INTERACTIVE CONTROL ENGINEERING
Principles and Applications

Volume 16 Control in Iron, Steel, and Metals Production

CHAPTER 8

Hot Strip Mill Simulator for Inferential Gauge Control With a Sign-Inverted Error

Pages 159–186

Edited by John D. Hedengren

How to cite this chapter. "Hot Strip Mill Simulator for Inferential Gauge Control With a Sign-Inverted Error," in *Control in Iron, Steel, and Metals Production*, John D. Hedengren, Ed., vol. 16 of *Interactive Control Engineering: Principles and Applications*. Advanced Process Solutions, LLC, 2026, ch. 8, pp. 159–186.

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

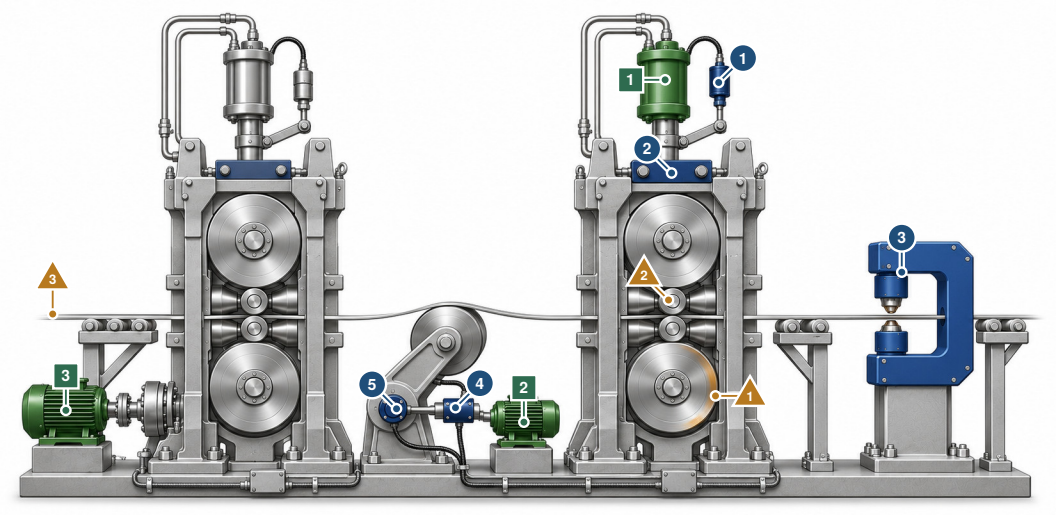
The app for this chapter
https://apmonitor.com/apps/hot-strip-mill/

The volume online: contents, other chapters and citation records
https://apmonitor.com/apps/books/v16-iron-steel-metals/#hot-strip-mill

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

CHAPTER 8

Hot Strip Mill Simulator for Inferential Gauge Control With a Sign-Inverted Error



The controller moves 1 capsule that sets the F7 roll gap 2 looper motor that holds the strip tension 3 F6 drive, trimmed to hold the looper angle

The controller reads 1 capsule position 2 roll force 3 X-ray gauge, 5 m downstream 4 strip tension 5 looper angle

What disturbs it 1 back-up roll out of round 2 work rolls growing as they heat 3 incoming strip, thick patches and cold bands

Notice that capsule position and roll force are read at the stand at once, while the X-ray gauge sees the strip only half a second later. Illustration, not to scale.

IN PLAIN TERMS **Rolling Hot Steel to an Exact Thickness**

The main idea. A hot strip mill squeezes red-hot steel between heavy rolls into a thin sheet. The squeeze stretches the frame that holds the rolls, so the set gap is not the thickness that comes out.

The challenge. The thickness can be worked out at once from how hard the rolls push, or read half a second late by a meter down the line. The quick answer is right when the steel changes but points the wrong way when the rolls warm and grow, so the late meter must correct it.

Why it matters. Strip that is too thick or too thin sells for less or is rolled again.

THREE WAYS TO READ THIS CHAPTER

Explore	12 min	The first page, the two history boxes of Section 8.2, the fact box of Section 8.3, the fact and try-it boxes of Section 8.7, and the key points
Study	60 min	The whole chapter with its pause points and the twelve exercises
Practice	20 min	The first page, Section 8.7, Section 8.8, Section 8.9 and its open challenge

The control question A finishing stand can infer the strip thickness at once from its roll gap and roll force, or measure it half a second later with an X-ray gauge downstream. For which disturbances is the fast inferred gauge exact, for which is it wrong the other way, and what must sit underneath it?

Learning objectives

- Derive the gap gain and the deadbeat gain of a stand from its mill modulus and the plastic modulus of the pass.
- Sort a disturbance by the gaugemeter error identity into one the inferred gauge reads exactly and one it inverts or biases.
- Explain from the length-indexed dead time why a loop on the X-ray gauge cannot carry the gain of a loop on the gaugemeter.
- Set up a monitor trim on the X-ray gauge, with its action and its speed relative to the gauge loop.
- Choose a gauge control architecture for a coil by asking which measurement can see each disturbance.

Prerequisites Sections 1.4.6 (mill spring and gaugemeter), 1.4.5 (delay as a length), 1.3 (loop action), 1.6.2 and 1.6.4.

The app https://apmonitor.com/apps/hot-strip-mill/

See also Volume 5, Chapter 3 for a loop closed through a slow analyzer, Volume 5, Chapter 10 for a soft sensor trimmed by a delayed laboratory and Volume 5, Chapter 13 for an estimate corrected by late fixes.

8.1 Introduction

A property that a plant is paid for is often measured somewhere other than where it is made. A composition leaves the column before the analyzer sees it, a coating weight is read a frame downstream of the applicator, and a thickness is read after the metal has left the rolls. The standard answer is to compute the property from instruments that sit at the point of manufacture, run the loop on that computed number, and correct it slowly from the honest measurement when the honest measurement finally arrives [1–3]. The arrangement is old, it is everywhere, and the usual way of describing its weakness is that the inferred value drifts.

Drift is the wrong word for what goes wrong. An inferred measurement is built out of a model, and the model is exact for whatever it accounts for and blind to whatever it does not. For the blind class the error is not a slow random walk. It is a deterministic function of the disturbance, and it can carry the opposite sign to the change in the true property and be several times larger. A loop closed on such a measurement then amplifies that disturbance while holding its own measurement exactly on setpoint, and the operator sees a controller working perfectly.

A hot strip mill finishing stand is the clearest instance of this in industry, because both halves of the argument are arithmetic rather than statistics. The stand is a spring that deforms under the load it applies, so the delivered thickness is the commanded roll gap plus the roll force divided by the mill modulus. At the design pass of this chapter's simulator the gap is 1.000 mm, the stretch is 2.000 mm and two thirds of the product is machine deflection. Read the other way the same relation is a thickness measurement, computed from a position transducer and a load cell that answer in milliseconds, which is what the British Iron and Steel Research Association gaugemeter has been since the 1950s [4, 5]. The X-ray gauge that measures the strip itself stands 5.0 m past the stand, which at 10.00 m/s is 0.500 s of pure dead time, and on an entry thickness wave the two measurements at identical tuning give 11.6 µm and 173.8 µm of root-mean-square deviation.

The inversion is what the simulator is built to show. Every term of the gaugemeter except the modulus is measured at the stand, so any disturbance that arrives with the metal, a cold band, a thick patch on the transfer bar or a change in interstand tension, appears in the roll force and the inferred value matches the truth once the force has settled, straying by at most 1.5 µm while it moves. Anything that moves the roll surface without moving the position transducer does the opposite. Work rolls growing thermally, back-up rolls that are not round and a modulus calibration taken on

a different roll set all close the gap, raise the force, and are read as a thicker strip when the strip is getting thinner. The error is the whole of the gap offset, while the true gauge moves by only 0.2104 of it, so the error is 4.75 times the size of the real change and points the other way. Measured open loop, 200 µm of thermal crown moves the delivered gauge −41.9 µm and the gaugemeter reading +156.7 µm, an error of +198.6 µm. With the gaugemeter loop closed at the tuning that wins every other comparison in this chapter, the delivered gauge ends 198.6 µm thin, which is the entire crown.

The class generalizes past steel. Any plant that runs a fast estimator under a slow reference measurement inherits the same division of labor and the same exposure, including analyzer-trimmed property models, laboratory-corrected soft sensors and inertial platforms corrected by satellite fixes [2, 3]. On a moving-material line it inherits a second property, because every transport delay is a length divided by a speed and every periodic disturbance is a wavelength in delivered material rather than a period in seconds [6]. However, a student handed a soft sensor as a transfer function block never meets an upset whose estimated error points the wrong way. This chapter works with a hot strip mill simulator in which the same inferred gauge is exact for one class of disturbance and inverted for another, so that the choice of measurement appears as a hundred micrometers of product rather than as a remark about model mismatch. The mill spring, the roll force model and monitor gauge control are all standard practice. What the chapter adds for the reader is an exact error identity for the inferred gauge, used to order four control architectures against six disturbances, an ordering in which the architecture decides the outcome and the tuning barely moves it. Section 1.4.6 introduced the mill spring and the gaugemeter error on the running example W-1, and this chapter measures them on a full stand.

Section 8.2 traces gauge control on a finishing train from the screwdown to monitor gauge control. Section 8.3 specifies the thirteen-state two-stand model down to its implicit stand solution, its instrument chain and its length-indexed transport delays, with every constant and initial value, and Section 8.3.8 states where each simplification fails. Section 8.4 derives the error identity and sets the inferred gauge against the delayed X-ray gauge. Section 8.5 describes the five loops, the architecture switch and the eccentricity estimator. Section 8.6 lists the fifteen scenarios, which open on a roll gap moved by hand and close on one coil carrying five disturbances and a speed change. Section 8.7 reports what each scenario measures, and Section 8.8 sets the loops in their real hardware. Sections 8.9, 8.10 and 8.11 give the limitations, the lessons that transfer outside steel and a summary, and Section 8.12 holds the exercises.

8.2 Gauge Control on a Finishing Train

Strip was rolled to thickness for half a century before anything on a mill measured thickness. A rolling schedule fixed a screwdown position for each stand from the incoming bar, the ordered gauge and experience of the grade, the operator adjusted it by hand between coils, and the product was checked with a micrometer on the cold coil hours later. That arrangement needs a large margin, because strip below the ordered thickness cannot be shipped. The theoretical apparatus arrived first, with roll pressure through the arc of contact put on a defensible footing by the slab analysis of the early 1940s [7] and reduced to a form a mill engineer could evaluate for a hot pass a decade later [8]. Both give the force. Neither gives the thickness, because the force deflects the machine.

The step that turned the deflection from a nuisance into an instrument was taken at the British Iron and Steel Research Association, and the two history boxes of this section tell how the gaugemeter and the hydraulic capsule came about.

The second measurement arrived in parallel. Radiation gauges, first isotope and then X-ray, made an absolute traceable thickness available on the runout table, and because the reading describes strip that has already been made they were put to work correcting the inferred gauge rather than closing the fast loop. Monitor gauge control, in which the fast gaugemeter carries the regulation and the slow absolute gauge trims its bias, was standard on hot mills through the 1960s and remains the structure a

modern finishing train runs [5, 9]. The two measurements are complementary in exactly the way the frequency content of the disturbances requires, and no later technique has displaced the arrangement.

> ⧖ HISTORY **The Gaugemeter and Lock-On (1950s)**
>
> Before 1960 the British Iron and Steel Research Association developed automatic gauge control based on the load cell, and with it a gauge meter formula that would drive the later progress of gauge control [10]. In its first form the formula computed the exit thickness at the head end of each strip from the roll gap and the roll force, and the result was stored, a step called lock-on. The roll gap was then adjusted so that the thickness of the rest of the strip agreed with the stored value, which removed the unevenness along the rolling direction. The target was the head end's own thickness. Close study of the mill modulus in the 1960s then opened the way to the absolute mode of automatic gauge control [10].

Three refinements were added to that skeleton and all three are in the simulator. Mass-flow control uses the fact that the volume rate through each stand must match, so the thickness ratio across two stands is the speed ratio their drives have to run at, which keeps the looper from being emptied every time a gauge controller moves a gap. Looper and interstand tension regulation became a subject of its own, because the looper angle is a pure integral of the mass-flow mismatch and the tension it sets feeds back into the roll force and therefore into the gauge [6]. Roll eccentricity compensation answered the one disturbance neither gauge loop can reach, by demodulating the force signal against the measured back-up roll angle and subtracting the once-per-revolution component before the force is divided by the modulus [11]. A century of rolling theory is reviewed in [10], and the modeling detail behind the schedules, the mill spring calibration and the automatic gauge control architectures is collected in [12, 13].

What did not change is the exposure. The gaugemeter is a measurement of the mill, not of the strip, and it is only as good as the mill's calibration was on the day it was taken. Stand stiffness governs achievable rolling accuracy and is quoted as a chord over the working range, because the force against stretch curve is not linear at low load [14]. A modulus measured on one roll set, one width and one oil temperature is used on all of them until somebody retakes it.

> ⧖ HISTORY **The Hydraulic Screwdown (Late 1960s)**
>
> A rolling mill with a hydraulic screw-down device was developed at the end of the 1960s. It could change the mill modulus, and it entered tandem strip mills in the 1970s to raise the precision of gauge control [10]. Its control used a mechanical servo at first and later an electro-hydraulic servo, which moved the screw-down eight times faster than the conventional one. The speed brought a new problem. A fast roll gap loop and the interstand tension loop now responded at different rates and disturbed each other, and that interference opened the way to non-interference control of gauge and tension [10]. The capsule and servo valve of this chapter's F7 stand are of that kind.

8.3 Process Model

The simulated plant is stands F6 and F7 of a seven-stand finishing train, the looper between them, and the runout table as far as the X-ray gauge. 1200 mm of C-Mn strip enters F6 at 4.600 mm, leaves it at 3.600 mm, and leaves F7 at 3.000 mm and 880 °C, running 10.00 m/s at the F7 exit and 8.333 m/s into it. F6 takes 13.697 MN on a 1.400 mm gap and F7 takes 12.000 MN on a 1.000 mm gap, against mill moduli of 6.226 and 6.000 MN/mm. Interstand tension is 10.0 MPa, held by 13.535 kN m of looper torque at a 25° arm angle. The gauge tolerance is ±25 µm and the mill is paid by the meter inside it. The machine, its instruments and each symbol appearing below are drawn in Figure 8.1.

What the five loops reach, and what they do not, is sorted in Table 8.1. Units are those of a rolling mill, so thickness, gap and contact length appear in mm, roll force in MN, mill modulus in MN/mm, flow stress and strip tension in MPa, the strip path and the line speed in m and m/s, looper torque in kN m and temperature in °C. The clock is in seconds, because the transport delay sets the time scale

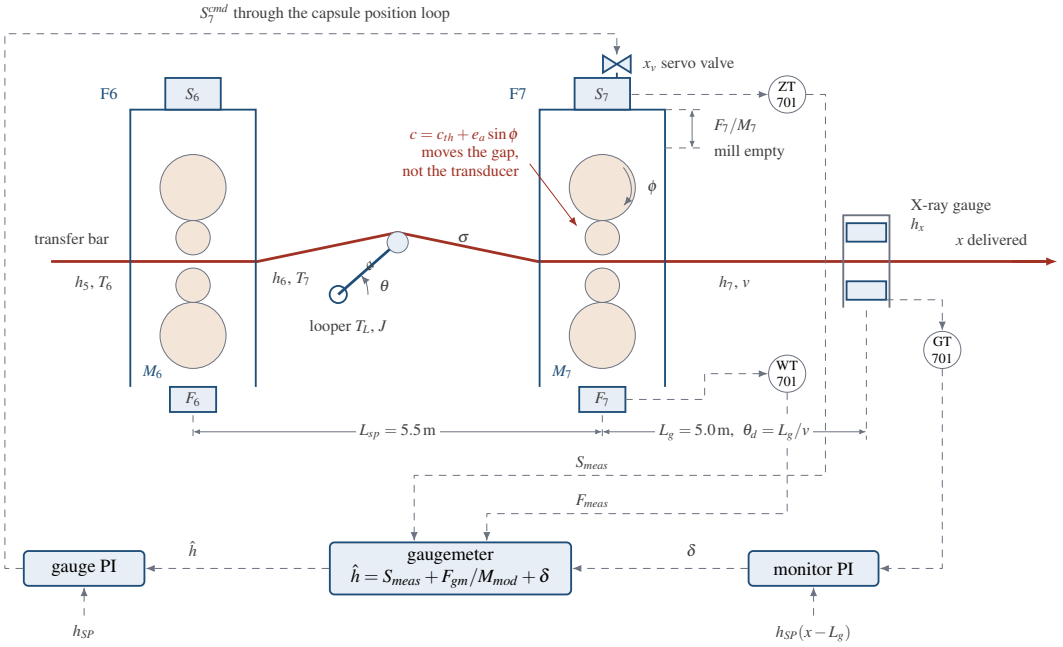


Figure 8.1 Stands F6 and F7 of a hot strip finishing train with the two rival thickness measurements. Each stand is a spring of modulus M_i standing F_i/M_i taller under load, so the delivered gauge h_7 is the commanded gap S_7 plus that stretch plus the gap offset c made by thermal crown c_{th} and runout $e_a \sin\phi$. Transducer ZT-701 and load cell WT-701 answer in milliseconds and the gaugemeter combines them into $\hat{h}$, while gauge GT-701 reads the strip itself L_g downstream for $\theta_d = L_g/v$ of transport delay. The looper sets the interstand tension σ, which subtracts from the roll force.

and the interesting events last tenths of a second. Thirteen states are integrated and every conversion between a meter and a millimeter is made once, at the point of use.

8.3.1 The Mill Spring and the Two Curves

A stand is an elastic structure loaded by the pass it is making. The housing, the screws, the chocks, the bearings and the oil column in the hydraulic capsule stretch together, and on a light hot gauge they stretch by more than the strip is thick. The delivered thickness is therefore the unloaded roll gap plus that stretch, and the roll force that sets the stretch is itself a function of the thickness, so the two have to be solved together as shown in Equation (8.1).

$$h_i = S_i + c_i + \frac{F_i(h_i)}{M_i}, \qquad i \in \{6,7\} \tag{8.1}$$

where h_i is the exit thickness in mm, S_i is the unloaded roll gap the position transducer reads, M_i is the mill modulus in MN/mm and c_i is everything that moves the roll surface without moving the transducer. F6 carries no such offset in this model, and F7 carries the one given in Equation (8.2).

$$c_7 = c_{th} + e_a \sin\phi \tag{8.2}$$

where c_{th} is roll thermal crown as a gap offset in mm, negative when the rolls grow into the strip, e_a is the back-up roll runout amplitude and ϕ is the back-up roll angle. Both terms are the subject of Section 8.4.

Equation (8.1) is the intersection of the elastic line of the stand with the plastic curve of the pass [4], and it cannot be marched. The plastic curve is nearly four times the steeper of the two, so

Table 8.1 Classification of Variables, With the Design Line-Up Every Scenario Opens On.

Symbol	Role	Nominal	Range
h_7	CV, unmeasured truth	3.000 mm	2.4–3.6 mm SP
$\hat{h}$, h_x	CV as inferred and as delayed	3.000 mm	—
S_7 (ZT-701)	CV of the capsule loop	1.000 mm	0.05–2.60 mm
σ	CV, measured	10.0 MPa	4–22 MPa SP
θ	CV, measured	25°	12–38° SP
S_7^{cmd}	MV of the gauge loop	1.000 mm	0.05–2.60 mm
x_v	MV of the capsule loop	0 %	±100 %
δ	MV of the monitor loop	0 mm	±0.25 mm
T_L	MV of the tension loop	13.535 kN m	0–32 kN m
r_6	MV of the angle loop	0 %	±5 %
v_{SP}	DV, operator set	10.0 m/s	4–14 m/s
S_6^{cmd}	DV, operator set	1.400 mm	0.6–2.4 mm
A_e, λ_e	DV, entry wave and its wavelength	0 %, 30 m	0–6 %, 8–120 m
ΔT_s, λ_s	DV, skid depth and pitch	0 °C, 125 m	0–60 °C, 40–300 m
e_a	DV, back-up roll runout	0 mm	0–0.080 mm
c_{tgt}	DV, thermal crown target	0 mm	−0.30–0.15 mm
ε_M	DV, modulus calibration error	0 %	±8 %
T_{bias}	DV, bulk temperature bias	0 °C	−50–40 °C
σ_x, σ_F	DV, gauge and load cell noise	0, 0	0–12 µm, 0–0.30 MN
x_v, S_7, S_6	states, valve and two capsules	0, 1.000, 1.400 mm	±1, 0–4 mm
v, r_6	states, the two drives	10.00 m/s, 0	0.3–14 m/s, ±0.05
x, ϕ	states, strip and roll angle	0, 0	m, 0–2π
θ, $\dot{\theta}$	states, looper arm	25°, 0	3–45°
σ, T_L	states, tension and torque	10.0, 13.535	0–60 MPa, 0–32
c_{th}, T_b	states, the two slow drifts	0, 0	as their disturbances

substituting the force back into the right-hand side amplifies an error by that factor on every pass and successive substitution diverges immediately. The application solves it by Newton iteration on the residual $r(h) = h - S - c - F(h)/M$, with the slope re-evaluated by a forward difference at every pass. The derivative $r'(h) = 1 + Q/M \approx 4.8$ is bounded away from zero and never changes sign, and the iteration stops at a residual of 10^{-13} mm. The solve is called eight times per physics step, four Runge-Kutta stages on two stands, and from the previous step's answer as a warm start it converges in two or three iterations.

> **DID YOU KNOW? Most of the Stretch Is in the Rolls**
>
> Of the elastic deformation of a rolling mill under load, the rolls take 40 to 70 %, the housing 10 to 16 % and the screw-down device 4 to 20 %. Bending and flattening of the rolls is by far the largest share [10].

8.3.2 Roll Force and Flow Stress

The separating force follows the slab treatment of the roll bite [7, 8], with the mean pressure taken as the plane-strain flow stress lifted by a Coulomb friction hill and lowered by both strip tensions, as shown in Equation (8.3).

$$F = w L_c \left[1.1547\, k_f - \tfrac{1}{2}(\sigma_b + \sigma_f) \right] Q_p, \quad L_c = \sqrt{R' \Delta h} \tag{8.3}$$

$$Q_p = \frac{h_m}{\mu L_c} \left(\mathrm{e}^{\mu L_c / h_m} - 1 \right) \tag{8.4}$$

where w is the strip width in mm, L_c is the contact arc length, $R' = R\,f_H$ is the work roll radius after Hitchcock flattening, Δh is the draft, h_m is the mean of the entry and exit thicknesses, μ is the roll-strip friction coefficient, and σ_b and σ_f are the back and front tensions. Millimeters times millimeters times N/mm^2 is newtons, so a factor of 10^{-6} converts the product to MN. At the design pass the arc is 16.72 mm against a mean thickness of 3.30 mm, which puts Q_p at 2.350, so the friction hill is most of the reason a 3 mm pass costs twelve meganewtons. The bracket in (8.3) is floored at 15 MPa so that a large tension cannot drive the computed pressure negative.

Flow stress is the reason a hot strip is a different material every second, as shown in Equation (8.5).

$$k_f = k_0\,\mathrm{e}^{-\beta(T-T_0)}\left(\frac{\max(1,\dot{\varepsilon})}{\dot{\varepsilon}_0}\right)^{m}, \quad \dot{\varepsilon} = \frac{\bar{v}}{L_c}\ln\frac{h_{in}}{h_{out}} \tag{8.5}$$

where k_0 is the flow stress at the reference temperature T_0 and reference strain rate $\dot{\varepsilon}_0$, β is the softening coefficient, m is the strain rate exponent, and $\bar{v}$ is the mean of the entry and exit strip speeds through the arc, converted to mm/s at the point of use. The temperature term carries the disturbance that matters. A skid mark 35 °C deep raises the flow stress by 16 %, and with both gaps frozen it puts 134.0 µm of extra thickness on the product.

8.3.3 The Two Plant Gains

Differentiating (8.3) at the operating point gives the plastic modulus, and combining it with (8.1) gives the two numbers every gauge argument in this chapter turns on, as shown in Equation (8.6).

$$Q = -\frac{\partial F}{\partial h}, \qquad \frac{\partial h}{\partial S} = \frac{M}{M+Q}, \qquad \frac{\partial h_{out}}{\partial h_{in}} = \frac{\partial F/\partial h_{in}}{M+Q} \tag{8.6}$$

At the design pass $Q = 22.513\,\mathrm{MN/mm}$ against a mill modulus of 6.000 MN/mm, so a gap move arrives at the product attenuated to 0.2104 of itself and an entry thickness variation walks through at 0.6534, both checked numerically by the application's verification suite. The reciprocal of the first, $(M+Q)/M = 4.752$, is the deadbeat gain, the gap move that cancels a gauge error in one step, and the shipped gauge loop is tuned just below it at 4.00. A mill is a soft spring against a stiff strip, which is why no hot mill has ever been run on gap position alone.

8.3.4 Actuators, Drives and the Looper

The F7 capsule is a velocity device. A servo valve spool follows the valve command through a short lag and the capsule integrates the spool position, so the position loop closes around an integrator and needs almost no reset action of its own, as shown in Equations (8.7) and (8.8).

$$\tau_v\frac{\mathrm{d}x_v}{\mathrm{d}t} = \frac{\mathrm{clamp}(u_v,\pm 100)}{100} - x_v \tag{8.7}$$

$$\frac{\mathrm{d}S_7}{\mathrm{d}t} = K_q\,x_v, \qquad \tau_{g6}\frac{\mathrm{d}S_6}{\mathrm{d}t} = S_6^{cmd} - S_6 \tag{8.8}$$

where x_v is the normalized spool position, u_v is the valve command in percent and K_q is the capsule slew rate at full valve in mm/s. F6 has no gauge control in this application, so its capsule is represented as a closed position loop with a single time constant. The drives, the delivered strip coordinate and the back-up roll angle follow Equation (8.9).

$$\begin{aligned} \tau_d\frac{\mathrm{d}v}{\mathrm{d}t} &= v_{SP} - v, & \tau_{s6}\frac{\mathrm{d}r_6}{\mathrm{d}t} &= \frac{r_6^{cmd}}{100} - r_6 \\ \frac{\mathrm{d}x}{\mathrm{d}t} &= v, & \frac{\mathrm{d}\phi}{\mathrm{d}t} &= \frac{v}{(1+f_s)R_b} \end{aligned} \tag{8.9}$$

where x is the strip delivered past F7 in meters, R_b is the back-up roll radius in meters and f_s is the forward slip at the roll exit. At 10.00 m/s the back-up roll turns at 13.87 rad/s, which is 2.207 Hz, and that frequency moves with the line speed while the runout amplitude does not.

The looper is a torque balance about its pivot and the interstand span is an elastic member, as shown in Equation (8.10).

$$\begin{aligned} J\,\ddot{\theta} &= 10^{3}\,(T_L - T_{grav} - T_{strip}) - c_L\,\dot{\theta} \\ \tau_T\,\frac{\mathrm{d}T_L}{\mathrm{d}t} &= T_L^{cmd} - T_L \\ \frac{\mathrm{d}\sigma}{\mathrm{d}t} &= K_s\left(v_{in,7} - v_6 + \frac{\mathrm{d}L}{\mathrm{d}\theta}\,\dot{\theta}\right) \end{aligned} \tag{8.10}$$

where J is the arm inertia in kg m^2, c_L is the hydraulic damping, $T_{grav} = W\cos\theta$ is the arm weight torque and T_{strip} is the torque the strip pull applies. The factor 10^3 converts the torques from kN m to N m for the inertia. The geometry is two straight runs from each roll bite up to a looper roll sitting $a\sin\theta$ above the pass line, as shown in Equation (8.11).

$$\begin{aligned} \ell &= \sqrt{a^2\sin^2\theta + (L_{sp}/2)^2}, \qquad L = 2\ell \\ T_{strip} &= \frac{2\,\sigma\,h_6\,w\,(a\sin\theta/\ell)\,a\cos\theta}{10^{3}} \end{aligned} \tag{8.11}$$

where K_s is the span stiffness in MPa/m of length mismatch, $v_{in,7} = v\,h_7/h_{in,7}$ is the strip speed into F7, with $h_{in,7}$ the F7 entry thickness of Equation (8.15), and $v_6 = v\,\rho_{mf}(1 + r_6)$ is the F6 exit speed. Between its 3° and 45° stops the arm stores 115 mm of strip, of which only 41 mm lies below the 25° design angle. At 10.00 m/s that is one hundredth of a second of production, which is why the looper angle is never left to an operator and why the F6 drive is slaved to the thickness ratio through the mass-flow lag ρ_{mf} of Equation (8.12).

$$\tau_{mf}\,\frac{\mathrm{d}\rho_{mf}}{\mathrm{d}t} = \mathrm{clamp}\left(\frac{h_7}{h_{in,7}},\,0.4,\,1.0\right) - \rho_{mf} \tag{8.12}$$

8.3.5 What Arrives, and What Drifts

Two disturbances arrive with the metal and both are written as functions of delivered strip length rather than of time, because that is the coordinate the X-ray gauge, the coiler and the customer's complaint all live in, as shown in Equation (8.13).

$$\begin{aligned} h_5(x) &= h_5^{0}\left[1 + 0.01\,A_e\sin\left(\frac{2\pi x}{\lambda_e}\right)\right] \\ T_7(x) &= T_0 + T_b - \Delta T_s\,\mathrm{e}^{\,\kappa[\cos(2\pi x/\lambda_s)-1]} \end{aligned} \tag{8.13}$$

where A_e is the transfer bar thickness variation in percent of 4.600 mm, λ_e and λ_s are wavelengths in delivered meters, ΔT_s is the skid mark depth and $\kappa = 8$ makes the wrapped Gaussian about three times narrower than its own pitch while staying smooth enough for a fixed-step integrator. F6 rolls material that will be delivered L_{st} meters later, so its entry thickness is evaluated at $x + L_{st}$ and its temperature is the F7 profile at the same place plus ΔT_{67}. A skid pitch of 1.5 m on the slab becomes 125 m on 3 mm strip, which arrives every 12.5 s at ten meters a second and is therefore far slower than roll runout and separable from it.

The two slow drifts are first-order approaches to their targets, as shown in Equation (8.14).

$$\tau_c\,\frac{\mathrm{d}c_{th}}{\mathrm{d}t} = c_{tgt} - c_{th}, \qquad \tau_b\,\frac{\mathrm{d}T_b}{\mathrm{d}t} = T_{bias} - T_b \tag{8.14}$$

The work roll thermal time constant τ_c of 40 s is what makes the drift slow enough for a slow loop to catch, which is the whole design of monitor gauge control.

8.3.6 Instruments and the Length-Indexed Delay

Four measurements feed the controllers. Each is a first-order lag stepped with the exact discrete update $z_k = \alpha z_{k-1} + (1-\alpha)u_k$, where $\alpha = \mathrm{e}^{-\Delta t_c/\tau}$, so no scan period can destabilize it. The load cell carries $\tau_{load} = 0.004\,\mathrm{s}$, the position transducer $\tau_{pos} = 0.002\,\mathrm{s}$, the X-ray gauge $\tau_{xray} = 0.030\,\mathrm{s}$ and the tensiometer $\tau_{sig} = 0.020\,\mathrm{s}$. The position transducer reads the capsule, which is above the whole roll stack, and that single fact is why roll runout is invisible to it. Gauge and load cell noise enter as first-order filtered Gaussian sequences, scaled to put the stationary standard deviation on the slider value. Their generator is seeded at 20260914, which makes a noisy run repeat bit for bit.

Two transport delays exist in the model, the X-ray gauge's and the span between the stands, both keyed to delivered strip length, as shown in (8.15).

$$h_x = h_7(x - L_g)\,, \quad \theta_d = \frac{L_g}{v}, \quad h_{in,7} = h_6(x - L_{st}) \tag{8.15}$$

Each delay line is a ring buffer keyed on x with linear interpolation between samples, and the span line starts full of strip rolled at the starting conditions. Indexed on time it would be right at 10.00 m/s and quietly wrong at every other speed the mill runs, whereas indexed on length the dead time follows every acceleration for free. Chapter 10 measures the same distinction for a predictor on a coating line. On the gap step of scenario 1 the application measures 0.525 s of lag at 10 m/s and 0.860 s at 6 m/s between the true and reported gauges crossing half of the settled step. That is L_g/v plus about 25 ms, which the 0.030 s averaging time of the gauge and the 0.004 s scan at which the delay line is written and read add to the transport delay.

8.3.7 Parameters, Initialization and Numerics

Table 8.2 holds every constant. The model dialog shows and adjusts all but seven of them, and those seven carry a hidden mark, meaning their values are never displayed and have to be worked out from behavior instead.

The mill modulus deserves a sentence of its own, because 6.00 MN/mm looks small beside the figures a rolling engineer carries in the hundreds. Those figures are the same number in different units, since 6.00 MN/mm is 612 tf/mm of force per millimeter of stretch, and the check that matters is the stretch the value produces. At 12.000 MN the stand stands 2.000 mm taller than when it is empty, which is two thirds of a 3.000 mm product and is why the commanded gap is 1.000 mm. A stiffer stand would pass a smaller fraction of the entry variation and would need a smaller deadbeat gain [14]. F6 carries 6.2261 MN/mm and stretches 2.200 mm under 13.697 MN.

No scenario starts away from equilibrium. Bisection on Equation (8.1) at the scheduled exit thicknesses gives the two roll gaps, returning 1.4000 mm at F6 and 1.0000 mm at F7, and the looper torque is initialized to the 13.535 kN m that balances the arm weight and the strip pull at 25°, of which 5.891 kN m lifts the arm itself. The application's verification suite re-solves those three constants and fails if they have drifted. States advance on a 0.002 s fixed step under classical Runge-Kutta of fourth order, controllers execute on a 0.004 s period, and the historian keeps 50 samples for each simulated second. The engine, its control period and its speed setting are those of Section 1.6.1, and a different speed rescales simulated time against the wall clock and nothing else.

Four conditions raise an annunciation without stopping the run, and each clears only after the condition has receded by a margin. Interstand tension below 1.5 MPa is a loop-out, above 28 MPa is necking, roll force above 20.0 MN is past the stand rating, and a roll gap driven to zero leaves the capsule with nothing to give. Tension is clamped to 0 and 60 MPa and the looper angle to its two stops, with the angular rate zeroed at a stop so the integrator cannot pass through it.

8.3.8 Assumptions

The model is detailed along the strip and absent across it, and each simplification fails somewhere.

One dimension. There is no strip width profile, no roll crown across the barrel, no flatness and no work roll bending, shifting or pair-crossing. On a real mill the shape across the strip takes at least as

Table 8.2 Model Parameters, in Two Panels. Hidden Constants Are Never Displayed and Have to Be Inferred From the Plant's Behavior.

Symbol	Meaning	Value	Unit	Shown
w	strip width	1200	mm	adjustable
R	work roll radius	340	mm	adjustable
f_H	roll flattening factor	1.37	—	adjustable
μ	roll-strip friction	0.30	—	adjustable
R_b	back-up roll radius	0.70	m	adjustable
F_{max}	stand force limit	20.0	MN	adjustable
L_g	F7 to the X-ray gauge	5.0	m	adjustable
L_{sp}	F6 to F7 stand centers	5.5	m	adjustable
h_5^0	F6 entry thickness	4.600	mm	adjustable
h_6^0	F6 exit thickness	3.600	mm	adjustable
h_7^0	F7 exit thickness	3.000	mm	adjustable
T_0	strip temperature at F7	880	°C	adjustable
ΔT_{67}	F6 hotter than F7	25	°C	adjustable
k_0	flow stress, T_0 and $\dot{\varepsilon}_0$	224.79	MPa	adjustable
β	flow stress softening	0.0045	1/°C	adjustable
m	strain rate exponent	0.13	—	adjustable
K_s	span tension stiffness	300	MPa/m	adjustable
M_7	F7 mill modulus	6.00	MN/mm	adjustable
M_6	F6 mill modulus	6.2261	MN/mm	adjustable
τ_v	servo valve response	0.006	s	adjustable
K_q	capsule slew, full valve	9.0	mm/s	adjustable
τ_{g6}	F6 capsule response	0.030	s	adjustable
τ_d	main drive response	0.50	s	adjustable

Symbol	Meaning	Value	Unit	Shown
τ_{s6}	F6 speed trim response	0.15	s	adjustable
τ_{mf}	mass-flow lag	0.25	s	adjustable
τ_T	looper torque response	0.08	s	adjustable
τ_c	roll thermal constant	40	s	adjustable
τ_b	bulk temperature lag	6.0	s	adjustable
J	looper inertia	600	kg m^2	adjustable
a	looper arm length	0.80	m	adjustable
c_L	looper damping	5000	N m s	adjustable
W	looper gravity torque	6.50	kN m	adjustable
τ_{load}	load cell response	0.004	s	adjustable
τ_{pos}	transducer response	0.002	s	adjustable
τ_{xray}	X-ray averaging time	0.030	s	adjustable
b	gauge tolerance	0.025	mm	adjustable
b_σ	tension tolerance	2.0	MPa	adjustable
σ_{lo}	strip loops out below	1.5	MPa	adjustable
σ_{hi}	strip necks above	28	MPa	adjustable
f_s	forward slip at the exit	0.03	—	hidden
L_{st}	F6 to F7, delivered	6.6	m	hidden
$\dot{\varepsilon}_0$	reference strain rate	100	1/s	hidden
σ_{56}	F5 to F6 back tension	10.0	MPa	hidden
θ_{min}	looper bottom stop	3	°	hidden
θ_{max}	looper top stop	45	°	hidden
τ_{sig}	tensiometer response	0.020	s	hidden

much control effort as the thickness along it, and the two interact through the roll stack, because a gauge correction changes the load, which changes the roll deflection, which changes the crown [13].

The mill spring is one linear modulus per stand. A real force against stretch curve is nonlinear at low load where the chocks are still bedding in, which is why mill modulus is quoted as a chord over the working range [14]. The modulus error disturbance makes the point that the modulus is a calibration, but the linearity is an idealization.

Hitchcock flattening is frozen at 1.37. Real flattening grows with load, so the model understates the plastic modulus far from the design point, although within the working range the two gains of (8.6) are accurate.

The strip is one temperature. Through the thickness and across the width the strip is isothermal, where in reality the surface is far below the core after descaling and the deformation reheats it. The profile is written in delivered meters rather than emerging from a heat balance down the mill.

No mill resonance. Stands, drives and the roll stack are first-order lags, so the third-octave chatter mode near 100 Hz to 200 Hz that a real gauge controller must stay clear of does not exist, and neither does torsional coupling between the two stands through the strip.

The mass-flow link is idealized. F6's speed follows the true thickness ratio through a 0.25 s lag, where a real scheme forms that ratio from the X-ray gauge and the upstream gaugemeters. Feeding the compensation a biased gaugemeter turns a gauge bias into a permanent speed error and parks the looper on a stop, which is not what the application is about.

The eccentricity estimator is handed a perfect roll angle. It assumes one sinusoid on one roll, where a real stand has four rolls at three frequencies plus harmonics and the estimator has to acquire the phase rather than being given it [11].

Two stands, not seven. There is no threading, no head end, no tail out and no coiler, so the transient a real gauge control system spends most of its effort on, the first meters of a coil, is absent.

Three disturbances are stylized. A thermal crown arriving on a clean first-order approach to exactly 200 µm is not what a real roll set does, but it isolates the gap-side disturbance and makes the steady state of Equation (8.17) readable off a trend. A modulus error applied as a step is equally artificial, because a calibration goes stale between roll changes rather than during a coil, although it separates a pure measurement bias from every dynamic effect in the plant. The skid marks are one wrapped Gaussian of a single pitch and depth, and that regularity is what makes their 12.5 s arrival interval measurable.

Repeatability itself is unrealistic. Freezing the step and the scan period yields an identical trajectory on any machine, whereas a distributed control system in a mill carries timing jitter and an execution order that shifts.

8.4 The Two Measurements

8.4.1 An Exact Error Identity

The gaugemeter rearranges Equation (8.1) into a thickness, using the position transducer, the conditioned load cell signal and a modulus that is a calibration rather than a constant of nature, as shown in Equation (8.16).

$$\hat{h} = S_{meas} + \frac{F_{gm}}{M_{mod}} + \delta, \qquad M_{mod} = M_7 \left(1 + \frac{\varepsilon_M}{100}\right) \tag{8.16}$$

where ε_M is the modulus calibration error in percent and δ is the monitor trim of Section 8.5. Subtracting the truth of (8.1) from (8.16) leaves the whole argument of the chapter in one line, as shown in Equation (8.17).

$$\hat{h} - h_7 = -c_7 + F_7 \left(\frac{1}{M_{mod}} - \frac{1}{M_7}\right) + \delta \tag{8.17}$$

Three properties of (8.17) carry the application. Nothing on the right depends on the strip, so every disturbance that arrives with the metal, a thick patch on the transfer bar, a skid mark, a cold slab or a tension change, leaves the error unmoved and the inferred gauge is exact. Measured open loop

Table 8.3 Open-Loop Response to Six Disturbances, Roll Gap Held. The Rule That Separates the Two Blocks Is Whether the Disturbance Reaches the Strip or the Roll Gap. The Skid and Cold-Strip Errors Are Read at the Band Bottom and in Steady State, and the Entry-Wave Error Is the Lag of the Force Filters.

Disturbance	Δh_7 (μm)	$\hat{h} - h_7$ (μm)	Sign
Arrives with the metal			
Skid mark, 35 °C deep	+134.0	−0.04	—
Entry wave, 3 % of 4.6 mm	±51.2	±1.6	—
Cold strip, −30 °C	+115.5	0.00	—
Moves the roll gap			
Thermal crown, −200 μm	−41.9	+198.6	inverted
Runout, 40 μm amplitude	±8.6	∓39.6	inverted
Modulus error, +5 %	0.0	−95.2	bias

at the bottom of a 134.0 μm skid mark the gaugemeter error is 0.04 μm, and on a whole strip 30 °C cold, which moves the product 115.5 μm, it is zero to the digits shown once the strip has settled. Only while the force is moving does the 14 ms of load cell and gaugemeter filtering open a gap, at most 1.6 μm on any strip-side row of Table 8.3. Second, a gap offset enters at full size with a minus sign while the true gauge moves by only $M/(M+Q)$ of it, so the ratio of error to real change is $-(M+Q)/M = -4.75$. Third, a modulus error is a pure bias proportional to the force, invisible at the stand and constant for the coil, which is what the monitor trim δ exists to cancel and to display.

Table 8.3 separates the two halves on the running model. Every row is open loop with the roll gap held, so the entries are properties of the plant and the instruments rather than of any controller. Each row is the scenario that carries the disturbance with Gauge control switched to manual before the run, and a steady disturbance is read once it has settled.

The crown row is the one worth checking against the algebra. Equation (8.17) says the error should be exactly the negative of the gap offset, and at 200 s after the step the crown state has reached −198.65 μm and the measured error is 198.65 μm. The true gauge has moved −41.91 μm, so the measured ratio is -4.741 against the -4.752 that (8.6) predicts, the small difference being the rise in Q as the pass gets heavier. The modulus row is (8.17) evaluated directly, since $12\,\text{MN}\,(1/6.30 - 1/6.00) = -95.2\,\mu\text{m}$.

🕮 **Pause.** Try Exercise 8.2 before reading on.

8.4.2 Where the Two Curves Cross

Figure 8.2 draws Equation (8.1) the way a mill engineer reads it, as the intersection of the elastic line of the stand with the plastic curve of the pass, and puts the crown disturbance on the same axes.

Reading the figure explains the mechanism in one sentence. Crown slides the mill line left, the intersection moves down the plastic curve to a thinner strip and a higher force, and the gaugemeter converts that higher force back into a thickness along the line the mill no longer sits on. A controller that believes the reported 3.1586 mm shuts the gap until the reported value is 3.000 mm again, which by Equation (8.17) is precisely the condition $h_7 = h_{SP} + c_7$. The steady state of a gaugemeter loop under a gap offset is therefore a delivered gauge low by the entire offset, exactly, and the measured result at 210 s is −198.6 μm against a crown of −198.65 μm.

🕮 **Pause.** Try Exercise 8.1 before reading on.

8.4.3 What the Five Meters Cost the Other Loop

The X-ray gauge has no such defect. It measures the product, it is absolute and traceable, and its error under every disturbance in Table 8.3 is zero once the strip reaches it. What it has instead

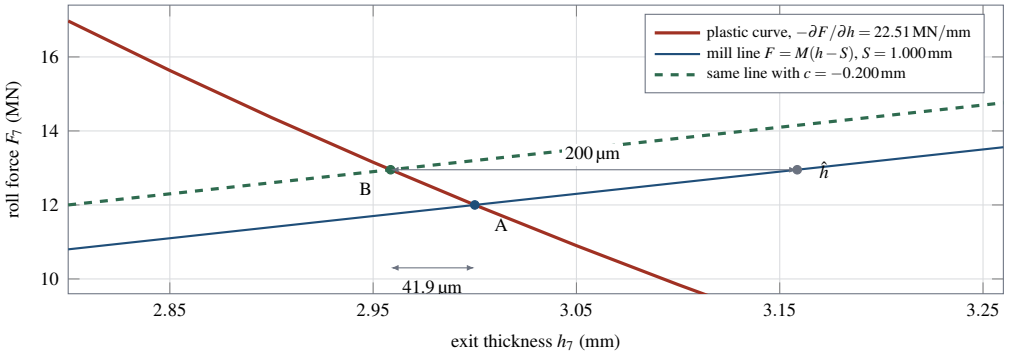


Figure 8.2 Mill line and plastic curve at the F7 design pass, and what 200 µm of thermal crown does to both. The plastic curve is the thick solid line, the mill line the thin solid line and the mill line under crown the dashed line. Crown displaces the mill line bodily to the left without moving the position transducer, so the intersection moves from A to B, 41.9 µm thinner at 12.951 MN, while the gaugemeter reads that force back along the undisplaced line and reports 3.1586 mm. The gap between B and $\hat{h}$ is the whole of the crown, the wrong way up.

is 0.500 s of transport delay at line speed, and dead time is the one plant property no tuning rule removes [15, 16]. A pure delay θ_d in series with the static gap gain of (8.6) reaches −180° of phase at a period of $2\theta_d = 1.0$ s where the loop gain is the reciprocal of the gap gain, so the ultimate gain is $1/0.2104 = 4.75$. Adding the capsule, the X-ray averaging filter and reset action moves both numbers, and at $K_c = 4.00$ and $\tau_I = 0.25$ s the delayed loop cycles at 1.374 s rather than 1.0 s, or at 2.09 s when 6 m/s makes θ_d 0.833 s. Neither is a property of the controller. Both are five meters of runout table divided by a speed.

That is the whole of the trade. One measurement is exact for the strip, blind to the mill and available in four milliseconds, and the other is true about everything and half a second late. Monitor gauge control runs both and splits the job by frequency, with the gaugemeter above about one radian a second where it is the only measurement that exists and the X-ray below it where it is the only measurement that is right.

8.5 Controllers as Implemented

8.5.1 The Shared Algorithm

All five loops run the velocity-form PID of Section 1.6.2, Equations (1.13) to (1.16), with back-calculation on the integral alone and the tracking constant left at its default of τ_I. The action a takes the value +1 on a reverse-acting loop and −1 on a direct-acting one, in the convention of Section 1.3. Here the scan is $\Delta t_c = 0.004$ s and the output bias is zero throughout. Every loop sets $\tau_D = 0$, which idles the derivative path and its measurement filter. Nor does any loop impose a rate limit, since the capsule slew K_q, the valve lag τ_v and the torque lag τ_T already supply the real ones from the machinery. Selecting manual solves for the integral against the operator's actuator value, which makes both transfers bumpless.

8.5.2 The Five Loops and Their Directions

Table 8.4 lists what each loop reads, what it writes and which way it acts. Two of the directions are worth stating in words, because they are the two that get reversed in practice.

Opening the gap thickens the strip, so the gauge loop is reverse acting. Its gain of 4.00 is set just below the deadbeat value of 4.752 from (8.6), with the remainder held back for the phase the capsule and the load cell filter cost. The monitor loop is direct acting, and this is the sign that gets flipped. Its output is a bias added to the gaugemeter's reading, and raising a measurement drives the loop

Table 8.4 The Five Implemented Loops, With the Tuning Every Scenario Except 4, 14 and 15 Opens On.

CV	MV	K_c	τ_I (s)	Limits	a
$\hat{h}$ or h_x	S_7^{cmd}	4.00	0.25	0.05–2.60	+1
h_x	δ	0.60	2.0	±0.25	−1
S_7	x_v	500	1.00	±100	+1
σ	T_L	0.300	0.80	0–32	+1
θ	r_6	0.120	4.0	±5	+1

underneath it to shut the gap and roll thinner, so the manipulated variable moves the process variable down. Getting that one backwards does not oscillate. It walks the gauge quietly off specification over half a coil, which is worse, and the application's verification suite tests the direction of all five loops by perturbing each and checking the sign of the response.

🕮 **Pause.** Try Exercise 8.3 before reading on.

The capsule loop closes around the integrator of (8.8), so it is nearly a pure proportional loop, and 500 %/mm against $K_q = 9.0\,\text{mm/s}$ at full valve places the closed capsule at 45/s. It is in service in manual mode as well, because it is the servo that makes the capsule a position device at all and what manual changes is only who writes its setpoint. The looper angle loop is in service in both modes for the same reason, since the angle is a pure integral of the mass-flow mismatch and one percent of F6 speed swings the arm 27 °/s. Manual mode releases only the tension loop and the gauge loop.

8.5.3 The Architecture Switch

The gauge loop reads one of three numbers, selected by a four-position control that is the whole argument of the application. In *manual* the gap is whatever the operator typed, in *X-ray only* the loop reads h_x, in *gaugemeter* it reads (8.16) without the trim, and in *gaugemeter with monitor* it reads (8.16) in full with the monitor loop in service. The monitor setpoint is h_{SP} delayed as in (8.15) and averaged as the X-ray gauge is. When the monitor loop is out of service its integral tracks a zero output, so switching it in is bumpless.

8.5.4 The Synchronous Eccentricity Estimator

Neither gauge loop can reach 2.207 Hz, which is fourteen radians a second against a delayed loop that stops at one and a fast loop that would have to chase a signal of the wrong sign. What the mill has instead is independent knowledge of where the back-up roll is, from a shaft encoder, and two integrators in the rotating frame turn that into an estimate of the once-per-revolution force component, as shown in Equation (8.18).

$$
\begin{aligned}
F_{ecc} &= e_a^c \cos\phi + e_b^s \sin\phi \\
\frac{\mathrm{d}e_a^c}{\mathrm{d}t} &= 2K_{ecc}\,(F_{sig} - \bar{F} - F_{ecc})\cos\phi \\
\frac{\mathrm{d}e_b^s}{\mathrm{d}t} &= 2K_{ecc}\,(F_{sig} - \bar{F} - F_{ecc})\sin\phi \\
F_{gm} &= F_{sig} - F_{ecc}
\end{aligned}
\tag{8.18}
$$

where F_{sig} is the load cell signal after the adjustable force filter τ_F, $\bar{F}$ is a first-order average with a 1.0 s time constant, removed first so the estimator locks onto the ripple rather than spending its first second chasing twelve meganewtons of steady load, and K_{ecc} is a convergence rate in reciprocal seconds. The factor of two undoes the mean of $\cos^2$, which is what makes K_{ecc} a rate rather than an arbitrary number. Setting K_{ecc} to zero switches the estimator off and zeroes both states. At $K_{ecc} = 3/\text{s}$ in scenario 12 the estimate reaches 90 % of its final amplitude 0.19 s after it is switched in and settles at 0.177 MN, against the 0.189 MN the runout is actually worth in the open-loop force.

The alternative offered alongside it is the force filter τ_F, and the comparison fails instructively, because a low-pass slow enough to remove 2.207 Hz also removes the entry thickness wave at 0.333 Hz that the whole scheme exists to catch. With $\tau_F = 0.20$ s in place of the estimator, a 3 % entry wave fired on top of the runout of scenario 12 reaches 63.6 µm of gauge error, against 43.9 µm with the estimator in service. The estimator removes energy at one frequency the mill has independent knowledge of and leaves the rest of the spectrum alone. The phasor canceller of Chapter 6 removes strand bulging in the same way, at a frequency computed from the casting speed and the roll pitch.

8.5.5 Three Models Kept Apart

Equations (8.1) to (8.15) are the simulated plant, thirteen integrated states carrying two implicit algebraic solves. Whatever model the controller holds is far smaller, because the gauge loop carries nothing beyond a gain and a reset time and its tuning encodes one plant number, the gap gain $M/(M+Q) = 0.2104$, whose reciprocal is the deadbeat gain. The gaugemeter is a third thing again, a measurement model consisting of one constant M_{mod} whose disagreement with the plant's M_7 is the modulus term in (8.17). Confusing the three is easy here, because the symbol M appears in the plant, in the measurement and in the tuning rule and means something different each time. The mill modulus M_7 is a property of the machine, M_{mod} is a number in a calculation, and $K_c = 4.00$ is a controller gain in millimeters of gap per millimeter of gauge error.

8.6 Scenarios

Table 8.5 lists fifteen scenarios ordered the way a laboratory session would run them. A scenario pins the opening mode, the architecture, the five tunings, every disturbance setting and any timed event, then puts a question before the run and a measurement after it. Three tuning sets recur. The shipped set is $K_c = 4.00$ and $\tau_I = 0.25$ s on the gauge loop, the detuned set that a delayed measurement forces is $K_c = 1.2$ and $\tau_I = 1.5$ s, and the deliberately poor set that opens the challenge is $K_c = 0.8$, $\tau_I = 2.5$ s, a monitor trim too slow to converge inside a coil and a force filter of 0.20 s.

Each scenario also carries a completion rule, condensed in the last column of Table 8.5, and credit, Reset and the challenge verdict work as Section 1.6.4 describes. A run that needs a different architecture or tuning is repeated from Reset. The challenge is judged on the coil as given, so changing a model parameter or a disturbance voids the attempt. The scenario names in the table are the ones the application shows, and all but scenarios 5 and 15 are experiments.

8.7 Results

Each figure below comes off the running application. The exported record sets the true gauge beside both measurements, the capsule position beside its command, the live plastic modulus and gap gain, and the running yield and deviation metrics, which is enough to redo any measurement here. Root-mean-square gauge deviation and meters off gauge accumulate with the length rolled, so both are reported together with the interval over which they were taken. Peak deviations, capsule positions and limit cycle periods carry no such dependence. Each run loads its scenario, makes any listed change before Run and runs from $t = 0$, and peaks, crossings and periods are read at the 2 ms physics step.

8.7.1 Two Thirds of the Product Is Mill Spring

Scenario 1 settles the question with no controller in the picture. Opening the F7 gap from 1.000 to 1.250 mm moves the delivered gauge from 3.0000 to 3.0546 mm, which is 54.6 µm of product for 250 µm of gap. Closing it to 0.750 mm gives 2.9474 mm, another 107.2 µm for 500 µm. The ratio is 0.218 on the opening step and 0.214 on the closing one, both a little above the linearized 0.2104, which describes only small moves, and the arithmetic is on the force trace, because opening the gap 250 µm lightens the pass, the force falls from 12.000 to 10.828 MN, the stretch falls from 2.000 to

Table 8.5 Scenario Exercises and Completion Requirements. Each Scenario Is Named as the Application Shows It, With the Mechanism It Exercises in *Italics*.

#	Scenario	Required to complete
1	The gap you set is not the gauge you get. *the mill spring, a gap gain of 0.21*	Run in Manual through the gap moves at 6, 16 and 26 s, then move the gap at least 5 µm off 1.000 mm and hold it for 3 s.
2	Measuring the mill modulus. *force against stretch, not gauge against gap*	Run in Manual through the three gap moves and hold the gap at 1.000 mm for 5 s after the strip arrives 30 °C colder at 26 s.
3	What the roughing mill sends you. *entry walk-through at 0.65*	Run 9 s in Manual on the 3 % wave, three of its periods, then 9 s more with the wave cut to 1 %.
4	Feedback from five meters downstream. *what half a second of dead time permits*	Run 30 s in Auto on the X-ray gauge with $K_c \leq 1.3$ and the 3 % wave.
5	Pushing the delayed loop. *the limit cycle at $2\theta_d$, a looper on its stops* (demonstration)	On the X-ray gauge at $K_c \geq 3.9$ let the gauge error pass ±100 µm, lower K_c to 3.0 or less, and hold the error within ±40 µm for 9 s.
6	The gauge you cannot measure, inferred. *an inferred measurement with no delay*	With gaugemeter and monitor, hold the true gauge within 5 µm of 2.950 mm for 2 s after the step at 8 s, and of 3.000 mm for 2 s after the step back at 20 s.
7	Skid marks. *the disturbance the gaugemeter is exact for*	Run 25 s of skid marks on each of gaugemeter with monitor at $K_c \geq 3.9$, no gauge control, and the X-ray gauge at $K_c = 1.2$ and $\tau_I = 1.5$ s.
8	The rolls grow into the strip. *the sign inversion, driven by the loop*	Keep the gaugemeter loop without trim in Auto for 110 s after the crown starts at 10 s.
9	The monitor trim. *splitting one CV by frequency*	Run 100 s past the crown onset with the trim at 0.60 and 2.0 s, then 60 s each with the trim gain at 0 and at 1.5 with $\tau_I = 0.3$ s.
10	A calibration that is a season old. *a bias no instrument at the stand can find*	Run the gaugemeter loop with the +5 % error to at least 18 s, switch the monitor in until the true gauge holds within 5 µm for 3 s, then set the error to −5 % and see the gauge error pass −20 µm.
11	Once per revolution. *amplification at 2.207 Hz*	Run 5 s each with gaugemeter and monitor (estimator off), no gauge control, the gaugemeter alone at $K_c \geq 3.9$, and the X-ray gauge at $K_c = 1.2$ and $\tau_I = 1.5$ s.
12	Telling the roll from the strip. *synchronous demodulation against a shaft angle*	Run to 20 s with the estimator in at $K_{ecc} \geq 2.9/\text{s}$ from 12 s, then 3 s with $K_{ecc} = 0$ and $\tau_F \geq 0.15$ s, then fire the 3 % entry wave and run 6 s more.
13	The looper, the tension and why the gauge cares. *the looper, mass flow and why the gauge cares*	Run 10 s past the F6 gap opening at 42 s, then Reset, open the gauge loop before 8 s and run 5 s past the 16 MPa tension step.
14	Slow down, and watch the dead time grow. *a plant that changes under a fixed tuning*	Run 15 s at 6 m/s on the X-ray gauge with $K_c \leq 1.3$, the same on gaugemeter and monitor, and, from the reloaded case, on the X-ray gauge at $K_c = 4.0$ and $\tau_I = 0.25$ s.
15	Challenge: One coil, everything at once. *five disturbances and a speed change at once* (challenge)	Hold the true delivered gauge within ±25 µm of 3.000 mm over the last third of the 90 s coil, with the setpoint, model parameters and disturbances as given.[a]

[a] With the shipped configuration (X-ray gauge, $K_c = 0.8$, $\tau_I = 2.5$ s) the worst error over the last third is 61.2 µm, 36.2 µm outside the band, at 65.4 % prime yield. With the suggested configuration (gaugemeter and monitor, $K_c = 7.0$ and $\tau_I = 0.15$ s, trim 1.2 and 1.2 s, $K_{ecc} = 2/\text{s}$) it is 20.3 µm, inside with 4.7 µm to spare, at 98.8 % prime yield.

1.805 mm, and 195 µm of the 250 µm given away comes straight back as a housing no longer standing as tall.

> **TRY IT IN THE APP 1 · The Gap You Set Is Not the Gauge You Get**
>
> Load scenario 1 and press Run in Manual. At 6 s the F7 roll gap opens by 250 µm. Watch the delivered gauge: it moves from 3.0000 to 3.0546 mm, only 54.6 µm, because the roll force falls from 12.000 to 10.828 MN and the stand stretches less.

Scenario 2 turns the same experiment into the calibration a mill crew takes on the head end of a coil. The modulus is not the slope of gauge against gap. It is the slope of force against stretch, and stretch is $h - S$, which is readable from the trends. At $+100\,\mu$m the force is 11.530 MN and the stretch 1.9216 mm, so $\Delta F/\Delta\delta = 0.470/0.0784 = 6.00\,\text{MN/mm}$, and at $-200\,\mu$m the same ratio is $0.472/0.0787$. The cold strip that follows moves the gauge 115.5 µm with the gap untouched and the modulus unchanged, because what changed was the load.

Scenario 3 runs a 3 % transfer bar variation, which is 138 µm of amplitude on the 4.600 mm entering F6, against no gauge control at all. F6 passes 57 % of what it is given and F7 passes 65 %, so the amplitude is 78 µm at the F6 exit and 51 µm at the product against a 25 µm tolerance. Read at 60 s the root-mean-square deviation is 36.2 µm and the prime yield is 32.5 %. A hot mill without gauge control does not make slightly worse strip. It makes mostly seconds.

8.7.2 What Five Meters Cost

Scenario 4 closes the loop on the X-ray gauge with the loop detuned to $K_c = 1.2$ and $\tau_I = 1.5$ s to suit the delay. Read at 60 s on the same disturbance it buys almost nothing, with 34.2 µm of root-mean-square deviation against 36.2 µm open loop and the prime yield moving from 32.5 to 35.1 %. The loop is working, and it is working a wavelength late, and a correction applied late to a periodic disturbance is as likely to add to it as to subtract from it.

Scenario 5 answers the obvious objection by restoring $K_c = 4.00$ and $\tau_I = 0.25$ s, the tuning the gaugemeter uses successfully in scenario 6, on a disturbance cut to a third. The loop limit-cycles at 1.374 s with a peak deviation of 328 µm and 164 µm of root-mean-square deviation read at 60 s, against 12.0 µm for doing nothing at all. The part that matters on a real mill is below the gauge trace, because the gap swings 2.31 mm, the looper cannot store the mass-flow swing that follows, the arm reaches both its 3° and 45° stops and the tension swings between 0.0 and 34.8 MPa. Loose strip between two stands at ten meters a second is a cobble and a wrecked roll set. Neither 4.75 nor 1.374 s is a tuning constant.

8.7.3 The Same Gain on a Measurement That Is Not Late

Scenario 6 closes the loop on $\hat{h}$ at the identical 4.00 and 0.25 s, with the monitor trim in service. A 50 µm setpoint step is inside 5 µm of the new setpoint 0.54 s after it takes effect, crosses it at 0.65 s and overshoots by 4.5 µm at 1.03 s. With the gaugemeter alone the same step settles inside 5 µm in 0.99 s with no overshoot, and the difference belongs to the trim. The step is no error to its delayed setpoint, but when the thinner strip reaches the X-ray gauge 0.51 s after the step the true gauge is still 11 µm short of the new setpoint, so the trim rises to 17.5 µm at 0.57 s and the gauge loop rolls that much thinner. The bare $S + F/M$ parts from the truth by at most 25.3 µm, 18 ms into the capsule's move, and by more than 5 µm for 60 ms, which is the load cell and its filter fourteen milliseconds behind a moving capsule. Twelve seconds after the step the inferred and true thicknesses agree to 0.20 µm, and that residue is the trim δ, because the uncorrected $S + F/M$ agrees with the truth to 0.002 µm, as Equation (8.17) requires with the mill unchanged.

Table 8.6 puts all four architectures on one disturbance at identical gains, which is the comparison the application exists to make.

Reading Table 8.6 across gives the result the chapter is built on. The delayed loop at the gain the fast loop carries comfortably is seven times worse than doing nothing, at 173.8 µm of root-mean-

Table 8.6 Four Architectures on the Two Strip-Side Disturbances, at Identical Gauge Loop Tuning Except Where Noted. The Entry Wave Block Is Scenario 3 With the Wave Cut to 2 %, Read From the Record Over $15 \le t \le 45$ s, Which Is 300 m of Strip, and the Skid Block Is Scenario 7 Read From the Scorecard at $t = 80$ s, Over 800 m.

Architecture	RMS (µm)	Peak (µm)	Prime (%)	Off gauge (m)
Entry wave 2 %, $K_c = 4.00$ and $\tau_I = 0.25$ s				
Manual, gap held	24.1	34.4	52.3	143
X-ray only	173.8	340.3	6.1	282
Gaugemeter	11.6	16.4	100.0	0
Gaugemeter + monitor	11.5	16.4	100.0	0
Skid marks 35 °C, X-ray detuned to 1.2 and 1.5 s				
Manual, gap held	41.3	134.0	80	159
X-ray only, detuned	32.2	107.8	81	151
Gaugemeter	11.6	32.7	90	82
Gaugemeter + monitor	10.2	32.7	95	42

square deviation against 24.1 µm, while the inferred loop holds 11.6 µm and puts no strip outside the band. Nothing about the two controllers differs, because the plant, the disturbance, the gain and the reset time are identical and the whole difference is the 0.500 s between the roll bite and the gauge. The monitor trim changes the result by less than 0.1 µm here, which is correct, because an entry wave leaves the gaugemeter no bias to remove.

> **✦ DID YOU KNOW? Same Gain, Opposite Outcomes**
>
> From the simulator: on the 2 % entry wave at identical tuning, the loop on the X-ray gauge puts 282 m of 300 m off gauge and the loop on the gaugemeter puts none, while a held gap leaves 143 m off gauge (Table 8.6). Half a second of delay turns the best controller into the worst.

8.7.4 The Disturbance the Gaugemeter Was Invented For

Scenario 7 puts cool bands on the strip. The slab rested on water-cooled skid rails in the reheat furnace, so it carries bands 1.5 m apart that the mill has stretched more than eighty times, arriving 125 m apart and therefore every 12.5 s. Thirty-five degrees is 16 % more flow stress, F6 feels it first and its exit thickens by 99.3 µm, F7 meets that strip and the cold band together 0.66 s later, and 134.0 µm of extra gauge reaches the coiler with nothing controlling it. The lower block of Table 8.6 runs the four architectures on that one disturbance.

The gaugemeter is right for this disturbance and not merely fast. A skid mark is a change in the strip, every effect of a change in the strip appears in the roll force, and the roll force is one of the two numbers the gaugemeter is built from, so Equation (8.17) does not move and the inferred error at the bottom of a 35 °C band is 0.04 µm against a 134.0 µm disturbance. The measurement does not have to be told about temperature or about the furnace.

8.7.5 The Three It Gets Backwards

Scenarios 8 to 12 are the other half of Table 8.3, now with the loops closed.

Scenario 8 is the headline. The gaugemeter loop holds its own measurement dead on setpoint and leaves the delivered gauge 198.6 µm thin, against 41.9 µm with the loop open, so closing the loop made the disturbance 4.7 times worse. At 210 s the capsule is commanded at 0.197 mm instead of 1.000 mm with 16.82 MN on a stand scheduled for 12.000 MN, and the prime yield over the run is 7.4 %. The controller worked perfectly and destroyed the coil.

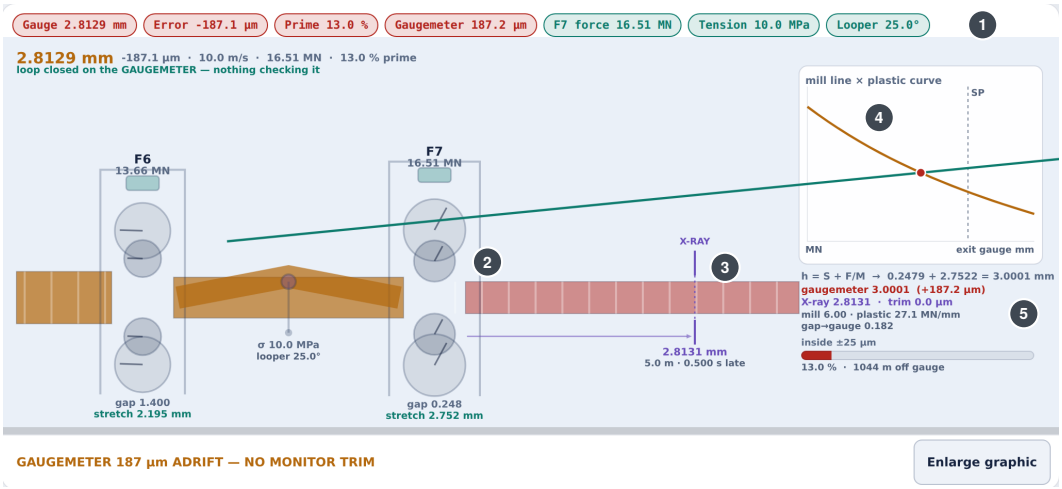


Figure 8.3 Scenario 8 at $t = 120$ s, with 200 µm of thermal crown 110 s into a 40 s time constant and the gauge loop closed on the gaugemeter alone, cropped to the plant view. The badge row reads the gaugemeter 187.2 µm adrift against a true error of −187.1 µm, positive by nearly the whole crown while the truth is negative. The capsule has been driven to 0.2479 mm against a scheduled 1.000 mm and the force to 16.51 MN against 12.000 MN. Callouts: (1) the badge row, (2) stand F7 with its force, gap and stretch, (3) the X-ray gauge and its late reading, (4) the mill line and plastic curve at the present pass, (5) the gaugemeter reading against the X-ray reading.

> **TRY IT IN THE APP** **8 · The Rolls Grow Into the Strip**
>
> Load scenario 8, leave Gauge control at Gaugemeter and run past two minutes. The work rolls start to grow at 10 s. The inferred gauge stays on 3.000 mm while the delivered gauge walks thin. By 210 s the capsule is commanded at 0.197 mm and the strip is 198.6 µm thin.

Scenario 9 adds one slow loop and nothing else. Comparing the X-ray gauge with the setpoint and adding the result as a bias to the gaugemeter's reading, at $K_c = 0.60$ and $\tau_I = 2.0$ s, roughly an eighth the speed of the loop it corrects, holds the peak gauge error to 12.5 µm and the steady error to 0.13 µm against 198.6 µm without it. The trim walks to −0.1985 mm, the capsule settles at 1.198 mm and the force returns to 12.00 MN, which is where the stand should have been all along. A faster trim, at 1.5 and 1.0 s, holds the peak to 3.1 µm, but at 1.5 and 0.3 s, or 3.0 and 1.0 s, the trim has become the delayed loop again and cycles at up to 201 µm.

Scenario 10 is the second job a monitor loop does. A new roll set went in and nobody retook the modulus, so the arithmetic divides by 6.30 MN/mm and the stand has 6.00 MN/mm. Equation (8.17) predicts −95.2 µm at a frozen force and the mill delivers +81.5 µm, because the force falls as the loop opens the gap. Either number is three times the tolerance band and constant for the coil, and no instrument at the stand can find it, because the position is right, the force is right and the arithmetic is right while the constant is wrong. Switching the monitor loop in at 30 s brings the error inside 5 µm 14.9 s later, inside 1 µm after 23.6 s, and parks the trim at +0.0952 mm, which is the calibration error displayed in millimeters. A trim that walks a long way is a maintenance request.

Scenarios 11 and 12 are the disturbance neither loop can reach. 40 µm of runout on a 0.70 m back-up roll at ten meters a second breathes the gap at 2.207 Hz while the position transducer, above the whole roll stack, never moves. The measured gauge ripple, half its peak-to-peak value, is 8.6 µm with no gauge control, 8.0 µm with the X-ray loop at its detuned 1.2 and 1.5 s and 21.0 µm with the gaugemeter, so the fast loop is two and a half times worse than doing nothing. Equation (8.6)

predicts the trends to within 0.5 µm, since $0.2104 \times 40\,\mu\text{m} = 8.4\,\mu\text{m}$ of true ripple runs against $-0.7896 \times 40\,\mu\text{m} = -31.6\,\mu\text{m}$ of reading, measured open loop as 8.6 µm and 31.1 µm. At the shipped 4.00 and 0.25 s the X-ray loop would only limit-cycle, as in scenario 5, which is why scenario 11 runs it detuned. The X-ray loop is untroubled only because 2.207 Hz is fourteen radians a second and the loop stops at one, so being slow is not always a defect. Switching the estimator of (8.18) in at $K_{ecc} = 3/\text{s}$ in scenario 12 takes the ripple from 20.3 µm, with gaugemeter and monitor, to 8.2 µm, just under the no-control value, while losing none of the gaugemeter's speed elsewhere.

8.7.6 The Looper, and a Plant That Changes

Scenario 13 shows what the machinery underneath is for. Tension subtracts from the pressure the rolls apply at 0.024 MN/MPa at frozen thickness, which is 0.83 µm of gauge per megapascal at F7 alone. With the gauge loop open the step from 10 to 16 MPa moves the delivered gauge −10.6 µm, or 1.77 µm per megapascal, because the same tension also pulls the F6 exit 8.6 µm thinner. Closed, the effect appears in the actuator instead of the product, because the gap moves 49 µm through the same step while the delivered gauge strays 1.3 µm and settles within 0.03 µm. The gaugemeter takes a tension change out without being told, because it is one more thing that arrives through the roll force. Opening the F6 gap by 80 µm thickens the F6 exit by only 22 µm and is also a mass-flow change that the looper absorbs and the F6 speed trim walks back, with the delivered gauge moving 9.0 µm at worst.

Scenario 14 changes the plant rather than the disturbance. The line drops from 10 to 6 m/s for the tail out, so the dead time grows from 0.500 to 0.833 s while the gauge has not moved. The slowdown is also a disturbance in its own right, because the strain rate and with it the flow stress fall with the speed, and with the gap held it leaves the strip 56.5 µm thin. The detuned loop removes that over about fifteen seconds with a worst excursion of 67.4 µm, and then survives the new dead time, with the worst excursion going from 34 to 39 µm and the root-mean-square deviation from 22.9 to 22.5 µm, read over the ten seconds before the change and over 25 to 60 s. What has gone is the room for integral action, because the integral gain a delayed loop can carry falls as the reciprocal of the delay, while the proportional limit set by the gap gain barely moves. The scenario 5 tuning, whose short integral time leans on that room, cycles in this case at 2.09 s with 381 µm peaks in the fifteen seconds after the change, against 1.374 s at full speed in scenario 5. What does not change is the disturbance wavelength, because the entry wave is 30 m long and arrives every five seconds instead of every three, slower in time and identical in strip. On the gaugemeter at its own 4.00 and 0.25 s the slowdown costs a 6.4 µm transient and no margin at all, because there is no length in the measurement to divide by anything. Left at the case's detuned 1.2 and 1.5 s, the gaugemeter removes the slowdown itself barely faster than the X-ray loop does, peaking at 50.2 µm.

8.7.7 The Graded Challenge

The challenge is ninety seconds of one coil carrying a 1.5 % transfer bar, 16 °C skid marks, 45 µm of runout, a modulus 4 % adrift, 150 µm of thermal crown arriving over the run and a speed change to 7 m/s, starting on the X-ray gauge with the poor tuning. Six disturbances, and each is defeated by a different line in the scheme. Table 8.7 measures what each line is worth. Figure 8.4 draws its five rows as a chart.

Table 8.7 summarizes the whole application. The two configurations that stay on the delayed measurement differ in every tuning, a gauge loop at 0.8 and 2.5 s with weak inner loops against 1.2 and 1.5 s with sound ones, and by almost nothing in result, at 61.2 against 60.9 µm of peak deviation, so tuning is not the variable. The configuration that runs the fast loop without a slow one behind it is the worst of the five at 65.9 µm and 29.9 % prime yield, because the crown and the stale modulus go straight into the product. Adding the monitor trim takes the peak to 28.6 µm and the prime yield to 89.5 %, and adding the estimator takes them to 20.3 µm and 98.8 % with 10 m off gauge in 767 m. Get the architecture right and almost any sensible tuning passes. Get it wrong and no tuning does.

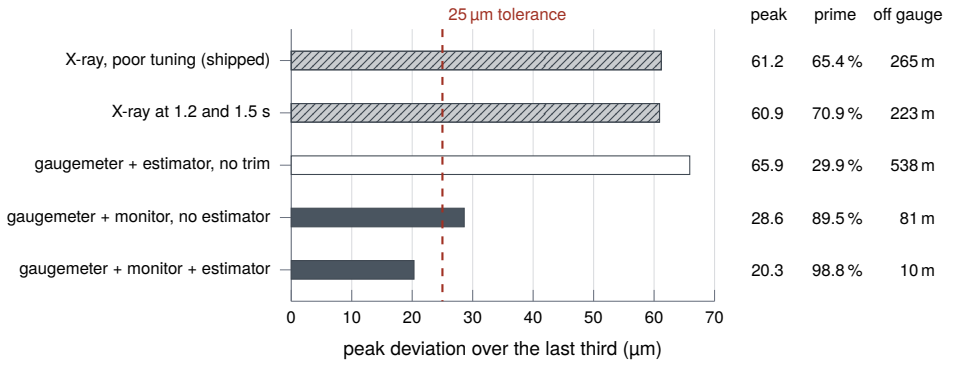


Figure 8.4 The five configurations of Table 8.7 as a chart: the peak deviation over the last third of the coil (bars and first column, in µm), with prime yield and meters off gauge in the other two columns. Hatched bars read the X-ray gauge, the white bar the gaugemeter without a trim, solid bars the gaugemeter with the monitor trim. Notice that the two hatched bars, whose tunings differ in every loop, are almost equal, and only the full scheme falls inside the dashed tolerance line.

Table 8.7 Five Configurations on the Same Ninety-Second Coil, Each After the First Built From the Suggested Settings With One Change. Peak Deviation Is the Application's Score, the Worst Logged Error Over the Last Third ($t \geq 59.4$ s), and RMS, Prime Yield and Meters off Gauge Are Cumulative Over the Whole Run of 767 m.

Configuration	Peak (µm)	RMS (µm)	Prime (%)	Off gauge (m)
Shipped: X-ray, poor tuning	61.2	27.0	65.4	265
X-ray at 1.2 and 1.5 s	60.9	24.0	70.9	223
Gaugemeter + estimator, no monitor trim	65.9	47.0	29.9	538
Gaugemeter + monitor, no estimator	28.6	17.1	89.5	81
Gaugemeter + monitor + estimator (suggested)	20.3	9.9	98.8	10

8.8 Technology in Practice

The scheme of Section 8.5 runs on three instruments, one hydraulic actuator and a layered control system. Figure 8.5 draws the gauge loops with that hardware in place of the schematic's symbols.

What Measures It

TECHNOLOGY **Sensor: The X-Ray Thickness Gauge**

An X-ray thickness gauge sends a beam through the strip from a source on one side to a detector on the other. The transmitted intensity falls exponentially with the thickness crossed, following Beer's law [17]. At low photon energies the attenuation grows with about the fourth to fifth power of the atomic number, so the reading depends on the alloy as well as on the thickness. Gauges of this family have been moving from radioactive sources to radiation generators [17]. Inside a finishing train there is no room for the gauge and the environment is hostile, so it stands only behind the final stand [18]. In this chapter it is GT-701, a 0.030 s averaging lag at the end of 5.0 m of runout table.

The gaugemeter is built from two instruments at the stand. The position transducer reads the capsule, and the load cell under the roll stack reads the roll force. A load cell is a steel element of precisely known modulus with strain gauges bonded to it, and the gauges change resistance as the element compresses or relaxes on a microscopic scale [19]. Stray electric currents through a load cell are read as a change of load, and the currents of arc welding equipment can damage it for good. In

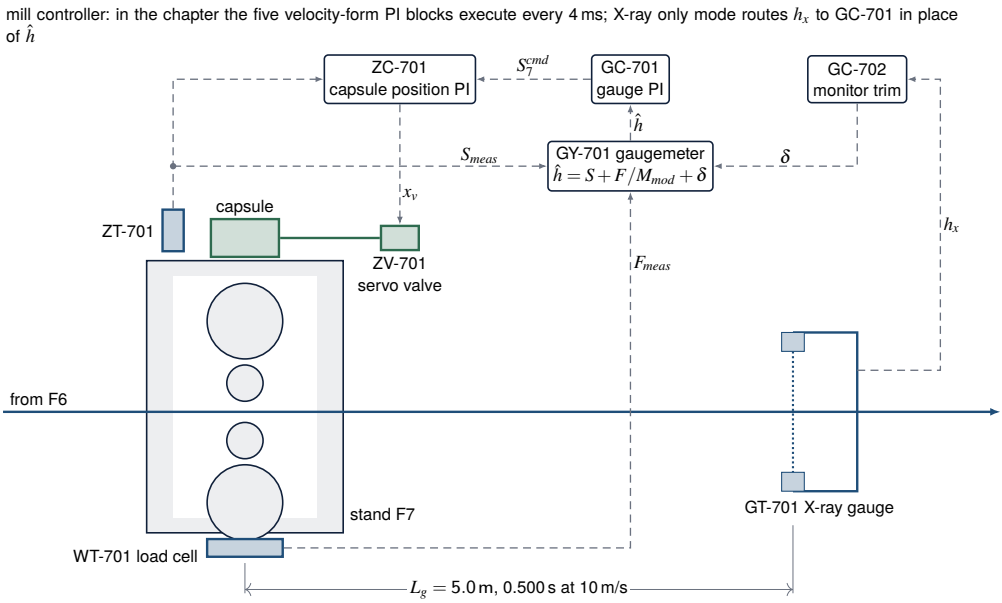


Figure 8.5 The F7 gauge loops drawn with their hardware. The position transducer ZT-701 on the capsule and the load cell WT-701 under the roll stack feed the gaugemeter GY-701 at once, and the gauge controller GC-701 sends the gap command to the capsule position loop ZC-701, which drives the servo valve ZV-701. The X-ray gauge GT-701 reads the strip itself 5.0 m downstream and feeds only the monitor trim GC-702, whose output δ biases the gaugemeter. Notice that no instrument at the stand reads the strip: the fast loop sees the mill, and only the slow loop sees the product.

the chapter's model the load cell WT-701 is a 0.004 s lag and the transducer ZT-701 a 0.002 s lag, which is why the inferred gauge is available at once.

What Moves It

The gap is moved by the capsule under the servo valve, inside a position loop that the chapter models as a proportional loop around the capsule's integrator. That loop, ZC-701 in the figure, is in service in manual and automatic alike, and the gauge loop only writes its setpoint.

> **TECHNOLOGY** **Final Element: The Hydraulic Capsule**
>
> A hydraulic gap control system moves the rolls and creates the rolling force with high-response servo valves, hydraulic cylinders and a precise measurement of the cylinder's movement, and accumulators hold its supply pressures constant [20]. Hydraulic oil hardly compresses, so a cylinder positions firmly and a modest piston generates a very large force. The oil must be purged of gas bubbles, the lines kept free of leaks, and the system stores almost no energy of its own [19]. A model of such a system carries the mill housing stretch and the compression of the oil as parameters [20]. Here the servo valve is a 0.006 s lag and the capsule an integrator of 9.0 mm/s at full valve.

Where the Controller Runs

A published model of a finishing mill stacks the work in layers: a setup calculation that fixes the gaps, speeds and tensions of each coil, a monitor thickness control that corrects from the X-ray gauge, the automatic gauge control, and the hydraulic gap control that moves the rolls [20]. That model's gap control samples every 1 ms. A more recent design names the programmable logic controller as the platform for a multivariable looper and gauge law [21]. Control loops of this kind run with cycle times from milliseconds to minutes [22]. The chapter's five loops all execute every 4 ms. The capsule

loop belongs to the gap control layer, the gauge loop with its architecture switch to the gauge control layer, and the monitor trim to the monitor layer above it.

> **IN PRACTICE** **A Gain Held Below the Limit**
>
> Inside the finishing train the thickness cannot be measured, so the roll force automatic gauge control works from the roll force alone, and its gain is usually set below the value a stable loop would allow. A smaller gain keeps the loop stable through the large variation of the mill modulus and the uncertainty that roll eccentricity adds, at the price of variation left uncorrected. The monitor control trims the roll gaps from an X-ray gauge that stands only behind the final stand. Skid marks from the reheating furnace are the major part of the entry thickness variation [18]. This chapter's shipped gain of 4.00 likewise sits below the deadbeat value of 4.752.

What Is Modernizing It

Three developments touch this chapter's loops. Adaptive feedforward gauge control extracts the long-term variation of the previous stand's exit gauge, mostly skid marks, with an adaptive filter and feeds it to the next stand. It runs on the conventional control system and was shown on real coil data [18]. Machine vision now measures the dimension this chapter leaves out. A shape meter at the finishing exit projects a pattern of light-emitting diode dots onto the strip and reads it with an area camera. Its angle error against a standard target stays within 0.45° at two sigma, it succeeds on more than 98.5 % of each coil, and it feeds flatness control through the work roll bender and leveling [23]. Diagnosis has moved onto the exit gauge itself: comparing the spectra of thickness defects relates each one to a roll, which supports a policy for changing and maintaining rolls, and a least-squares estimate shortens the eccentricity tests [24].

8.9 Limitations

What the application demonstrates soundly is inferential control and the way it fails, and what it is not is a rolling model. One dimension, a linear mill spring, frozen Hitchcock flattening and an absent shape controller between them forbid any statement about the schedule, crown, flatness or chatter of a real mill. The absent second dimension is the larger omission, because work roll bending, shifting and pair-crossing, the roll crown across the barrel and the flatness the coiler shape meter reads all interact with thickness through the roll stack, and on many mills coordinating gauge against shape is a harder problem than either alone [13]. The thermal crown here is a scalar gap offset standing in for a profile.

Three further boundaries matter for anyone reading the numbers as mill data. The strip is isothermal, so the flow stress of (8.5) is an average no instrument on a real mill measures and the skid depth is a modeling input rather than a temperature anybody reads. The interstand span is one lumped elastic state whose stiffness is dominated by roll-bite slip, and slip is not modeled explicitly. There is no threading, no head end and no tail out, so the transient that consumes most of a real gauge control system's effort, the first meters of a coil before the X-ray gauge has anything to read, is outside the model entirely [9]. Five features were kept, the mill spring, the two plant gains, the length-indexed delay, the gap-side offset and the modulus calibration, and each was kept because a gauge controller must be designed against it. Section 8.3.8 states where each simplification fails.

> **OPEN CHALLENGE** **Gauge, Looper and Tension as One Problem**
>
> Gauge, looper angle and strip tension interact, because every gap move made by the gauge control changes the mass flow the looper must absorb. Practice still runs separate proportional-integral loops that do not take these dynamic relations into account [21]. A multivariable gauge and tension control based on optimal regulator theory was proved in simulation and experiments in 1993 [25]. An integrated design of 2019 removed the interaction in simulation and was described as a prospect for use on site [21]. Progress would be an integrated controller proved on a working mill across strip sizes.

8.10 Discussion

The measured result that transfers outside steel is the ordering in Table 8.7. Three of the five configurations differ only in architecture and span 20.3 to 65.9 µm of peak deviation, while the two that differ only in tuning span 60.9 to 61.2 µm. A loop built on an inferred measurement is therefore not usefully described by its tuning until the disturbance classes have been enumerated and each one assigned to a part of the scheme. Equation (8.17) is the tool for that enumeration, and its analogue exists wherever an estimator is used in a loop. A soft sensor built from temperatures and flows is exact for whatever those measurements carry and structurally wrong for whatever changes the relationship between them [2, 3].

The frequency split is the other transferable piece. The crossover between the fast inferred measurement and the slow true one is set by the dead time of the true one and by nothing else, which is why it moves with the line speed here and why a monitor trim tuned at 10 m/s is carrying less margin at six. The same reasoning fixes where an analyzer trim, a laboratory correction or a satellite fix belongs in a cascade, and the reset time of the slow loop is the design variable that places it. Volume 5, Chapter 10 corrects a moisture soft sensor from a delayed laboratory in this way, and Volume 5, Chapter 13 corrects a dead-reckoned position from late acoustic fixes. The application exposes that choice directly, because the monitor trim reset time is a slider and the two failure modes at its ends are both reachable inside a scenario.

Yield is the economic weight behind all of it. A hot strip mill is paid by the meter inside ±25 µm, and the configurations in Table 8.7 put between 10 and 538 m of one 767 m coil outside that band. Strip rolled off gauge is downgraded or rerolled and the energy spent reheating the slab is spent either way, and scenario studies of the European iron and steel industry put process efficiency and yield alongside fuel switching among the levers available to it [26].

Methods other than the five loops fit the same written model. Capsule travel, a stand force rating and two looper stops make the thirteen-state plant a constrained optimization, and the constraint that binds in scenario 8 is the roll force rather than the gauge. An estimator running on these equations over a horizon could tell the gap offset c_7 apart from the modulus error ε_M instead of lumping both into one trim, the distinction a maintenance crew needs and a single bias cannot make.

8.11 Conclusion

The last two stands of a hot strip finishing train make a compact laboratory for inferential control. Thirteen states are integrated across the two stands, the looper and the runout table, with each stand’s exit thickness solved implicitly against its plastic curve at every integration stage, two transport delays indexed on delivered strip length and two slow drifts that make an inferred gauge wrong. At the design pass the stand stretches 2.000 mm under 12.000 MN, so a 1.000 mm gap delivers 3.000 mm of strip and only 0.2104 of a gap move reaches the product. Rearranged, the same relation is a thickness measured at the stand, and its difference from the truth is exactly the negative of the unseen gap offset plus a term in the modulus calibration error, which divides every disturbance in two. On a 2 % entry wave at identical gains the inferred loop holds 11.6 µm of root-mean-square deviation with no strip off gauge while the delayed loop gives 173.8 µm and 282 m off gauge in 300 m, and on 200 µm of thermal crown the same inferred loop leaves the gauge 198.6 µm thin against 41.9 µm with the loop open. A slow X-ray trim at an eighth the speed of the loop it corrects takes that to a 12.5 µm peak and 0.13 µm at 210 s, and a synchronous estimator locked to the back-up roll angle takes the 2.207 Hz ripple from 20.3 to 8.2 µm. A fuller model would add the cross-strip dimension, a head end and a tail out, several eccentric rolls at different frequencies and a mill resonance to stay clear of. Within its limits the simulator shows where an inferred measurement is exact, where it is inverted, and what to put underneath it.

KEY POINTS

- A finishing stand stretches under the force of rolling, so the gap it is set to is not the thickness it delivers.
- The gaugemeter is exact for every disturbance that arrives with the metal and inverted for anything that moves the roll surface without moving the position transducer.
- On a 2 % entry wave at identical gains the inferred loop holds 11.6 µm of root-mean-square deviation while the delayed loop gives 173.8 µm.
- A slow X-ray trim takes the crown error from 198.6 µm to a 12.5 µm peak, and the architecture, not the tuning, decides the challenge coil.

Where else this applies. A fast estimate trimmed by a late true measurement also corrects the moisture soft sensor of Volume 5, Chapter 10 and the dead-reckoned position of Volume 5, Chapter 13, and the roll-angle canceller has a twin in the strand bulging canceller of Chapter 6.

References

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

[2] C. Brosilow and M. Tong, "Inferential control of processes: Part II. the structure and dynamics of inferential control systems," *AIChE Journal*, vol. 24, no. 3, pp. 492–500, 1978.

[3] 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.

[4] W. C. F. Hessenberg and R. B. Sims, "Principles of continuous gauge control in sheet and strip rolling," *Proceedings of the Institution of Mechanical Engineers*, vol. 166, no. 1, pp. 75–90, 1952.

[5] G. W. Alderton and W. C. F. Hessenberg, "The continuous measurement and control of gauge in strip rolling," *International Materials Reviews*, vol. 1, no. 1, pp. 239–269, 1956.

[6] I. S. Choi, J. A. Rossiter, and P. J. Fleming, "Looper and tension control in hot rolling mills: A survey," *Journal of Process Control*, vol. 17, no. 6, pp. 509–521, 2007.

[7] E. Orowan, "The calculation of roll pressure in hot and cold flat rolling," *Proceedings of the Institution of Mechanical Engineers*, vol. 150, no. 1, pp. 140–167, 1943.

[8] R. B. Sims, "The calculation of roll force and torque in hot rolling mills," *Proceedings of the Institution of Mechanical Engineers*, vol. 168, no. 1, pp. 191–200, 1954.

[9] T. S. Bilkhu, "Hot strip mill gauge control," in *UKACC International Conference on CONTROL 2010*, 2010, pp. 150–155.

[10] M. Ataka, "Rolling technology and theory for the last 100 years: The contribution of theory to innovation in strip rolling technology," *ISIJ International*, vol. 55, no. 1, pp. 89–102, 2015.

[11] A. Kugi, W. Haas, K. Schlacher, K. Aistleitner, H. M. Frank, and G. W. Rigler, "Active compensation of roll eccentricity in rolling mills," *IEEE Transactions on Industry Applications*, vol. 36, no. 2, pp. 625–632, 2000.

[12] W. L. Roberts, *Hot Rolling of Steel*. New York, NY: Marcel Dekker, 1983.

[13] V. B. Ginzburg, *Steel-Rolling Technology: Theory and Practice*. New York, NY: Marcel Dekker, 1989.

[14] K. Wiedemer, "Significance of mill stand stiffness in relation to rolling accuracy," *Metals Technology*, vol. 1, no. 1, pp. 181–185, 1974.

[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] K. J. Åström and R. M. Murray, *Feedback Systems: An Introduction for Scientists and Engineers*, 2nd ed. Princeton, NJ: Princeton University Press, 2021.

[17] International Atomic Energy Agency, "Technical data on nucleonic gauges," International Atomic Energy Agency, Vienna, IAEA-TECDOC 1459, 2005, iSBN 92-0-107805-6.

[18] Y. K. Lee, Y. J. Jang, and S. W. Kim, "Adaptive feed-forward automatic gauge control in hot strip finishing mill," *ISIJ International*, vol. 47, no. 10, pp. 1444–1451, 2007.

[19] 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.

[20] P. Bhowal and S. K. Mukherjee, "Modeling and simulation of hydraulic gap control system in a hot strip mill," *ISIJ International*, vol. 36, no. 5, pp. 553–562, 1996.

[21] J. Sun, Y. Hu, F. Yin, Y. Hu, W. Peng, and D. Zhang, "Looper-gauge integrated control in hot strip finishing mill using inverse linear quadratic theory," *ISIJ International*, vol. 59, no. 9, pp. 1562–1572, 2019.

[22] 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.

[23] Y. Isei, T. Kato, M. Osugi, and T. Ohta, "Development of a shape meter employing the LED dot pattern projection method for a hot strip finishing mill," *ISIJ International*, vol. 60, no. 12, pp. 2876–2885, 2020.

[24] D. F. Garcia, J. M. Lopez, F. J. Suarez, J. Garcia, F. Obeso, and J. A. Gonzalez, "A novel real-time fuzzy-based diagnostic system of roll eccentricity influence in finishing hot strip mills," *IEEE Transactions on Industry Applications*, vol. 34, no. 6, pp. 1342–1350, 1998.

[25] K. Kimura, S. Nakagawa, A. Haraguchi, and H. Miura, "Advanced gauge control technology for hot strip mill," *Tetsu-to-Hagané*, vol. 79, no. 3, pp. 352–359, 1993, in Japanese with an English synopsis.

[26] N. Pardo and J. A. Moya, "Prospective scenarios on energy efficiency and CO_2 emissions in the European iron and steel industry," *Energy*, vol. 54, pp. 113–128, 2013.

8.12 Exercises

8.1 ○ Concept check · 2 min

In scenario 8 the gauge loop is closed on the gaugemeter alone, with no monitor trim, at $K_c = 4.00$ and $\tau_I = 0.25\,\mathrm{s}$. The work rolls grow 200 µm into the strip on a 40 s time constant. Long after the crown has settled, where is the delivered gauge?

(a) On the setpoint, because the integral action removes any steady error

(b) About 42 µm thin, the open-loop change, because the loop cannot see the crown

(c) About 200 µm thick, because the gaugemeter reads the higher force as thicker strip

(d) About 200 µm thin, because the loop holds the gaugemeter reading on setpoint

8.2 ○ Concept check · 1 min

Which of these disturbances does the gaugemeter read exactly once the roll force has settled?

(a) A strip that arrives 30 °C colder than the last one

(b) Work rolls that grow 200 µm into the strip as they heat

(c) A back-up roll that is out of round by 40 µm

(d) A mill modulus calibration that is 5 % too high

8.3 ○ Concept check · 2 min

The monitor trim compares the X-ray gauge with the delayed setpoint and adds its output δ to the gaugemeter reading. Why does it run with the action $a = -1$, direct acting in the convention of Section 1.3?

(a) The X-ray gauge sits downstream, and a transport delay reverses the sign of a gain

(b) A larger trim raises the reading the gauge loop holds, so the gap shuts and the strip thins

(c) A trim is a bias on a measurement, and every bias loop in the apps runs with -1

(d) A thermal crown makes the gaugemeter error positive, so the trim has to come out negative

8.4 ▭ Theory and calculation · 6 min

The F7 stand has a mill modulus $M_7 = 6.00\,\mathrm{MN/mm}$ and rolls 3.000 mm strip under 12.000 MN, with no gap offset. The plastic modulus of the pass is $Q = 22.513\,\mathrm{MN/mm}$.

(a) Find the mill stretch and the unloaded roll gap. *Checkpoint:* 2.000 mm and 1.000 mm.

(b) Find the gap gain $\partial h/\partial S$ and the deadbeat gain of the gauge loop. *Checkpoint:* 0.2104 and 4.752.

(c) A back-up roll out of round by 40 µm moves the roll surface without moving the position transducer. Find the amplitude of the ripple in the true gauge and in the gaugemeter reading, and say whether the two move together. *Checkpoint:* 8.4 µm and 31.6 µm, in opposite directions.

8.5 Theory and calculation · 5 min

The gaugemeter divides the F7 force by $M_{mod} = 6.30\,\mathrm{MN/mm}$, a calibration 5 % high, while the stand has $M_7 = 6.00\,\mathrm{MN/mm}$. There is no gap offset.

(a) Find the gaugemeter error at the design force of 12.000 MN. *Checkpoint:* −95.2 µm.

(b) The gauge loop on the gaugemeter alone opens the gap until the force has fallen to 10.27 MN. Find the steady error of the delivered gauge. *Checkpoint:* +81.5 µm.

(c) The monitor trim is switched in and returns the force to 12.000 MN. Find where the trim parks, and where it would park for a calibration 5 % low, $M_{mod} = 5.70\,\mathrm{MN/mm}$. *Checkpoint:* +0.0952 mm and −0.1053 mm.

8.6 Theory and calculation · 6 min

The X-ray gauge stands $L_g = 5.0\,\mathrm{m}$ past F7. Model the loop closed on it as the gap gain 0.2104 in series with a pure delay θ_d and a proportional controller.

(a) Find θ_d at 10 and at 6 m/s. *Checkpoint:* 0.500 s and 0.833 s.

(b) Find the ultimate gain and the period of oscillation at the ultimate gain at each speed. *Checkpoint:* 4.75 at both speeds, with periods of 1.0 s and 1.67 s.

(c) The app limit-cycles at $K_c = 4.00$ and $\tau_I = 0.25\,\mathrm{s}$, with periods of 1.374 s at 10 m/s and 2.09 s at 6 m/s. Name what the simple model leaves out, and say whether the simple model predicts a limit cycle at this gain. *Checkpoint:* The capsule, the X-ray averaging filter and the reset action. It does not.

8.7 App challenge · 8 min

Load scenario 1, The gap you set is not the gauge you get, and change no setting. The F7 roll gap command steps from 1.000 to 1.250 mm at 6 s and to 0.750 mm at 16 s. Run at a speed of 1 with the default seed and pause once between 12 and 15.9 s and once between 22 and 25.9 s, where the readings are steady, or read the rows at 15 s and 25 s in the exported CSV file (columns `gauge_mm`, `F7_MN` and `stretch7_mm`).

(a) In the first pause, read the Gauge error, Roll force and Mill stretch tiles, and find the gap gain of the opening step. *Checkpoint:* 54.6 µm, 10.83 MN, 1.805 mm, and a gap gain of 0.218.

(b) In the second pause, read the same tiles and find the gap gain of the closing step from 1.250 to 0.750 mm. *Checkpoint:* −52.6 µm, 13.18 MN, 2.197 mm, and a gap gain of 0.214.

(c) From the readings of (a) and (b), estimate the mill modulus. *Checkpoint:* 6.00 MN/mm, within 0.05.

Accept tile readings within one unit of their last digit.

8.8 App challenge · 9 min

Load scenario 8, The rolls grow into the strip. Its crown event starts at 10 s. Use the default seed and a speed of 2.

(a) Run the scenario as shipped, with Gauge control at Gaugemeter, to 120 s. Read the Gauge error, Gaugemeter error and Roll force tiles. *Checkpoint:* −187.1 µm, 187.2 µm and 16.51 MN.

(b) Reload the scenario, set Gauge control to Manual before Run, and run to 120 s again. Read the same three tiles. *Checkpoint:* −39.5 µm, 187.2 µm and 12.89 MN.

(c) Find the ratio of the two gauge errors, and explain why the gaugemeter error is the same in both runs. *Checkpoint:* 4.7.

Accept each reading within 0.2 µm or 0.02 MN.

8.9 App challenge · 9 min

Load the challenge, One coil, everything at once. Each run covers the whole 90 s coil at a speed of 2 with the default seed. The app pauses by itself at the end of the coil, when the clock reads 90.0 s. For each run, read the challenge result, the worst delivered gauge error over the last third of the coil, and the Prime yield tile at the end.

(a) Run the challenge as shipped, on the X-ray gauge with the poor tuning. *Checkpoint:* 61.2 µm and 65.4 %.

(b) Reload, press Reveal suggested tuning, and run again. *Checkpoint:* 20.3 µm and 98.8 %.

(c) Open the advanced controls, set Eccentricity estimator rate to 0, press Reset and run again. Name the disturbance that now spoils the coil. *Checkpoint:* 28.6 µm and 89.5 %, the back-up roll runout.

Accept each value within 0.2 µm or 0.2 %.

8.10 Application area · 4 min

A mill builder offers a stiffer F7 housing with $M_7 = 12.0\,\mathrm{MN/mm}$. Keep the pass of the chapter: 12.000 MN, $Q = 22.513\,\mathrm{MN/mm}$, and the entry sensitivity $\partial F/\partial h_{in}$ implied by the walk-through factor of 0.6534 on the present stand.

(a) Find the stretch, the gap gain and the deadbeat gain of the stiffer stand. *Checkpoint:* 1.000 mm, 0.348 and 2.88.

(b) Find the fraction of an entry thickness variation that reaches the product. *Checkpoint:* 0.540, against 0.6534.

(c) Find the gaugemeter error for 200 µm of crown, and its ratio to the true change. *Checkpoint:* 200 µm, and 2.88 times the true change.

8.11 Application area · 3 min

The slab rested on water-cooled skids 1.5 m apart in the reheat furnace, and the cool bands arrive 125 m apart on the 3.000 mm strip. Neglect any change in width.

(a) Find the elongation from slab to strip and the slab thickness it implies. *Checkpoint:* 83.3 and 250 mm.

(b) Find the time between bands at 10 and at 7 m/s, the tail-out speed of the challenge. *Checkpoint:* 12.5 s and 17.9 s.

8.12 Review the agent · 5 min

An AI assistant is asked how to deal with the back-up roll runout and replies as follows. “Skip the eccentricity estimator and set the force signal filter to $\tau_F = 0.20\,\mathrm{s}$. Its corner is at 0.80 Hz, so it cuts the 2.207 Hz runout by a factor of 2.9, while the 0.333 Hz entry wave keeps 92 % of its amplitude. The filter needs no shaft encoder and no tuning, so it is the more robust choice and costs nothing at the entry wave.” Find the flaw and say how you would check it.