margin from a valve position.
- Compare two slow PI loops with one predictive plan on the purge and the trim, in integrated error and in margin.

Prerequisites Sections 1.4.2 (equilibrium-limited beds), 1.4.3 (inert balance, recycle gain, feed ratio error, throughput manipulator), 1.6.2 (velocity form) and 1.6.4 (scenarios and exported files). Section 1.6.3 is needed only for the predictive controller.

The app https://apmonitor.com/apps/ammonia-loop/

See also Volume 6, Chapter 7 for the snowball effect of a reactor and column recycle, and Volume 6, Chapter 8 for a gas recycle with an inert and a purge. Volume 7, Chapter 2 for ignition and extinction folds in an exothermic reactor. Volume 9, Chapter 4 for offset-free predictive control with a Kalman disturbance estimate.

2.1 Introduction

A recycle is normally drawn as an arrow that leaves the separator and returns to the front of the block diagram, and it is normally treated as bookkeeping. An ammonia synthesis loop makes that treatment untenable. The extent of reaction over the iron catalyst is 0.0386 per pass at the design point of the simulator, so 6.00 kmol/s of gas circulates against 1.01 kmol/s of makeup, a recycle ratio of 5.95 (circulator flow over makeup, as the application reports it), and the properties a student would predict from the converter alone are not the properties the plant has. The converter answers a temperature change in minutes and the loop answers a composition change in five hours. A front-end error of five percent in the hydrogen to nitrogen ratio arrives inside the loop as an error of forty percent. The misconception the simulator corrects is that a recycle plant is a unit operation with a return line, when in fact the return line sets the gains, the time constants and the place a disturbance finally appears.

DID YOU KNOW? World Ammonia Output, 2024

The world's ammonia plants produced an estimated 150 million metric tons in 2024, counted as the nitrogen the ammonia contains. In the United States 37 plants ran that year, natural gas was the dominant feedstock, and about 88 percent of the output went to fertilizer [1].

The shape of the problem is general and it does not require a catalyst. Whenever a component enters a recycle loop and can leave it only through a small purge, the steady-state loop concentration is the ratio of two small flows and the loop amplifies any persistent imbalance in proportion to the inverse of the purged fraction. Reactor and separator systems joined by a recycle carry that amplification as a documented structural property [2], the snowball effect of Volume 6, Chapter 7, and the loop inventory divided by the purge is the time constant that comes with it. A second and separate question follows from conservation of mass. One valve on a recycle loop sets throughput, and the plant designer decides which variable that valve is allowed to hold, which fixes where every disturbance lands and which measurements are informative [3, 4]. Both questions are the ordinary subject matter of plantwide control rather than special features of an ammonia plant [5], and

hydrotreaters, reforming units, polymerization loops and gas plants with recovery all present them in the same form.

Ammonia loop models in the literature answer process questions. The converter has been simulated at steady state and in transient since 1965 [6, 7], the loop variables have been mapped by steady-state simulation [8], and a full control structure has been designed for the process by self-optimizing methods [9]. However, a student who is handed the loop as a set of transfer functions never performs the experiment that shows a five percent front-end error becoming forty percent, because the amplification lives in a mass balance that a transfer function has already absorbed into a gain. This chapter works with an interactive ammonia synthesis loop whose per-pass conversion is small enough that the recycle governs the plant, so that recycle amplification and the choice of what the makeup valve holds become experiments rather than assertions. The simulator carries fourteen states, a three-bed quench converter solved against the Gillespie-Beattie equilibrium, a feed and effluent exchanger that gives the vessel two steady states, a refrigerated separator, a sampled loop analyzer, a low selector on the makeup valve and a linear model predictive controller for the two slow handles. Thirteen scenarios run the same plant through a converter optimum, an extinction, two inventory pairings, three slow composition disturbances and a day on which two PI loops and one plan meet the same pair of drifts. The plant runs in a browser and exports every variable it computes. The equilibrium correlation, the kinetic form and the plantwide arguments are standard and are cited to their sources. Chapter 1 treats the same balances on its running example in Sections 1.4.2 and 1.4.3.

Section 2.2 traces how the loop came to be controlled. Section 2.3 carries every equation and constant the plant is built from. Section 2.4 derives the steady inert level and its time constant, measures the amplification of a front-end ratio error and states the inventory question. Section 2.5 describes the four PI loops, the selector and the predictive controller. Section 2.6 reports the scenarios, from the converter outward, and compares the two slow loops with one plan. Section 2.7 names the hardware and the practice behind the model. Sections 2.8 and 2.9 state where the model stops holding and what it supports beyond the implemented controllers, and Section 2.10 collects the measured results.

2.2 The Synthesis Loop and How It Came To Be Controlled

Nitrogen fixation at industrial scale began with the high-pressure catalytic synthesis commercialized in 1913, and the process is credited with sustaining roughly half of the present world population [10, 11]. The reaction that made that possible is also the reason the plant is a loop. Ammonia synthesis is exothermic, so its equilibrium constant falls steeply with temperature, while the rate over promoted iron is negligible until roughly 400 °C. The catalyst therefore has to be run where the equilibrium is poor, and the equilibrium ammonia fraction at any temperature the catalyst tolerates is small enough that the unconverted gas has to be recovered and returned [12]. Recovery by refrigerated condensation and return by a circulator is the arrangement that every modern plant uses, and it is what turns a reactor design problem into an inventory problem.

> **HISTORY A Loop Under Constant Pressure, 1905 to 1920**
>
> In his Nobel lecture of 1920, Fritz Haber traced the loop to his own laboratory. In work set aside in 1905 he had found that at red heat no catalyst gives more than a trace of ammonia at normal pressure, and had described a circulation that brought the gas to the hot metal and then washed the ammonia out, converting the gas stage by stage. Taking the problem up again in 1908 with Robert le Rossignol, he worked at 200 atmospheres, the limit of his compressor. Their small plant made and removed ammonia in one cycle at constant high pressure and led the heat of the outgoing gas to the fresh gas, so the process supplied its own heat [13]. Industry, he said, kept both features: the circulator and the feed and effluent exchanger of Figure 2.1.

The engineering attention followed the vessel first. Converter modeling was posed as a steady-state simulation of an autothermal reactor whose feed is preheated by its own effluent, and the transient behavior of the same reactor was analyzed in the same year [6, 7]. The structural fact those studies inherited had been established a decade earlier. An autothermal process whose feed is heated by its product has a heat generation curve and a heat removal line that can intersect more than once, so it possesses an ignited state, an extinguished state and an ignition condition between them [14]. The cooled stirred tank of Volume 7, Chapter 2 has the same ignition and extinction folds, and the adiabatic beds of Chapter 4 have an ignition floor of the same origin. The practical consequence is that a converter which has gone out does not come back when the operator returns the quench valve to where it was, and that the same feedback can produce sustained oscillation on an industrial machine rather than a clean approach to either branch [15].

Pause. Try Exercise 2.2 before reading on.

Loop questions were answered later and mostly by simulation. The purge rate, the inert level, the loop pressure and the hydrogen to nitrogen ratio were mapped against each other by steady-state calculation, which established the now familiar picture of a purge that buys conversion with hydrogen [8]. Before an on-line analyzer existed, the loop composition was a laboratory result rather than a measurement, the purge was a valve set by a standing instruction and revised on a shift report, and the converter was managed by the temperature profile because the temperature profile was the only thing measured continuously. Gas chromatography on the loop replaced the standing instruction with a controller, and a controller pointed at a five-hour plant is a different design problem from the two temperature loops next to it.

> **HISTORY The Leaks Were the Purge, 1932**
>
> Carl Bosch's Nobel lecture of 1932 shows how the high-pressure plants he built watched their loops. An automatic recorder of the gas density followed the composition of the circulating gas, and it also controlled the proportion of hydrogen in the mixture. Of the thousands of flanges and valves he wrote that everything had to be gas-tight, and that the measures had been so successful that the gas losses were just enough to keep the argon from the air and the accumulating methane from rising too high in the circulating system [16]. The leaks did the work of a purge. In the chapter's plant that work belongs to valve FV-102, watched by chromatograph GC-101.

The current arrangement became standard when plantwide control was posed as a subject in its own right rather than as a collection of loops. The recycle amplification of a persistent imbalance was named and quantified for reactor and separator systems [2], the choice of which valve sets throughput was identified as the decision that fixes the rest of the control structure [3], and the requirement that every inventory in a plant be held by a flow on one side or the other of it was stated as a consistency condition [4, 5]. The Tennessee Eastman process of Volume 6, Chapter 8, with its gas recycle, its inert and its purge, became the common test of those ideas. Applied to this process the result is four loops, a pressure controller and a converter temperature controller that answer in minutes, and an inert controller and a synthesis gas ratio controller that answer in hours, with the throughput manipulator

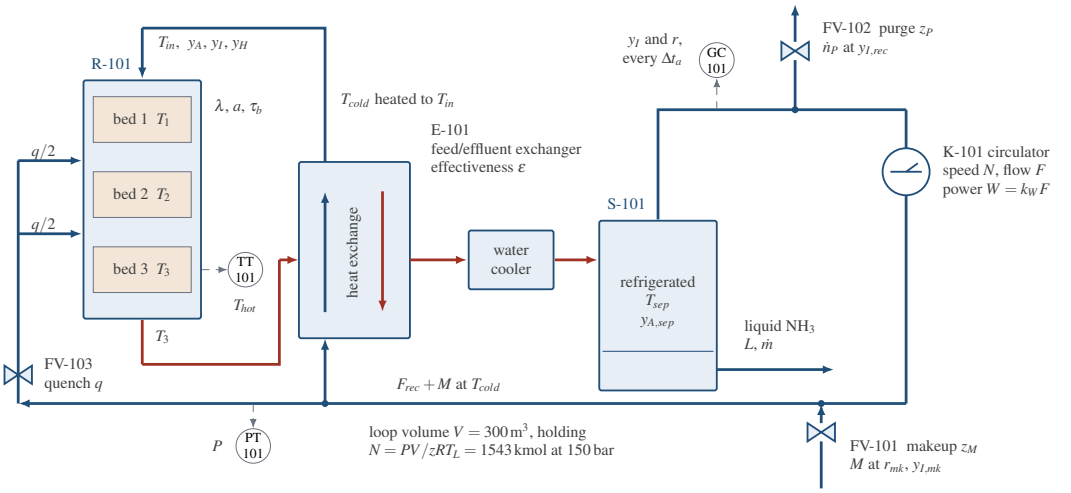


Figure 2.1 Synthesis loop L-101 with its three-bed quench converter, the feed and effluent exchanger that preheats the converter charge with the converter product, the refrigerated separator and the circulator. Gas at T_{cold} leaves the circulator, is heated to T_{in} in E-101 by the bed three effluent at T_3, and passes down through beds at T_1, T_2 and T_3 with a fraction q of the cold gas split evenly between the two quench points. Cooled effluent enters S-101, where liquid ammonia L is condensed at T_{sep} and leaves as product $\dot{m}$. The overhead gas splits into purge $\dot{n}_P$ through FV-102 and recycle F_{rec} back to the circulator, where makeup M at ratio r_{mk} and inert fraction $y_{I,mk}$ joins it through FV-101. Composition y_A, y_I and y_H is carried at the converter inlet mixing point, and the loop holds N kilomoles of gas in volume V at pressure P. PT-101 and TT-101 are continuous, and chromatograph GC-101 reports the loop inert fraction and the hydrogen to nitrogen ratio once every Δt_a.

chosen deliberately rather than by default [9]. Interest in ammonia as an energy carrier rather than only as a fertilizer intermediate has renewed attention to the efficiency of exactly this loop [17], and the efficiency number the loop owns is the syngas consumed per tonne of product, 128.6 kmol/t at the design purge here against a stoichiometric 117.4 kmol/t.

2.3 Process Model

Loop L-101 circulates 6.00 kmol/s of synthesis gas at 150 bar through a three-bed quench converter R-101, a feed and effluent exchanger E-101, a water cooler, and a refrigerated separator S-101 that condenses the ammonia out at –4.7 °C. Makeup syngas enters through valve FV-101 on the circulator discharge, and the purge leaves through FV-102 on the separator overhead. Quench valve FV-103 takes cold gas around the exchanger and injects half of it between beds one and two and half between beds two and three. The circuit, its instrument tags and every symbol carried by the equations below are drawn in Figure 2.1.

Table 2.1 lists everything the control system touches. Pressure is in bar, temperature in °C, flow in kmol/s, valve position in percent, and the clock is in minutes, so every rate below is per minute and every flow that enters one is converted with the factor of 60 written out. Fourteen quantities are integrated. Eight of them are the plant, and six are first-order actuator positions.

2.3.1 Loop Inventory and Pressure

The gas held in the loop is an ideal inventory with a compressibility correction at a fixed mean temperature, so pressure and moles are the same variable, as shown in Equation (2.1).

$$N = \frac{PV}{zRT_L}, \qquad \frac{\mathrm{d}P}{\mathrm{d}t} = 60\,\frac{P}{N}\,(M + F_{rec} - F) \tag{2.1}$$

Table 2.1 Classification of Variables. Nominal Values Are Those of the Solved Design Point at Which Every Scenario Opens.

Symbol	Role	Nominal	Range
P loop pressure	CV, measured	149.97	120–168 bar SP
T_{hot} hot spot	CV, measured	505.2	465–525 °C SP
y_I loop inerts	CV, sampled	10.07	5–18 % SP
r loop H_2:N_2	CV, sampled	3.000	2.6–3.4 SP
z_M makeup valve	MV	59.3	0–100 %
z_P purge valve	MV	37.4	0–100 %
q quench valves	MV	10.0	0–45 %
r_{mk} front-end trim	MV	3.00	2.4–3.6
N circulator speed	MV, operator	100	60–110 %
refrigeration	MV, operator	63	0–100 %
a catalyst activity	DV, unmeasured	1.00	0.50–1.05
makeup inert factor	DV, unmeasured	1.00	0.30–3.00
Δr_{mk} ratio drift	DV, unmeasured	0	±0.30
makeup availability	DV, measured	1.00	0.40–1.05
exchanger condition	DV, unmeasured	1.00	0.85–1.03
refrigeration shortfall	DV, unmeasured	0	−10–20 °C
P (bar)	state	149.97	20–260
y_A ammonia	state	0.0316	0.0001–0.45
y_I inerts	state	0.1007	0–0.60
y_H hydrogen	state	0.6508	0.05–0.92
T_1, T_2, T_3 (°C)	states	497/505/498	100–700
T_{sep} (°C)	state	−4.65	−40 to 60
six actuator positions	states	as above	as above

where N is the loop holdup in kmol, V is the loop gas volume, z is the compressibility, T_L is the mean loop gas temperature, M is the makeup flow, F_{rec} is the recycle returning from the separator and F is the circulator throughput, all in kmol/s. At the design point $N = 1543$ kmol. Pressure is inventory, and nothing else on this plant is, which is the fact the whole control structure is arranged around.

The circulator is a fixed volumetric machine, so its molar throughput rises with speed and with density, and its shaft power is proportional to throughput, as shown in Equation (2.2).

$$F = F_0 \frac{N_{sp}}{100} \frac{P}{P_{ref}}, \qquad W = k_W F \tag{2.2}$$

where N_{sp} is the speed demand in percent, F_0 is the rated throughput and k_W is the power coefficient. Equation (2.2) couples pressure back into contact time, because a loop running at higher pressure passes gas through the beds faster.

2.3.2 Composition at the Mixing Point

The loop is treated as one well-mixed gas volume whose composition is taken at the converter inlet, where the recycle and the makeup meet. Writing the species balance as $\mathrm{d}(N y_k)/\mathrm{d}t$ and substituting the inventory balance of (2.1) removes the total flow terms and leaves Equation (2.3).

$$\frac{\mathrm{d}y_k}{\mathrm{d}t} = \frac{60}{N}\left[M\left(y_{k,mk} - y_k\right) + F_{rec}\left(y_{k,rec} - y_k\right)\right] \tag{2.3}$$

where y_k is the mole fraction of ammonia, inerts or hydrogen at the mixing point, $y_{k,mk}$ is that species in the makeup and $y_{k,rec}$ is that species in the recycle. Nitrogen is the balance, $y_N = 1 - y_A - y_I - y_H$, and the loop ratio is $r = y_H / y_N$. Makeup carries no ammonia, so $y_{A,mk} = 0$, and the other two follow

from the front-end specification, $y_{I,mk}$ set by the reformer and $y_{N,mk} = (1 - y_{I,mk})/(1 + r_{mk})$ with $y_{H,mk} = r_{mk} y_{N,mk}$.

Equation (2.3) is the whole recycle argument in one line. A species that the converter does not consume and the separator does not condense appears in the makeup term and in the recycle term with almost the same concentration, so its only net removal is the purge, and it accumulates until the two balance.

2.3.3 Equilibrium and One Adiabatic Bed

The equilibrium constant is the Gillespie-Beattie correlation for the Haber equilibrium, written for ammonia over the roots of its reactants and returning K_p in inverse atmospheres, as shown in Equation (2.4) [18].

$$\begin{aligned} \log_{10} K_p = {} & \frac{2250.322}{T} - 0.85340 - 1.51049 \log_{10} T \\ & - 2.58987 \times 10^{-4}\, T + 1.48961 \times 10^{-7}\, T^2 \end{aligned} \tag{2.4}$$

where T is in kelvin. Because the reaction is exothermic, K_p falls by a factor of 1.49 for every 30 °C of rise at 400 °C and by 1.36 at 500 °C, and its square, the constant for the whole reaction, by about two, so every degree that speeds the reaction up worsens the destination it is heading for. With $y_N = (1 - y_A - y_I)/(1 + r)$ and $y_H = r y_N$, the equilibrium condition $y_A = K_p P\, y_N^{1/2} y_H^{3/2}$ collapses to a quadratic in the unreacted fraction, which is solved in closed form rather than by bisection, as shown in Equation (2.5).

$$y_A^{eq} = s - u, \quad u = \frac{-1 + \sqrt{1 + 4\kappa s}}{2\kappa}, \quad \kappa = \frac{K_p P\, r^{3/2}}{(1 + r)^2} \tag{2.5}$$

where $s = 1 - y_I$ and u is the combined hydrogen and nitrogen fraction at equilibrium. It is the same kind of closed form as Equation (1.2) of Chapter 1, written in mole fractions instead of conversion. Equation (2.5) is evaluated inside the bed solve, which is evaluated inside the derivative, which the fixed-step integrator calls four times per step, so the closed form matters for the speed of the steady-state solver.

A bed is adiabatic, so its state moves along a straight line in the ammonia against temperature plane, and the extent it can reach is where that line meets the equilibrium curve, as shown in Equation (2.6).

$$\begin{aligned} T(x) &= T_{in} + \lambda x, \qquad y_A(x) = \frac{y_A + 2x}{1 - 2x} \\ x_{eq} &: \; y_A(x) = y_A^{eq}\big(T(x), P, r, \tfrac{y_I}{1-2x}\big) \\ x &= \eta\, x_{eq}, \qquad \eta = 1 - \exp(-k(T_m)\, \tau_b^*\, a) \end{aligned} \tag{2.6}$$

where x is the nitrogen extent per mole of bed feed, λ is the adiabatic temperature rise for complete conversion of the nitrogen present, η is the fractional approach to equilibrium, a is the catalyst activity multiplier and $\tau_b^* = \tau_b F_0 / \max(0.3, F)$ is the contact time scaled by the actual circulation. The equilibrium extent is found by bisection in forty steps between zero and $0.98(1 - y_A - y_I)/(1 + r)$, which is the nitrogen the feed actually contains. The rate group is one Arrhenius term evaluated at the mean bed temperature $T_m = T_{in} + \tfrac{1}{2}\lambda x_{eq}$ with a square-root hydrogen coverage, as shown in Equation (2.7).

$$k(T_m) = k_0\, \mathrm{e}^{-E_a / R T_m} \sqrt{y_H / 0.652} \tag{2.7}$$

where the reference fraction 0.652 is the design loop hydrogen content, so the coverage term is unity at the design ratio. A single lumped group stands in for a Temkin-Pyzhev form with a diffusion correction [19], and it is calibrated so the first bed reaches 82.8 % of its equilibrium at the design inlet while beds two and three reach 98.2 and 98.1 %. Bed one converts $x_1 = 0.0263$, bed two 0.0101 and bed three 0.0050.

Quench mixes cold gas into the effluent between beds, half of the total fraction at each of the two points, and the converter extent is the flow weighted sum, as shown in Equation (2.8).

$$\begin{aligned} T^{+} &= \frac{T + \tfrac{q}{2} T_{cold}}{1 + \tfrac{q}{2}}, \qquad y_A^{+} = \frac{y_A^{out} + \tfrac{q}{2} y_A}{1 + \tfrac{q}{2}} \\ x_{tot} &= \frac{x_1 + \left(1 + \tfrac{q}{2}\right) x_2 + (1+q)\, x_3}{1+q} \end{aligned} \tag{2.8}$$

where q is the quench fraction of circulator throughput and the weights account for the quench gas that never saw the beds ahead of it. At the design line-up $x_{tot} = 0.03862$, the converter outlet carries 11.79 % ammonia against 3.16 % at the inlet, and the hot spot is 505.2 °C against a sintering limit of 530 °C.

2.3.4 The Converter Heats Its Own Feed

One algebraic equation gives this vessel two steady states, and it is Equation (2.9).

$$T_{in} = T_{cold} + \varepsilon\, c_{hx}\, (T_3 - T_{cold}) \tag{2.9}$$

where ε is the exchanger effectiveness and c_{hx} is a condition multiplier that represents a leaking bypass. At design, (2.9) lifts gas leaving the circulator at 40 °C to 419.1 °C. The operating point is a fixed point of the loop formed by (2.6) and (2.9) rather than a consequence of a valve position, which is the structure that produces an ignited branch and an extinguished branch [14, 15]. The beds carry thermal capacity so the gas answers immediately and the catalyst does not, as shown in Equation (2.10).

$$C_{cat} \frac{\mathrm{d}T_i}{\mathrm{d}t} = 60\, F\, c_{pg} \left(T_i^{out} - T_i\right), \qquad i = 1,2,3 \tag{2.10}$$

where C_{cat} is the catalyst and vessel heat capacity per bed in kJ/°C, c_{pg} is the molar heat capacity of the gas in kJ per kmol per °C and T_i^{out} is the algebraic bed outlet of (2.6). Both sides carry kJ per minute, and the resulting bed time constant is close to two minutes.

2.3.5 Separator, Purge and Makeup

The separator condenses ammonia to its vapor pressure with a departure factor, and its temperature follows the refrigeration demand through a first-order lag, as shown in Equations (2.11) and (2.12).

$$y_{A,sep} = k_{sep}\, \frac{\mathrm{e}^{A_s - B_s/(T_{sep} + 273.15)}}{P} \tag{2.11}$$

$$\tau_{sep} \frac{\mathrm{d}T_{sep}}{\mathrm{d}t} = -\frac{R_{duty}\, z_{ref}}{100} + 30 + \Delta T_{cool} - T_{sep} \tag{2.12}$$

where A_s and B_s are vapor pressure constants, k_{sep} is the departure from ideal separation at loop pressure, z_{ref} is the refrigeration demand in percent, R_{duty} is the refrigeration span and ΔT_{cool} is a shortfall disturbance. Equation (2.11) is clamped to the interval $[10^{-4}, 0.5]$. At –4.65 °C it leaves 3.79 % ammonia in the recycle, where an ideal Raoult calculation would predict about 2.4 %, and the ammonia going back round is the strongest single brake on an equilibrium-limited converter.

The material balance round the converter and separator gives the recycle and the product, as shown in Equation (2.13).

$$\begin{aligned} \Xi &= F x_{tot}, \qquad F_{out} = F - 2\Xi, \qquad y_A^{out} = \frac{F y_A + 2\Xi}{F_{out}} \\ L &= \frac{F_{out}\left(y_A^{out} - y_{A,sep}\right)}{1 - y_{A,sep}}, \qquad F_{gas} = F_{out} - L \\ F_{rec} &= \max\left(0,\, F_{gas} - \dot{n}_P\right), \qquad \dot{m} = 17.03\, L \end{aligned} \tag{2.13}$$

where Ξ is the nitrogen converted in kmol/s, F_{out} accounts for four moles of reactant becoming two of product, L is the liquid ammonia in kmol/s and $\dot{m}$ is the product in kg/s, reported on the scorecard as $86.4\,\dot{m}$ in t/d. Recycle compositions are the converter outlet corrected for the mole change and for the ammonia removed, as shown in Equation (2.14).

$$\begin{aligned} y_{I,rec} &= \frac{y_I}{1-2x_{tot}}\,\sigma, \qquad y_{H,rec} = \frac{y_H - 3x_{tot}}{1-2x_{tot}}\,\sigma \\ \sigma &= \frac{1-y_{A,sep}}{1-y_A^{out}} \end{aligned} \tag{2.14}$$

The purge is a square-root valve on the loop pressure and the makeup is a linear valve whose delivery falls as the loop it feeds gets harder to push into, as shown in Equation (2.15).

$$\begin{aligned} \dot{n}_P &= k_P \frac{z_P}{100}\sqrt{P} \\ M &= M_{max}\frac{z_M}{100}\,\alpha_{mk}\,\mathrm{clamp}\left(1 - 0.35\frac{P-P_{ref}}{P_{ref}},\, 0.3,\, 1.4\right) \end{aligned} \tag{2.15}$$

where k_P is the purge valve capacity in kmol/s per $\sqrt{\text{bar}}$, M_{max} is the makeup at a fully open valve and α_{mk} is the availability the front end is actually delivering. The pressure term in (2.15) is what gives the loop any self-regulation at all, because without it the only outlet that responds to pressure is the purge. Separating α_{mk} from z_M is what allows the front end to be taken away independently of the valve, and scenario 10 turns on exactly that distinction.

Each of the six actuators follows its command through a first-order lag, $\tau_j\,\mathrm{d}u_j/\mathrm{d}t = u_j^{cmd} - u_j$, with the makeup, purge, quench, speed, refrigeration and front-end trim constants listed in Table 2.2. Every state is clipped after each step, the pressure to 20–260 bar, the bed temperatures to 100–700 °C and the actuator positions to their travel.

2.3.6 Initialization and Numerics

Every scenario opens at a solved equilibrium. The application marches the derivative set from a fixed starting guess at a step of 0.25 min for 16 000 iterations, which is 67 hours of plant time, and successive substitution converges because every pass through the converter removes part of the reactant inventory. Evaluated at that line-up, the largest relative state derivative is 9.6×10^{-12}, which the application's own verification suite checks. The engine of Section 1.6.1 integrates by the classical fourth-order Runge-Kutta method at a fixed 0.05 min (3 s) step, the controller executes every 0.25 min (15 s), and the historian logs four samples per simulated minute. The exported file holds 104 columns, among them the true inert level beside the last analysis and each commanded valve position beside the actual one. The challenge scenario is the one exception to the solved start, because it is handed over lit and already drifting with a purge valve at 22 %, and solving that line-up would present a converter that is already out.

2.3.7 Assumptions

The model is detailed on the loop and thin on the vessel internals, and each simplification fails in a stated place.

Three lumped adiabatic beds. Each bed carries one temperature and one extent rather than an axial profile, so the hot spot is a bed outlet, while a real hot spot sits perhaps a third of the way into the bed and moves as the catalyst ages. Mixing at the quench points is instantaneous.

One kinetic group instead of a rate expression. The approach to equilibrium in (2.6) is a single Arrhenius term with a square-root hydrogen coverage, calibrated to a bed-one approach of 82.8 %. A Temkin-Pyzhev rate integrated down the bed with a diffusion correction is the honest form [19], and it carries a stronger ammonia inhibition than this one does.

Table 2.2 Model Parameters. Every Constant Above the Second Rule Is Shown in the Application and Adjustable by the User. The Six Below It Are Hidden.

Symbol	Meaning	Value	Unit
λ	adiabatic rise per unit extent	2958	°C
k_0	rate pre-exponential	1.46×10^{12}	—
E_a	apparent activation energy	165 000	J/mol
τ_b	contact time per bed	0.55	s
T_{max}	catalyst sintering limit	530	°C
T_{cold}	gas leaving the circulator	40	°C
ε	exchanger effectiveness	0.828	—
c_{pg}	loop gas heat capacity	35.5	J/mol/°C
V	loop gas volume	300	m^3
F_0	circulator rated throughput	6.00	kmol/s
P_{ref}	design loop pressure	150	bar
P_{max}	loop design pressure	172	bar
M_{max}	makeup at a full-open valve	1.70	kmol/s
k_P	purge valve capacity	0.01850	$\text{kmol s}^{-1}\text{bar}^{-1/2}$
$y_{I,mk}$	inerts in the makeup	0.0100	—
r_{mk}	makeup H_2:N_2	3.00	—
k_{sep}	separator departure factor	1.60	—
τ_{sep}	separator thermal lag	2.5	min
τ_{mk}	makeup valve travel	0.40	min
τ_{pg}	purge valve travel	0.25	min
τ_q	quench valve travel	0.30	min
τ_N	circulator speed response	0.50	min
τ_{ref}	refrigeration response	3.0	min
τ_{rmk}	front-end response to a trim	12	min
$y_{I,max}$	inert operating limit	0.16	—
W_{max}	circulator driver rating	12 500	kW
k_W	power per unit throughput	1650	kW per kmol/s
τ_{Tm}	thermocouple response	0.25	min
τ_{Pm}	pressure transmitter response	0.10	min
Δt_a	loop gas analysis cycle	15	min
—	ammonia value	0.42	$/kg
—	makeup syngas cost	2.50	$/kmol
—	electricity cost	0.075	$/kWh
Hidden constants, never displayed in the application			
C_{cat}	bed heat capacity	26 000	kJ/°C
z	loop compressibility	1.12	—
T_L	mean loop gas temperature	40	°C
A_s	vapor pressure constant	11.76	—
B_s	vapor pressure constant	2817	K
R_{duty}	refrigeration span	55	°C

Mole fractions rather than fugacities. Equation (2.4) is a fit to real gas data, but it is applied here to mole fractions. At 150 bar the fugacity coefficients depart from unity by tens of percent, and this is the largest single approximation in the model, so the equilibrium ammonia fractions are optimistic by a few points.

One well-mixed loop volume. Equation (2.3) places the whole inventory at one point. A real loop is a train of vessels with their own holdups and the analyzer looks at one place in it, so the true composition response carries a transport delay this model does not have.

Four components and one lumped inert. Argon and methane are combined into a single species that neither reacts nor condenses. The two behave differently in the separator, and methane in particular is

not entirely insoluble in liquid ammonia at loop conditions.

A separator that is a vapor pressure and a factor. Equation (2.11) is not a flash calculation, ammonia dissolved in the liquid drawn off is ignored, and k_{sep} is fitted rather than derived.

A circulator with no map. Equation (2.2) has no head curve, no surge line and no efficiency variation, so the machine cannot be operated into a region a real one would refuse.

Catalyst deactivation as a slider. Activity is one multiplier applied uniformly. Real deactivation is worst at the front of bed one, takes years, and changes the shape of the profile rather than only its height.

No front end and no hydrogen recovery. Everything upstream of the makeup valve is a supply and a ratio, and the purge goes to fuel rather than through the recovery unit most modern plants install [12]. The recovery unit is left out because it hides the economics the application is about.

2.4 What the Recycle Does

2.4.1 Two Small Flows Set the Steady Composition

Setting (2.3) to zero for the inert species and using the fact that the purge is drawn from the recycle stream gives the steady inert level directly, as shown in Equation (2.16).

$$y_{I,rec}^{\infty} = \frac{M\, y_{I,mk}}{\dot{n}_P}, \qquad \tau_I = \frac{N}{\dot{n}_P} \tag{2.16}$$

where τ_I is the inert time constant in seconds. This is Equation (1.4) of Chapter 1 written for the purge stream. Both are ratios of two small numbers. At the design line-up $M y_{I,mk}/\dot{n}_P = 1.0082 \times 0.01/0.08473 = 0.1190$, and the application reports 11.90 % inerts in the purge stream against 10.07 % at the converter inlet, the difference being the makeup dilution and the mole change through the beds. The time constant is $1543/0.08473 = 18\,209$ s, which is 303 min or 5.1 h. Cutting the purge valve from 37.4 to 20 % moves the inert level from 10.07 toward 15.34 % and covers 63 % of that distance in 5.12 h, so the measured response runs on the five-hour clock of (2.16). The smaller purge alone would set a slower clock, but the makeup falls with production as the inerts rise, and the smaller inert feed speeds the approach.

The same structure governs the hydrogen to nitrogen ratio, which is why both of the loop's slow variables are slow for the same reason. What the loop cannot consume in the proportion it is supplied has to leave through the same small purge, and the purge removes it at the loop concentration rather than at the concentration it arrived in.

🕮 **Pause.** Try Exercise 2.3 before reading on.

2.4.2 Amplification, and Why It Is Not Symmetric

Figure 2.2 is the settled loop ratio against the ratio the front end delivers, measured by running the plant to steady state with the pressure controller holding 150 bar. A once-through plant would lie on the identity line. This one does not come close to it.

Defining the amplification as the relative loop error divided by the relative makeup error, a front end 2.0 % low is amplified 9.9 times, one 5.0 % low is amplified 8.0 times and one 10.0 % low is amplified 5.9 times. The gain falls as the error grows because a loop that is already nitrogen rich purges nitrogen faster, so the removal term catches up with the imbalance sooner. The curve is also not symmetric. The same relative errors high are amplified 13.5, 17.2 and 23.9 times, because hydrogen leaves through a purge less readily than nitrogen does once the loop is hydrogen rich. Production is a maximum at the stoichiometric ratio, 679 t/d at $r_{mk} = 3.00$, falling to 553 t/d at 10 % low and 516 t/d at 10 % high.

Section 1.4.3 derives the settled excess of one reactant after a feed ratio error, Equation (1.6), which is the mechanism behind this curve. The amplification takes as long as it is large. Following the loop ratio after the front end drifts from 3.00 to 2.85, the settled value is 1.804 and the trajectory covers 63 % of that change in 3.2 h, 95 % in 10.7 h and 97.8 % in 14 h. Production falls from 678 to

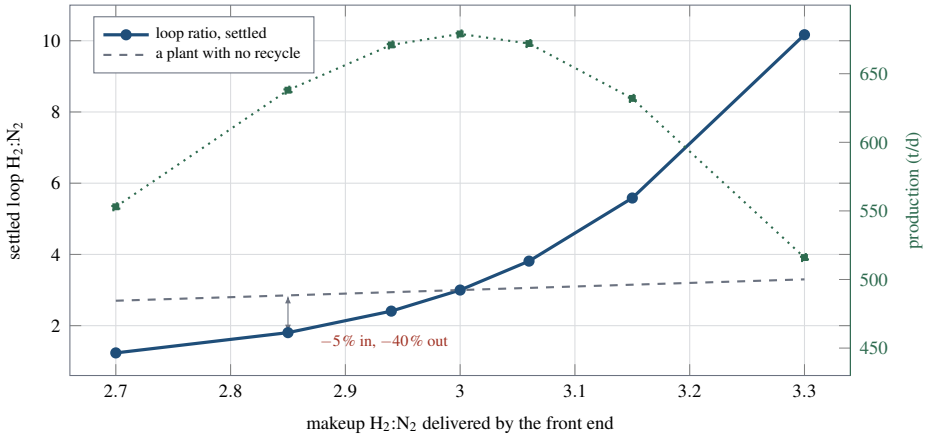


Figure 2.2 Recycle amplification of a front-end composition error, measured at steady state with the pressure controller holding 150 bar and the purge valve fixed at 37.4 %. The solid curve is the loop ratio, the dashed line is where a plant without a recycle would sit, and the dotted curve read against the right axis is ammonia production. Against the 679 t/d made at the stoichiometric ratio, a makeup ratio 5 % low settles the loop at 1.80 and costs 41 t/d, while the same error high settles it at 5.58 and costs 47 t/d.

638 t/d over the same period and the loop pressure never leaves 150.0 bar, because a controller is holding it. The whole disturbance therefore arrives as a production number that drifts down across half a day on a plant where nothing has alarmed.

> **TRY IT IN THE APP** **6 · A Small Error at the Front End**
>
> Load scenario 6 and press Run. From 30 min the front end delivers 2.85 to 1 instead of 3 to 1. Watch the loop H_2:N_2 trend head for 1.80, covering 63 % of the way in 3.2 h, while production falls from 678 toward 638 t/d and the pressure trend stays flat at 150 bar.

2.4.3 Which Variable Is the Inventory

At steady state the makeup entering the loop equals the reactant leaving it as ammonia plus the reactant leaving it through the purge, so the makeup valve is the throughput manipulator and there is no other candidate. Section 1.4.3 poses the same choice for the running example. The design decision is which variable it is allowed to hold, and the application offers both answers as a segmented control. With the pressure controller on the makeup valve, pressure is the controlled inventory and production is whatever holding it requires. With the operator on the makeup valve, production is set and pressure is whatever the converter needs in order to consume the gas, with the pressure controller retained behind a low selector as a high limit. Each arrangement is consistent in the sense that the loop inventory is regulated by a flow entering it [4], and each is used in practice [3, 9].

The two differ in where a disturbance appears, which Section 2.6 measures. What matters for the operator is that only one of the two failure modes announces itself. A plant losing catalyst activity with the pressure controller in charge quietly makes less ammonia, and the same plant with the operator in charge quietly consumes the margin to its relief valves.

Pause. Try Exercise 2.1 before reading on.

Table 2.3 The Four Loops as Implemented, With the Default Tunings Every Automatic Scenario Loads

CV	MV	K_c	τ_I (min)	Action	Plant
P	z_M	2.4	9	reverse	minutes
T_{hot}	q	0.14	8	direct	minutes
y_I	z_P	2.5	300	direct	5.1 h
r	r_{mk}	0.09	260	reverse	3.2 h

2.5 Controller Design

2.5.1 The Implemented PI Algorithm

All four loops use the shared velocity-form law of Section 1.6.2, Equations (1.14) to (1.17), executed every $\Delta t_c = 0.25$ min. The action is reverse on a loop where a rising measurement calls for a falling output and direct where it calls for a rising one, and the measurement passes through a first-order filter advanced by its exact zero-order-hold recursion. The clamp acts on the emitted output only, back-calculation unwinds the integral with a tracking constant T_t that defaults to τ_I, and Manual operation solves the integral so that the output equals the operator's valve position, which makes transfer in both directions bumpless. All four loops run with $\tau_D = 0$, so the derivative path is present in the implementation and inactive in every shipped tuning. No loop carries an output rate limit, because the actuator lags of Table 2.2 already limit how fast a valve moves.

2.5.2 Four Loops on Two Time Scales

Table 2.3 lists the loops, their default tunings, their direction and the plant each has in front of it. The pressure and hot-spot loops read continuous transmitters filtered at 0.10 and 0.25 min. The inert and ratio loops read one chromatograph on the loop that reports both numbers every 15 min and holds them between analyses, so both slow loops move in steps. The sample period is not the difficulty on those two loops. A 15 min analysis on a 5.1 h plant is twenty times faster than the process it measures, and the difficulty is the process.

Gains carry the units of their pairing, 2.4 % of makeup valve per bar, 0.14 % of quench per °C, 2.5 % of purge valve per percent of inerts, and 0.09 of makeup ratio per unit of loop ratio. Output limits are 0 to 100 % on the makeup and purge valves, 0 to 45 % on the quench, and 2.4 to 3.6 on the front-end trim. The two reset times in the hundreds of minutes are conservative by a wide margin, and scenario 9 measures how wide. Running the doubled makeup inerts of scenario 9 with the shipped $K_c = 2.5$ and $\tau_I = 300$ min leaves the inert level at 10.52 % after 48 h and still approaching its 10 % setpoint, while $K_c = 10$ with $\tau_I = 60$ min first reaches 10.00 % 6.3 h after the change and moves the purge valve through 79 percentage points of travel against 55. Neither tuning oscillates, and even the slider maxima of $K_c = 30$ and $\tau_I = 20$ min settle within eight hours with no sustained cycle, after one undershoot to 9.81 %, 12 min with the purge valve wide open and 125 points of travel. The reason is in the model rather than in the tuning rule. The inert level is a single lag of 5.1 h behind a 15 min sample hold, so the loop has almost no phase lag beyond first order and gain alone cannot destabilize it. What limits the purge controller in practice is the cost of overshooting, because every excess percent of valve position is hydrogen leaving through the stack for hours before the analyzer reports that it was not needed.

2.5.3 Four Structures and One Selector

The application offers manual operation, a basic-loops structure in which the pressure and hot-spot controllers are in service and nobody is watching the composition, a supervised structure that adds the inert loop and, on a separate selection, the ratio loop, and a predictive structure in which one plan replaces the two slow loops (Section 2.5.4). The basic-loops setting is where most of the scenarios

live, because a plant with the fast loops closed and the slow ones open works perfectly for a shift and poorly for a month.

When the operator owns the makeup valve, the pressure controller is retained as an override against the high limit rather than against the normal setpoint, and a low selector passes whichever signal asks for less gas, as shown in Equation (2.17).

$$z_M = \min\left(z_M^{op},\ \mathrm{PI}\left(P_{hi}, \hat{P}\right)\right) \tag{2.17}$$

where z_M^{op} is the operator demand and $P_{hi} = 165\,\mathrm{bar}$ is the high limit against a design pressure of 172 bar. Whenever the operator demand is selected, the pressure controller is tracked to the applied output through the same expression used for manual operation, so it takes over without a jump. Pointing the override at the normal setpoint instead of the high limit is the mistake worth naming, because it makes the production structure into a slower version of the pressure structure. The override then retakes the valve the moment the pressure reaches the number the other structure was already holding, and the two arrangements become indistinguishable.

2.5.4 Model Predictive Control

Section 1.6.3 sets out the predictive option of the framework in general, and this subsection gives what is specific to the loop. Under the predictive structure the pressure and hot-spot loops run exactly as in Table 2.3, and a quadratic program chooses the purge valve z_P and the front-end trim r_{mk} together [20, 21]. A plan is made every $T_s = 15\,\mathrm{min}$, on each new loop analysis at the declared cycle, and its filter update reads the last analysis reported. The horizon is $N_h = 48$ periods (12 h, adjustable from 6 to 24 h), grouped into twelve blocks of $\{1,1,1,1,2,2,3,4,6,9,15,3\}$ periods. Every decision is expressed as a change from the commands in force, and the first block alone is sent to the valves and kept until the next analysis.

The plan predicts with the equations of Section 2.3, evaluated at the declared parameter values of Table 2.2 and the nominal disturbances of Table 2.1, closed around the pressure and hot-spot PI loops at their shipped tuning with the 0.10 and 0.25 min transmitter lags, and linearized at the design line-up by central differences. Its eighteen states are the fourteen of the plant, the two transmitter filters and the two loop integrals, and an exact zero-order hold over T_s gives Equation (2.18).

$$\begin{aligned} \tilde{x}_{k+1} &= A_d \tilde{x}_k + B_d\left(\tilde{u}_k + d_k\right) + E_d \tilde{w}_k, \qquad d_{k+1} = d_k \\ \tilde{y}_k &= C\tilde{x}_k + D_w \tilde{w}_k \end{aligned} \tag{2.18}$$

where tildes denote departures from the design point, $\tilde{u} = (z_P, r_{mk})$ holds the decided inputs, d one unmeasured disturbance per decided input, $\tilde{w}$ the known inputs held over the horizon (the pressure and hot-spot setpoints, the operator's makeup demand, the circulator speed and the refrigeration demand), and $\tilde{y} = (y_I, r, z_M)$ the loop inerts in percent, the loop ratio and the makeup valve command, which alone reads a known input directly through D_w. Two versions are built, one for each owner of the makeup valve. Step-response fits were not used, because a step made in manual opens the two fast loops that stay in service under the plan.

The settled gains are -0.194 percent of inerts per percent of purge (-0.244 with the operator on the makeup valve) and 11.48 of loop ratio per unit of front-end ratio. With the pressure loop on the makeup valve the inerts cover 45.5, 78.4 and 97.5 % of a purge move's settled effect after 2, 5 and 12 h, close to a 3.3 h time constant. Against the engine's mean response to ± 1 percent of purge and ± 0.01 of trim from the line-up of scenario 4 the model errs by at most 0.26 % of the largest response between 15 and 720 min, while a single 5 percent opening lowers the inerts by 0.866 % in 12 h against the model's 0.945 %, the nonlinearity the model leaves out.

A steady-state Kalman filter estimates $\tilde{x}$ and d at every decision under every structure, so a switch to the plan starts warm, and it reads only what the slow loops read, the analyzed y_I and r, with assumed errors $\delta_I = 0.01\,\%$ and $\delta_r = 0.002$. Each disturbance carries process noise $\varphi\,\delta_m^2 / G_m^2$, with $\varphi = 10^4$ and G_m the settled gain of its input on its own analysis, which makes the prediction offset

free [22]. The makeup valve's offset from the model, b_M, is read from the pressure loop's command at each decision. Replacing the filter by a constant output bias, on a diagnostic copy that is not part of the application, raises the inert IAE of scenario 12 from 56.6 to 539.3 % · min and its ratio IAE from 8.5 to 68.5 min, because a bias corrects a drift only as fast as the five-hour plant reveals it.

Each plan minimizes Equation (2.19), written in dollars per decision period.

$$\begin{aligned} J = & \sum_{k=1}^{N_h} \left[q_r \left(\frac{r_k - r_{\mathrm{SP}}}{0.05} \right)^2 + c_P z_{P,k} \right] \\ & + 50 \lambda_u \sum_{j=1}^{12} \left[\left(\frac{\Delta z_{P,j}}{5} \right)^2 + \left(\frac{\Delta r_{mk,j}}{0.05} \right)^2 \right] \\ & + \sum_{i \in \{I,M\}} \left(\rho_i s_i^2 + \mu_i s_i \right) \end{aligned} \tag{2.19}$$

Here $q_r = 20$, $\lambda_u = 5$, Δ is the move of block j, and $c_P = \$5.10$ per percent of purge prices the synthesis gas the valve vents in one period at the declared \$2.50 per kmol and the design pressure, so the plan purges only as much as the inert ceiling demands. The slacks s_I and s_M of the soft rows of Equation (2.20) cost $\rho_I = 10^5$ per square percent and $\mu_I = 10^4$ per percent of inerts, and $\rho_M = \mu_M = 10^4$ for the makeup valve, so a reachable limit costs nothing.

$$\begin{aligned} & 0 \le z_{P,k} \le 100, \qquad 2.4 \le r_{mk,k} \le 3.6 \\ & y_{I,k} \le y_{I,\mathrm{SP}} + \max(0,\, \hat{y}_{I,0} - y_{I,\mathrm{SP}})\, \mathrm{e}^{-kT_s/\tau_f} + s_I \\ & z_{M,k} + b_M \le 98 + s_M, \qquad s_I,\ s_M \ge 0 \end{aligned} \tag{2.20}$$

where $\hat{y}_{I,0}$ is the filter's present inert level and $\tau_f = 120\,\mathrm{min}$ lets a level already above its setpoint come back along a lag instead of pricing the whole horizon. The input rows are hard and are the limits the slow loops clamp to. The makeup row and a purge hold act only with the pressure controller on the makeup valve, and while the pressure reads more than 2 bar under its setpoint the plan may close the purge but not open it. The model has the pressure loop on its setpoint, and a loop being drained concentrates its inerts, which the filter would otherwise take for purge that has lost its worth. The compressor limit is a soft row rather than a hard constraint, because the makeup valve is an output of the pressure loop, so a hard row on it would be infeasible in every front-end cut. The only hard rows are the input bounds, which always contain the present commands, so a plan cannot be infeasible.

If the solver returns a plan that is suboptimal, nonfinite or beyond the valve limits, the supervised inert and ratio loops, which have followed the applied commands throughout, take over both handles bumplessly and give them back at the next acceptable plan. Nothing in the plan comes from the run's parameters or from the plant's states. Its inputs are the analysis, the pressure transmitter, the makeup command the pressure loop sends, the setpoints, the operator's settings and the commands applied, so an edit a learner makes to the plant is, for the plan, model error that the filter absorbs into its disturbance estimates. A dual active-set method [23] solves a plan in a mean of 0.77 ms in a browser at the 12 h horizon and 1.19 ms at 24 h over three runs, under 0.01 percent of one core at real time.

2.6 Results

Thirteen scenarios are arranged from the converter outward, listed in Table 2.4. Each fixes the starting mode, the control structure, the four tunings, the disturbance settings and any timed events, and each carries the briefing and the completion rule of Section 1.6.4. The exercises at the end of the chapter run several of them.

2.6.1 The Converter Has an Optimum and a Cliff

Scenario 2 runs the quench by hand. Closing it from 10 to 4 % takes the beds from 497/505/498 to 528/528/525 °C an hour later and the ammonia leaving bed three falls from 11.79 to 10.33 %.

Table 2.4 Scenario Sequence and the Mechanism Each One Exercises

#	Scenario	What it exercises
1	Four valves by hand	two actuators answer in minutes and two in hours
2	The converter has a best temperature	rate against equilibrium on one vessel
3	Blow it out and get it back	two steady states from (2.9)
4	Pressure and hot spot on automatic	a held variable stops carrying information
5	Which variable is the inventory	throughput manipulator pairing
6	A small error at the front end	recycle amplification over 14 hours
7	Inerts and the five hours they take	the time constant of (2.16)
8	What a purge valve is buying	conversion against hydrogen loss
9	The makeup gets dirty	a controller that converts one cost into another
10	The front end cuts back	self-regulation until an actuator saturates
11	Five years of catalyst	constraint margin spent by the first disturbance
12	Five hours ahead: Supervised or MPC	one plan against two PI loops on two slow drifts
13	A day on the loop (challenge)	structure first, then the slow loops early

Opening it to 14 % runs the beds at 471/485/479 °C after the next hour and the outlet rises to 12.55 %. Hotter is faster and worse, because the equilibrium curve of (2.4) has come down further than the rate has climbed.

Scenario 4 puts the pressure controller on the makeup valve and the hot-spot controller on the quench and then moves the temperature setpoint. Stepping the hot spot from 505 to 498 °C half an hour into the run takes production from 678 to 726 t/d two hours later, an improvement of 7.0 %, while the loop pressure stays within 0.4 bar of 150 bar. Waiting a day gives 719 t/d, because the extra production draws more makeup and the inert level rises from 10.07 to 10.52 % at a fixed purge valve, which is the two time scales appearing inside a single step response. Continuing downward in five-degree steps, each setpoint held from the start of its own run and read at 1440 min, gives 738, 748, 753 and 754 t/d at 493, 488, 483 and 478 °C, with the quench valve opening from 10.10 to 16.26 %, so the converter's own optimum is near 480 °C and worth 11 % of production.

Nobody operates there, and scenario 3 is the reason. Opening the quench to 22 % for seventy minutes and then returning it to exactly 10 % does not relight the converter, which goes out 6.6 min after the valve opens. The inlet temperature falls from 419 to 90 °C and stays there, while the loop pressure runs from 150 to 214 bar within twenty minutes of the opening, and to the 260 bar ceiling of the model by 51 min, because the makeup gas is still arriving and nothing is consuming it. Held from the design state in manual, a quench of 15.3 % keeps the converter lit and 15.4 % puts it out, so past that point the converter is on the wrong side of the ignition condition, and the recovery is a startup heater rather than a valve position.

2.6.2 Where the Disturbance Lands

Scenario 5 runs the same plant twice. With the pressure controller on the makeup valve the loop sits at 149.99 bar and makes 678 t/d. Handing the valve to the operator and asking for 66 % takes production to 749 t/d and lets the pressure float to 155.5 bar, which is what the converter needs in order to consume the extra gas. Dropping the catalyst activity to 0.75 then separates the two arrangements. Production holds at 744 t/d and the pressure carries the loss instead, reaching 157.6 bar two hours after the change and settling at 158.3 bar, 6.7 bar inside the 165 bar high limit at which (2.17) would take the valve back. Under the other pairing the same disturbance appears as a production number and nothing else, a loss of 17 t/d two hundred minutes after the change with the pressure on 150 bar.

Scenario 10 is the self-regulation the arrangement buys. Removing 45 % of the makeup gas availability with the pressure controller in charge drives the makeup valve from 59.4 to 100.0 % within ten minutes, where it stops being a controller. The pressure falls to 140.6 bar eight minutes after the cut, recovers only to 146 bar, and production settles at 631 t/d against 679 t/d. Nothing trips

Table 2.5 Measured Outcomes of the Composition and Constraint Scenarios. Every Reading States the Time Since the Disturbance at Which It Was Taken, Because a Loop Composition Is Still Moving Hours After the Disturbance.

#	Reading	y_I (%)	r	Rate (t/d)
—	design line-up	10.07	3.000	677
6	front end 3.00 to 2.85, after 4.25 h	10.01	2.128	659
6	the same, after 13.5 h	9.73	1.834	639
6	the same, settled after 24 h	9.65	1.807	638
7	purge 37.4 to 20 %, after 0.5 h	10.55	3.000	675
7	the same, after 3 h	12.40	3.000	629
7	the same, after 5.5 h	13.54	3.000	598
7	the same, after 9.5 h	14.51	3.000	572
8	purge shut to 8 %, after 11.2 h	20.81	3.000	471
8	purge open to 85 %, after 11.7 h	5.44	3.000	791
9	makeup inerts doubled, after 24 h	11.38	3.000	647
9	the same, after 72 h	10.21	3.000	677
11	activity to 0.65, after 4.5 h	9.88	3.000	650
11	exchanger −8 %, 20 min later	8.92	3.000	0

and nothing alarms. The hot spot dips 6.9 °C below setpoint as the pressure falls and is back within 1 °C twenty minutes later, and it overshoots by 9.5 °C when the gas returns at 400 min and recovers as quickly, because a recycle absorbs a large feed disturbance by changing its own inventory. The useful part of that behavior ends exactly where the valve travel does.

2.6.3 Slow Composition and the Cost of a Purge

Table 2.5 collects the measured outcomes of the scenarios that turn on composition, each with the observation window it was read over. Scenario 7 halves the purge and is the clearest of them. Thirty minutes after the cut the inert level has moved from 10.07 to 10.55 %, which on a trend is indistinguishable from noise, and three hours after the cut it is 12.4 %. Production falls from 678 to 629, then 598, then 572 t/d while the pressure stays within 0.6 bar of 150 bar and the circulator power, after a rise of 0.04 MW in the first twenty minutes, stays within 0.01 MW of 9.90 MW. Sixteen percent of the plant output has gone and the only instrument that reports it is the analyzer.

Scenario 8 takes both extremes of the purge valve. Nearly shut at 8 %, the inert level reaches 20.81 % after 11.2 h and production falls to 471 t/d, so a fifth of the loop is gas that does nothing but occupy partial pressure and get compressed, and the saving is 72 kg/h of hydrogen against 389 kg/h at design. Wide open at 85 %, the loop is clean at 5.44 % inerts and makes 791 t/d, and the stack is carrying 939 kg/h of hydrogen while the makeup valve has gone from 59.4 to 75.0 %. That is 26 % more syngas bought for 17 % more ammonia.

Table 2.6 sweeps the valve properly, with each point run to a settled state after moving the valve five minutes into a design-start run. This protocol follows the lit converter branch; a cold start at the lowest purge can instead converge to an extinguished converter. Read down the table the trade is visible in two places at once. Syngas per tonne of product rises monotonically with purge, from 123.7 to 139.3 kmol/t, so feedstock efficiency wants the valve shut, while production rises monotonically from 484 to 792 t/d, so plant output wants it open.

Which of the two wins depends on the prices and pairing, a consequence of Section 2.4. With pressure control on makeup and the illustrative syngas price of $2.50 per kmol, the settled margin peaks near 42 % purge at $2 049 per hour, compared with $2 041 per hour at the design 37.4 % and $1 628 per hour at 85 %. Sampled valve positions from 37.5 to 47.5 % remain within 1 % of the

Table 2.6 Purge Valve Sweep on the Lit Branch, Each Move Made After a Design Start and Settled Over 96 h With the Pressure Controller Holding 150 bar. Margin, Quoted in the Text, Is the Application's Own Metric, Ammonia Sold Less Makeup Gas and Circulator Power, Evaluated at the Settled State.

z_P (%)	y_I (%)	Rate (t/d)	z_M (%)	H_2 purged (kg/h)	Syngas (kmol/t)
15	18.15	484	40.8	140	123.7
25	13.32	599	51.4	249	125.9
37.4	10.08	679	59.4	389	128.6
42	9.27	699	61.7	441	129.7
55	7.58	740	66.8	591	132.5
85	5.43	792	75.1	939	139.3

sampled best. Both extremes cost money: dilution loses production at low purge, while the additional syngas bought to hold pressure costs more than its additional ammonia earns at high purge. The flat optimum explains why a valve setting can look correct on every screen and still lose money over a year. The application's price slider extends to \$5 per kmol; reducing it to \$0.145 per kmol instead makes the margin rise as the purge opens. This optimum belongs to the stated prices and pairing, rather than to a universal purge setting.

> **+ DID YOU KNOW? Venting More Gas Makes More Ammonia**
>
> From the simulator: opening the purge valve from 15 to 85 % sends 939 kg/h of hydrogen out of the loop instead of 140 kg/h, yet production rises from 484 to 792 t/d, because the loop then carries 5.43 % inerts instead of 18.15 % (Table 2.6).

Scenario 9 supplies the cleanest argument for treating the inert setpoint as an economic decision. Doubling the methane slipping through into the makeup with the inert controller in service takes the purge valve from 37.6 to 69.9 % in the twenty-four hours after the change and to 81.4 % after seventy-two, still opening, while the hydrogen leaving through it rises from 391 to 845 kg/h and the syngas consumed per tonne rises from 128.7 to 141.9 kmol/t for 677 t/d of product against 678 before the change. The controller does the right thing and the plant is worse off anyway, because an inert controller does not solve an inert problem. It converts an upstream problem into a hydrogen bill, and the answer is at the reformer.

2.6.4 Margin, Measured Before It Is Spent

Scenario 11 ages the catalyst to 0.65 and the plant absorbs it. The beds cool, the hot-spot controller closes the quench from 10.08 to 8.45 % to put the temperature back on 505 °C, and production settles at 650 t/d. Every indication on the display is normal. Four and a half hours later the feed and effluent exchanger loses 8 % of its effectiveness, and the quench controller now needs more travel than the 8.45 % it has. The valve is shut four minutes after the event, T_{in} falls, and the converter is out six minutes after it, with the loop pressure passing 160 bar two minutes later. Fresh catalyst survives the same exchanger loss, with the quench driven shut and the inlet holding at 392 °C, which is the comparison the scenario asks for after a Reset.

The first disturbance did not damage the plant. It removed the plant's ability to survive the second one, and no controlled variable said so, because the hot spot was on setpoint at every instant until it was not. Valve position is therefore the trend worth its own strip on this plant, and the application reports the quench command beside the hot spot for that reason.

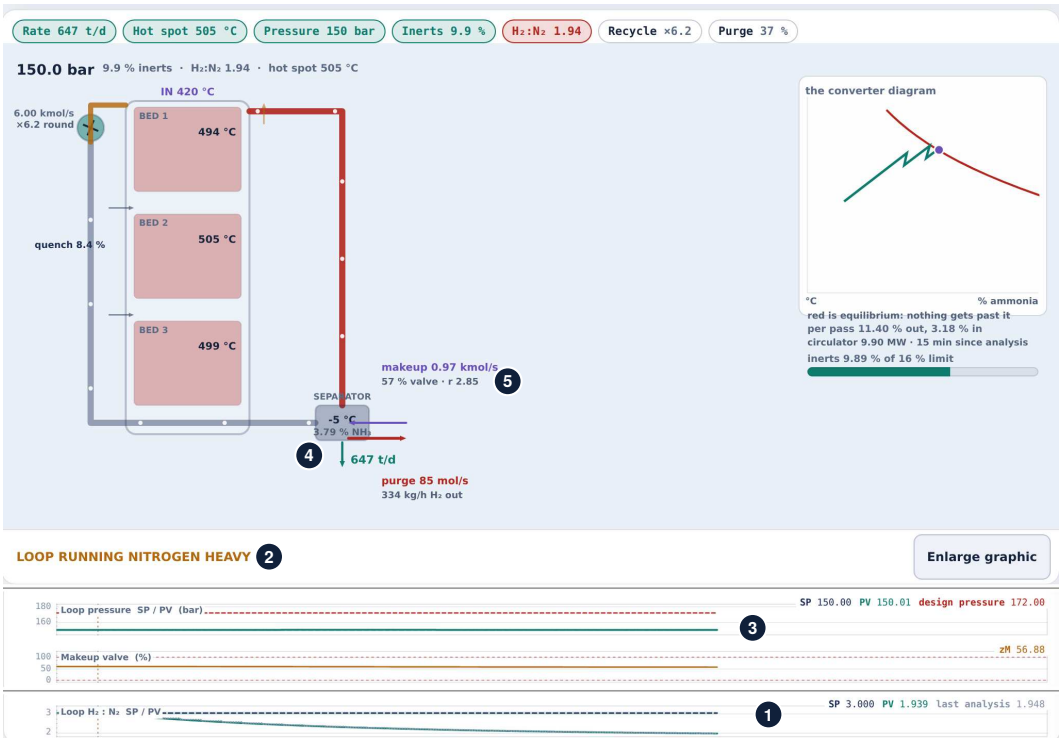


Figure 2.3 Scenario 6 at $t = 480$ min on the 12 h trend window, seven and a half hours after the front-end ratio drifted from 3.00 to 2.85, with the pressure and hot-spot loops in service and the composition loops open, cropped to the badges, the plant view and the loop pressure, makeup valve and loop H_2:N_2 strips of the trends. (1) The loop ratio has fallen to 1.939 against a last analysis of 1.948 that is 14.75 min old, and it is still moving toward 1.80, which is what the status line (2) calls a loop running nitrogen heavy. The point of the view is the pressure trend (3) read against production (4). Loop pressure is 150.01 bar against a 150 bar setpoint and has stayed within 0.04 bar of it since the drift, while production has drifted from 678 to 647 t/d and the makeup valve (5) has closed from 59.4 to 56.9 % to keep the pressure there. The hot spot is 504.9 °C on a 505 °C setpoint, the quench sits at 8.41 %, the loop carries 9.89 % inerts, and the recycle badge reads 6.2, up from 5.94 before the drift, because the same circulation is now serving less makeup.

> **TRY IT IN THE APP** **11 · Five Years of Catalyst**
>
> Load scenario 11 and press Run. At 30 min the catalyst activity falls to 0.65 and the quench closes from 10.08 to 8.45 % to hold the hot spot. Watch the quench strip at 300 min, when the exchanger loses 8 %: the valve is shut four minutes later, and the converter is out six minutes after the event.

2.6.5 The Challenge

The graded scenario hands over a fifteen-hour shift with the purge valve already at 22 %, makeup inerts 30 % above specification and rising to 80 % above at two hours, a front end running 0.06 low on ratio, catalyst at 0.90, and a reformer that removes 45 % of the makeup availability from the fifth hour to the eighth and a half. The case ships with weak tunings on the basic loops, and with those the converter goes out at 302.7 min, 2.7 min into the front-end cut. That run fails, with a worst pressure error of 3.15 bar over the last five hours, 7.7 % inerts and a ratio of 2.70 at the end. The target, judged at 900 min on the case as given, is the loop pressure within ±3 bar of 150 over the last five hours, the

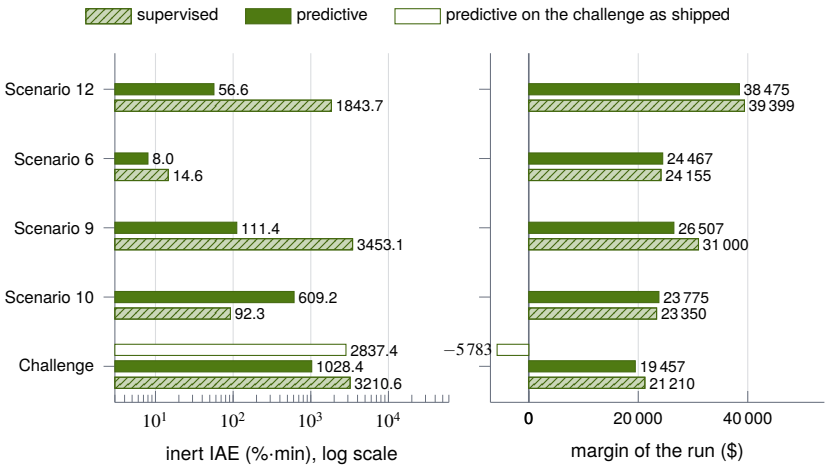


Figure 2.4 Table 2.7 as bars, digit for digit: the inert IAE on a logarithmic axis (left) and the margin of each run (right), with the supervised loops hatched, the predictive controller solid and the predictive controller on the challenge as shipped open. Notice that the plan cuts the inert error in every case but scenario 10, yet earns less margin than the supervised loops on scenario 12, scenario 9 and the challenge, and that on the challenge as shipped it ends with a negative margin.

hot spot under 530 °C with the converter alight all day, and at the end inerts under 13 % and the ratio within 0.15 of three. The suggested configuration, the default tunings of Table 2.3 with both slow loops in service from the start and the purge and quench stations left where the shift found them, meets it with a worst pressure error of 0.04 bar, a peak hot spot of 512.7 °C, 12.47 % inerts, 0.53 point inside the limit, and a ratio of 2.99, and the application verifies that on itself.

The answer is a structure and a schedule rather than a tuning constant. On the default tunings the two fast loops hold the plant for the whole shift, alight and within 0.004 bar of setpoint over the last five hours, but on their own they let the inerts cross the 16 % operating limit at 3.1 h and finish at 21.0 % with the ratio at 2.22. The inert loop and the ratio loop decide where the plant finishes, and both need three time constants to get there. With the fast loops tuned from the start, putting the slow loops in service at twenty minutes finishes inside the target at 12.89 % inerts and a ratio of 2.93, waiting two hours finishes at 13.21 % and fails, and waiting six hours finishes at 14.73 %. The front-end cut in the middle is a distraction and treating it as one is correct, because the loop absorbs it by depressuring to 141.7 bar and recovers on its own before the scored window opens, while a purge valve moved for a three-hour disturbance is still moving five hours after the disturbance has gone. Figure 2.3 shows the application during scenario 6, which is the same lesson without the clock.

2.6.6 Two Slow Loops Against One Plan

Table 2.7 runs the supervised structure, with its shipped tuning, the ratio loop on the analyzer and the suggested settings on the challenge, against the predictive structure with the settings of scenario 12 on every case and only the structure changed. Both face the same events, analyzer, transmitters and duration, the application has no measurement noise, and neither receives a preview. Figure 2.4 draws the table’s inert error and margin columns side by side.

Scenario 12 starts on the design line-up and at 60 min raises the methane slipping into the makeup by half and turns the front end 0.10 nitrogen heavy, two drifts that each take hours to reach the loop. On the first analysis that shows them, at 75 min, the plan opens the purge from 37.9 to 56.2 %, reaching 59.5 by 120 min, while it takes the trim to 3.12 by 120 min and eases it back to the 3.10 the drift needs by the end of the day (Figure 2.5). The inerts peak at 10.36 % at 81.5 min and stay

Table 2.7 Supervised Loops Against Model Predictive Control

Case (run length)	Structure	IAE y_I (% · min)	Peak, end y_I (%)	Purge travel (%)	IAE r (min)	Trim travel	IAE P (bar · min)	H_2 purged (t)	NH_3 (t)	Margin ($)	Plans, mean (µs)
Scenario 12 (1440 min)	supervised	1843.7	11.98, 10.73	22.8	269.95	0.157	26.0	11.74	648.6	39 399	—
	predictive	56.6	10.36, 10.00	22.0	8.47	0.147	13.3	14.58	680.2	38 475	96, 529
Scenario 6 (720 min)	supervised	14.6	10.07, 10.01	0.3	321.60	0.185	5.0	4.45	337.0	24 155	—
	predictive	8.0	10.07, 10.00	0.3	13.50	0.243	1.4	4.72	340.1	24 467	48, 525
Scenario 9 (1440 min)	supervised	3453.1	13.65, 11.42	42.5	0.00	0.000	43.4	14.28	623.3	31 000	—
	predictive	111.4	10.70, 10.00	46.0	0.03	0.000	26.5	20.44	679.9	26 507	96, 764
Scenario 10 (720 min)	supervised	92.3	10.64, 10.07	8.3	0.00	0.000	1680.3	4.61	328.5	23 350	—
	predictive	609.2	11.41, 9.64	117.8	0.08	0.002	708.8	4.56	331.9	23 775	48, 502
Challenge (900 min)	supervised, suggested	3210.6	15.11, 12.47	40.7	164.31	0.111	370.7	6.59	368.9	21 210	—
	predictive, suggested	1028.4	13.10, 10.10	183.4	39.34	0.285	862.8	9.70	400.5	19 457	60, 489
	predictive, as shipped	2837.4	13.10, 5.59	122.8	37.52	0.274	33 197.0	4.81	150.7	−5 783	60, 455

Errors and travel accumulate from $t = 0$ as on the scorecard, end values are read at the run length, no plan was rejected, and solve times are wall-clock time in Node. The last row switches only the structure on the challenge as shipped.

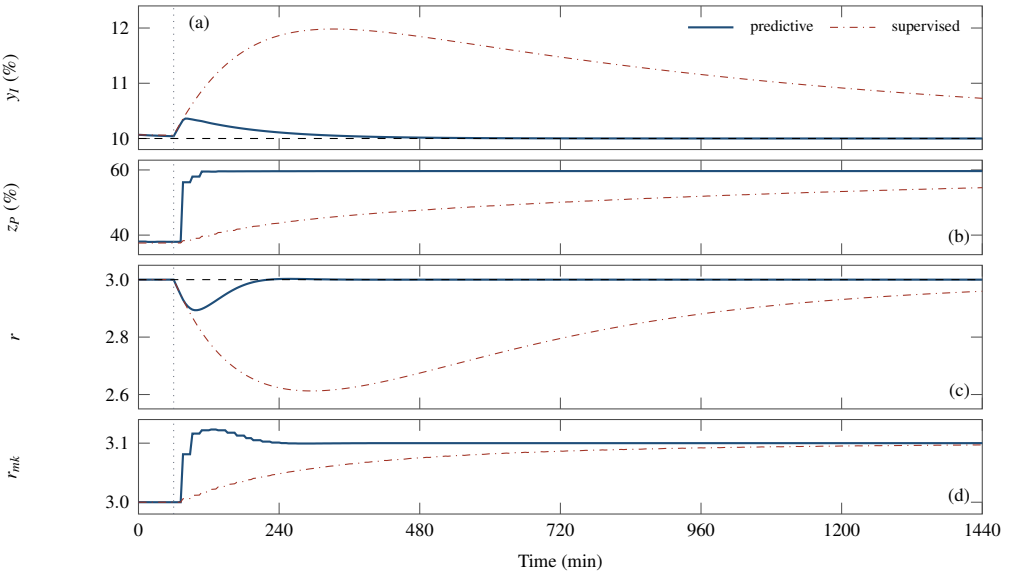


Figure 2.5 Scenario 12 over its 1440 min day under the predictive structure (thick solid lines) and the supervised structure (dash-dotted lines), with methane slip raised by half and the front end turned 0.10 nitrogen heavy at 60 min (dotted). (a) True loop inerts against the 10 % setpoint (dashed). (b) Purge valve command. (c) True loop H_2:N_2 against its setpoint of 3 (dashed). (d) Front-end trim command.

within 0.1 % of the setpoint from 252 min, and the ratio, lowest at 2.89, stays within 0.02 of 3 from 178.5 min. Under the supervised loops the purge creeps from 37.6 to 54.6 % over the day, the inerts peak at 11.98 % at 333 min and still read 10.73 % at 1440 min, and the ratio falls to 2.61. The plan cuts the inert IAE from 1843.7 to 56.6 % · min and the ratio IAE from 269.95 to 8.47 min with less travel, 22.0 against 22.8 percent of purge and 0.147 against 0.157 of trim, because it moves each handle sooner rather than further. The bill is the purge. Holding 10 % on dirtier gas vents 14.58 against 11.74 t of hydrogen, and the day's margin is $924 lower. Neither the makeup ceiling nor the purge hold acts on this day, with the makeup valve at most at 62.7 % and the pressure never below 149.4 bar, and switching the ceiling off changes no plant variable.

Scenario 6, the front end alone 0.15 nitrogen heavy, is a clear win on the ratio, with a ratio IAE of 13.50 against 321.60 min and a largest error of 0.163 against 0.596, for 0.243 against 0.185 of trim

travel and \$312 more margin.

Scenario 9 is a loss in money. The plan holds the doubled makeup inerts at 10.00 %, against 11.42 % under supervision at 1440 min, with an inert IAE of 111.4 against 3453.1 % · min and 56.6 t more ammonia, but purging 16 039 against 11 547 kmol of gas costs \$4 493 of margin. The loss is the price of the 10 % setpoint rather than of the controller. Supervised with a fast inert loop, $K_c = 10$ and $\tau_I = 60$ min, which lies outside the shipped tuning, also ends the day at 10.00 % and earns \$26 753, \$246 more than the plan, with 79.1 against 46.0 percent of purge travel, and the shipped loop is cheaper only because it is still 0.52 % above its setpoint after two days.

Scenario 10 is a loss in inerts. The plan holds the pressure better, with an IAE of 708.8 against 1680.3 bar · min and \$425 more margin, and the lowest pressure is the same 140.6 bar at 48.25 min under both, before any plan has closed the purge. Its inert IAE is 609.2 against 92.3 % · min and its purge travel 117.8 against 8.3 percent, because while the makeup valve is at its stop, from 50.25 to 400.25 min, the makeup ceiling and the hold close the purge, to 20.2 % at 90 min, trading inerts, which can wait five hours, for pressure, which cannot. A replay outside the application, through the basic loops, of the supervised purge with the plan's purge substituted only while the makeup is at its stop gives an inert IAE of 645.1 and a pressure IAE of 715.0, so the closure during the cut is the whole trade.

The graded challenge keeps its scoring and is given here for information only. Switching only the structure does not meet it. The plan moves only the purge and the trim, and the weak fast tunings still lose the converter in the front-end cut, at 334.6 min against 302.7 min for the shipped line-up. The loop then spends 548.4 min above the 172 bar design pressure, peaking at 228.8 bar. With the suggested settings the plan meets the target with a worst pressure error of 0.06 bar, 10.1 % inerts and a ratio of 3.00, and it breaks no unscored limit, with a peak pressure of 160.9 bar and no minute over temperature, over the design pressure or over 16 % inerts. Its pressure IAE is 862.8 against 370.7 bar · min and its margin \$1 753 lower, because inerts held near 10 % mean more purge and more makeup during the cut.

2.7 Technology in Practice

What Measures It

Figure 2.6 draws the loops of Section 2.5 with their devices. Loop pressure and the hot spot are read continuously. The two slow variables come from one chromatograph on the loop gas, which separates the species of a sample in time, so its cycle takes minutes and some dead time is inevitable [24]. GC-101 reports every 15 min and holds its last result in between, and the model's single well-mixed loop leaves out the transport delay of a real sample line (Section 2.8).

> **TECHNOLOGY** **Sensor: The Loop Gas Chromatograph**
>
> A process chromatograph injects a small fixed sample into a long column, where each species travels at its own speed, so the detector at the end need only tell the carrier gas from anything else. A thermal conductivity detector senses the passing gas by the cooling of a heated element, and the other common type, the flame ionization detector, cannot see ammonia. Most process units recalibrate themselves from bottles of calibration gas on a schedule. The sample valve matters most for accuracy: if the injected amount is not repeatable, the readings change from cycle to cycle while the gas does not, and a leaking valve shifts the baseline or corrupts the analysis [24].

What Moves It

Four handles move the loop: the makeup valve, the purge valve, the quench valve and the trim sent to the front end, which sets the ratio of the gas the reformer delivers. The synthesis gas is raised to loop pressure by centrifugal compressors, usually driven by steam turbines on steam raised from the plant's surplus heat, and the catalyst is divided into layers cooled between them, either directly by

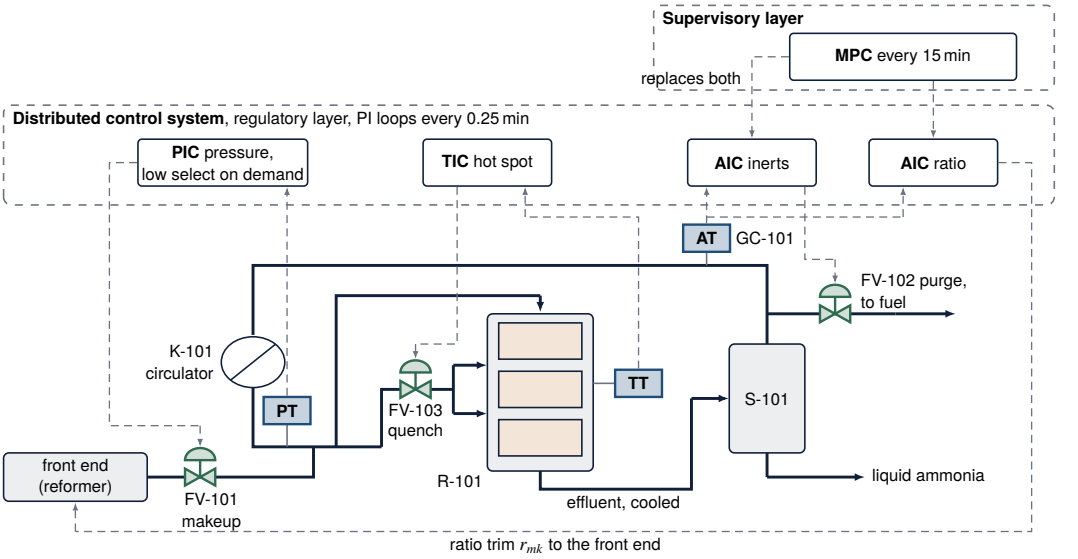


Figure 2.6 The loops of the synthesis loop with their hardware. Lettered boxes are sensing devices: the loop pressure transmitter (PT), the bed temperature element (TT) and the loop gas chromatograph (AT), which reports the inerts and the ratio together. Valve symbols are final elements, and dashed arrows are signals. Notice that one analyzer feeds both slow loops, that the predictive controller in the supervisory layer replaces those two loops on the same two handles, and that the operator's makeup demand passes a low selector against the pressure override. The exchanger, the cooler and the circulator's drive are left out. Not to scale, and not a piping diagram.

cold synthesis gas or by raising steam [25]. The model keeps the direct form, one quench valve split evenly between two bed inlets, and a circulator with no map, which cannot surge.

> **TECHNOLOGY** **Final Element: A Control Valve With a Positioner**
>
> A positioner compares the stem position of a pneumatic control valve with the controller's signal and adjusts the actuator air until the two agree, a cascade inside the valve. A smart positioner also measures the stem position, the actuator and supply air pressures and the temperature. Without a positioner, or after poor maintenance, friction makes the stem move in stair steps and ignore small reversals of the signal. On loss of air an air-to-open valve fails closed, and the risk of the process should decide which way each valve fails [24]. The model's actuator lags of 0.25 to 0.40 min stand in for this hardware.

Where the Controller Runs

A plant runs such loops in a distributed control system, in which a supervisory level oversees local controllers and sends them set points, loops are sometimes nested, and cycle times range from milliseconds to minutes [26]. The app executes its four PI loops every 0.25 min and its predictive controller every 15 min, on each new analysis, so both sit inside that range. The PI loops and the low selector belong to the regulatory layer. The plan belongs to the supervisory layer above it and writes to the purge and the trim, the same two handles the slow loops hold when it is not in service.

IN PRACTICE **Purge, Startup and the Long Run**

A 2007 European reference document on ammonia plants describes the practice the model leaves out. A continuous purge keeps argon and methane from accumulating, and its size holds the loop at roughly 10 to 15 percent inerts. A converter is brought up to temperature with a fired startup heater, and a full startup takes one or two days. Plants run for long periods, and shutdowns for technical failure average 5.7 a year [25]. The design point here, 10.07 % inerts, sits inside that band, and the converter of scenario 3 stays out because the model has no such heater.

What Is Modernizing It

Three developments bear on this loop. Recent designs recover the hydrogen from the purge by cryogenic separation, membranes or pressure swing adsorption and return it to the loop, where older designs burned the purge as reformer fuel [25]. Recovery changes the price of the purge that Table 2.6 trades against production, which is why the model leaves it out. Model-based advanced control with weighted and hierarchical optimization was reported in service on an ammonia plant in 2004 [25], the industrial counterpart of the plan of Section 2.5.4. And a loop fed with hydrogen from electrolysis will have to change its rate with the renewable supply, which calls for plants that start, stop and adjust production quickly [27]. The open challenge after the Limitations takes this up.

2.8 Limitations

The model is not a design tool and no number in it should be used to size anything. The kinetics are one fitted group rather than a mechanism, the equilibrium of (2.4) is applied to mole fractions at a pressure where fugacity coefficients are not unity, and the separator is a correlation with a factor on it. The three-bed converter cannot show an axial profile, so it cannot show where a hot spot actually sits or how deactivation moves it. The loop is one well-mixed volume, so the transport delay a real analyzer sees between the converter and the sample point is absent, and the measured composition responses here are therefore slightly optimistic. There is no compressor map, so nothing in the application can surge. Startup, catalyst reduction and the fired heater are outside the model, which means a blown-out converter in this application stays out and the recovery that a real plant performs cannot be practiced. Economics are three illustrative prices, rather than a market quotation or a complete plant cost model. At the shipped \$2.50 per kmol, the design feedstock costs about \$322 per tonne against \$420 of revenue. Hydrogen recovery, purge fuel credit, upstream production costs and capital are absent, so the economic optimum is an exercise in the stated assumptions. Disturbances arrive as commanded steps and timed events where a plant drifts. Fixing the integration step, the controller period and the random seed reproduces one trajectory on every machine, and no distributed control system is that repeatable. Five features survived all of that trimming, the recycle amplification of (2.3), the two steady states of (2.9), the two time scales, the sampled composition measurement and the selector of (2.17), and each one is present because it changes the controller that has to be built. Section 2.3.7 names where each remaining assumption stops holding.

The plan's model is linear about a single design point, with the fast loops closed at their shipped tuning, so it leaves out the weaker effect of a large purge move (0.866 % of inerts for a 5 percent opening against its 0.945 %), a retuned fast loop, as in the challenge, and a makeup valve at its stop, which only the priced ceiling and the measured hold stand in for. It prices the purge at the declared gas price whatever the slider says, and it holds whatever inert setpoint it is given, which is why it gives up \$4 493 of margin on scenario 9.

OPEN CHALLENGE **A Loop That Follows the Wind**

Ammonia made from hydrogen split out of water with renewable electricity must follow an intermittent supply. A 2020 review finds that processes with chains of high-pressure compressors, heat integration and sensitive catalysts cannot leave steady state without large stores of hydrogen and electricity [27]. A 2018 study of

a three-bed autothermal converter found its widest steady range in the inert fraction and the hydrogen to nitrogen ratio, the two slow variables of this chapter, and noted that a pilot plant had been studied only at steady state. Outside that range the converter goes out or overheats its catalyst [28]. Progress would be a loop that changes rate on demand, proven in dynamic operation, without losing its converter.

2.9 Discussion

The organizational split around this plant matches the technical one. The converter belongs to the reaction engineer, the front end belongs to a different unit and often a different department, and the two variables that decide whether the loop is economic, the inert level and the hydrogen to nitrogen ratio, belong to nobody in particular and are trended by nobody in particular. Both of them are five-hour variables measured every fifteen minutes and neither shows on the pressure trace, and a loop left for a month at the 1.80 ratio of scenario 6 gives away 40 t/d, which is 1200 tonnes of ammonia that no alarm reported. The quantity worth putting in front of an operator is not on the standard display at all, and it is the purge-corrected efficiency, 128.6 kmol of syngas per tonne against a stoichiometric 117.4, with the difference attributable to the purge.

The documented model supports methods well beyond the four implemented loops and the linear plan over two of their handles. Recycle amplification is exactly the structure that steady-state real-time optimization, Section 1.4.7 and Chapter 11, exists to handle, because the optimum purge rate of Table 2.6 moves with the makeup inert content, the catalyst activity and the price of hydrogen, and it is flat enough that an operator will not find it by hand. A nonlinear predictive controller on the same equations would carry the sintering limit, the design pressure and the quench travel as explicit constraints and would coordinate the makeup and purge valves as one two-input problem rather than as two loops with a five-hour interaction between them [9]. Moving horizon estimation on the fourteen states would recover the catalyst activity and the exchanger condition from the temperature and pressure records, which turns the invisible margin loss of scenario 11 into a soft sensor that reports before the second disturbance arrives rather than after. The equations transfer directly to an algebraic modeling language, with the six hidden constants of Table 2.2 recast as estimated parameters.

2.10 Conclusion

This chapter has used an ammonia synthesis loop as an interactive laboratory for the behavior a recycle imposes on a plant. The simulator integrates fourteen states and solves rather than assumes its design point, 149.97 bar and a 505.2 °C hot spot with 10.07 % inerts at a hydrogen to nitrogen ratio of 3.000, making 677 t/d from 1.01 kmol/s of makeup against 6.00 kmol/s of circulation, at which line-up the largest relative state derivative is 9.6×10^{-12}. Single-pass extent is 0.0386, the recycle ratio is 5.95, and the purge is 8.4 % of the makeup, which fixes the inert level at 11.90 % in the purge stream and the composition time constant at 5.1 h. A front end delivering hydrogen and nitrogen 5 % below three to one settles the loop at 1.80, an amplification of eight that costs 40 t/d against the 678 t/d the plant was making, and it takes 14 h to arrive with the loop pressure never leaving 150.0 bar. The same error high is amplified 17 times, because hydrogen leaves through a purge less readily than nitrogen does. Handing the makeup valve to the operator rather than to the pressure controller changes nothing about the plant and everything about the result, since a catalyst aged to 0.75 costs 17 t/d of production under one pairing and 2.3 bar of pressure margin under the other. A hot-spot setpoint seven degrees lower is worth 7.0 % of production, and the quench travel that buys it is the same travel a converter needs in order not to go out, and scenario 11 spends that travel four and a half hours before it is needed. A linear model predictive controller over the purge and the trim holds the inerts within 0.1 % of setpoint from 192 min after a two-drift upset that two PI loops have not corrected a day later, for $924 of margin in purged hydrogen. Further work would add a hydrogen recovery unit on the purge, a compressor map with a surge line, a fired startup heater so that an extinguished converter can be recovered, and an axial converter profile. The synthesis loop is

an accessible browser platform for learning what a recycle does to a small error and what it costs to choose the wrong inventory.

KEY POINTS

- An ammonia loop sends most of its gas around again, so a small, steady error in the makeup gas grows inside the loop over hours.
- The inert level and its five-hour time constant are both set by the small purge.
- Holding pressure turns a catalyst loss into lost production, and holding production turns it into lost pressure margin.
- The quench valve position, not the hot spot, shows the margin left.
- A predictive controller on purge and trim corrects slow drifts sooner than two PI loops and pays in purged hydrogen.

Where else this applies. Volume 6, Chapter 7 shows the same amplification, and Volume 6, Chapter 8 a gas recycle with an inert and a purge.

References

[1] U.S. Geological Survey, *Mineral Commodity Summaries 2025*. Reston, VA: U.S. Geological Survey, 2025, nitrogen (fixed)–ammonia, pp. 128–129.

[2] W. L. Luyben, "Snowball effects in reactor/separator processes with recycle," *Industrial & Engineering Chemistry Research*, vol. 33, no. 2, pp. 299–305, 1994.

[3] R. M. Price, P. R. Lyman, and C. Georgakis, "Throughput manipulation in plantwide control structures," *Industrial & Engineering Chemistry Research*, vol. 33, no. 5, pp. 1197–1207, 1994.

[4] E. M. B. Aske and S. Skogestad, "Consistent inventory control," *Industrial & Engineering Chemistry Research*, vol. 48, no. 24, pp. 10 892–10 902, 2009.

[5] J. J. Downs and S. Skogestad, "An industrial and academic perspective on plantwide control," *Annual Reviews in Control*, vol. 35, no. 1, pp. 99–110, 2011.

[6] R. F. Baddour, P. L. T. Brian, B. A. Logeais, and J. P. Eymery, "Steady-state simulation of an ammonia synthesis converter," *Chemical Engineering Science*, vol. 20, no. 4, pp. 281–292, 1965.

[7] P. L. T. Brian, R. F. Baddour, and J. P. Eymery, "Transient behaviour of an ammonia synthesis reactor," *Chemical Engineering Science*, vol. 20, no. 4, pp. 297–310, 1965.

[8] L. D. Gaines, "Ammonia synthesis loop variables investigated by steady-state simulation," *Chemical Engineering Science*, vol. 34, no. 1, pp. 37–50, 1979.

[9] A. Araújo and S. Skogestad, "Control structure design for the ammonia synthesis process," *Computers & Chemical Engineering*, vol. 32, no. 12, pp. 2920–2932, 2008.

[10] V. Smil, *Enriching the Earth: Fritz Haber, Carl Bosch, and the Transformation of World Food Production*. Cambridge, MA: MIT Press, 2001.

[11] J. W. Erisman, M. A. Sutton, J. Galloway, Z. Klimont, and W. Winiwarter, "How a century of ammonia synthesis changed the world," *Nature Geoscience*, vol. 1, no. 10, pp. 636–639, 2008.

[12] M. Appl, "Ammonia, 2. Production processes," in *Ullmann's Encyclopedia of Industrial Chemistry*. Weinheim, Germany: Wiley-VCH, 2011.

[13] F. Haber, "The synthesis of ammonia from its elements," Nobel Lecture, 2 June 1920, 1920, nobelprize.org/uploads/2018/06/haber-lecture.pdf, read November 2026.

[14] C. van Heerden, "Autothermic processes," *Industrial & Engineering Chemistry*, vol. 45, no. 6, pp. 1242–1247, 1953.

[15] J. C. Morud and S. Skogestad, "Analysis of instability in an industrial ammonia reactor," *AIChE Journal*, vol. 44, no. 4, pp. 888–895, 1998.

[16] C. Bosch, "The development of the chemical high pressure method during the establishment of the new ammonia industry," Nobel Lecture, 21 May 1932, 1932, nobelprize.org/uploads/2018/06/bosch-lecture.pdf, read November 2026.

[17] D. R. MacFarlane, P. V. Cherepanov, J. Choi, B. H. R. Suryanto, R. Y. Hodgetts, J. M. Bakker, F. M. Ferrero Vallana, and A. N. Simonov, "A roadmap to the ammonia economy," *Joule*, vol. 4, no. 6, pp. 1186–1205, 2020.

[18] L. J. Gillespie and J. A. Beattie, "The thermodynamic treatment of chemical equilibria in systems composed of real gases. I. an approximate equation for the mass action function applied to the existing data on the Haber equilibrium," *Physical Review*, vol. 36, no. 4, pp. 743–753, 1930.

[19] D. C. Dyson and J. M. Simon, "Kinetic expression with diffusion correction for ammonia synthesis on industrial catalyst," *Industrial & Engineering Chemistry Fundamentals*, vol. 7, no. 4, pp. 605–610, 1968.

[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] D. Q. Mayne, J. B. Rawlings, C. V. Rao, and P. O. M. Scokaert, "Constrained model predictive control: Stability and optimality," *Automatica*, vol. 36, no. 6, pp. 789–814, 2000.

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

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

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

[25] European Commission, "Reference document on best available techniques for the manufacture of large volume inorganic chemicals: Ammonia, acids and fertilisers," European Commission, Integrated Pollution Prevention and Control, Tech. Rep., Aug. 2007, eippcb.jrc.ec.europa.eu, read November 2026.

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

[27] C. Smith, A. K. Hill, and L. Torrente-Murciano, "Current and future role of Haber–Bosch ammonia in a carbon-free energy landscape," *Energy & Environmental Science*, vol. 13, no. 2, pp. 331–344, 2020.

[28] I. I. Cheema and U. Krewer, "Operating envelope of Haber–Bosch process design for power-to-ammonia," *RSC Advances*, vol. 8, no. 61, pp. 34 926–34 936, 2018.

2.11 Exercises

2.1 ⌕ Concept check · 2 min

The loop runs on the basic loops with the pressure controller on the makeup valve and the hot-spot controller on the quench. The catalyst activity falls from 1.00 to 0.75. What does the operator see on the display once the loops have settled?

(a) The loop pressure climbs until the high limit takes the makeup valve back
(b) The hot spot falls below its setpoint and stays there with the quench shut
(c) Production falls while the loop pressure stays on its setpoint
(d) The loop inerts climb, because the converter consumes less of the makeup

2.2 ⌕ Concept check · 1 min

In scenario 3 the quench is opened to 22 % with the plant in Manual, and the converter goes out. At 90 min the quench returns to exactly 10 %, the value at which the converter ran lit at the start. What happens next?

(a) The converter relights within minutes, as the exchanger restores its inlet temperature
(b) The converter relights slowly, on the five-hour clock of the loop composition
(c) The converter relights once the pressure has risen far enough to speed up the reaction
(d) The converter stays out, because a dead bed cannot heat its own feed back to ignition

2.3 ⌕ Concept check · 2 min

The loop pressure answers a makeup valve move in minutes, while the loop inert level answers a purge valve move over hours. What sets the inert time scale?

(a) The loop gas holdup divided by the small purge flow that is the only exit for inerts
(b) The 15 min cycle of the loop gas analyzer that measures the inerts
(c) The 300 min integral time of the inert controller in its shipped tuning
(d) The heat capacity of the catalyst beds, which delays every change in conversion

2.4 Theory and calculation · 8 min

At the design line-up the loop pressure is 149.97 bar, the purge valve is at 37.4 % with capacity $k_P = 0.01850\,\mathrm{kmol/s}$ per $\sqrt{\mathrm{bar}}$, and the makeup is 1.0082 kmol/s carrying 1.00 % inerts. The loop holds $V = 300\,\mathrm{m^3}$ of gas with $z = 1.12$ at a mean $T_L = 40\,^{\circ}\mathrm{C}$. Use $R = 8.314\,\mathrm{J/mol/K}$.

(a) Find the purge flow and the share of the makeup it removes. *Checkpoint:* 0.08473 kmol/s, 8.4 % of the makeup.

(b) Find the loop holdup N. *Checkpoint:* 1543 kmol.

(c) Find the steady inert fraction in the purge stream and the inert time constant, in seconds and in hours. *Checkpoint:* 11.90 %, about 18 210 s or 5.06 h.

(d) The purge valve is cut to 20 % with the pressure held at 150 bar. Compute $N/\dot{n}_P$ at the new purge. Scenario 7 measures 63 % of the change in 5.12 h. Why is the loop faster than this ratio? *Checkpoint:* About 9.5 h. The makeup falls with production.

2.5 Theory and calculation · 6 min

The equilibrium constant of Equation (2.4) is used in Equation (2.5) with the pressure written as the model writes it, $P = 150$.

(a) Evaluate K_p at 400 °C and at 430 °C and find the factor by which it falls over those 30 °C. *Checkpoint:* 1.291×10^{-2} and 8.671×10^{-3}, a factor of 1.49.

(b) For a loop gas with $r = 3$ and $y_I = 0.10$ at 150 bar, find κ, u and the equilibrium ammonia fraction at 450 °C. *Checkpoint:* $\kappa = 0.3295$, $u = 0.7262$, $y_A^{eq} = 0.1738$.

(c) Repeat (b) at 505 °C, the design hot spot, and say what the two results mean for a bed that runs hotter. *Checkpoint:* $\kappa = 0.1770$, $y_A^{eq} = 0.1103$.

2.6 Theory and calculation · 4 min

At the design line-up the circulator moves $F = 6.00\,\mathrm{kmol/s}$, the converter extent is $x_{tot} = 0.03862$ per mole of feed, the separator returns $y_{A,sep} = 0.0379$ in its overhead gas, the purge is 0.08473 kmol/s drawn from that gas, and the makeup is 1.0082 kmol/s. Ammonia is 17.03 kg/kmol.

(a) Find the ammonia made in the converter and the liquid ammonia drawn from the separator at steady state, and express the product in tonnes per day. *Checkpoint:* 0.4634 kmol/s made, 0.4602 kmol/s drawn, 677 t/d.

(b) Find the recycle ratio, the circulator flow over the makeup as the app reports it, and the makeup gas consumed per tonne of product. *Checkpoint:* 5.95 and 128.6 kmol/t.

2.7 App challenge · 8 min

Load scenario 6, A small error at the front end, and change no setting. At 30 min the front-end ratio drifts 0.15 low, from 3.00 to 2.85, with the pressure and hot-spot loops in service. Run at 60x with the default seed. Pause between 720 and 740 min, set Front-end ratio drift to +0.15, and run 240 min more. Export the CSV file, whose time column is in hours.

(a) Read the true loop ratio, column `ratio`, and the production at 720 min, and the largest departure of the true pressure, column `press_bar`, from 150 bar between 30 and 720 min. *Checkpoint:* 1.853 within 0.005, 640 t/d within 1 t/d, and under 0.05 bar.

(b) Taking the amplification as the relative loop ratio error divided by the relative makeup error, find its value at 720 min, and compare it with the settled value of 1.804 that Figure 2.2 reports. *Checkpoint:* 7.6 at 720 min, against 8.0 settled.

(c) Read the loop ratio 240 min after the drift was set to +0.15. How far is it from the settled value of Figure 2.2 for a makeup ratio of 3.15? *Checkpoint:* 3.40 within 0.01, against 5.58 settled.

2.8 App challenge · 8 min

Load scenario 11, Five years of catalyst, and change no setting. The catalyst activity falls to 0.65 at 30 min and the feed and effluent exchanger loses 8 % of its effectiveness at 300 min. Run to 400 min at 60x with the default seed and export the CSV file, whose time column is in hours.

(a) Read the quench command, column `quench_cmd_pct`, and the production at 299.75 min, the last row before the exchanger loss. *Checkpoint:* 8.45 % within 0.02 %, and 650 t/d.

(b) Find the first row after 300 min at which the quench command reaches zero and the first at which the hot spot, column `hotspot_C`, is below 380 °C. *Checkpoint:* 304.25 min and 306.0 min.

(c) Press Reset. Once the activity event has fired at 30 min, and before 300 min, set Catalyst activity to 1.00. Run to 360 min. Does the converter stay lit? Read the converter inlet, column `tin_C`, and the quench command at 360 min. *Checkpoint:* It stays lit, with the inlet at 392.0 °C within 0.3 °C and the quench at zero.

2.9 App challenge · 8 min

Load scenario 12, Five hours ahead: Supervised or MPC, and change no setting. At 60 min the methane slip rises by half and the front end turns 0.10 nitrogen heavy. Run to 1440 min at 180x with the default seed under the MPC structure the scenario loads, and export the CSV file, whose time column is in hours. Then press Reset, set Control structure to Supervised before pressing Run, run the same 1440 min, and export again.

(a) From the MPC run, read the peak of `inert_pct` and its time, and the purge command, column `purge_cmd_pct`, at 76 min. *Checkpoint:* 10.36 % at 81.5 min, and 56.2 %.

(b) From the supervised run, read the peak of `inert_pct` and its time, and its value at 1440 min. *Checkpoint:* 11.98 % at 333 min, and 10.73 %.

(c) Read the IAE of the Loop inerts row of the scorecard and the column `metric_profit` at 1440 min for both runs. What does the plan's better control cost? *Checkpoint:* IAE 56.6 against 1843.7, within 1, and $38 475 against $39 399, $924 less margin.

2.10 Application area · 3 min

A plant manager has to choose which variable the makeup valve holds. Plant A sells to a customer who takes a fixed tonnage every day. Plant B sends its product into a pipeline that takes whatever it makes. Which pairing suits each, and what must the operators of each watch that the other need not?

2.11 Application area · 4 min

In scenario 9 the methane slip from the reformer doubles the makeup inerts. The inert controller holds its setpoint by opening the purge, and once settled the gas per tonne rises from 128.7 to 141.9 kmol/t at about 677 t/d. At the shipped syngas price of $2.50 per kmol, estimate what the dirtier gas costs per day. Where does the fix belong, and what can the loop engineer do in the meantime?

2.12 Review the agent · 5 min

An AI assistant reviews Table 2.6 and writes: "Production rises monotonically with the purge valve, from 484 to 792 t/d. Opening the purge from the design 37.4 % to 85 % adds 17 % to production with the loop still at 150 bar, so the plant should run with the purge at 85 %." Find the flaw and say how you would check it.