INTERACTIVE CONTROL ENGINEERING
Principles and Applications

Volume 5 Delays, Slow Measurements, and State Estimation

CHAPTER 4

Kraft Digester Simulator That Cooks Four Hours Ahead of Its Kappa Analyzer

Pages 82–104

Edited by John D. Hedengren

How to cite this chapter. "Kraft Digester Simulator That Cooks Four Hours Ahead of Its Kappa Analyzer," in *Delays, Slow Measurements, and State Estimation*, John D. Hedengren, Ed., vol. 5 of *Interactive Control Engineering: Principles and Applications*. Advanced Process Solutions, LLC, 2026, ch. 4, pp. 82–104.

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

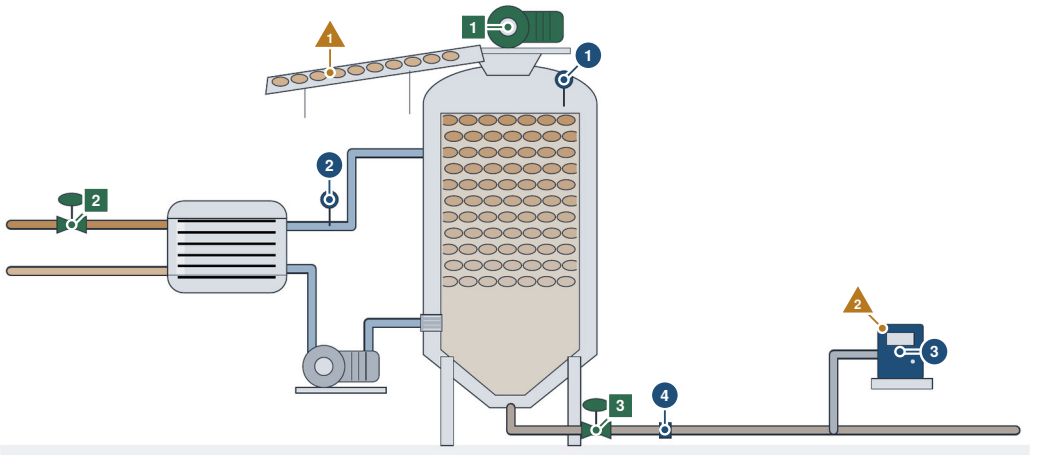
The app for this chapter
https://apmonitor.com/apps/kraft-digester/

The volume online: contents, other chapters and citation records
https://apmonitor.com/apps/books/v05-delays-estimation/#kraft-digester

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

CHAPTER 4

Kraft Digester Simulator That Cooks Four Hours Ahead of Its Kappa Analyzer



The controller moves 1 chip meter that feeds the chips 2 steam valve of the circulation heater 3 blow that sets the production rate

The controller reads 1 chip level 2 heater outlet temperature 3 online kappa analyzer 4 pulp flow at the blow

What disturbs it 1 a new wood or wetter chips 2 the analyzer drifting

Chips heated at the top reach the blowline about four hours later, and only then can the analyzer far down the line grade them. Illustration, not to scale.

IN PLAIN TERMS **Cooking Wood Chips by a Late Report**

The main idea. A digester cooks wood chips in a hot chemical bath that frees the fibers for paper. The chips sink slowly down a tall vessel, and a heater near the top sets how far they cook.

The challenge. The test that grades the pulp comes about four hours after the heater, and its instrument can slowly drift. The controller must choose what to hold steady between results, and how far to trust each late one.

Why it matters. A poor choice of what to hold can waste more than a hundred dollars of wood, chemicals and steam per tonne of pulp.

THREE WAYS TO READ THIS CHAPTER

Explore	11 min	the first page, the history segment and the fact on kraft pulping (Section 4.2), the try-it on the production step (Section 4.6), the fact on the held variable and the try-it on the drifting analyzer (the same section) and the key points
Study	45 min	the whole chapter with its three pause points and the twelve exercises at its end
Practice	15 min	the first page, the results (Section 4.6), Technology in Practice (Section 4.7), Limitations (Section 4.8) and the two open challenges

The control question A continuous digester cooks its chips for four hours past the heater, and the kappa number that grades the cook arrives late from an analyzer that can drift and later still from a laboratory. What should the heater hold between quality reports, and how should the late and biased reports correct it?

Learning objectives

- Compute the cooking residence, Vroom's rate and the H-factor change that a heater step or a production change brings.
- Explain why a heater step reaches the blowline intact four hours later, and why no heater setting rescues the chips already past the heater.
- Derive the heat-balance and H-factor feedforwards for a production change.
- Compare held variables by their steady loss, and explain why holding the residual alkali makes a new wood worse.
- Reconcile a drifting analyzer with the laboratory, and judge when a better estimate is a worse feedback signal.

Prerequisites Sections 1.4.2, 1.4.4, 1.4.5, 1.4.6, 1.4.8 and 1.6.2. Chapter 2 for a delay in the path of the manipulated input.

The app https://apmonitor.com/apps/kraft-digester/

See also Volume 3, Chapter 6 for a dose computed from a measured upstream variable ahead of a three-hour delay, and Volume 3, Chapter 2 for matching loop structure to where a disturbance enters. Volume 9, Chapter 10 for choosing the setpoint a fast loop holds under a slow quality, and Volume 9, Chapter 8 for an optimum that sits on a constraint.

4.1 Introduction

The digester simulated here makes 1500 t/d of unbleached pulp from 132.98 t/h of dry chips at a design yield of 47.0 %. The chips spend an hour in an impregnation zone at 115 °C and four hours in a cooking zone at about 160 °C, and they leave at kappa 30. When the cooking temperature is raised by 2 °C half an hour into a run, the pulp at the blowline does not change for 4.125 h, then its true kappa falls to 27.1–27.6 and stays there, and the online analyzer reports the new value at 5.08 h on the clock. No measurement of the cook exists during those hours, because nothing inside the vessel reports the lignin left in the chips.

A plant whose product is graded by a property measured long after the product was made is common well beyond pulping. Kappa itself cannot be measured on line inside a batch digester and has been inferred from liquor measurements [1], and the general problem of fusing outputs that arrive at different rates and with different delays has been posed as multirate state estimation for the chemical process industry [2]. A continuous digester adds two features to that class. The vessel is a plug-flow reactor, so a change made at the heater travels down with the chips and arrives intact one residence time later instead of spreading as a first-order lag, and the cook can be steered by two handles, temperature and alkali charge, that reach one quality through different paths. Between quality reports the temperature has to hold some measured variable constant, and the choice of that variable is a question of controlled-variable selection [3, 4].

The simulator corrects two beliefs that a block diagram leaves intact. The first is that a measured variable close to the quality is a safe thing to hold. Holding the residual alkali at the extraction screens with temperature looks reasonable, but on an alkali-hungry wood the residual falls, the heater cools to restore it and kappa rises to 34.8 against 32.7 for the H-factor held on the same wood. The second is that a better estimate of the quality is a better feedback signal. The reconciled kappa removes a drifting analyzer's bias, yet on the same disturbance day without the drift, at a 20 % charge, the H-factor held kappa within ± 2 on the raw analyzer at 11 of 11 seeds and on the reconciled estimate at none, because the estimate is built to wait out short upsets.

The published digester models and controllers are research tools built around fundamental models and a control benchmark [5, 6], and the industrial practice reported for kappa control reduces the

vessel to a first-order plus dead-time response [7]. Neither lets a learner see the H-factor of each slug of chips, the analyzer's and the laboratory's clocks and the truth on one axis, and the cost of each candidate held variable side by side. This chapter works with a kraft digester simulator that runs in a browser, described completely enough for a second implementation, and measures on it how feedforward, reconciliation and the choice of held variable act on a quality that answers hours late. The H-factor, the effective-alkali balance, the bias updates and the loss definition are established methods, and the results are measurements on one teaching model at stated seeds. The model carries the chip column as forty Lagrangian slugs whose lignin and alkali ledgers close to rounding, its estimator and controllers read nothing but declared instruments, and every rate law, constant, tuning and default is stated. Nine graded cases run from a silent blowline to a day with four simultaneous upsets.

Section 4.2 follows kraft pulping and its control. Section 4.3 gives the vessel, its slugs, its instruments and its constants, and Section 4.4 measures how a heater step and a production step travel through it. Section 4.5 sets out the control layers: three PI loops, the H-factor feedforward, the held variable chosen by its loss, the estimator that reconciles three clocks, the kappa trim and the economic optimizer. Section 4.6 measures the nine cases, verifies the code and compares the kinetics with published data, Section 4.7 sets the loops on plant hardware, and Sections 4.8 and 4.9 state the limitations and conclude.

4.2 Cooking to a Number: Kraft Pulping and Its Control

Kraft pulping dissolves the lignin that binds wood fibers in a hot solution of sodium hydroxide and sodium sulfide, and how far the cook has gone is graded by the kappa number of the pulp, a permanganate test whose value is about seven times the lignin left on the pulp in percent [8]. Mills steer a cook by one index of time and temperature, the H-factor, whose origin the history segment below tells. Kinetic models of the cook were developed alongside the index, among them a theoretical model of the kraft process published in 1983 [9], and measured rate constants of the bulk and residual delignification phases for several hardwoods and a pine remain an active subject [10, 11].

> ⧖ **HISTORY** **1957: One Number for Time and Temperature**
>
> The kraft process was patented by Dahl in 1884, and the first kraft mill went into operation in 1890. The rate of delignification rises steeply with temperature. In 1957 Vroom combined cooking time and temperature into one number, the H-factor, the time integral of a reaction rate relative to that at 100 °C, with an activation energy of 134 kJ/mol, so that one hour at 100 °C counts 1 and one hour at 171 °C about 1000. Mills have since cooked to an H-factor target, which predicts the kappa of a cook only while the alkali and the liquor-to-wood ratio stay fixed. Laboratory cooks of one eucalyptus hybrid at eight temperatures from 140 to 175 °C at one H-factor gave kappa numbers with no significant difference, and the yield fell at the hotter cooks [8].

The continuous digester moved the cook from a batch vessel to a column that the chips descend for hours, and the control problem moved with it. A fundamental model of a continuous digester, extended from the Purdue model with a lumped-parameter description of the flow transport, was published for simulation and control in 1997 [5]. Because kappa was not measured inside the vessel, the control literature of the 1990s was inferential. A batch digester was controlled through an extended Kalman filter on liquor measurements [1], a continuous digester through a model identified by subspace methods [12], and the effective alkali of a continuous digester became a controlled variable of its own [13]. A pulp mill benchmark problem for control was posed in 2004 [6], and nonlinear inferential control of kappa at several locations of a continuous digester was formulated with multirate measurements [14]. Prediction and control of the blow-line kappa number were reported in 2005 [15], and an industrial model predictive controller that treats the digester as a first-order plus dead-time process reduced the standard deviation of kappa by 40 % on a dual-vessel continuous digester [7].

Table 4.1 Classified Variables and Their Design Values

Variable	Role	Design	Range
κ	CV, blowline kappa (truth)	30.0	SP 20–45
T_h	CV of the heater loop	160 °C	140–170 °C
h	CV, chip level	5.0 m	SP 1–9 m
v_s	MV, steam valve	58.0 %	0–100 %
n_{cm}	MV, chip meter speed	100 %	0–150 %
T_{SP}	MV of the held variable	160 °C	140–170 °C
AC	MV, alkali charge on wood	20 %	14–26 %
P	production setpoint	1500 t/d	1000–1900
f_B	DV, wood B share	0	0–100 %
m	DV, chip moisture	50 %	40–62 %
s_{wl}	DV, liquor strength	100 %	70–110 %
L_f, L_s	states, lignin pools	0.224, 0.056	kg/kg wood
C, W	states, carbohydrate, acids	0.680, 0.030	kg/kg wood
E	state, effective alkali	50.0 g/L at entry	—
T, H	states, temperature, H-factor	115 °C, 0 at entry	—
E_r	output, residual alkali	10.1 g/L	floor 8 g/L
Y	output, yield	47.0 %	—

✦ DID YOU KNOW? Kraft Makes Most of the Pulp

A 2005 profile for the United States Department of Energy put about 98 % of the country's chemical pulp capacity, or 80 % of all its pulping capacity, on the kraft process. It also found that continuous digesters predominate, being more space-efficient and easier to control than batch digesters [16].

The question of which variable to hold constant has its own history. Optimizing control structures, in which a measured variable is held so that the plant stays near its optimum as disturbances arrive, were formalized for chemical plants in 1980 [3]. The self-optimizing form judges a candidate by its loss, the extra cost of holding it at a constant setpoint against re-optimizing after each disturbance [4], and combinations of measurements extend the candidate set [17]. A digester offers the textbook candidates on one vessel: the cooking temperature, the H-factor, the residual alkali and a model's predicted kappa.

4.3 The Digester as Modeled

The simulated digester, shown in Figure 4.1, is a generic teaching design with no mill as its source. Chips enter the top separator through a chip meter, white liquor is ratioed to them at an alkali charge AC, and the chip column descends through an impregnation zone and a cooking zone. The circulation heater at the top of the cooking zone sets the temperature of the chips that pass it, and they carry that temperature down, cooling slowly to the surroundings. The spent liquor is drawn through the extraction screens, where a titrator reads the residual alkali E_r, and the pulp is blown below them, where an online analyzer and the laboratory read its kappa. The chip level h between the chip meter and the blow is the vessel's one inventory.

The variables are listed in Table 4.1 with their role. Time is in hours, wood flows in tonnes of dry wood per hour, alkali in grams per liter as NaOH and temperature in °C. Five lumped states are integrated together with ten states for each of the 40 slugs. The lumped states are the heater outlet temperature T_h, the chip level h, the column position ϕ in slugs, and two ledgers for the steam used and the heat given to the chips.

4.3.1 One Slug of Chips

Each slug is a closed batch of dry wood and the liquor that travels with it. Its state per kilogram of original dry wood is the bulk lignin L_f, the residual lignin L_s, the carbohydrate C, the alkali-consuming acid groups W (in kilograms of NaOH they consume), the effective alkali E in the liquor,

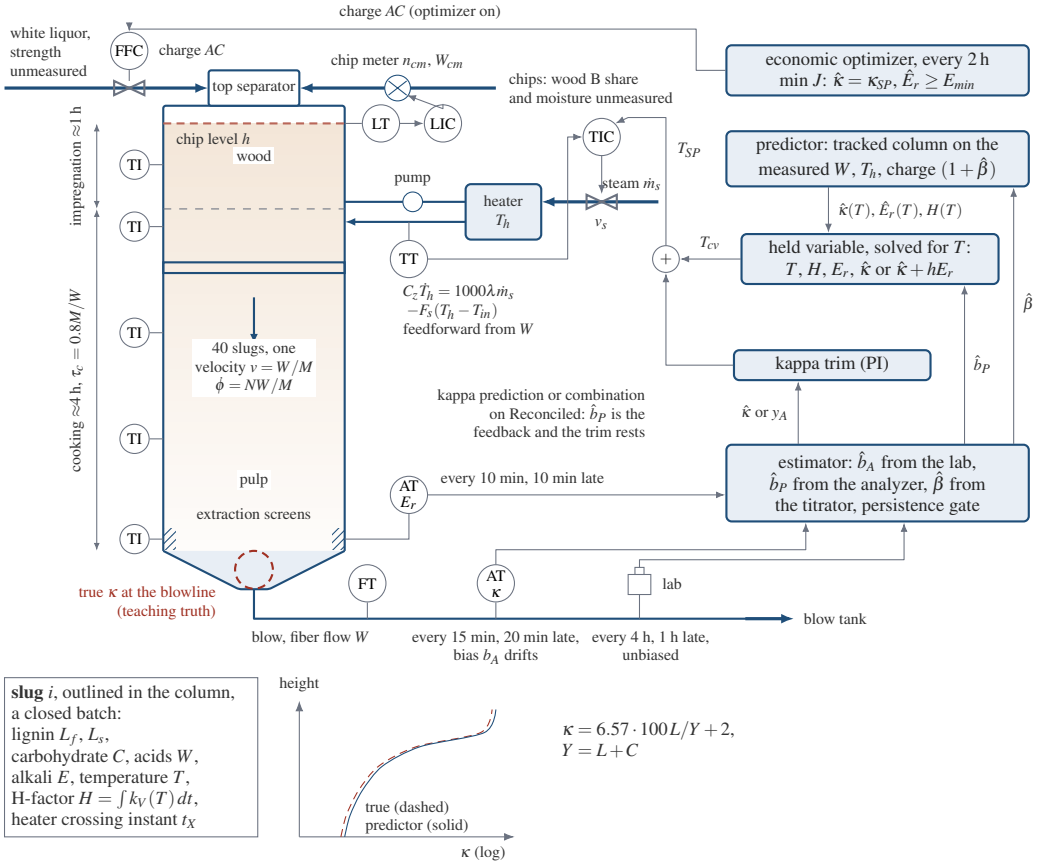


Figure 4.1 Continuous kraft digester as modeled. The column of $N = 40$ slugs is painted from wood to pulp. One slug, outlined, carries its lignin pools L_f and L_s, carbohydrate C, acid groups W, alkali E, temperature T, H-factor H and heater crossing instant t_X. The column moves at one velocity $v = W/M$ set by the blow's fiber flow W. The circulation heater sets T_h with the steam valve v_s under a temperature controller whose setpoint T_{SP} is the held variable's temperature plus the kappa trim. Measurements are open circles, the truth at the blowline and the true profile are dashed (teaching truth), and the inputs of the predictor are tagged rather than wired. Process model only, not to scale.

the temperature T in kelvin and the H-factor H, and its rates in hours are given by Equation (4.1).

$$
\begin{aligned}
k_V(T) &= \exp\bigl(43.2 - 16113/T\bigr), \qquad \frac{\mathrm{d}H}{\mathrm{d}t} = k_V(T) \\
\frac{\mathrm{d}L_f}{\mathrm{d}t} &= -A_L \rho_s k_V E^{0.5} L_f, \qquad \frac{\mathrm{d}L_s}{\mathrm{d}t} = -\frac{A_L}{r_s}\rho_s k_V E^{0.5} L_s \\
\frac{\mathrm{d}C}{\mathrm{d}t} &= -A_C k_C(T)\, E^{0.75} \max(0, C - C_\infty) \\
\frac{\mathrm{d}W}{\mathrm{d}t} &= -A_W k_W(T)\, E\, W \\
\frac{\mathrm{d}E}{\mathrm{d}t} &= -\frac{1000}{LW}\bigl(b_L r_L + b_C r_C + r_W\bigr) \\
\frac{\mathrm{d}T}{\mathrm{d}t} &= -\frac{T - T_a}{\tau_{loss}}
\end{aligned}
\tag{4.1}
$$

where r_L, r_C and r_W are the magnitudes of the lignin, carbohydrate and acid-group rates, ρ_s is the wood's reactivity, LW the liquor-to-wood ratio in liters per kilogram, and k_C and k_W are Arrhenius factors equal to 1 at 100 °C with activation temperatures $E_C/R = 17\,500\,\mathrm{K}$ and $E_W/R = 6000\,\mathrm{K}$. Vroom's rate is 1.019 at 100 °C with the constant 43.2 and 403.6 per hour at 160 °C, and $16113\,R$ is 134.0 kJ/mol. The carbohydrate loss has the higher activation energy (145.5 kJ/mol), so a hotter cook loses more yield at the same H-factor, as the laboratory cooks showed [8]. Kappa and yield follow from the slug's lignin and carbohydrate, as shown in Equation (4.2).

$$\kappa = 6.57 \cdot 100\,\frac{L}{Y} + 2, \qquad L = L_f + L_s, \qquad Y = L + C \tag{4.2}$$

where 6.57 kappa units per percent of lignin on pulp is the commonly quoted proportionality (about seven [8]) and the 2 is a teaching allowance for the hexenuronic acids the test also oxidizes.

A fresh slug enters with $L_f = 0.8L_0$, $L_s = 0.2L_0$, $C = C_0$, $W = W_0$, $T = T_{imp}$, $H = 0$ and $E_0 = 1000\,AC\,s_{wl}/LW$, where $LW = 3.0 + m/(1-m)$, so that 50 % moisture gives 4.0 L/kg and a 20 % charge gives 50 g/L. Wood A, the design wood, has $L_0 = 0.280$, $C_0 = 0.680$, $W_0 = 0.030$ and $\rho_s = 1$. Wood B has $L_0 = 0.290$, $C_0 = 0.670$, $W_0 = 0.050$ and $\rho_s = 0.95$, so it carries more lignin, consumes more alkali and cooks a little slower, and a mix is taken linearly in the share f_B. Each slug also keeps ledgers of the lignin and carbohydrate dissolved and the alkali consumed, which make $L_f + L_s + D_L = L_0$, $C + D_C = C_0$ and $E\,LW/1000 + A = E_0 LW/1000$ linear invariants that the integrator preserves.

4.3.2 A Column of Slugs

The column is carried by the method of characteristics. All 40 slugs ride it at one velocity, chips and free liquor together, and its position advances at the blow's fiber flow, as shown in Equation (4.3).

$$\begin{aligned} \frac{\mathrm{d}\phi}{\mathrm{d}t} &= \frac{N W_{col}}{M}, \qquad \frac{\mathrm{d}h}{\mathrm{d}t} = \frac{W_{cm} - W_{col}}{\rho_A} \\ C_z \frac{\mathrm{d}T_h}{\mathrm{d}t} &= 1000\,\lambda\,\dot{m}_s - F_s\,(T_h - T_{in}) \\ F_s &= \max\!\left(F_c,\ 0.1 F_{c,0}\right) \\ F_c &= 1000\,W_{col}(c_w + LW\,c_l) \end{aligned} \tag{4.3}$$

where $M = 664.9\,\mathrm{t}$ is the dry wood held in the two zones, W_{col} the column's descent in tonnes of wood per hour (the blow's flow, zero while the column hangs), W_{cm} the chip meter's flow, $\rho_A = 13.35\,\mathrm{t/m}$ the wood per meter of column, T_{in} the temperature of the slug arriving from impregnation, $\dot{m}_s = v_s\,\dot{m}_{s,max}\,p_s$ the steam in tonnes per hour with $\dot{m}_{s,max} = 90\,\mathrm{t/h}$ and p_s the header pressure as a fraction of design, $c_w = 1.35\,\mathrm{kJ/(kg\,K)}$, $c_l = 4.0\,\mathrm{kJ/(L\,K)}$, $\lambda = 2000\,\mathrm{kJ/kg}$, and $C_z = \tau_h F_{c,0}$ with $\tau_h = 0.1\,\mathrm{h}$ and $F_{c,0} = 2307.2\,\mathrm{MJ/(h\,K)}$ at design. The floor on F_s is a declared stand-in for the heat the circulation trades with a hung column. When ϕ passes 1 the bottom slug is blown with the instant it crossed, every slug moves down one slot, a fresh slug of the present feed enters the top, and the slug that moves into the first cooking slot takes T_h and records its crossing instant t_X. The cooking zone holds $0.8M$, so its residence is $\tau_c = 0.8M/W$, 4.000 h at design and 3.478 h at 1725 t/d, and one slug passes every 7.5 min at design. In this chapter τ_c is this residence and not the closed-loop time constant of Chapter 1. Because the column advances with the fiber flow, the delay of every slug is kept in wood rather than in time, the transport delay defined in volume of Section 1.4.4. Chapter 7 follows its material the same way along a line of well-mixed cells.

A hang-up holds the whole column for 0.5 h while the chip meter keeps feeding the top, and it then drops the distance it missed at once. The slugs in the cooking zone keep cooking while they hang. Every such event, and the weak-liquor hour at a strength of 85 %, the analyzer's drift and its recalibration, is applied once, between integration steps, and marked on the trends, so that no event hides inside the rate equations.

4.3.3 What the Instruments Report

The controllers read only a measurement packet. The heater outlet carries colored noise of 0.15 °C with a 0.02 h correlation time, five thermocouples at fixed heights read 0.3 °C noise, the blow's fiber flow 0.5 %, the level 3 cm and the steam flow 1 %. Three instruments report samples, as shown in Equation (4.4).

$$\begin{aligned} y_E(t_s + \tfrac{1}{6}) &= E_r(t_s) + v_E, & t_s &\in \{0, \tfrac{1}{6}, \tfrac{2}{6}, \dots\} \\ y_A(t_s + \tfrac{1}{3}) &= \kappa(t_s) + b_A(t_s) + v_A, & t_s &\in \{0, \tfrac{1}{4}, \tfrac{1}{2}, \dots\} \\ y_L(t_s + 1) &= \kappa(t_s) + v_L, & t_s &\in \{0, 4, 8, \dots\} \\ \frac{\mathrm{d}b_A}{\mathrm{d}t} &= \delta \text{ after the drift event} \end{aligned} \tag{4.4}$$

where $\kappa(t_s)$ and $E_r(t_s)$ are the true kappa and residual alkali of the slug at the blowline when the sample is drawn, the noise standard deviations are 0.2 g/L, 0.4 and 0.3, and $\delta = 0.3$ kappa per hour. The titrator samples every 10 min and reports 10 min later, the analyzer every 15 min and 20 min later, and the laboratory every 4 h and an hour later, without bias, on bottles drawn at the instants of analyzer samples. A reading holds until the next report, no sample is drawn while the column hangs, and every report is queued with the instant its sample left the blowline. The wood, the moisture, the liquor's strength and the analyzer's bias are not measured. The age of each held result, and the dead time it adds to a loop, follow Section 1.4.5.

4.3.4 Constants, Start-Up and Integration

Table 4.2 gives the constants with their defaults. The kinetic constants are teaching values. A_L and A_C were fitted so that the design slug leaves at kappa 30.000 and yield 47.00 %, with a residual alkali of 10.12 g/L and an H-factor of 1505.5 after cooling 1.72 °C over the cook. Each case starts lined out: one slug of the present feed is marched down the column at the present heater temperature and flow and copied into each slot at its age. The engine of Section 1.6.1 steps 10 s, integrating the lumped states and every slug by the classical fourth-order Runge-Kutta method, and the controller executes and a record is stored once per simulated minute.

4.3.5 Assumptions

The free liquor rides with the chips and there is no axial mixing, so the model has no counter-current wash zone and the cook stops at the extraction screens. Effective alkali stands for the hydroxide and hydrosulfide together, so sulfidity is not a variable. The kinetics are teaching values in the form of a two-pool lignin model and are not any wood's. The heater sets the temperature of each slug that crosses it, so radial gradients, channeling and the steaming of chips before the vessel are outside the model, and the reaction heat is neglected. A wood mix is taken as one wood with averaged composition and reactivity, while a real mix cooks each chip by its own kinetics. The residual alkali floor and the temperature limit are teaching values. A mill sets its floor from trials that keep dissolved lignin from redepositing on the fibers and its limit from the vessel's rating and the pulp's strength, and the model does not simulate redeposition, so the floor is judged as a declared limit. A hang-up holds the whole column, while a real column hangs above some point and leaves a void below it.

4.4 Hours Between the Heater and the Blowline

4.4.1 A Step That Arrives Intact

Case 1 holds the heater at 160 °C and the learner raises the setpoint to 162 °C at 0.5 h. The first slug to cross the heater after the step does so at 0.625 h, the next slug boundary plus the heater's lag, and reaches the blowline at 4.625 h. The slugs that follow leave at a median of 27.32 from 5 to 10 h, within 27.08–27.60 from 6 to 10 h, and the analyzer reports the change at 5.08 h from a sample drawn at 4.75 h. At 160 °C the logarithm of Vroom's rate rises 8.59 % per kelvin, so two degrees multiply the cooking H-factor by 1.186. The plug flow keeps the step's shape. A tanks-in-series grid of 40

Table 4.2 Model Constants. Those Marked Fixed Have No Control in the Simulator

Symbol	Meaning	Value	Unit
N, M	slugs, dry wood in the zones	40, 664.9	—, t
—	impregnation share (fixed)	0.2	—
ρ_A	dry wood per meter	13.35	t/m
T_{imp}	impregnation temperature	115	°C
—	liquor added per kg wood	3.0	L/kg
c_w	heat capacity of wood	1.35	kJ/(kg K)
c_l	heat capacity of liquor	4.0	kJ/(L K)
$\lambda, \dot{m}_{s,max}$	steam heat, valve capacity	2000, 90	kJ/kg, t/h
τ_h	heater time constant	0.1	h
A_L	bulk delignification (fixed)	9.5553×10^{-4}	$\mathrm{h}^{-1}\,(\mathrm{g/L})^{-0.5}$
r_s, f_s	residual pool slowness, share (fixed)	6, 0.20	—
A_C	carbohydrate loss (fixed)	4.0771×10^{-5}	$\mathrm{h}^{-1}\,(\mathrm{g/L})^{-0.75}$
C_∞	surviving carbohydrate (fixed)	0.30	kg/kg
A_W	acid-group consumption (fixed)	0.06	$\mathrm{h}^{-1}\,(\mathrm{g/L})^{-1}$
b_L, b_C	alkali per lignin, carbohydrate (fixed)	0.10, 0.45	kg/kg
T_a, τ_{loss}	surroundings, cooling time (fixed)	30, 300	°C, h
—	delignification multiplier	1.0	—
E/R	plant activation temperature	16113	K
—	analyzer period, delay, noise	0.25, 0.33, 0.4	h, h, —
δ	analyzer drift when on	0.3	h^{-1}
—	laboratory period, delay, noise	4, 1, 0.3	h, h, —
—	hang-up duration	0.5	h
—	weak liquor strength	85	%
p_s, p_a, p_w	steam, alkali, wood price	20, 150, 70	$/t
q	off-target kappa cost	5	$/t per kappa^2
E_{min}	residual alkali floor	8	g/L
T_{max}	cooking temperature limit	170	°C

cells, which an early design of the simulator specified, smeared the design kappa by 2.64 % and failed the refinement test, which is why the column is carried as slugs.

4.4.2 Chips Caught in Transit

A production increase is the disturbance that a fixed cooking temperature cannot absorb. In case 2 the production rises 15 % at 1 h, the residence falls from 5.00 to 4.348 h and the cooking time to 3.478 h, so at a fixed temperature every slug gathers 0.8696 of the design cooking H-factor. The first slugs are 0.3 off by 1.54 h, kappa reaches 32.56 at 5.46 h, and it stays within 32.04–32.61 from 8 to 12 h. The chips already in the cooking zone when the step arrives have crossed the heater, and no later move of the heater reaches them. A sweep of eight heater settings over the worst slug blown from 1 to 8 h measures the consequence. Temperature alone reaches +2.62, temperature with an analyzer trim at gains of 0.15, 0.3, 0.5 and 1.0 °C per kappa reaches +2.75, +3.05, +3.45 and +4.49, the H-factor +2.46, the H-factor with the analyzer trim +2.75, and the kappa prediction on the reconciled estimate +2.59. In seven of the eight the worst slug crossed the heater at the step. No setting of the heater holds ±2 through this step, which is the honestly infeasible variant of the case.

🕮 **Pause.** Try Exercise 4.2 before reading on.

4.4.3 The Premise of the Feedforward

The H-factor is worth feeding forward only if one H-factor gives one kappa across production rates. At an H-factor of 1505 and a 20 % charge, an independent bisection on the heater gives 158.45, 160.00 and 161.53 °C at 1300, 1500 and 1725 t/d, and the slugs leave at kappa 29.842, 30.005 and 30.166. The spread of 0.32 is the impregnation's share of the alkali consumption, which shortens with the rate. The engine with the H-factor held and the noise off gives 29.848, 30.010 and 30.170 from 10 to 14 h. The premise holds to a third of a kappa unit on this model, and the laboratory cooks support it within their scatter [8].

4.5 The Control Layers as Coded

The controller runs every minute on the measurement packet. Production is the blow's fiber flow, $W = P/(24 \cdot 0.47)$ in tonnes of wood per hour, so the throughput is set at the bottom of the vessel and the chip level is held at the top.

4.5.1 The Three PI Loops

The heater, level and kappa-trim loops are proportional-integral loops in the velocity form of Section 1.6.2, written for these loops in Equation (4.5).

$$\begin{aligned} e_k &= a\,(r_k - y_k), \qquad s_k = s_{k-1} + K_c(e_k - e_{k-1}) \\ \hat{u}_k &= u_b + s_k + I_{k-1}, \qquad u_k = \min\bigl(u_{max}, \max(u_{min}, \hat{u}_k)\bigr) \\ I_k &= I_{k-1} + \Bigl[\frac{K_c}{\tau_I} e_k + \frac{u_k - \hat{u}_k}{\tau_I}\Bigr]\Delta t \end{aligned} \tag{4.5}$$

where $a = \pm 1$ is the action, u_b a bias and $\Delta t = 1\,\text{min}$. The symbols map onto Equations (1.15) to (1.18) as $s_k = S_k$ and $\hat{u}_k = v_k$, with no derivative term and a tracking time constant equal to τ_I. The heater loop drives the steam valve from 0 to 100 % at $K_c = 2.5\,\%$ per °C and $\tau_I = 6\,\text{min}$. Its bias is a heat-balance feedforward, the valve that lifts the chips from impregnation to T_{SP} at the measured fiber flow, as shown in Equation (4.6).

$$u_b = \frac{100}{\dot{m}_{s,max}} \, \frac{W_f\,(c_w + 4.0\,c_l)(T_{SP} - T_{in,0})}{\lambda} \tag{4.6}$$

where W_f is the measured fiber flow through a 0.02 h filter, limited to 1.5 times the production setpoint and to no less than a tenth of the design flow, and $T_{in,0} = 30 + (T_{imp} - 30)\exp(-0.2M/(300\,W_f))$, in °C, is the declared heat balance's estimate of the arriving chips, 114.72 °C at design. It holds 58.0 % of the valve at design and follows a production move within minutes, before the heater's outlet has fallen. The level loop moves the chip meter from 0 to 150 % of the design wood flow at 10 % per meter with $\tau_I = 2\,\text{h}$, an averaging tuning that lets the level absorb the chip meter's moves. The kappa trim acts on the cooking temperature at 0.15 °C per kappa with $\tau_I = 3\,\text{h}$, reverse acting, within $\pm 12\,^{\circ}\text{C}$. The setpoint of the heater loop is $T_{SP} = \min(170, \max(140, T_{cv} + \text{trim}))$, where T_{cv} is the held variable's temperature.

4.5.2 The H-Factor Feedforward

The held variable H-factor chooses the heater temperature that gives the chips entering the cooking zone now their target H-factor at the measured rate, as shown in Equation (4.7).

$$\begin{aligned} H_{SP} &= H_{imp} + \int_0^{\tau_c} k_V\bigl(T_a + (T_h - T_a)\mathrm{e}^{-t/\tau_{loss}}\bigr)\,dt \\ \tau_c &= 0.8M/W \end{aligned} \tag{4.7}$$

where H_{imp} is the H-factor the predictor's slug gathers in impregnation. The integral is taken by Simpson's rule on 24 intervals. Because $\ln H$ is nearly linear in T_h, Newton's method on $\ln H$ with the slope $(E/R)/(T_h + 272.65)^2$, kept inside a bisection bracket on 140–170 °C, stops when the H-factor is within 10^{-6} of the target or after 30 iterations. A target that the bracket cannot reach returns the bound and is flagged, and the solve never extrapolates. On a fresh instance it reached 162.2698 °C in three iterations against a brute-force scan's 162.2700 °C.

4.5.3 Which Variable the Heater Holds

Five candidates are offered, the temperature itself, the H-factor, the residual alkali E_r, the predicted kappa $\hat{\kappa}$, and the combination $\hat{\kappa} + hE_r$. The last three are computed by the controller's predictor, the declared kinetics run on the design wood and moisture, the commanded charge times $(1 + \hat{\beta})$ and the

Table 4.3 Steady Loss of Each Held Variable ($/t Pulp) at a 20 % Charge, With the Steady Kappa on Wood B

Held variable	Wood B	Moisture	Liquor	Rate	κ on B
temperature	29.45	0.00	12.83	20.12	32.75
H-factor	29.45	0.00	12.83	0.19	32.75
residual alkali	119.68	0.47	190.07	0.00	35.21
kappa prediction	9.90	0.05	0.50	0.63	31.74
combination	7.27	0.08	2.27	0.71	31.54

measured flow, for the slug about to enter the cooking zone. Each is solved for the temperature that holds it at its setpoint by the Illinois variant of regula falsi on 140–170 °C, stopping within 10^{-3} or after 16 iterations, and the solves agreed with 0.001 °C scans to 5.6×10^{-4} °C.

A candidate is scored by its loss, as shown in Equation (4.8).

$$\begin{aligned} L(c,d) &= J\big(T_h(c=c_s,d),d\big) - \min_{T_h} J(T_h,d) \\ J &= \frac{p_w + p_a AC_d + p_s\, s(T_h)}{Y} + q\,(\kappa - \kappa_{SP})^2 \end{aligned} \tag{4.8}$$

where J is the cost in dollars per tonne of pulp of the plant's steady state, AC_d the charge delivered, s the steam per tonne of wood, and c_s the candidate's value at the design point. The candidate is measured as the controller would see it at steady state, with $\hat{\beta}$ fitted so that the predictor's residual equals the plant's. The disturbances are wood B at 60 %, moisture at 56 %, liquor at 90 % strength and production up 15 %. The ideal is found by golden section on 148–172 °C, the held temperature by the candidate's crossing on a 0.25 °C grid with the plant evaluated at that temperature, and the weight h by golden section on $[-2,2]$ minimizing the summed loss over wood B and the weak liquor. At the 20 % charge $h = -0.0869$ kappa per g/L, and the table is given in Table 4.3. Holding the residual alkali with temperature is the worst choice on the wood and the liquor, because both lower the residual and the heater then cools. The kappa prediction and the combination recover most of the loss on the wood, and the H-factor and the residual alkali are the candidates that ignore a production change and a moisture change respectively. Volume 9, Chapter 10 asks the same question of a dissolved oxygen setpoint held under an effluent ammonia that answers in hours.

4.5.4 Reconciling Three Clocks

The estimator keeps its own column of 40 slugs, advanced on the measured fiber flow, given the measured heater temperature as each slug crosses, and cooked with the declared kinetics on the commanded charge, so its predicted kappa $\hat{\kappa}_P$ exists for every slug in the vessel and for the slug leaving now. Each update pairs two sources for the same sample, aligned by the instant the sample left the blowline, as shown in Equation (4.9).

$$\begin{aligned} \hat{b}_A &\leftarrow \hat{b}_A + \lambda_A\big(y_A(t_s) - y_L(t_s) - \hat{b}_A\big) \\ \hat{b}_P &\leftarrow \hat{b}_P + \lambda_P\big(y_A - \hat{b}_A - \hat{\kappa}_P(t_s) - \hat{b}_P\big) \\ \hat{\beta} &\leftarrow \hat{\beta} + g_\beta\left(\beta_s + \frac{y_E - \hat{E}_P(t_s)}{\partial E/\partial \ln AC} - \hat{\beta}\right) \\ \hat{\kappa}(t) &= \hat{\kappa}_P(t) + \hat{b}_P \end{aligned} \tag{4.9}$$

where the first update runs on each laboratory report, the second on each analyzer report and the third on each titrator report, $\lambda_A = 0.8$, $\lambda_P = 0.1$, $g_\beta = 0.2$, β_s is the factor the tracked slug was given when it entered five hours earlier, and $\partial E/\partial \ln AC = 31\,\mathrm{g/L}$ at design. Averaging the factor each report implies for its own slug, rather than integrating the raw residual, keeps the five-hour delay out of the update. Each line is a bias update of the form of Equation (1.12), compared with the value stored for the same sample as in Equation (1.14). The laboratory is the one unbiased instrument, so it alone

makes the analyzer's bias observable (Section 1.4.6). Chapter 10 corrects a soft sensor from a delayed laboratory in the same way. The biases are clamped to ± 15 kappa and $\hat{\beta}$ to ± 0.4, and an update whose partner is missing is skipped. A persistence gate guards the two predictor updates. A titrator report whose implied factor lies more than 0.05 from $\hat{\beta}$, or an analyzer report more than 1.5 kappa from its reconciled prediction, is held back until seven titrator or five analyzer reports on the same side agree, which is more than an hour of reports. On the titrator a run must also stay within 0.05 of the factor that started it. An upset that passes within an hour, such as a weak-liquor hour, therefore never enters the slugs that follow it, while a new wood is believed after its first hour.

4.5.5 The Kappa Trim and the Economic Optimizer

The kappa feedback is Off, a PI on the raw analyzer, or Reconciled. With the temperature, the H-factor or the residual alkali held, Reconciled feeds $\hat{\kappa}$ to the PI. With the kappa prediction or the combination held, Reconciled adds $\hat{b}_P$ to the predicted kappa inside the solve, so the bias is the feedback and the trim rests at zero. The economic optimizer runs every 2 h when switched on and chooses the charge, as shown in Equation (4.10).

$$\begin{aligned} \min_{AC,\,T_h} \; & \frac{p_w + p_a AC(1+\hat{\beta}) + p_s\, s(T_h)}{\hat{Y}} \\ \text{s.t. } & \hat{\kappa}(AC,T_h) + \hat{b}_P = \kappa_{SP}, \quad \hat{E}_r(AC,T_h) \geq E_{min} \\ & 140 \leq T_h \leq T_{max}, \quad 0.15 \leq AC \leq 0.26 \end{aligned} \tag{4.10}$$

The equality removes T_h by the regula falsi solve, a bisection of 22 halvings finds the least charge whose predicted residual stays on the floor, and a golden-section search of 24 steps minimizes the cost along the remaining range, 51 predictor evaluations in all. The optimizer holds the floor on its own predictor with no back-off. On a fresh instance it found 21.237 % at 163.25 °C and \$242.4918/t, against a brute-force search's 21.240 % and \$242.5105/t, which cannot land on the floor.

4.5.6 Three Models Kept Apart

The plant model is Section 4.3 with the true wood, moisture, liquor and analyzer bias. The controller's model is the predictor and the estimator's column, the declared kinetics on the design wood, which know none of the four. The loss table is a designer's offline analysis on the declared model with the rate multiplier at 1 and Vroom's activation temperature. A grep of the controller, estimator, predictor and optimizer sources finds no read of any truth field.

4.6 Nine Scenarios, Measured

The cases appear in Table 4.4. Every case starts lined out at the design point with the defaults above and departs from them as its row says. The numbers below come from the simulator's own engine run without a page at the case's seed (20261030), advanced step by step exactly as the page would advance it, with a learner's setting entered after the run starts where the case asks for it, and for the challenge also at seeds 1–10. The judged quantity is the true kappa of each slug as it is blown, a quantity no instrument reports. The exported file holds the analyzer, laboratory, reconciled and predicted kappa beside the true kappa, whose column name marks it as teaching truth (Section 1.6.4).

4.6.1 Feedforward Through the Rate Step

With the H-factor held in case 3 the heater setpoint jumps from 160.00 to 161.53 °C at the step and the predicted H-factor of the entering chips stays at 1505.0. The chips already past the heater lose their time all the same and peak at 32.46 at 4.48 h, which is the last slug outside ± 1 of 30, and from 7 to 12 h kappa stays within 29.88–30.42. The feedforward ends the upset one cooking residence after the step, while temperature alone never ends it. Volume 3, Chapter 6 makes the same move on a coagulant dose, computed from a measured upstream variable with the slow product measurement demoted to a trim.

Table 4.4 The Nine Scenarios

#	Case and events	Teaching point
1	Temperature held, learner steps T_{SP}, 10 h	the four-hour silence
2	Production +15 % at 1 h, temperature held, 12 h	chips caught in transit, infeasible variant
3	As 2 with the H-factor held	feedforward
4	Wood B 60 % at 1 h, each held variable, 12 h	self-optimizing loss
5	Analyzer drift from 1 h, kappa prediction, analyzer trim, 24 h	multirate reconciliation
6	Wood B at 1 h, H-factor, analyzer trim, 36 h	feedback through the delay
7	Kappa prediction, reconciled, optimizer at 1 h, 12 h	steam against alkali
8	Hang-up at 2 h, H-factor, 10 h	level and over-cooked slugs
9	Challenge: rate, drift, wood B and a weak hour, 24 h	the whole problem

TRY IT IN THE APP **3 · The Same Step With H-Factor Feedforward**

Load the case and press Run at 0.1 h/s. Watch the cooking temperature setpoint jump at the production step at 1 h, then the true kappa (dashed). The chips already past the heater still peak at 32.46 at 4.48 h, and from 7 to 12 h kappa stays within 29.88–30.42.

4.6.2 A New Wood and the Held Variable

Case 4 makes 60 % of the chips wood B at 1 h. Holding the H-factor, kappa reaches 32.59–33.03 from 10 to 14 h with the heater setpoint unmoved at 160.00 °C. Holding the residual alkali, the heater cools to 158.35 °C and kappa reaches 32.50–35.09, because the wood's acid groups consume alkali and the heater tries to put it back. The kappa prediction reads the lower residual as less alkali, heats to 160.73 °C and holds 31.63–32.79, and the combination heats to 160.87 °C and holds 31.44–32.81. The steady loss table ranks them in the same order, \$119.68, \$29.45, \$9.90 and \$7.27 per tonne. None reaches 30, because the part of wood B that is slower rather than hungrier reaches no instrument but kappa. Holding the kappa prediction at the design charge, the titrator reads 7.68 g/L at 14 h, under the 8 g/L floor, so the heater can hold kappa but only the charge can hold the floor.

Pause. Try Exercise 4.1 before reading on.

DID YOU KNOW? **The Closest Measurement Is the Worst to Hold**

From the simulator: on wood B, holding the residual alkali measured inside the cook costs \$119.68 per tonne of pulp against \$29.45 for the H-factor, and its steady kappa is 35.21 against 32.75 (Table 4.3). The heater cools to restore the alkali the new wood consumed.

4.6.3 A Drifting Analyzer and the Laboratory

In case 5 the analyzer's bias drifts upward at 0.3 kappa per hour from 1 h, and the measured traces are shown in Figure 4.2. With the trim on the raw analyzer the loop holds the analyzer a few units above 30 while the true kappa falls to 26.52 at 24 h, the bias then being 6.90. With feedback through the reconciled estimate the true kappa stays within 28.68–29.51 from 12 to 24 h. The laboratory pins $\hat{b}_A$ once every four hours, at 5.56 after the report of 21 h, and the 1.34 it has not yet seen at 24 h leaks into the predictor bias. With no feedback at all the predictor alone holds 29.65–30.49, because nothing in this case changes what the predictor knows.

TRY IT IN THE APP **5 · The Analyzer Drifts and the Laboratory Corrects It**

Press Run at 0.5 h/s on the Clocks view. The trim on the raw analyzer holds the analyzer a few units above 30 while the true kappa (dashed) falls to 26.52 at 24 h. Set Kappa feedback to Reconciled, press Reset and run again. From 12 to 24 h the true kappa now stays within 28.68–29.51.

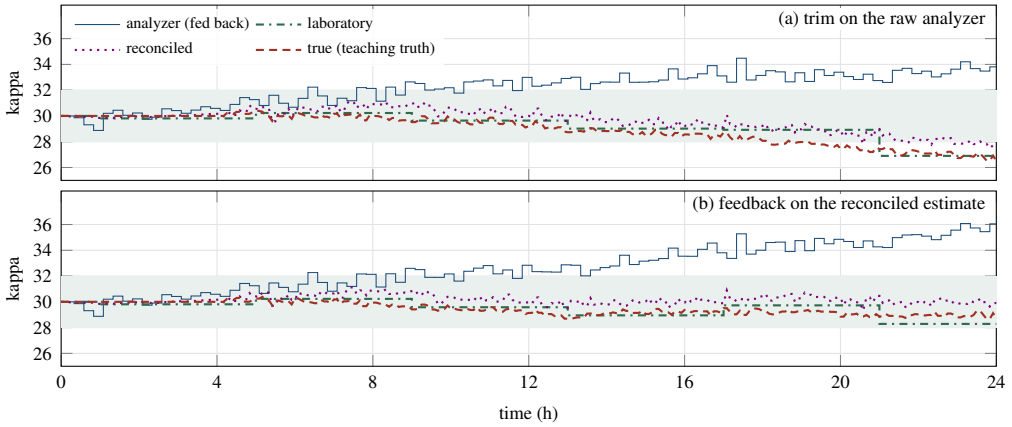


Figure 4.2 Kappa on three clocks in case 5, the analyzer drifting from 1 h. (a) The trim on the raw analyzer holds the analyzer a few units above 30 while the true kappa falls to 26.52. (b) Feedback through the reconciled estimate keeps the true kappa within 28.68–29.51 after 12 h. The analyzer (thin solid) and the laboratory (dash-dotted) hold their last report, the reconciled estimate is dotted, the truth is dashed, and the shaded band is ±2.

4.6.4 Feedback Through Five Hours

Case 6 holds the H-factor on wood B with the analyzer trim and sweeps its gain. At 0.15 and 0.3 °C per kappa the last slug outside ±1 leaves at 24.12 and 12.25 h. At 0.5 kappa undershoots to 27.88 at 14.50 h and is not settled until 33.25 h, and at 0.8 it rings and grows to 20.22–41.10 from 30 to 34 h. The vessel's delay sets how fast the feedback can be, whatever the gain, which is the bandwidth limit of Section 1.4.2.

4.6.5 Steam Against Alkali

Case 7 turns the optimizer on at 1 h. It lowers the charge from 20.00 to 18.83 % and the heater setpoint rises from 159.85 to 160.61 °C. From 7 to 12 h the alkali falls from 425.0 to 396.3 kg per tonne of pulp, the steam rises from 0.814 to 0.817 t/t and the yield from 46.97 to 47.60 %, with 41 of 41 slugs within ±2. The optimum sits on the residual floor of its predictor, and the true residual from 7 to 24 h is 7.72–8.39 g/L, under the floor in 49 % of the steps. A constraint held at its limit on a model is broken about half the time by the model's error, which is the reason a mill backs off from it. As in the data hall of Volume 9, Chapter 8, the cheapest feasible point sits on a constraint.

4.6.6 A Hang-Up

In case 8 the column hangs at 2 h. The level climbs from 5.0 to 8.86 m while the chip meter feeds a column that does not move, and it is 3.87 m at the drop at 2.50 h. Four slugs leave at the drop at kappa 28.56, 29.02, 29.49 and 29.99, the lowest with an H-factor of 1636.0 against 1505.5 for an untouched slug, each equal to the declared kinetics cooked for its own 4.375, 4.250, 4.125 and 4.000 h from crossing to exit. Four hours later the chips that fell past the heater in the drop arrive under-cooked at 31.36 at 6.12 h, and the slug that crossed while the heater settled at 161.00 °C leaves at 28.63 at 6.62 h. The heater, whose load returns at once when the column drops, swings between 154.14 and 161.56 °C after the drop.

4.6.7 The Challenge

The challenge runs 24 h. At 1 h production rises 15 % and the analyzer starts to drift, at 2 h the chips become 60 % wood B, and from 3 h the liquor runs at 85 % strength for an hour. The target is that every slug blown from 12 to 24 h leaves within ±2 of kappa 30 and with a true residual alkali at

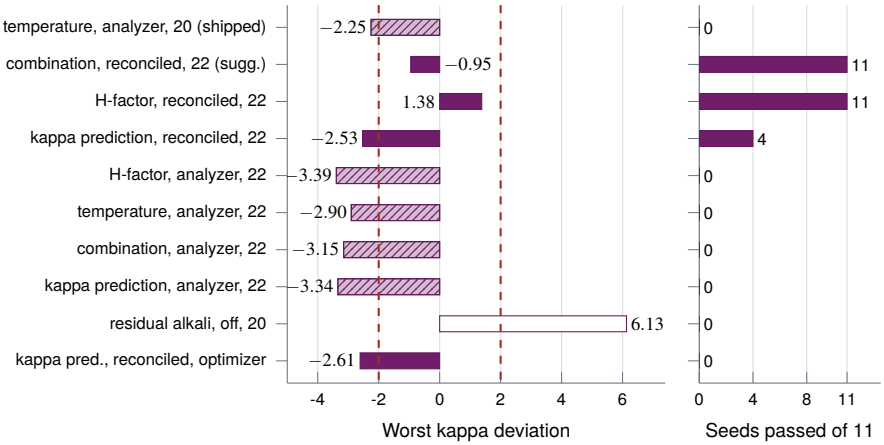


Figure 4.3 The challenge day by structure, from Table 4.5: the worst deviation of the true kappa from 30 over the eleven seeds (left, dashed lines at ± 2) and the seeds that held kappa within ± 2 and the residual alkali on its floor (right). Rows name the held variable, the feedback and the charge in percent. Dark bars are feedback through the reconciled estimate, hatched bars the raw analyzer, the white bar no feedback. Notice that every structure on the raw analyzer falls below -2 because of the drift, and that only the combination and the H-factor on the reconciled estimate pass every seed.

Table 4.5 The Challenge at the Case's Seed and Seeds 1–10. Worst Kappa Deviation and Lowest Residual Over the Eleven, Steam and Alkali per Tonne of Pulp at the Case's Seed

Held, feedback, charge	Pass	Worst	E_r low (g/L)	Steam (t/t)	Alkali (kg/t)
temperature, analyzer, 20 (shipped)	0	−2.25	5.39	0.942	434
combination, reconciled, 22 (sugg.)	11	−0.95	10.07	0.915	487
H-factor, reconciled, 22	11	1.38	10.08	0.914	485
kappa prediction, reconciled, 22	4	−2.53	9.08	0.949	496
H-factor, analyzer, 22	0	−3.39	8.55	0.970	498
temperature, analyzer, 22	0	−2.90	8.85	0.951	492
combination, analyzer, 22	0	−3.15	8.70	0.961	495
kappa prediction, analyzer, 22	0	−3.34	8.58	0.971	498
residual alkali, off, 20	0	6.13	8.72	0.792	404
kappa pred., reconciled, optimizer	0	−2.61	4.83	0.952	465

or above 8 g/L, with every parameter as shipped and the held variable, the feedback and the charge left to the learner. The results over eleven seeds are listed in Table 4.5 and drawn in Figure 4.3. The shipped temperature loop with the analyzer trim keeps all 110 slugs of the case's seed within ± 2 only because the drift's over-cooking cancels wood B's under-cooking, and every one of them is under the floor. The suggested setting, the combination on the reconciled estimate at a 22 % charge, passes at every seed, as does the H-factor on the reconciled estimate. Every structure on the raw analyzer fails because of the drift, and at 20 % 12 of 121 structure-seed runs hold kappa while none also holds the floor. The kappa prediction on the reconciled estimate passes 4 of 11, because its analyzer gate, which counts reports on one side only, admits the weak hour's under-cooked slugs into $\hat{b}_P$ when they reach the analyzer an hour after wood B's, both reading high. Turning the optimizer on with that structure takes the true residual to 5.00 g/L, because the weak hour and wood B reach the titrator five hours after the charge was chosen. The steam and alkali per tonne are reported rather than judged, because holding kappa and the floor on wood B costs more of both than the shipped settings' off-floor pulp. The suggested setting at 16 h is shown in Figure 4.4.

The estimator's gate is also why a better estimate is not always a better feedback signal. On the

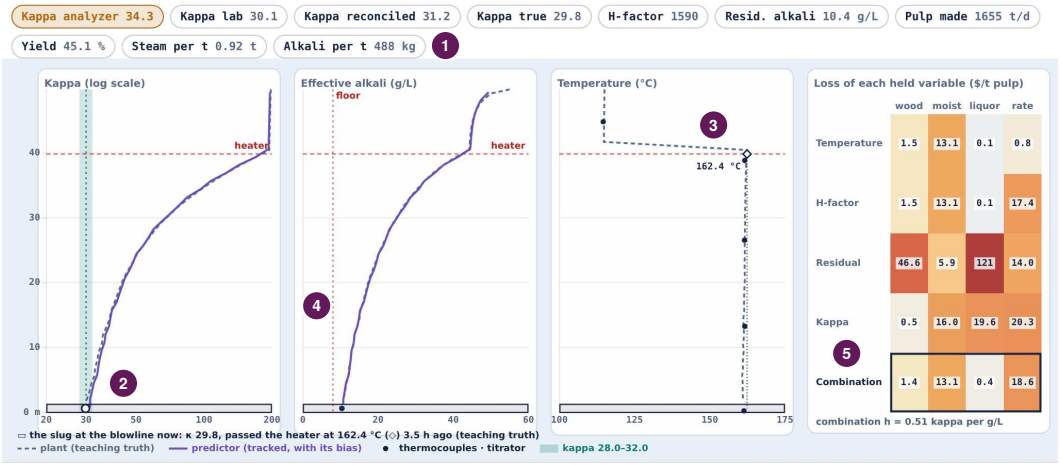


Figure 4.4 The challenge at 16 h with the suggested setting (combination on the reconciled estimate, 22 % charge), Profiles view, cropped to the readings and the profiles. (1) The readings: the analyzer reads 34.3 with its drift, the laboratory 30.1, the reconciled estimate 31.2 and the true kappa 29.8, the chips entering are predicted to gather an H-factor of 1590, the titrator reads 10.4 g/L, and 1655 t/d of pulp leaves at 45.1 % yield, 0.92 t of steam and 488 kg of alkali per tonne. (2) The slug at the blowline, which passed the heater at 162.4 °C 3.5 h earlier. (3) The heater crossing on the temperature profile. (4) The residual alkali floor on the alkali profile. (5) The loss table at the 22 % charge, which boxes the combination's row ($h = 0.51$ kappa per g/L).

same day without the drift, at 20 %, the H-factor held kappa within ±2 on the raw analyzer at 11 of 11 seeds and on the reconciled estimate at none. The reconciled estimate holds back a large new offset for 75 minutes and then moves a tenth of the way per report, which is right for a predictor's bias and slow as the input of a PI loop that already waits five hours for the vessel.

🕮 **Pause.** Try Exercise 4.3 before reading on.

4.6.8 Code Verification

The ledgers close to rounding. Over the suggested challenge, 180 slugs that entered and left close their lignin to 1.4×10^{-15}, their alkali to 1.1×10^{-15} and their carbohydrate to 2.0×10^{-15} kg per kg of wood, the delivered charge equals 22 % times the liquor's strength for every slug including the nine of the weak hour, and the vessel's lignin closes to -2.3×10^{-13} of 1230.3 t. Through the hang-up's drop the vessel's lignin closes to -3.4×10^{-13} of 563.2 t. Each slug's H-factor agrees with an independent five-point Gauss-Legendre quadrature of Vroom's rate along its recorded temperature to 1.3×10^{-15}. A slug that crossed the heater after a step in case 1 reached the blowline 4.000 000 h later, and each of 32 slugs at 1725 t/d took 3.4783 h from heater to blowline. Doubling the slugs to 80 moves the suggested challenge's blowline kappa by at most 0.088 % with the noise off, halving the step to 5 s by 0.052 %, and neither changes a verdict at the seeds checked. The loss table, recomputed by an independent bisection and re-optimization, agrees to $0.01/t. A measuring script reproduces on the engine all 249 numbers quoted in the simulator's texts, catalog entries and change record.

4.6.9 Agreement With Published Data

Vroom's rate as coded is 1.01 % above the form with the constant 43.19 at 100, 150, 160, 170 and 171 °C, because of the rounded constant, and one hour at 171 °C gives 1014 against the 1000 quoted [8]. The bulk delignification rate at 35 g/L is 0.0158, 0.0380 and 0.0880 per minute at 150, 160 and 170 °C, inside the ranges 0.0117–0.0371, 0.0160–0.0631 and 0.0244–0.1125 measured for

six woods [11]. Its activation energy is Vroom's 134 kJ/mol, while the measured bulk values are 57–91 kJ/mol, so the model's rate rises 5.58 times from 150 to 170 °C against 2.09–3.22 measured, and its residual pool shares that activation energy. The kinetics are therefore teaching values, close to the published rates at the design temperature and steeper than them away from it. The trends of Section 4.4, constant kappa at constant H-factor and lower yield in a hotter cook, agree in direction with the laboratory cooks [8].

4.6.10 Runtime

At one simulated hour per second in a browser window of 1366 by 768 pixels, three runs of the challenge each took 24.0 s, the longest case's 36 h took 36.0–36.1 s, the largest of 12951 control calls took 10.2 ms, and the optimizer took 4.3–9.6 ms over 15 calls. These are bounds on one machine.

4.7 Technology in Practice

What Measures It

The cook is graded at the blowline, far in time from the moves that decide it. A 2003 report on continuous digesters put the dead time between the quality variables and manipulated inputs such as temperature at 4 to 8 h, and named the variability of the natural feedstock as the other difficulty [18]. Kappa is measured either by an online analyzer or in the laboratory [19]. The chapter keeps both, the analyzer every 15 min with a 20 min delay and a bias that can drift, and the laboratory every 4 h an hour late, in Equation (4.4).

> **TECHNOLOGY Sensor: The Online Kappa Analyzer**
>
> The kappa number is defined by a laboratory test: the volume of 0.1 N permanganate solution consumed by one gram of dry pulp under standard conditions, about seven times the lignin left on the pulp in percent [8]. An online analyzer automates the measurement. It collects its pulp samples itself and judges them by their optical properties through a calibration curve. Mills also keep models that infer kappa from process measurements, to warn the operators when the estimate and the analyzer disagree widely and to watch the analyzer's calibration [19]. A drifting calibration shows on a trend as the analyzer walking slowly away from the laboratory, the fault of case 5.

What Moves It

The cook has two handles. The alkali charge is set by the white liquor ratioed to the chip meter, and the temperature by the circulation heater, whose outlet the heater loop holds with the steam valve under the heat-balance feedforward of Equation (4.6). In a plant control system such loops are often nested or cascaded, the setpoint of one taken from the output of another [20]. Figure 4.5 draws the loops of Section 4.5 on field devices.

> **TECHNOLOGY Final Element: The Steam Valve of the Heater**
>
> A continuous digester is heated by steam injected directly or by circulating the cooking liquor through a heat exchanger [16], the arrangement of this chapter. Excess friction in the stem of a pneumatic control valve makes it move in jerky steps and ignore small reversals of its signal, a fault called stiction that is severe in poorly maintained valves without positioners, and an up-and-down test of the output in manual mode reveals it. An air-to-open valve fails closed on loss of air, and the safer failure for the process decides which action is chosen [21]. The model's valve has no friction, and the heater answers it through one lag of 0.1 h.

Where the Controller Runs

A digester's loops run in a distributed control system, which places a supervisory level over field controllers, with cycle times from milliseconds to minutes [20]. In the 2003 project a fundamental digester model ran beside a mill's system on an application node that exchanged data with it over an open communication protocol [18]. The application's controller executes once a simulated minute

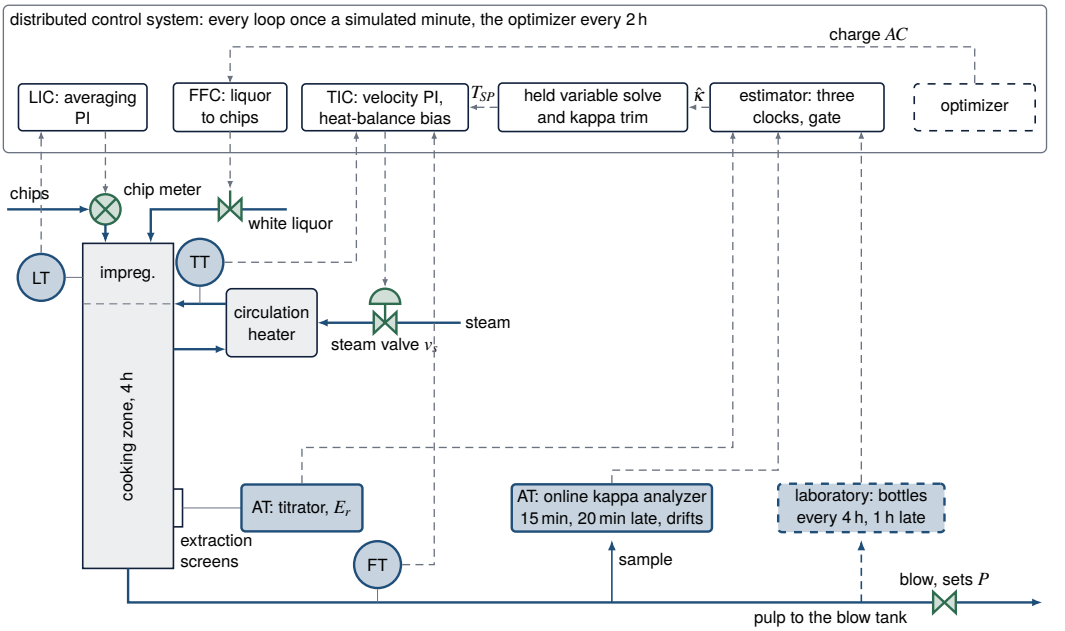


Figure 4.5 The loops of the digester on plant hardware. The heater loop (TIC) holds the circulation heater's outlet (TT) with the steam valve, the level loop (LIC) moves the chip meter, and the white liquor follows the chips at the charge *AC*. The held variable's solve writes the heater setpoint T_{SP} from the estimator, which reconciles the titrator at the extraction screens, the online kappa analyzer and the laboratory. Notice that every quality instrument sits at or below the screens, so each reads chips that passed the heater hours earlier.

on a plant integrated every 10 s, short beside a slug that passes every 7.5 min and a cook of 4 h. The velocity-form PI loops are the base layer. The held variable's solve, the estimator and the kappa trim sit above them as the supervisory layer that writes T_{SP}, and the economic optimizer runs above that every 2 h.

What Is Modernizing It

Three developments bear on this vessel. The first is predictive control. The 2003 project tested a predictive controller on step-response models against five PI loops, both sampled every 10 min, with the fundamental model standing in for the digester, and the PI loops needed a large tuning filter to avoid sustained oscillations [18]. The second is the adaptive soft sensor. A 2013 review reports a recursive regression model adapted to frequent process changes in a pulp digester, whose estimates fed back to a PID controller improved control [22]. The third is performance monitoring. The same project built a toolbox that judges loops from routine operating data, detecting oscillations and estimating dead time, with methods kept simple enough to supervise hundreds of loops [18]. Case 6, where a larger trim gain rings, is the kind of loop such a tool would flag.

> **IN PRACTICE** **Timing a Change to the Chip Front**
>
> A swing between softwood and hardwood is one of the transitions that digester operators and engineers must manage. One simple policy ramps each operating variable to its new value after a delay that roughly synchronizes the change with the moving front of the chips [18]. Mills align their records the same way. One study compared a chip sample taken at the chip bin with the kappa of pulp blown three and a half hours later [19]. The H-factor feedforward here times its moves by the slug entering the cooking zone.

4.8 Limitations

The vessel is a teaching design whose kinetics are fitted to one design point. Their bulk rate falls within published ranges at the design temperatures, but their temperature sensitivity is Vroom's index rather than a measured one, so a temperature move of several degrees is answered more strongly than the published rates suggest. Plug flow without a wash zone, effective alkali without sulfidity, a heater that sets the chips' temperature at once and an averaged wood mix leave out the radial, chemical and chip-scale structure that fundamental digester models carry [5]. The floor and the temperature limit are declared values, and the model does not simulate the lignin redeposition the floor stands for.

> ⚠ **OPEN CHALLENGE** **A Raw Material No Instrument Reads**
>
> An industry plan of 1999 gave priority to new sensors because the main raw material is naturally variable and because process operations often prevent timely and accurate measurement. Its priorities included cost-effective online sensors for pulp and fiber properties, and control systems that keep performing without deteriorating over time [23]. A 2003 report on continuous digesters added that real-time kappa measurement along the vessel is not generally feasible [18]. In this chapter the part of wood B that cooks slower reaches no instrument but kappa. A sensor that read the reactivity of the chips before the heater would be progress.

The estimator's analyzer gate counts reports on one side of the estimate, so an hour-long upset that follows a lasting change on the same side enters the predictor bias, as the kappa prediction's 4 of 11 shows. The optimizer holds the floor on its predictor without a back-off, so the true residual sits on the floor on average. The loss table is a steady-state analysis at declared prices, and the combination's weight changes with the charge, from -0.0869 at 20 % to 0.514 at 22 %. Every run is one seed of instrument noise at a time, and the challenge's statements rest on eleven.

> ⚠ **OPEN CHALLENGE** **When Should an Estimate Believe a Change?**
>
> Soft sensors must be updated as a process changes, and a 2013 review sets out the trade-off. Recursive updates adapt gradually, can adapt excessively when a process runs in a narrow range for a while, and cannot follow abrupt changes. In one plant a model updated daily from the laboratory kept a bias after each change of operating condition [22]. The estimator of this chapter meets the same choice in its persistence gate. The gate ignores an hour-long upset but admits the weak hour that follows a new wood, so the kappa prediction passes 4 of 11 seeds. A rule that told a lasting change from a passing one would be progress.

4.9 Conclusion

This chapter studied a continuous kraft digester, a plug-flow reactor whose quality is measured hours after it is decided, on a simulator that carries the chip column as 40 Lagrangian slugs whose lignin, alkali and H-factor close to rounding. A two-degree heater step reaches the blowline 4.125 h later and intact, and a 15 % production step leaves chips past the heater that no heater setting can rescue (worst slugs $+2.46$ to $+4.49$). The H-factor feedforward restores kappa one cooking residence later, holding 29.88–30.42. On a new wood the residual alkali is the worst held variable at \$119.68/t and the combination of predicted kappa and residual alkali the best at \$7.27/t, and the reconciled estimate holds kappa at 28.68–29.51 under a drifting analyzer that drives the raw loop to 26.52. On the graded day the combination on the reconciled estimate at a 22 % charge passes 11 of 11 seeds. The kinetics are teaching values whose temperature sensitivity exceeds the measured one. A fuller model would add a wash zone, sulfidity, a back-off for the optimizer's floor and a level test on both gate channels. The simulator shows that on a long plug-flow vessel the decision is what to hold between measurements, not how hard to chase the last one.

KEY POINTS

- In a continuous digester every heater move rides down with the chips and reaches the blowline hours later with its shape intact.
- No heater setting rescues chips already past the heater when production rises, and the H-factor feedforward restores kappa one cooking residence later.
- On a new wood the residual alkali is the worst variable to hold, at $119.68/t, and the combination of predicted kappa and residual alkali the best, at $7.27/t.
- The reconciled estimate removes a drifting analyzer's bias, yet without the drift the raw analyzer is the better feedback signal, because the estimate waits out short upsets.

Where else this applies. Volume 9, Chapter 10 chooses the setpoint a fast loop holds under a quality that answers in hours, and Chapter 10 corrects a soft sensor from a delayed laboratory in the same way.

References

[1] J. H. Lee and A. K. Datta, "Nonlinear inferential control of pulp digesters," *AIChE Journal*, vol. 40, no. 1, pp. 50–64, 1994.

[2] R. D. Gudi, S. L. Shah, and M. R. Gray, "Adaptive multirate state and parameter estimation strategies with application to a bioreactor," *AIChE Journal*, vol. 41, no. 11, pp. 2451–2464, 1995.

[3] M. Morari, Y. Arkun, and G. Stephanopoulos, "Studies in the synthesis of control structures for chemical processes: Part I: Formulation of the problem. process decomposition and the classification of the control tasks. analysis of the optimizing control structures," *AIChE Journal*, vol. 26, no. 2, pp. 220–232, 1980.

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

[5] P. A. Wisnewski, F. J. Doyle, and F. Kayihan, "Fundamental continuous-pulp-digester model for simulation and control," *AIChE Journal*, vol. 43, no. 12, pp. 3175–3192, 1997.

[6] J. J. Castro and F. J. Doyle, "A pulp mill benchmark problem for control: problem description," *Journal of Process Control*, vol. 14, no. 1, pp. 17–29, 2004.

[7] W. Gilbert and B. Allison, "Control of continuous digester kappa number using generalized model predictive control," *TAPPI Journal*, vol. 23, no. 9, pp. 467–476, 2024.

[8] T. E. S. Segura, M. A. Vivian, E. A. Bonfatti Júnior, C. Sarto, F. Schmidt, and F. G. d. Silva Júnior, "Effectiveness of the H-factor for controlling eucalyptus kraft pulping," *Scientia Forestalis*, vol. 47, no. 124, pp. 791–798, 2019.

[9] R. R. Gustafson, C. A. Sleicher, W. T. McKean, and B. A. Finlayson, "Theoretical model of the kraft pulping process," *Industrial & Engineering Chemistry Process Design and Development*, vol. 22, no. 1, pp. 87–96, 1983.

[10] D. Almeida, R. B. Santos, P. W. Hart, and H. Jameel, "Hardwood pulping kinetics of bulk and residual phases," *TAPPI Journal*, vol. 14, no. 10, pp. 652–662, 2015.

[11] D. Almeida, H. Jameel, P. Hart, and R. Santos, "Hardwood kraft pulping kinetics," in *7th International Colloquium on Eucalyptus Pulp*, Vitória, Espírito Santo, Brazil, May 2015. [Online]. Available: https://www.eucalyptus.com.br/artigos/18_Diego+Almeida1.pdf

[12] R. Amirthalingam and J. H. Lee, "Subspace identification based inferential control applied to a continuous pulp digester," *Journal of Process Control*, vol. 9, no. 5, pp. 397–406, 1999.

[13] B. Gough and J. T. Kay, "Kraft continuous digester effective alkali control," *IEEE Industry Applications Magazine*, vol. 3, no. 4, pp. 48–53, 1997.

[14] N. Padhiyar, A. Gupta, A. Gautam, S. Bhartiya, F. J. Doyle, S. Dash, and S. Gaikwad, "Nonlinear inferential multi-rate control of Kappa number at multiple locations in a continuous pulp digester," *Journal of Process Control*, vol. 16, no. 10, pp. 1037–1053, 2006.

[15] R. Rantanen, T. Ahvenlampi, M. Tervaskanto, U. Kortela, and A. Korhonen, "Prediction and control of blow-line kappa number," *IFAC Proceedings Volumes*, vol. 38, no. 1, pp. 495–500, 2005.

[16] M. Miller, M. Justiniano, and S. McQueen, "Energy and environmental profile of the U.S. pulp and paper industry," Energetics, Incorporated, for the U.S. Department of Energy, Industrial Technologies Program, Columbia, MD, Tech. Rep., Dec. 2005.

[17] V. Alstad and S. Skogestad, "Null space method for selecting optimal measurement combinations as controlled variables," *Industrial & Engineering Chemistry Research*, vol. 46, no. 3, pp. 846–853, 2007.

[18] F. J. Doyle III, "Final report: Model-based approach to soft-sensing and diagnosis for control of a continuous digester," University of Delaware, Department of Chemical Engineering, Newark, DE, Technical Report DOE-ID13882-F, Jun. 2003.

[19] F. M. Correia, J. V. H. d'Angelo, G. M. Almeida, and S. A. Mingoti, "Predicting kappa number in a kraft pulp continuous digester: A comparison of forecasting methods," *Brazilian Journal of Chemical Engineering*, vol. 35, no. 3, pp. 1081–1094, 2018.

[20] 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.

[21] 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.

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

[23] American Forest & Paper Association, "Agenda 2020: The path forward, an implementation plan," American Forest & Paper Association, Washington, DC, Tech. Rep., 1999.

4.10 Exercises

4.1 ⌕ Concept check · 2 min

In case 4, 60 % of the chips become wood B at 1 h. Holding the H-factor, the true kappa settles near 32.7. Holding the residual alkali at the extraction screens, it climbs to 34.8 by 14 h. Why does the held variable measured in the cook itself do worse?

(a) The titrator reports 10 min late, so the heater oscillates around the residual alkali target

(b) Wood B cooks more slowly, which lowers the H-factor, and only the residual alkali reveals it

(c) The residual alkali target of 10.1 g/L drives the heater to its lower limit of 140 °C

(d) Wood B's acid groups consume alkali, the residual falls, and the heater cools to put it back

4.2 ⌕ Concept check · 1 min

In case 2 the production rises 15 % at 1 h. The chapter finds that no heater setting, with feedforward, analyzer feedback or both, keeps every slug within ±2 of kappa 30 through the step. What limits every structure?

(a) The chips already in the cooking zone have crossed the heater, and no later heater move reaches them

(b) The steam valve reaches full opening when the heat load of the chips rises by 15 %

(c) The H-factor stops predicting kappa once the production rate has changed

(d) The analyzer reports 20 min late, which is too slow for the kappa trim

4.3 ⌕ Concept check · 2 min

On the challenge day without the analyzer drift, at a 20 % charge, the H-factor with the trim on the raw analyzer holds every slug within ±2 at 11 of 11 seeds, and the same structure on the reconciled estimate at none. Why is the better estimate the worse feedback signal here?

(a) The reconciled estimate carries the laboratory's noise, and the trim amplifies it

(b) The estimator holds back a large new offset for more than an hour, then moves a tenth of the way per report

(c) The laboratory's delay of an hour is added to the loop's dead time and makes the trim unstable

(d) With the H-factor held, Reconciled moves the bias into the solve and the trim rests at zero

4.4 Theory and calculation · 6 min

Vroom's rate is $k_V(T) = \exp(43.2 - 16113/T)$ per hour with T in kelvin. The cooking zone holds $0.8M$ with $M = 664.9\,\mathrm{t}$ of dry wood, and the column descends at $W = P/(24 \times 0.47)$ tonnes of wood per hour for a production P in tonnes of pulp per day.

(a) Find k_V at 100 °C and 160 °C, the fractional rise of k_V per kelvin at 160 °C, and the factor by which a step from 160 to 162 °C multiplies the rate. *Checkpoint:* 1.019 and 403.6 per hour, 8.59 % per kelvin, and 1.186.

(b) Find the cooking residence at 1500 and 1725 t/d, and the share of the design cooking H-factor that a slug gathers at the higher rate and a fixed temperature. *Checkpoint:* 4.000 h and 3.478 h, a share of 0.8696.

(c) From the slope of part (a), estimate the heater rise that restores the cooking H-factor at 1725 t/d. The H-factor feedforward of case 3 computes 161.53 °C. Say why it asks for less. *Checkpoint:* About 1.63 °C. The shorter cook cools the chips less.

4.5 Theory and calculation · 6 min

The heater's bias is the heat-balance feedforward of Equation (4.6), with $c_w = 1.35\,\mathrm{kJ/(kg\,K)}$, $c_l = 4.0\,\mathrm{kJ/(L\,K)}$, $\lambda = 2000\,\mathrm{kJ/kg}$, $\dot{m}_{s,max} = 90\,\mathrm{t/h}$, $T_{imp} = 115\,^\circ\mathrm{C}$ and $M = 664.9\,\mathrm{t}$.

(a) Find $T_{in,0}$ and the valve u_b at the design flow, 132.98 t/h, and $T_{SP} = 160\,^\circ\mathrm{C}$. *Checkpoint:* 114.72 °C and 58.0 %.

(b) Find u_b after the production step of case 3, at 152.93 t/h and $T_{SP} = 161.53\,^\circ\mathrm{C}$. *Checkpoint:* About 68.9 %.

(c) Without the feedforward, how far must the heater outlet fall below its setpoint before the proportional term alone, at $K_c = 2.5$ % per °C, opens the valve by the extra amount that part (b) needs? *Checkpoint:* About 4.4 °C.

4.6 Theory and calculation · 7 min

In case 5 the analyzer's bias drifts upward at 0.3 kappa per hour from 1 h. Laboratory bottles are drawn at 0, 4, 8, … hours and reported an hour later, each paired with the analyzer sample drawn at the same instant, and the analyzer bias estimate is updated with $\lambda_A = 0.8$ as in Equation (4.9).

(a) Using Equation (1.8), find the mean age of a held analyzer result and of a held laboratory result. *Checkpoint:* 0.458 h (27.5 min) and 3 h.

(b) Show that, with the noise left out, the estimate lags a ramp that rises by r between bottles by $r(1-\lambda_A)/\lambda_A$ once the update has settled. Find the lag and the estimate after the report at 21 h. *Checkpoint:* A lag of 0.30 and an estimate of 5.40, against a bias of 5.70 in the bottle drawn at 20 h.

(c) Find the true bias at 24 h and the part of it the estimator has not yet seen. Compare with the chapter's 5.56 and 1.34. *Checkpoint:* 6.90, and about 1.50 unseen.

4.7 App challenge · 7 min

Load scenario 1, Open loop: a temperature step and the four-hour silence, with the default seed. Press Run at a speed of 0.1 h/s, press Pause when the clock is between 0.5 and 0.6 h, and raise Cooking temperature setpoint from 160 to 162 °C. Set the speed to 1 h/s and press Run. The case pauses at 10 h. Export the CSV file.

(a) Find the first row after your step at which `truth_kappa_blowline` is below 29, and the time from your step to that row. *Checkpoint:* Between 4.0 and 4.2 h, for example the row at 4.633 h for a step at 0.50 h.

(b) Find the first row at which `kappa_analyzer` is below 28, and the instant its sample was drawn, column `kappa_analyzer_sample_h`. *Checkpoint:* Reported at 5.083 h from a sample drawn at 4.75 h.

(c) Read `truth_kappa_blowline` at 10 h, and find the change in kappa per kelvin of cooking temperature. *Checkpoint:* 27.28, within 0.05, so about -1.36 kappa per kelvin.

4.8 App challenge · 9 min

Load scenario 4, The wood changes: which variable to hold, with the default seed and Held variable as loaded, H-factor. Run at 1 h/s. The case pauses at 12 h. Press Run again and pause at 14 h. Export the CSV file and read `cook_T_sp_C` and `truth_kappa_blowline` in the row at 14 h. Set Held variable to Residual alkali, press Reset and repeat. Set Held variable to Kappa prediction, press Reset and repeat, and also read the titrator, column `EA_residual_gL`.

(a) Report the setpoint and the true kappa for each held variable. *Checkpoint:* H-factor 160.00 °C and 32.69. Residual alkali 158.35 °C and 34.76. Kappa prediction 160.73 °C and 31.80. Within 0.05 °C and 0.05.

(b) Report the titrator reading under the kappa prediction, and say what it means for the alkali charge. *Checkpoint:* 7.68 g/L, within 0.05 g/L, under the 8 g/L floor.

(c) Explain the ranking with the loss table, and why none of the three reaches 30. *Checkpoint:* \$119.68, \$29.45 and \$9.90 per tonne.

4.9 App challenge · 8 min

Load scenario 5, The analyzer drifts and the laboratory corrects it, with the default seed and the settings as loaded: Held variable Kappa prediction and Kappa feedback Analyzer. Run to 24 h at 1 h/s and export the CSV file.

(a) In the row at 24 h, read four columns: the true kappa `truth_kappa_blowline`, the analyzer reading `kappa_analyzer`, the true bias `truth_analyzer_bias` and the estimated bias `bias_analyzer`. *Checkpoint:* 26.52, 33.81, 6.90 and 5.56, within 0.02.

(b) The kappa trim has integral action, yet the analyzer stays more than three units above 30. Estimate the steady error of the trim on a ramp, with $K_c = 0.15\,^{\circ}\mathrm{C}$ per kappa, $\tau_I = 3\,\mathrm{h}$ and the gain of Exercise 4.7. *Checkpoint:* About 4.4 kappa.

(c) Set Kappa feedback to Reconciled, press Reset and run to 24 h again. Find the lowest and highest `truth_kappa_blowline` from 12 to 24 h. *Checkpoint:* 28.68 and 29.51, within 0.02.

4.10 Application area · 4 min

In case 7 the economic optimizer lowers the charge from 20.00 to 18.83 %. From 7 to 12 h the alkali falls from 425.0 to 396.3 kg per tonne of pulp and the steam rises from 0.814 to 0.817 t/t, and the yield of the slug blown at 12 h rises from 46.97 to 47.60 %. The declared prices are \$70 per tonne of dry wood, \$150 per tonne of alkali as NaOH and \$20 per tonne of steam.

(a) Find the saving per tonne of pulp and per day at 1500 t/d. *Checkpoint:* About \$6.22 per tonne, about \$9300 per day.

(b) The true residual alkali then runs at 7.72 to 8.39 g/L, under the 8 g/L floor in 49 % of the steps. Estimate the back-off on the predicted floor that keeps this range above 8 g/L, and its cost, taking $\partial E/\partial \ln AC = 31\,\mathrm{g/L}$. *Checkpoint:* About 0.3 g/L, about 3.6 kg of alkali and \$0.54 per tonne.

4.11 Application area · 3 min

With the challenge's shipped settings, temperature held, the trim on the raw analyzer and a 20 % charge, all 110 slugs blown from 12 to 24 h at the default seed are within ± 2 of kappa 30, and all 110 are under the 8 g/L floor, the lowest at 5.6 g/L.

(a) Why does kappa look good? *Checkpoint:* The drift's over-cooking cancels wood B's under-cooking.

(b) Which instrument shows the problem during the day, and why does a mill keep a residual alkali floor at all? *Checkpoint:* The titrator at the extraction screens, below 8 g/L from about 7.7 h. A floor keeps dissolved lignin from redepositing on the fibers.

4.12 Review the agent · 5 min

An AI assistant is asked for a setting that passes the challenge. It replies: "Hold the Kappa prediction with Kappa feedback on Reconciled and a 22 % charge. It uses the best estimate in the app, and it passes: 110 of 110 slugs within ± 2, worst -1.9 at 15.1 h, lowest residual 9.5 g/L. The loss table's combination only adds a heuristic weight on the residual alkali and cannot do better." Find the flaw, and say how you would check it.