INTERACTIVE CONTROL ENGINEERING
Principles and Applications

Volume 11 Control in Petroleum Refining

CHAPTER 13

A Claus Sulfur Recovery Simulator for Feedforward, Dead Time and an Optimum Turned Into a Zero

Pages 308–328

Edited by John D. Hedengren

How to cite this chapter. "A Claus Sulfur Recovery Simulator for Feedforward, Dead Time and an Optimum Turned Into a Zero," in *Control in Petroleum Refining*, John D. Hedengren, Ed., vol. 11 of *Interactive Control Engineering: Principles and Applications*. Advanced Process Solutions, LLC, 2026, ch. 13, pp. 308–328.

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

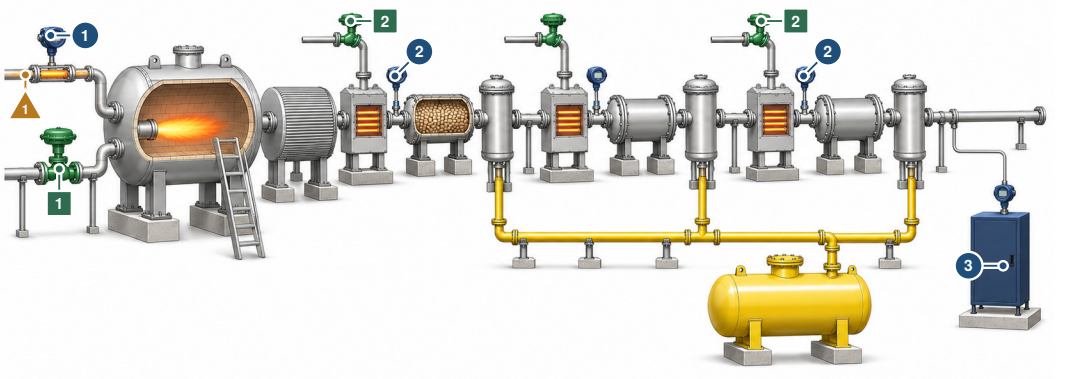
The app for this chapter
https://apmonitor.com/apps/sulfur-recovery/

The volume online: contents, other chapters and citation records
https://apmonitor.com/apps/books/v11-refining/#sulfur-recovery

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

CHAPTER 13

A Claus Sulfur Recovery Simulator for Feedforward, Dead Time and an Optimum Turned Into a Zero



The controller moves 1 combustion air valve, set by the feedforward and the trim
2 reheater valves that set the three bed inlet temperatures

The controller reads 1 acid gas flow meter 2 bed inlet temperatures 3 tail gas analyzer that reads the air demand

What disturbs it 1 the acid gas: its flow, its hydrogen sulfide and a hydrocarbon nobody measures

Notice the long sample line at the far right: the analyzer that sets the air reads gas that left the furnace minutes earlier, after three condensers. Illustration, not to scale.

IN PLAIN TERMS **Just Enough Air to Make Sulfur**

The main idea. A sulfur plant burns a third of a poisonous gas, hydrogen sulfide. The rest reacts with what burned to make liquid sulfur. An instrument at the end of the plant reads a value that is zero when the air is just right and shows which way it is wrong.

The challenge. That reading arrives minutes after the air has changed, so no tuning can shrink the first error when the gas changes. Measuring the gas coming in and setting the air from it ahead of time does most of the work.

Why it matters. One percent too little air can push a plant over its pollution permit.

THREE WAYS TO READ THIS CHAPTER

Explore	12 min	This page, the Did you know? box in the Introduction, the History boxes in Section 13.2, the Try it boxes and the second Did you know? box in Section 13.5, and Key points in Section 13.9. No equations.
Study	45 min	The whole chapter in order, stopping at the three pause points, then Exercises 1 to 12.
Practice	20 min	This page, Section 13.5 with Figure 13.2, Section 13.6 and Section 13.7 with its Open challenge box.

The control question Sulfur recovery has a maximum at one air rate, and a maximum has no sign, yet an ordinary PI loop holds the air of every Claus unit. What measurement makes that possible, and with the analyzer minutes downstream of the furnace, what do tuning, feedforward and a Smith predictor each do to the error a feed change leaves?

Learning objectives

- Explain how the air demand $y_{H_2S} - 2y_{SO_2}$ turns the recovery maximum into a zero a PI loop can hold, and price a percent of air from the stoichiometry.
- Compute the stoichiometric air of a feed, and the ratio the trim must find after an unmeasured change or cannot reach after a measured one.
- Show by a run that the first peak after a composition step is the same for every tuning, and say what shortens it.
- Compare no feedforward, flow feedforward and flow and analysis on the same disturbances, and read the result from the scorecard and tiles.
- Judge what a Smith predictor on the trim can and cannot remove, and why a drifting analyzer must be checked from outside the loop.

Prerequisites Sections 1.4.6, 1.4.7, 1.6.2 and 1.6.4, with Section 1.4.4 for the beds. Chapter 8 for a late analyzer in this volume.

The app https://apmonitor.com/apps/sulfur-recovery/

See also Volume 3, Chapter 4 for an oxygen trim on a fuel-air ratio and Volume 3, Chapter 6 for a feedforward with a slow ratio trim around an optimum. Volume 5, Chapter 3 for a loop on a slow analyzer and Volume 5, Chapter 5 for a Smith predictor built from a process model.

13.1 Introduction

A sulfur plant is judged on the second decimal place of its recovery, and the recovery has a maximum. One third of the hydrogen sulfide is burned to sulfur dioxide in a reaction furnace and the rest reacts with it, two to one, to make sulfur. Any other air rate leaves one of the two in excess and that excess cannot react with anything. Measured on the unit modeled here, recovery is 97.619 % at an air ratio of 1.000, 96.640 % at 0.990 and 96.640 % again at 1.020: one percent short of air costs exactly what two percent of surplus costs, because too little air strands hydrogen sulfide one for one while too much strands only half as much sulfur dioxide.

> **DID YOU KNOW? Sulfur From 109 Plants in 2005**
>
> In 2005 the United States produced 8.8 million metric tons of elemental sulfur, recovered at 109 petroleum refineries, natural gas processing plants and coking plants. About 90 % of the sulfur consumed was used as sulfuric acid [1]. About 90 to 95 % of recovered sulfur is made by the Claus process [2].

A maximum has no sign, and a fixed-sign controller cannot sit on one. That is the argument Volume 12, Chapter 6 makes at length, where a chloride optimum defeats an otherwise competent PID. Yet the air on every Claus unit in the world is on a PI loop and it works. The reason is not a better controller. It is a better measurement. The tail gas analyzer reports the air demand

$$D = y_{H_2S} - 2y_{SO_2}$$

which is zero at the recovery maximum, positive when the furnace is short of air and negative when it has too much. The optimization has been turned into a regulation by changing what is measured, and that move is worth more than any amount of search machinery. Section 1.4.6 sets this measurement beside the volume's other analyzers, and Section 1.4.7 among the few refinery optima that lie inside the range rather than on a limit.

The misconception the simulator corrects is that an optimum is an optimization problem. Faced with an objective that has a maximum, the first question is not what optimizer to build but whether

some measurable quantity is zero there: a difference, a ratio against a stoichiometric number, a derivative somebody can instrument. Where one exists the problem collapses to a loop.

Finding it does not make the problem easy. The analyzer sits at the far end of the plant behind three condensers, 3.2 min of transport and cycle away from the furnace it is controlling. Measured, a step in acid gas composition produces a peak air demand error of 2.67 mol% whether the loop is sluggish, correctly tuned or oscillating, because the error arrives with the step and for the next 3.2 min the controller has no information that anything has changed. On a dead-time-dominant loop, tuning changes how long an error lasts and not how big it gets [3]. The things that shorten the peak are feedforward and a faster measurement, which is engineering rather than tuning.

This chapter works with an interactive Claus unit whose stoichiometry, analyzer delay and dew point correlation are hidden, so that the reader discovers the air demand relationship, measures the division of labor between feedforward and feedback rather than taking it on trust, and watches what a drifting analyzer does. The Claus chemistry, feedforward and air demand control are all established practice. Section 13.2 gives the history of the process and of its control. Section 13.3 writes out the three stage Claus train, with its furnace, its equilibrium limited converters, its sulfur dew point fouling and its analyzer faults, in the detail needed to rebuild it. Section 13.4 describes the controllers, from the feedforward and its trim to the dew point override and the optional Smith predictor. Section 13.5 reports the scenario sequence, in which the same disturbance is answered about 168 times better by a ratio station than by any tuning, and the graded challenge. Section 13.6 sets the hardware of a real plant beside the model. Sections 13.7 and 13.8 state the limitations and the three results that carry past sulfur plants, and Section 13.9 sums up.

13.2 Background

Acid gas from an amine regenerator is mostly hydrogen sulfide and carbon dioxide, and something must be done with it. The Claus process burns part of the hydrogen sulfide and reacts the remainder with the sulfur dioxide produced, first thermally in a reaction furnace and then over alumina in a series of catalytic converters with a condenser after each. Modern plants recover between 95 and 98 percent of the sulfur before tail gas treating, and the balance leaves as sulfur dioxide from a stack under a permit.

> ⧖ HISTORY **1883 and 1937: One Catalytic Step Becomes Two**
>
> Carl Friedrich Claus patented the original process in 1883. It partially oxidized hydrogen sulfide over a catalyst at 200 to 350 °C, making sulfur and water in a single step. In 1937 the process was divided into two parts. A flame burned one third of the hydrogen sulfide completely to sulfur dioxide, and that sulfur dioxide then reacted over a catalyst with the remaining hydrogen sulfide at 170 to 370 °C [4]. Most modern plants that treat acid gas are built on this modified process [4], and the furnace and three converters of this chapter follow its two parts.

Three technical literatures bear on the control problem. The reaction furnace itself is the least understood part of the unit and has been reviewed as such, because the chemistry at 1200 °C in under a second involves far more than the two nominal reactions [5]. The catalytic section is an equilibrium problem, and process modeling of it has been used to evaluate modifications aimed at raising recovery and reducing emissions [6]. Alternative configurations, including high temperature air combustion and routes that recover hydrogen as well as sulfur, continue to be proposed [7, 8]. The environmental driver behind all of it is direct: sulfur dioxide from refinery sulfur plants is a measured contributor to local air quality and the subject of its own assessments [9].

Control developed around the analyzer. Before continuous tail gas analysis the air was set from the acid gas flow with a ratio a supervisor adjusted after a laboratory result, and the unit ran deliberately rich because the penalty for being lean was larger. The air demand analyzer changed that. Once the difference between the hydrogen sulfide and twice the sulfur dioxide could be measured continuously,

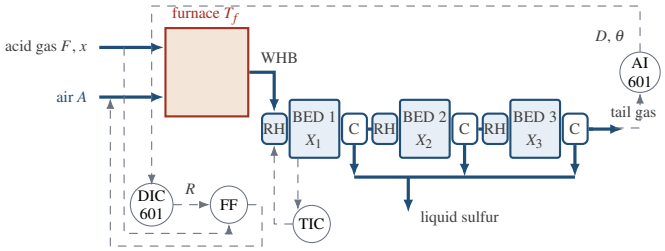


Figure 13.1 Claus unit U-601. Acid gas F of hydrogen sulfide fraction x burns with air A in a reaction furnace at T_f, followed by a waste heat boiler WHB. Three reheaters RH and catalytic beds of conversion X_i each precede a condenser C that removes liquid sulfur. AI-601 measures the tail gas air demand D after a dead time θ. The feedforward FF computes the air from the acid gas flow and composition and multiplies it by the ratio R that the trim controller DIC-601 sets. A temperature loop TIC, drawn for bed 1, holds each bed inlet on its reheater.

the whole problem became a trim loop on a stoichiometric feedforward, and the structure has not changed since. The oxygen trim on the air-to-fuel ratio of the fired heater in Volume 3, Chapter 4 has the same structure. What has changed is the feedforward. Plants with an acid gas chromatograph compute the air demand from the measured composition rather than from a commissioning number, and plants without one do not.

> ⧖ HISTORY **1957 and 1977: The Analyzer Takes Over the Air**
>
> A patent filed in 1957 found the ratios of a sulfur plant hard to hold, largely because the composition of its feed gas varied. Its remedy was to analyze a sample of the process gas for the ratio of hydrogen sulfide to sulfur dioxide and let the result reset the flow controller of the furnace air. The analyzer was a chromatograph: a timer injected the sample into a packed column at intervals, and the two gases left the column at different times [10]. A patent filed in 1977 had the structure of this chapter. A PI controller on the tail gas ratio multiplied the nominal air for a known volume of acid gas, and that product was multiplied by the measured acid gas flow to give the total air [11].

The dead time was and remains the difficulty. A sample point closer to the furnace shortens it and costs a hot, wet, corrosive sample line that plugs, so the trade is an instrumentation decision rather than a control one. Dead-time compensators exist and are surveyed [3], and on a plant whose disturbances are mostly composition rather than setpoint changes the feedforward is usually the better investment. Above the regulatory layer, the constraint and economic arbitration that a large unit needs is provided by multivariable control [12], as in the layers of Section 1.4.7, and the plantwide question of which unit owns which specification is the same one the fractionation train of Volume 10, Chapter 10 poses [13].

13.3 Process Model

Unit U-601 takes 8.0 kmol/min of acid gas at 85 mol% hydrogen sulfide with one percent hydrocarbon, and 17.33 kmol/min of air for a ratio of exactly one. The flame sits at 1226 °C, the thermal stage makes 63.5 % of the sulfur, and three converters with inlets at 245, 210 and 200 °C take the recovery to 97.619 %, leaving 7.46 t/d of sulfur to the stack against a permit of 9.0 t/d. Figure 13.1 shows the arrangement, and Table 13.1 classifies its variables.

13.3.1 Stoichiometry, and Where It Leaves the Plant

With R the ratio of air fed to air demanded, the furnace burns $R/3$ of the hydrogen sulfide and leaves $1 - R/3$. The Claus reaction wants those two in a ratio of two to one, which happens at $R = 1$ and

Table 13.1 Classification of Variables at the Design Operating Point

Symbol	Role	Nominal	Range
D (AI-601)	CV, analyzer	0.00 mol%	−1–1.5 SP
$T_{in,1..3}$	CV, continuous	245, 210, 200 °C	180–275
R	MV	1.000	0.90–1.10
A	MV	17.33 kmol/min	0–40
reheaters	MV	245, 210, 200 °C	150–285
F	DV, measured	8.0 kmol/min	3–13
x	DV, analyzed	85.0 mol%	45–95
hydrocarbon	DV, unmeasured	1.00 mol%	0–4
ammonia	DV, unmeasured	0 mol%	0–6
analyzer bias	DV, unmeasured	0	±1.5
T_f, $T_{out,i}$, activity	states	—	—

nowhere else, so whichever is in excess is sulfur that cannot react, Equation (13.1).

$$\text{loss} = \begin{cases} 1 - R, & R < 1 \\ \tfrac{1}{2}(R - 1), & R > 1 \end{cases} \tag{13.1}$$

The asymmetry in Equation (13.1) is worth noticing, because every operator's instinct to run rich follows from it, before adding that excess oxygen sulfates the catalyst.

The air the furnace actually demands is a stoichiometric calculation, Equation (13.2).

$$A_{st} = \frac{F}{0.21}\left(\tfrac{1}{2}y_{\mathrm{H_2S}} + \tfrac{3}{4}y_{\mathrm{NH_3}} + 3\,y_{HC}\right) \tag{13.2}$$

Only the first two terms of that are measured on most plants, and the third never is.

13.3.2 Equilibrium Limited Beds

Each converter is adiabatic and equilibrium limited, which makes it the opposite of the hydrocracker beds of Chapter 7. More conversion is more heat and more heat is a worse equilibrium, so the bed limits itself, Equations (13.3) and (13.4).

$$T_{out} = T_{in} + K\,\dot{n}\,X, \qquad X = X_{eq}\left(\tfrac{T_{in}+T_{out}}{2}\right) a \tag{13.3}$$

$$X_{eq}(T) = X_0 - k_X\,(T - 200\,^\circ\mathrm{C}) \tag{13.4}$$

with $\dot{n}$ the reactable sulfur reaching the bed in kmol/min and a the bed activity. Because $\mathrm{d}X_{eq}/\mathrm{d}T < 0$ the residual has exactly one root and it is stable, so a Claus converter cannot run away. In the terms of Section 1.4.4, its runaway index is negative and its steady gain from inlet to outlet is less than one.

13.3.3 The Sulfur Dew Point

What a converter can do is drown. Below the sulfur dew point liquid sulfur condenses in the pores, Equation (13.5), and the activity decays, Equation (13.6).

$$T_{dew} = T_0 + k\,\log_{10}\left(\frac{y_S}{0.01}\right) \tag{13.5}$$

$$\frac{\mathrm{d}a}{\mathrm{d}t} = \begin{cases} -k_f\,(T_{dew} - T_{in})\,a, & T_{in} < T_{dew} \\ +k_r\,(1 - a), & \text{otherwise} \end{cases} \tag{13.6}$$

Recovering the activity takes a hot gas soak measured in days, so the damage is delayed and large while the benefit of running colder is immediate and small.

Table 13.2 Model Parameters

Symbol	Meaning	Value	Class
thermal	furnace conversion, net of COS	65.5 %	adjustable
X_0	equilibrium at 200 °C	0.800	adjustable
k_X	equilibrium lost per °C	0.0030	adjustable
K	rise per kmol/min converted	45.0 °C	adjustable
T_0, k	dew point correlation	182, 30.0	adjustable
k_f	activity lost below the dew point	0.055/(°C h)	adjustable
COS made	sulfur as COS and CS_2	3.0 %	adjustable
T_h	half hydrolysis, bed 1 outlet	305 °C	adjustable
θ	analyzer dead time	3.2 min	adjustable
T_{NH_3}	ammonia destruction	1250 °C	adjustable
permit	sulfur to the stack	9.0 t/d	visible

13.3.4 Carbonyl Sulfide and the First Bed

Carbon dioxide ballast in the acid gas makes carbonyl sulfide and carbon disulfide in the furnace, and only the first bed hydrolyzes them back, Equation (13.7).

$$h = \frac{1}{1 + \exp[-s\,(T_{out,1} - T_h)]} \tag{13.7}$$

That single term is why the first bed wants to run hot while the other two want to run cold, and it is the subject of scenario 10.

13.3.5 Assumptions

The thermal stage converts a fixed fraction of whatever the stoichiometry allows. Real furnace chemistry depends on residence time, burner mixing and configuration, and is the least settled part of the unit [5].

All sulfur species are lumped as sulfur atoms. The real vapor is an equilibrium between several allotropes and the dew point depends on which.

Each converter is one lumped adiabatic equilibrium stage, so there is no axial profile and no mass transfer limit.

Condensers are perfect and there is no tail gas treating unit, which is what takes a 97 percent Claus unit to 99.9 percent overall [6].

Carbonyl sulfide and carbon disulfide are one lumped species hydrolyzed only in the first bed. Real hydrolysis also occurs in the second bed on the right catalyst.

The analyzer dead time is a pure transport delay followed by a first-order average of 0.8 min, with no dispersion in the sample line, and the instrument drifts but never fails outright. Real sample systems plug, and a plugged sample system reports the last thing it saw.

13.4 Controller Design

13.4.1 Velocity Form PID

Every loop uses the shared velocity-form PID of Section 1.6.2, with the derivative on the measurement, executed every $\Delta t = 0.05$ min. Equations (13.8) to (13.11) restate it in the symbols used below.

$$e_k = a\,(\mathrm{SP}_k - \mathrm{PV}_k) \tag{13.8}$$

$$P_k = K_c e_k, \qquad D_k = -K_c \tau_D\, a\, \frac{\mathrm{PV}_k^f - \mathrm{PV}_{k-1}^f}{\Delta t} \tag{13.9}$$

$$s_k = s_{k-1} + (P_k - P_{k-1}) + (D_k - D_{k-1})$$

$$c_k = s_k + I_{k-1} \tag{13.10}$$

$$
\begin{aligned}
u_k &= \mathrm{clamp}(c_k, u_{\min}, u_{\max}) \\
I_k &= I_{k-1} + \frac{K_c}{\tau_I} e_k \Delta t + \frac{\Delta t}{\tau_t} (u_k - c_k)
\end{aligned}
\tag{13.11}
$$

where a is the controller direction, s_k accumulates the proportional and derivative increments without limit, c_k is the output the algorithm demands, and the integral used at scan k was accumulated through scan $k-1$. These are Equations (1.16) to (1.19) with $s_k = S_k$ and $c_k = v_k$. The integral state I also carries the starting output, so the bias term u_b is zero. Anti-windup is the back-calculation term of Equation (13.11), with tracking time $\tau_t = \tau_I$ [14]. Every loop here runs with $\tau_D = 0$, so in practice all four are PI. The trim DIC-601 has $a = -1$ and limits of 0.90 and 1.10 on the ratio, and the reheater loops have $a = +1$ and limits of 150 and 285 °C on the reheater outlets. A loop in manual sets $I_k = u_{\text{man}} - s_k$, so the switch to automatic is bumpless. The 0 to 40 kmol/min limit on the air flow is applied after the trim's ratio has been multiplied into the feedforward and is not seen by the trim's anti-windup, but it is reached only near the top of the flow and composition ranges. Integration is fixed-step fourth-order Runge-Kutta at 0.01 min, on the engine clock of Section 1.6.1.

13.4.2 Feedforward, and What It Is Allowed to Know

The air flow is computed from Equation (13.2) and multiplied by the ratio the trim controller sets. Three settings differ in what the calculation uses. *None* holds the selected base air flow, which the feedback trim still multiplies. The base does not follow the trimmed output, avoiding a repeated multiplication on each scan. *Flow* uses the measured acid gas rate with the composition the unit was commissioned with. *Flow and analysis* uses the composition an acid gas chromatograph reports through a first-order lag of 2.0 min. In none of the three does the calculation measure the ammonia or the hydrocarbon. It uses commissioning values for both, set in the app's advanced feedforward settings, because no plant has an on-line analyzer for amine carryover. Everything the feedforward is wrong about is what the trim loop is for.

13.4.3 The Trim Loop and Its Authority

DIC-601 has $a = -1$, direct acting in the convention of Section 1.6.2, because more air is more sulfur dioxide and therefore a lower air demand. The process gain is about −36 mol% of air demand per unit of ratio, so the suggested gain of 0.022 is a loop gain of 0.8, which is about as far as a loop with 3.2 min of dead time and integral action can be pushed [3, 15].

The trim range is 10 % either side on purpose, because an analyzer that fails high should not be allowed to swing the air by a quarter. The same limit that stops the analyzer doing harm stops the trim doing good, which is the subject of scenario 8.

🕮 **Pause.** Try Exercise 13.1 before reading on.

13.4.4 The Dew Point Override

A one-directional override raises a bed inlet setpoint to the sulfur dew point plus a guard of 10 °C whenever the setpoint is closer to the dew point than that, up to the 275 °C top of the setpoint range. It can only push up, because there is no amount of conversion worth a wet catalyst bed, and because the damage is delayed and the benefit is immediate.

13.4.5 Smith Compensation of the Air Demand Trim

The optional Smith predictor compensates the modeled response to commanded air while retaining the stoichiometric feedforward and the three reheater loops. Its feedback variable remains signed air demand D in mol%. A Smith correction separates the command response from the measured residual [3, 16]. The nominal command-to-analyzer model is Equation (13.12).

$$
\widehat{G}_0(s) = \frac{K_S}{(0.25s+1)(0.8s+1)}, \qquad \widehat{G}_\theta(s) = e^{-\widehat{\theta} s} \widehat{G}_0(s)
\tag{13.12}
$$

Here s has units of inverse minutes, $K_S = -36.3728042448$ mol% per unit ratio, and $\widehat{\theta} = 3.2$ min. The 0.25 min air lag precedes transport, and the 0.8 min analyzer lag follows it. The delay-free branch retains both lags.

The gain is a centered finite secant at ratios 1 ± 10^{-5}, using declared chemistry at 8 kmol/min, 85 mol% hydrogen sulfide, 1 mol% hydrocarbon, zero ammonia, unit activities and bed inlets 245, 210 and 200 °C. Its limiting derivative is −36.3729014, a difference of 2.67 parts per million. Applied to in-run ratio steps of −0.005 and +0.005, the executed frozen model has root mean square analyzer errors of 0.000489 and 0.000 937 mol% over 40 min. A separate optimized second-order-plus-dead-time fit gives $K = -36.3251$, lags 0.800793 and 0.247 454 min, effective delay 3.227 694 min, and RMS 0.000 009 76 mol%. The effective fitted delay includes sample-grid phase. That smaller residual describes the fitted model, not the implemented one.

The model advances on the original 0.05 min scan using the command actually delivered over the preceding interval. Its input u is applied air divided by the controller's own feedforward base, or the held base when feedforward is off. It is distinct from the physical ratio $R = A/A_{st}$. For model air state x_k, initial ratio u_0, and incremental demand v_k, the normal-scan recursions are Equation (13.13).

$$\begin{aligned} x_k &= u_{k-1} + (x_{k-1} - u_{k-1})e^{-\Delta t/0.25} \\ v_k &= K_S(x_k - u_0), \qquad b = e^{-\Delta t/0.8} \\ \widehat{D}_{0,k} &= b\widehat{D}_{0,k-1} + (1-b)v_k \\ \widehat{D}_{\theta,k} &= b\widehat{D}_{\theta,k-1} + (1-b)v(t_k - \widehat{\theta}) \\ D_{c,k} &= D_{m,k} + \widehat{D}_{0,k} - \widehat{D}_{\theta,k} \end{aligned} \tag{13.13}$$

The timestamped history interpolates v and supplies zero incremental demand before its first sample. The two analyzer states initialize to zero, while x starts at u_0. Here D_m is the actual analyzer reading and D_c is the corrected feedback. The residual $D_m - \widehat{D}_\theta$ carries unknown loads into the PI after detection. Neither the hidden feed nor the actual stoichiometric requirement enters the predictor.

A separate PI with the same action $a = -1$ uses D_c in the velocity recursion already given, with the same 0.90–1.10 trim bounds. An aggregate-lag internal model control (IMC) tuning approximation [17] sets $\tau_I = 0.25 + 0.8 = 1.05$ min and $K_c = 1.05/(|K_S|\lambda)$. The fixed $\lambda = 2$ min gives $K_c = 0.014433861$ ratio per mol%. This approximates the two lags for tuning and does not exactly invert Equation (13.12). Integral action removes measured offset and cannot identify sensor bias.

While Smith acts, ordinary PI tracks delivered trim using actual analyzer history. Smith entry and recovery reseed its own corrected-measurement history and hold the first move. The independent 0 kmol/min to 40 kmol/min air bound remains authoritative. A nonfinite packet, state or prediction first holds delivered trim within its bounds, then transfers to tracked ordinary PI with a recorded fault. Numerically valid model mismatch does not trigger this transfer. If previous delivered trim is below 0.90 under severe air clipping, its bounded hold cannot preserve the same command. All histories are local to a run and are rebuilt on restart.

13.5 Results

Thirteen scenarios run from a manual ratio sweep to a graded challenge. Table 13.3 lists them, and Section 1.6.4 describes their completion rules, the scorecard and the exported log.

13.5.1 What a Percent of Air Costs

Table 13.4 reports the air ratio swept either side of stoichiometric.

One percent short of air costs exactly what two percent of surplus costs, and the air demand measurement passes through zero at the maximum with a gain of about −36 mol% per unit of ratio. The permit is 9.0 t/d, so one percent of air is the difference between a plant inside its license and a plant outside it. Run as scenario 1, the steps to 0.990 at 20 min and to 1.020 at 60 min settle on

Table 13.3 Scenario Sequence

#	Scenario	Teaching point
1	Two to one	what a percent of air costs
2	A measurement with a sign	optimum into a zero
3	Three minutes of nothing	dead time and peak error
4	Tuning against dead time	the limit cycle
5	Feedforward on the flow	measured disturbances
6	Feedforward on the composition	analyzed disturbances
7	The disturbance nobody measures	what feedback is for
8	When the trim runs out of room	authority
9	The analyzer that lies	sensor drift
10	Three beds, three windows	middle against end
11	A wet bed	slow constraints
12	A trim that remembers what it sent	Smith compensation
13	Challenge	graded run

Table 13.4 Recovery and Stack Against Air Ratio

R	Recovery (%)	Stack (t/d)	D (mol%)
0.980	95.661	13.60	+0.736
0.990	96.640	10.53	+0.366
1.000	97.619	7.46	0.000
1.010	97.130	8.99	−0.360
1.020	96.640	10.53	−0.712

Table 13.5 Three Tunings on One Composition Step

K_c	τ_I (min)	Peak $\lvert D\rvert$	IAE	Ratio band
0.010	20	2.67	153	settles slowly
0.022	12	2.67	42	settles
0.040	8	2.67	335	cycles, 0.064 wide

the values of Table 13.4, and the stack permit alarm comes up 0.19 min after each of them. It clears only when the stack falls below 7.65 t/d, 85 % of the permit, which happens for 0.18 min as the ratio passes 1.000 at 60 min and then not until 100.71 min, after the return to stoichiometric.

13.5.2 Tuning Does Not Reduce the Peak

Table 13.5 reports three tunings against a step in acid gas composition from 85 to 78 mol% at 20 min with the flow feedforward running. The runs are scenario 3 with the gains set before the run, and the last row is also scenario 4 as loaded.

The peak is 2.67 mol% for all three, and it arrives with the step, at 20.01 min, because the composition moves the true air demand at once. It then holds to within 0.01 mol% until 23.3 min, because the analyzer reading first moves at 23.26 min and until then no gain and no integral time has anything to act on. A faster measurement or a composition feedforward does not lower that first value either, but both shorten it. One minute after the step the error is 1.70 mol% with the dead time set to zero and 1.87 mol% with the composition analyzer in the feedforward, against 2.67 mol% for every tuning at 3.2 min. The integrated absolute error in mol% min, taken over the 380 min following the step at the 0.01 min model step, separates the tunings: correct tuning is 3.6 times better than sluggish tuning and eight times better than tuning that cycles. The sluggish loop is still 0.010 mol% from setpoint at 400 min. The gain of 0.040 never settles at all, swinging the air ratio between 0.900, the trim's lower stop, and 0.964, a band of 0.064 that is about seven percent of the air.

Table 13.6 Integrated Absolute Air Demand Error Over 380 min Following Each Disturbance

Disturbance	None	Flow	Flow and analysis
Flow 8.0 kmol/min to 9.6 kmol/min	1227	7.3	7.3
Composition swings, 4 h	—	163	69
Hydrocarbon 1.0 mol% to 1.6 mol%	—	21.6	21.6

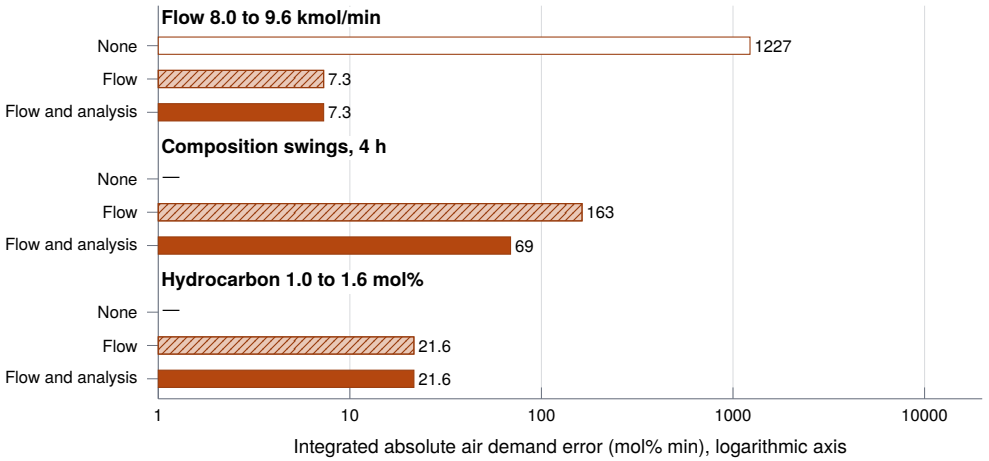


Figure 13.2 The integrated absolute air demand errors of Table 13.6 on a logarithmic axis, by disturbance and feedforward setting: none (white bar), flow (hatched) and flow and analysis (solid). A dash marks a case the table does not report. Notice that any feedforward takes the flow step from 1227 to 7.3, that the composition analyzer helps only on the composition swings, and that no setting changes the hydrocarbon slug, which only the feedback trim answers.

Pause. Try Exercise 13.2 before reading on.

13.5.3 Feedforward Against Feedback

Table 13.6 reports the integrated absolute error in mol% min over the same window for three disturbances and three feedforward settings, from scenarios 5, 6 and 7 with the feedforward chosen before the run, and Figure 13.2 draws it on a logarithmic axis. Each window starts at the first event of its scenario, and the error is integrated at the model step. The application scorecard, which integrates its 0.05 min log, agrees except on the brief flow transient, where it reads 7.2.

About 168 times less integrated error on the flow step. With no feedforward the base air remains at 17.33 kmol/min and the trim multiplies it up to 19.06 kmol/min at its high stop of 1.10. Its ten percent authority cannot supply the twenty percent increase needed by the load. The true air ratio settles at 0.9165 and recovery at 89.8 %, with 9.71 t emitted over the six hours after the step. The flow ratio brings the air up within about a minute, but the analyzer reports that minute 3.2 min later, and the trim answers an error the feedforward has already removed by taking the ratio to 1.032. The stack, over the permit from the step to 21.1 min, goes back over it four more times before 38.4 min. The permit alarm raised at 20.02 min therefore latches, because the stack settles at 7.73 t/d, inside the permit but above the 7.65 t/d at which the alarm resets.

The composition swings separate the second and third settings. Over the 380 min from the first of six moves, four hours of an amine unit taking the acid gas between 79 and 88 mol% hydrogen sulfide and two of settling, the flow feedforward alone leaves an integrated error of 163 and 2.67 t of sulfur to the stack, and adding the composition analyzer gives 69 and 2.26 t. Four hundred kilograms, about one and a half tonnes a day at that rate, bought with one chromatograph.

The hydrocarbon slug is invisible to both, because nothing on the plant measures it. Six tenths of a percent more hydrocarbon raises the true air demand by four percent, and the trim finds every bit of it, walking the ratio to 1.040 and recovering 97.62 %. That is the honest division of labor: feedforward handles what is measured and feedback handles what is real.

> **TRY IT IN THE APP** **7 · The disturbance nobody measures**
>
> Load the case and press Run. At 20 min the hydrocarbon in the acid gas rises from 1.0 to 1.6 mol%, and the feedforward cannot see it. Watch the trend of the air ratio the trim asks for: it walks to 1.040, and the recovery comes back to 97.62 %. Feedback finds the four percent of extra air demand that nothing on the plant measures.

13.5.4 The Analyzer That Lies

A tail gas analyzer drifting 0.5 mol% high leads the trim to add air until the analyzer reads its setpoint, which puts the true demand at −0.500 mol%. The plant runs 1.4 % rich with the air demand indication exactly on target, recovery falls to 96.94 % and the stack goes to 9.60 t/d, over the 9.0 t/d permit. The air demand alarm stays quiet, because the measured error is near zero and even the true error of 0.5 mol% is inside its 1.2 mol% band. The one warning is the stack permit alarm, raised at 21.77 min from the stack figure, which the analyzer does not influence. It latches: the stack drops back under the permit twice while the trim settles, from 25.0 to 30.4 min and from 34.4 to 38.1 min, but the alarm clears only below 7.65 t/d and so stays up for the rest of the run, with the air demand reading sitting on its false target throughout. That failure mode is why air demand analyzers are the most carefully maintained instruments on a sulfur plant, and none of the maintenance is control engineering [18].

Pause. Try Exercise 13.3 before reading on.

> **TRY IT IN THE APP** **9 · The analyzer that lies**
>
> Load the case and press Run. At 20 min the tail gas analyzer starts to read 0.5 mol% high, and every loop stays in Auto. Watch the air demand bar under the plant: the open marker, the analyzer, stays on two to one while the filled marker, the true demand, settles near −0.5 mol%. The stack permit alarm comes up at 21.77 min and stays.

13.5.5 Three Beds, Three Windows

Measured at steady state on bed 3, recovery is 97.79 % with the inlet at 185 °C and 97.27 % at 230, with the dew point at 180. Colder is better, exactly as the equilibrium says, with a phase change as the floor. Measured on bed 1, recovery is 96.94 % at 210 °C of inlet and 98.24 % at 275, which the equilibrium flatly denies. Hydrolysis of carbonyl sulfide and carbon disulfide runs from 19 to 76 % across that range, and everything bed 1 gives up in conversion is recovered by beds 2 and 3, while nothing is downstream of bed 3 to recover its losses. Scenario 10 makes the same moves on a smaller scale, bed 1 to 210 and then 265 °C and bed 3 to 225 °C, and settles at 96.94, 98.07 and 97.74 %. A stage in the middle of a train has different economics from the stage at the end.

13.5.6 The Graded Challenge

The challenge runs fifteen hours with two flow changes, three composition swings and a hydrocarbon slug, starting with the feedforward off and the trim tuned at half its correct gain. It is met when the air demand is held within 0.25 mol% of setpoint through the last third, every bed stays above its dew point, and the shift emits less than five tonnes of sulfur. The run as loaded makes 199.5 t of sulfur and emits 13.29 t. Configured as suggested it makes 208 t and emits 4.5 t, and the stack permit alarm still comes up at every one of the six events and clears between them.

The verdict is judged at 900 min on the true air demand at every logged sample from 594 min, the last third of the run, with the model parameters and disturbances as given. Changing either voids

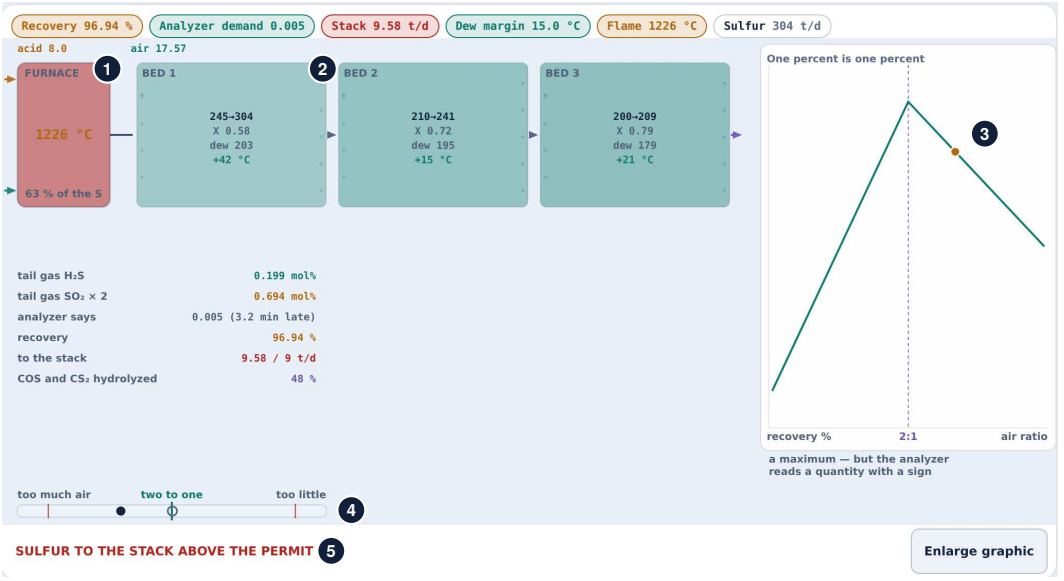


Figure 13.3 The Claus unit application running scenario 9, the analyzer bias case, at $t = 120$ min, 100 min after the +0.5 mol% drift, with every loop in automatic, cropped to the plant view with numbered callouts. (1) The furnace and (2) the three catalytic beds are drawn on one line, each bed carrying its inlet and outlet temperatures, its conversion, its sulfur dew point and its margin above it. (3) Beside them a sketch of recovery against air ratio marks the operating point at 1.014 and 96.94 %, on a curve recalculated from the current feed, bed temperatures and catalyst activities using the same model as the operating point. (4) The bar under the scene is the air demand, with two to one marked and the alarm limits either side. The true value of –0.495 mol% is the filled marker and the analyzer reading of 0.005 mol% the open one, on the optimum. (5) The latched stack permit alarm is on the status line below.

the attempt. As loaded, with no feedforward, $K_c = 0.010$, $\tau_I = 20$ min, bed inlets at 225, 205 and 205 °C and the override off, the worst error is 0.332 mol%, above the required band, and 13.29 t is emitted, although the dew point margin never falls under 11.1 °C. With the suggested settings, flow and analysis feedforward, $K_c = 0.022$, $\tau_I = 12$ min, inlets at 258, 208 and 194 °C and the override on, the worst error is 0.144 mol%, with 0.106 mol% to spare, the margin never falls under 11.6 °C, and 4.48 t is emitted. The same settings chosen before Run or carried through Reset give identical trajectories. Over seven other analyzer noise seeds the worst error ranges from 0.134 to 0.180 mol% and the emission from 4.45 to 4.49 t, all passing.

13.5.7 Measured Comparison of PI and Smith Trim

Scenario 12, the command-memory case, runs 300 min on flow feedforward. Hydrogen sulfide changes from 85 to 82 mol% at 20 min and returns at 150 min. Unmeasured hydrocarbon changes from 1.0 to 1.3 mol% at 90 min and returns at 210 min. The model contains none of these future events. Each paired experiment uses seed 20260922, identical instruments, noise, feedforward and event times. PI retains each scenario's ordinary tuning, including the unchanged suggested challenge settings. Smith uses $\lambda = 2$ and $\widehat{\theta} = 3.2$ throughout, except for explicitly labeled model-delay experiments. No tuning is adjusted by case.

Table 13.7 integrates true signed-demand error from $t = 0$ through each stated duration at the 0.05 min logging interval. These windows include startup, unlike the earlier event-relative tables. Travel sums absolute requested-ratio increments. Rail time uses the scorecard tolerance of 2×10^{-7}

Table 13.7 Fixed-Setting Trim Comparisons Including Startup

Case	Trim	Duration (min)	Peak $\lvert D \rvert$ (mol%)	IAE (mol% min)	Travel (ratio)	Rail (min)	Permit (min)	Crossings	Emitted (t)
Command-memory case	PI	300	1.209	52.239	0.281	0.00	60.65	10	1.778
	Smith	300	1.177	19.953	0.105	0.00	22.70	4	1.635
	Stronger PI	300	1.177	24.125	0.204	0.00	25.55	6	1.654
	Same-fast PI	300	3.849	541.683	7.333	32.65	262.70	44	3.942
	Smith, $\hat{\theta} = 1.6$	300	1.177	25.497	0.202	0.00	29.10	7	1.660
Composition step	PI	400	2.666	41.958	0.196	0.00	29.85	2	2.141
	Smith	400	2.666	14.549	0.077	0.00	7.45	1	2.063
Flow step, FF off	PI	400	6.710	1226.928	0.100	375.95	380.00	1	10.346
	Smith	400	6.708	1227.458	0.100	375.55	380.00	1	10.349
Composition swings	PI	420	3.581	162.604	1.023	0.00	171.70	17	2.882
	Smith	420	3.750	70.667	0.374	0.00	43.35	6	2.479
Hydrocarbon step	PI	400	1.366	21.578	0.097	0.00	37.90	2	2.202
	Smith	400	1.366	7.482	0.040	0.00	7.90	1	2.117
Stale-composition FF	PI	400	6.052	1198.411	0.100	375.75	380.00	1	5.391
	Smith	400	6.052	1198.979	0.100	375.30	380.00	1	5.393
Analyzer bias	PI	400	0.500	184.011	0.038	0.00	369.10	3	2.619
	Smith	400	0.500	188.985	0.014	0.00	377.55	1	2.634
Actual delay 6.4 min	PI	300	1.229	60.633	0.423	0.00	68.80	10	1.815
	Smith	300	1.174	65.042	0.367	0.00	76.05	12	1.833
	Smith, $\hat{\theta} = 1.6$	300	4.026	505.368	3.590	29.75	260.10	25	3.803
	Smith, $\hat{\theta} = 6.4$	300	1.177	32.232	0.105	0.00	35.55	4	1.690
Challenge, loaded	PI	900	5.915	1159.652	0.286	458.50	842.65	29	13.294
	Smith	900	5.890	1094.027	0.367	499.25	842.60	73	12.997
Challenge, suggested	PI	900	5.487	93.929	1.414	0.00	63.30	16	4.481
	Smith	900	5.440	66.411	0.815	0.00	37.20	11	4.312
Flow step, flow FF	PI	400	6.183	7.188	0.132	0.00	10.45	5	2.183
	Smith	400	6.183	3.881	0.049	0.00	3.75	2	2.165
Swings, analysis FF	PI	420	3.760	68.618	0.897	0.00	71.05	24	2.469
	Smith	420	3.732	49.312	0.327	0.00	52.75	11	2.385

at either trim bound. Permit time counts samples above the current stack limit, and crossings count upward transitions rather than hysteretic alarm resets. Stack rates and integrated mass are sulfur-equivalent quantities, not species-specific sulfur dioxide measurements. None of these comparisons reaches the independent air cap.

On this case, IAE is 52.239 for ordinary PI and 19.953 for Smith, with emitted sulfur 1.778 and 1.635 t. Permit time is 60.65 against 22.70 min, with 10 against 4 upward crossings. A fixed stronger PI ($K_c = 0.018$, $\tau_I = 4$ min) gives IAE 24.125, so part of the shipped-PI difference is tuning. Ordinary PI at the exact Smith tuning ($K_c = 0.014433861$, $\tau_I = 1.05$ min) gives 541.683 and 32.65 min on a trim rail. Compensation permits that faster integral setting in this experiment, without establishing global optimality.

At 32 min, ordinary PI has true demand -0.611391 and analyzer -0.310293 mol%. Smith has true demand -0.016843, analyzer -0.134111 and corrected feedback -0.029178 mol%, as shown in Figure 13.4. The correction reflects previously commanded air. On the isolated unknown-composition step, PI and Smith have exactly the same sampled peak of 2.665 970 mol%. Different whole-sequence peaks in Table 13.7 include different residuals from earlier events and do not demonstrate anticipation of unknown feed.

Smith loses when the feedforward leaves insufficient authority: IAE is 1227.458 against 1226.928 with feedforward off and 1198.979 against 1198.411 with stale composition. Analyzer bias gives 188.985 against 184.011. With actual transport 6.4 min, the nominal predictor gives 65.042 against PI's 60.633. An explicit matched model delay of 6.4 min lowers that IAE to 32.232, while a wrong 1.6 min model gives 505.368 and 29.75 min on a ratio rail. These losses remain numerically valid predictions, without fallback.

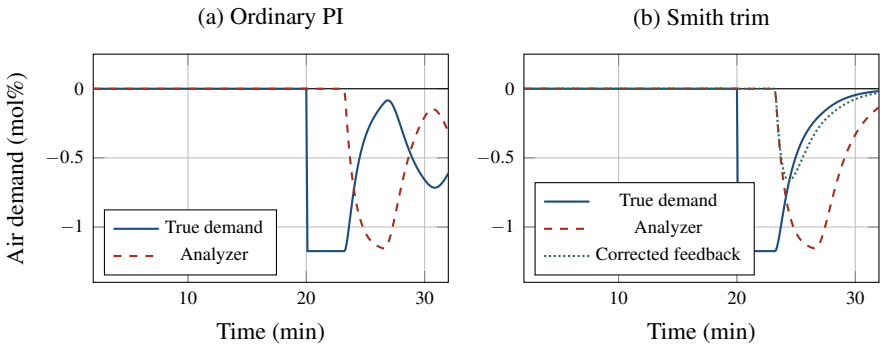


Figure 13.4 First unknown composition change under ordinary PI and Smith trim. Both panels show the same 30 min window ending at 32 min. True demand is the solid line, the analyzer the dashed line and, in panel (b), the corrected feedback the dotted line. Corrected feedback combines the analyzer with two incremental model branches and is distinct from true demand.

> ✦ DID YOU KNOW? **The Predictor Does Not Help Against a Lying Analyzer**
>
> From the simulator (Table 13.7). With the analyzer biased, the Smith trim integrates an error of 188.985 against 184.011 mol% min for ordinary PI, and spends 377.55 of the 400 min over the permit against 369.10. The predictor corrects for the air it sent, but the analyzer's bias stays in the residual it feeds back.

Across the 61 measured variants, active Smith runs execute 6000–18000 model/PI calls and reject zero samples. On the Node host of the measurements, the case's mean module time is 0.908 µs per call, including both model branches and PI. This timing depends on host load, JIT compilation and clock resolution, and does not measure plant integration or rendering.

13.6 Technology in Practice

Figure 13.5 draws the air demand trim, its feedforward and one of the bed inlet loops of Section 13.4 with the hardware a refinery installs.

What Measures It

The trim reads one instrument at the far end of the plant, and the feedforward reads the acid gas flow and, in the flow and analysis setting, a chromatograph on the acid gas. Chromatographs are relatively slow, and a 1989 patent records why a feed analysis corrected by the tail gas still leaves the plant off ratio: it does not see the hydrocarbons until the gas they leave reaches the tail gas analyzer, and the heavier hydrocarbons are usually the ones that upset the air [19]. Scenario 7 is that case.

> TECHNOLOGY **Sensor: The Tail Gas Analyzer**
>
> A chromatograph separates the species of a sample in time, so its result arrives a cycle of minutes after the sample and some dead time is inevitable [20]. A non-dispersive photometer measures a gas continuously by the infrared or ultraviolet light it absorbs, and each compound absorbs in its own pattern [20]. Either kind drifts: analytical instruments drift more than those for pressure, level, temperature or flow, so they are checked with span gases of known concentration, often by a self-calibration system on a schedule [20]. A drift looks calm on a trend. The chapter's analyzer has 3.2 min of transport delay and a 0.8 min average, and it drifts but never fails outright.

What Moves It

The trim moves the air, and the reheaters move the bed inlets. Each catalytic stage of a Claus plant is a reheater, a catalyst chamber and a condenser, and the beds must run above the sulfur dew point so

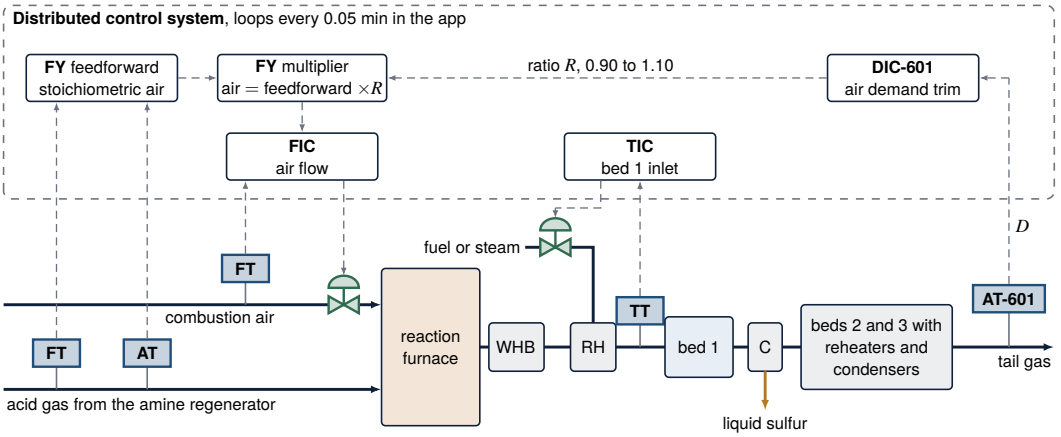


Figure 13.5 The air demand trim, its feedforward and the bed 1 inlet loop with their hardware. Lettered boxes are sensing devices: the acid gas flow meter (FT), the acid gas chromatograph used by the flow and analysis setting (AT), the bed inlet thermocouple (TT) and the tail gas analyzer (AT-601). Valve symbols are final elements, and dashed arrows are signals. Notice the length of the path from the air valve through the furnace and three condensers to the analyzer, which is the 3.2 min of the chapter, and that the trim multiplies the feedforward instead of adding to it. Beds 2 and 3, their loops, the dew point override and the Smith option are left out. Not to scale, and not a piping diagram.

that liquid sulfur does not deactivate the catalyst [2]. Plants reheat by mixing in hot gas bypassed from the waste heat boiler, by steam exchangers, by exchangers against hot gas from a converter, or by fired heaters that burn fuel or acid gas short of air, so that no oxygen reaches the catalyst [7]. The model sets each reheater outlet between 150 and 285 °C and leaves the method open.

> **TECHNOLOGY** **Final Element: Main and Trim Air Valves**
>
> The 1977 patent split the furnace air into a main line and a trim line, each with its own flow controller and control valve. The trim valve took the faster and finer adjustment over the short term, and a slow PI controller moved the main air so that the trim air returned to a chosen steady flow over the long term [11]. A pneumatic valve takes a set position when its air fails, chosen from the risks of the process, and friction in its stem makes it move in steps and deaf to small reversals [20]. The chapter has one air flow, which follows the computed demand through a 0.25 min lag between 0 and 40 kmol/min.

Where the Controller Runs

A refinery runs such loops in a distributed control system, in which a supervisory level oversees local controllers, loops are sometimes nested, and cycle times range from milliseconds to minutes [21]. The app executes every loop each 0.05 min, inside that range. The trim and the three reheater loops are PI blocks, the feedforward is a calculation block whose output the trim's ratio multiplies, and the dew point override raises a setpoint through a high select. The 1977 patent built the same calculation on an analog computer and noted that a digital computer could easily be programmed to do it [11].

> **IN PRACTICE** **The Stack Monitor Sits Outside the Loop**
>
> In the United States, a refinery sulfur plant whose construction or modification began after May 14, 2007, with a capacity above 20 long tons a day, burning ambient air and incinerating its tail gas, must keep its sulfur dioxide below 250 ppm by volume, dry and at zero excess air. It must monitor and record that concentration

continuously, with an oxygen monitor to correct for excess air, and each rolling 12-hour average above the limit is reported as excess emissions [22]. Emissions analyzers are often given daily calibration checks [20]. The stack monitor reads what leaves, not the air demand, so it is the check from outside the loop that the drifting analyzer of scenario 9 calls for.

What Is Modernizing It

Three developments touch this loop. The first is the inferred analyzer. A soft sensor trained on other plant measurements can act as a back-up for a failed hardware sensor, and it needs maintenance of its own because the data drift [18]. One such sensor estimated the hydrogen sulfide in the tail gas of a refinery sulfur unit, designed to replace the analyzer during its maintenance [23]. The second is dead-time compensation. The 1977 patent already compared a delayed analyzer with an analog model of the process run without its dead time [11], the idea of the Smith option of Section 13.4. The third is predictive control, with more than 4600 applications reported and refining the largest share [12], which can arbitrate the constraints a large unit shares with its neighbors.

13.7 Limitations

The thermal stage is a fixed conversion. The reaction furnace is the part of a Claus unit that a review describes as least understood [5], and this model does not attempt it.

There is no tail gas treating unit, which is what takes a modern plant from 97 to 99.9 percent overall [6, 8], so the application teaches the Claus section and not the plant.

Sulfur species are lumped and the dew point correlation is fitted, so the absolute dew point values carry less meaning than their changes.

Ammonia is represented only through its air demand and a furnace temperature alarm. Real ammonia chemistry in a Claus furnace is complicated and dominates the sour water debottlenecking literature.

The analyzer drifts but never plugs. A plugged sample system reports the last good value and looks calm, which is a different and harder fault than the bias implemented here.

The Smith model freezes local gain and both lags at the declared design point. Ratio normalization does not capture changes in the feedforward base, chemistry, catalyst condition or instrument dynamics. Unknown composition is learned through the delayed measured residual, and analyzer bias is preserved in that residual. The authority and bias rows in Table 13.7 show the resulting losses. Agreement of individual model recursions does not establish closed-loop robustness over the full adjustable plant range. The 6.4 min transport experiment illustrates the sensitivity to dead-time error studied in compensator design [16].

Finally, the constraint and economic arbitration a large unit needs [12] is absent, as is the plantwide question of which upstream unit owns the acid gas composition [13].

OPEN CHALLENGE **A Furnace Model Good Enough to Steer By**

The air a furnace needs depends on what burns in it, and no model predicts that reliably. A 2001 thesis found no comprehensive kinetic model of the reaction furnace, and found the equilibrium models in use failing worst when carbon dioxide and hydrocarbons are present [4]. A 2018 review describes detailed schemes of 130 species and more than 1500 reactions, too large and stiff for flow simulation, and a reduced scheme with mean absolute errors of 7.62 and 7.88 % against plant data [5]. The difficulty is a flame zone above 2000 °C lasting milliseconds, followed by up to two seconds without oxygen [5]. Progress would be a model fast and accurate enough to compute the air for a feed whose hydrocarbons change.

13.8 Discussion

Three results generalize past sulfur plants.

The first is the move that makes this optimum controllable. Recovery has a maximum and cannot be held by feedback. The difference between the hydrogen sulfide and twice the sulfur dioxide is zero

there and can. Where an objective has a maximum, the first question is whether some measurable quantity changes sign across it, and finding one converts a search into a loop.

The second is that on a dead-time-dominant loop the peak error is set by the disturbance, and how long it lasts by the delay, rather than by the tuning. Measured here, the peak is 2.67 mol% for every tuning from sluggish through correct to oscillating, and it holds for the 3.2 min of dead time. Effort on such a loop belongs in the measurement and the feedforward, and the useful question about a slow loop is where the delay is and whether it can be moved.

The third is the division of labor. Feedforward answered a flow disturbance about 168 times better than feedback, cut the error on a composition swing from 163 to 69 when it was given an analyzer, and answered a hydrocarbon slug not at all. Every feedforward is a model, every model is incomplete, and the feedback loop underneath stands in for the rest.

13.9 Conclusion

A Claus unit is an optimum that behaves like a setpoint, and the reason is a measurement rather than a controller. In the simulated unit of this chapter recovery is 97.619 % at an air ratio of one, 96.640 % at 0.990 and 96.640 % at 1.020, so one percent short of air costs exactly what two percent of surplus costs, and the stack figure moves from 7.46 t/d to 10.53 t/d against a 9.0 t/d permit for the same one percent. The air demand difference is zero at that maximum and carries a gain of −36 mol% per unit of ratio, which is what makes an ordinary PI loop sufficient. The analyzer is 3.2 min downstream, so the peak error after a composition step is 2.67 mol% for every tuning tried and lasts as long as the delay, and the remedy is feedforward rather than tuning: measured over the 380 min after a twenty percent flow step, the integrated air demand error is 7.3 with a flow ratio station and 1227 without one. A tail gas analyzer drifting half a mole percent takes the plant to 9.60 t/d of stack sulfur with the air demand reading on target, and only the latched stack permit alarm says so. A fuller model would add a tail gas treating unit, a reaction furnace with kinetics rather than a fixed conversion, ammonia chemistry, and a plugged sample system as a second analyzer fault. Even without them, the sulfur plant is a compact case for learning what feedforward is for.

The optional Smith trim reduces IAE from 52.239 to 19.953 on the fixed 300 min comparison, while the stronger ordinary PI reaches 24.125 and model-delay error, sensor bias and exhausted authority retain their measured limitations.

KEY POINTS

- The air demand reading is zero at the best air rate and has a sign on either side, so an ordinary PI loop holds an optimum like a setpoint.
- One percent short of air costs exactly what two percent of surplus costs: recovery is 96.640 % at 0.990 and at 1.020 against 97.619 % at 1.000.
- With the analyzer 3.2 min downstream, the first peak after a composition step is 2.67 mol% for every tuning, and the remedy is feedforward rather than tuning.
- An analyzer drifting half a mole percent takes the stack to 9.60 t/d with the reading on target, and only the stack permit alarm says so.

Where else this applies. Volume 3, Chapter 4 holds an oxygen trim on its air-to-fuel ratio with the same structure, and Volume 12, Chapter 6 shows a fixed-sign controller failing on an optimum that has no such measurement.

References

[1] U.S. Geological Survey, “Sulfur,” Mineral Commodity Summaries 2006, pp. 164–165, Reston, VA, 2006, usgs.gov.

[2] U.S. Environmental Protection Agency, “AP-42, fifth edition, volume I, chapter 8, section 8.13: Sulfur recovery,” Compilation of Air Pollutant Emissions Factors, Research Triangle Park, NC, April 2015, 2015, epa.gov.

[3] J. E. Normey-Rico and E. F. Camacho, “Dead-time compensators: A survey,” *Control Engineering Practice*, vol. 16, no. 4, pp. 407–428, 2008.

[4] A. E. Pollock, “Kinetic modeling of the modified Claus reaction furnace,” Ph.D. dissertation, University of Calgary, Calgary, Alberta, Jun. 2001.

[5] S. Zarei, “A review of investigations on claus reaction furnace of sulfur recovery unit,” *COJ Technical & Scientific Research*, vol. 1, no. 3, 2018.

[6] H. Ghahraloud, M. Farsi, and M. Rahimpour, “Modification of claus sulfur recovery unit by isothermal reactors to decrease sulfur contaminant emission: Process modeling and optimization,” *Chemical Product and Process Modeling*, vol. 14, no. 2, 2018.

[7] M. Sassi and A. K. Gupta, “Sulfur recovery from acid gas using the claus process and high temperature air combustion (hitac) technology,” *American Journal of Environmental Sciences*, vol. 4, no. 5, pp. 502–511, 2008.

[8] A. Taghizadeh Damanabi and F. Bahadori, “A new approach for hydrogen production in claus sulfur recovery process,” *Journal of Sulfur Chemistry*, vol. 40, no. 2, pp. 137–148, 2018.

[9] S. Delfi, M. Mosaferi, A. Khalafi, and K. Zoroufchi Benis, “Sulfur dioxide emissions in iran and environmental impacts of sulfur recovery plant in tabriz oil refinery,” *Environmental Health Engineering and Management*, vol. 5, no. 3, pp. 159–166, 2018.

[10] F. W. Karasek, “Sulfur production,” U.S. Patent 3 026 184, Mar. 20, 1962.

[11] D. E. Smith, “Control system for a sulfur process,” U.S. Patent 4 100 266, Jul. 11, 1978.

[12] 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.

[13] M. L. Luyben, B. D. Tyreus, and W. L. Luyben, “Plantwide control design procedure,” *AIChE Journal*, vol. 43, no. 12, pp. 3161–3174, 1997.

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

[15] S. Skogestad, “Simple analytic rules for model reduction and PID controller tuning,” *Journal of Process Control*, vol. 13, no. 4, pp. 291–309, 2003.

[16] J. E. Normey-Rico and E. F. Camacho, “Unified approach for robust dead-time compensator design,” *Journal of Process Control*, vol. 19, no. 1, pp. 38–47, 2009.

[17] D. E. Rivera, M. Morari, and S. Skogestad, “Internal model control: PID controller design,” *Industrial & Engineering Chemistry Process Design and Development*, vol. 25, no. 1, pp. 252–265, 1986.

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

[19] R. L. Reed and D. A. Holdeman, “Method for controlling Claus furnace with variable hydrocarbon feed composition,” U.S. Patent 4 836 999, Jun. 6, 1989.

[20] 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.

[21] 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.

[22] U.S. Environmental Protection Agency, “Standards of performance for petroleum refineries for which construction, reconstruction, or modification commenced after may 14, 2007,” Code of Federal Regulations, Title 40, Part 60, Subpart Ja, sections 60.102a and 60.106a, edition of July 1, 2025, 2025, govinfo.gov.

[23] A. di Bella, L. Fortuna, S. Graziani, G. Napoli, and M. G. Xibilia, “Soft sensor design for a sulfur recovery unit using genetic algorithms,” in *2007 IEEE International Symposium on Intelligent Signal Processing*, 2007, pp. 1–6.

13.10 Exercises

13.1 ○ Concept check · 2 min

Recovery on the Claus unit has a maximum at an air ratio of 1.000, yet the trim DIC-601 is an ordinary PI loop with a fixed direction. Why can it hold the plant at the maximum?

(a) The recovery curve is flat at its top, so a small air error costs too little for the loop to need a direction

(b) The analyzer reads the air demand, which is zero at the maximum and changes sign across it, so the error has a direction

(c) The trim's limits of 0.90 and 1.10 keep the ratio close enough to the maximum that the sign of the error never matters

(d) The stoichiometric feedforward computes the optimum air exactly, so the PI loop has only analyzer noise left to remove

13.2 Concept check · 2 min

In scenario 3 the acid gas steps from 85 to 78 mol% hydrogen sulfide at 20 min, on flow feedforward. Trim gains of 0.010, 0.022 and 0.040 all leave the same peak air demand error, 2.67 mol%. Why?

(a) The trim sits on its 0.90 stop through the excursion, so none of the three gains is acting

(b) The integral time is too long in all three tunings for the integral to act before the peak

(c) The error arrives with the step and the analyzer reports nothing for about 3.2 min, so no tuning has anything to act on

(d) The flow feedforward cancels whatever part of the error depends on the tuning of the trim

13.3 Concept check · 2 min

In scenario 9 the tail gas analyzer drifts 0.5 mol% high at 20 min, with every loop in Auto and the feedforward on flow and analysis. Which indication on the panel tells the operator that something is wrong?

(a) The air demand alarm, because the true air demand settles 0.5 mol% away from its setpoint

(b) The trim ratio, which runs to its 1.10 stop as the trim adds air to correct the reading

(c) The Analyzer air demand tile, which settles 0.5 mol% away from the setpoint

(d) The stack permit alarm, which is raised from the stack figure that the analyzer does not influence

13.4 Theory and calculation · 7 min

Use the stoichiometric loss of Equation (13.1) and the measured recoveries of Table 13.4.

(a) Compute the loss at air ratios of 0.98, 0.99, 1.01 and 1.02. How many times more does a unit of air shortfall cost than a unit of surplus? *Checkpoint:* 0.02, 0.01, 0.005 and 0.01; twice as much.

(b) From the rows at 1.000 and 0.990, find the drop in recovery per unit of loss. Use it to predict the recoveries at 0.980 and 1.010 and compare them with the table. *Checkpoint:* 97.9 points per unit of loss; 95.661 and 97.130 %, as measured.

(c) The stack carries 7.46 t/d at 97.619 % recovery. Find the sulfur fed, the recovery the 9.0 t/d permit allows at least, and the band of air ratio that keeps the stack inside the permit. *Checkpoint:* About 313 t/d of sulfur, a floor of 97.13 %, and a band from about 0.995 to 1.010.

13.5 Theory and calculation · 7 min

Use the air demand of Equation (13.2). The design feed is 8.0 kmol/min of acid gas at 85 mol% hydrogen sulfide, 1.00 mol% hydrocarbon and no ammonia.

(a) Compute the stoichiometric air at the design point. *Checkpoint:* 17.33 kmol/min.

(b) The hydrocarbon rises to 1.60 mol%, as in scenario 7. The feedforward on flow and analysis still assumes 1.00 mol%. By how much does the demand rise, and what ratio must the trim find? *Checkpoint:* 18.02 kmol/min, 3.96 % more; a ratio of 1.040.

(c) In scenario 5 the flow rises to 9.6 kmol/min with the feedforward off, so the trim multiplies the held base of 17.33 kmol/min and stops at 1.10. Find the true air ratio and the ratio the trim

would need. A shortfall costs 0.979 points of recovery per percent of air, and at the new flow the unit recovers 97.94 % at a ratio of one. Predict the recovery. *Checkpoint:* 0.9165 against 1.200 needed; about 89.76 %.

13.6 Theory and calculation · 7 min

Take the trim's process gain as $K_S = -36.37$ mol% of air demand per unit of ratio.

(a) Compute the loop gain $|K_S|K_c$ for the three tunings of Table 13.5. Which exceeds one, and what does the table show for it? *Checkpoint:* 0.36, 0.80 and 1.45; the gain of 0.040 cycles.

(b) The challenge holds the air demand within 0.25 mol% of setpoint. What band of air ratio is that, and how many points of recovery can sit at its lean and its rich edge? Use 97.9 points per unit of loss. *Checkpoint:* ± 0.0069 of ratio; 0.67 points lean and 0.34 rich.

(c) The Smith trim sets $K_c = 1.05/(|K_S|\lambda)$ and $\tau_I = 1.05$ min. Compute K_c for $\lambda = 2$ and 4 min, and the loop gain at $\lambda = 2$. Why can the Smith PI run this short integral time when ordinary PI at the same setting rails? *Checkpoint:* 0.01443 and 0.00722; loop gain 0.525.

13.7 App challenge · 8 min

Load scenario 1, "1 · Two to one (MANUAL)", as shipped (default seed) and run 150 min at the default speed of 5 min per second, about 30 s. The air ratio steps to 0.990 at 20 min, to 1.020 at 60 min and back to 1.000 at 100 min.

(a) Pause once between 30 and 59 min, once between 75 and 99 min and once after 130 min. Each time read the Recovery, To the stack and Analyzer air demand tiles. *Checkpoint:* 96.64 %, 10.53 t/d, 0.366 mol%; then 96.64, 10.53, −0.712; then 97.62, 7.46, 0.000.

(b) From the two off-design readings, estimate the gain of air demand on air ratio. Is it the same on both sides? *Checkpoint:* About −36 mol% per unit of ratio: −36.6 lean, −35.6 rich.

(c) Watch the alarm banner. After which steps does the stack permit alarm come up, and when does it clear for good? *Checkpoint:* Just after 20 and 60 min; it clears for good about 0.7 min after the return to 1.000.

13.8 App challenge · 8 min

Load scenario 3, "3 · Three minutes of nothing", as shipped (Feedforward on Flow only, Air demand trim gain 0.022, Air demand integral time 12.0 min, default seed). Run 60 min at 5 min per second and export the CSV.

(a) Find the largest |`demand_pct`|, the true air demand, after the step at 20 min, and the logged time at which `demand_meas`, the analyzer, first differs from zero by more than 0.01 mol%. Read the scorecard IAE at 60 min. *Checkpoint:* 2.666 mol% at 20.05 min; the analyzer at 23.30 min; IAE 34.43 (± 0.1).

(b) Pause, press Reset, set the gain to 0.040 and the integral time to 8.0 min, and run 60 min again. Repeat with 0.010 and 20.0 min. Record the peak and the IAE at 60 min each time. *Checkpoint:* Peak 2.666 mol% both times; IAE 48.01 and 66.45 (± 0.1).

(c) Table 13.5 ranks the three tunings 42, 153 and 335 over 380 min. Why does the gain of 0.040 rank second, not last, at 60 min? *Checkpoint:* Its cycle has run only half an hour; the sluggish loop's error is large early.

13.9 App challenge · 9 min

Load scenario 5, "5 · Feedforward on the flow", as shipped (Feedforward None, default seed). The acid gas rises from 8.0 to 9.6 kmol/min at 20 min. Run 400 min at 20 min per second, about 20 s, and export the CSV.

(a) When does `ratio_cmd` reach the 1.10 stop? At 400 min, read the true ratio on the status line of the Air demand card, the Recovery tile and the scorecard IAE. *Checkpoint:* About 24.1 min; 0.9165, 89.76 %, IAE 1226.9 (±1).

(b) Pause, press Reset, switch Feedforward to Flow only and run 400 min again. Read the scorecard IAE, the largest `ratio_cmd` and its time, the To the stack tile and whether the stack permit alarm is still on. *Checkpoint:* IAE 7.19 (±0.1); 1.032 at 23.70 min; 7.73 t/d; the alarm is still on.

(c) Why does the trim move the ratio to 1.032 at all, when the flow feedforward has already answered the step? *Checkpoint:* It acts on the transient the analyzer reports 3.2 min late.

13.10 Application area · 3 min

A plant engineer proposes moving the tail gas sample point closer to the furnace to shorten the analyzer delay, at the cost of a hot, wet, corrosive sample line.

(a) What would a shorter delay do to the first peak after a composition step, and what to the error one minute after it? *Checkpoint:* Nothing to the peak; it cuts the error at one minute.

(b) What does the chapter suggest as the better investment on a plant whose disturbances are mostly composition, and what fault of such a sample line does the model leave out? *Checkpoint:* Composition feedforward; a plugged sample line.

13.11 Application area · 4 min

Measured at steady state, bed 3 gives 97.79 % recovery at 185 °C of inlet and 97.27 % at 230, with its dew point at 180. Bed 1 gives 96.94 % at 210 °C and 98.24 % at 275.

(a) Why does bed 3 want to run cold and bed 1 hot? *Checkpoint:* Equilibrium favors cold; bed 1 alone hydrolyzes COS and CS_2.

(b) Why is the dew point protected by a one-directional override rather than an alarm, and in which direction do the challenge's suggested inlets of 258, 208 and 194 °C move from the loaded 225, 205 and 205 °C? *Checkpoint:* The damage is delayed and the gain immediate; bed 1 hotter, bed 3 colder.

13.12 Review the agent · 5 min

An AI assistant recommends: "Switch the Air demand trim to Smith predictor. It takes the 3.2 min dead time out of the loop, so the trim sees a composition change at once. The first peak after an unknown feed change falls, a drifting analyzer is corrected as well, and the comparison table shows the IAE falling from 52.239 to 19.953." Find the flaw and say how you would check it in the app.