INTERACTIVE CONTROL ENGINEERING
Principles and Applications

Volume 16 Control in Iron, Steel, and Metals Production

CHAPTER 10

Galvanizing Line Simulator for Coating Weight Control With a Delay Fixed in Length

Pages 211–231

Edited by John D. Hedengren

How to cite this chapter. "Galvanizing Line Simulator for Coating Weight Control With a Delay Fixed in Length," 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. 10, pp. 211–231.

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

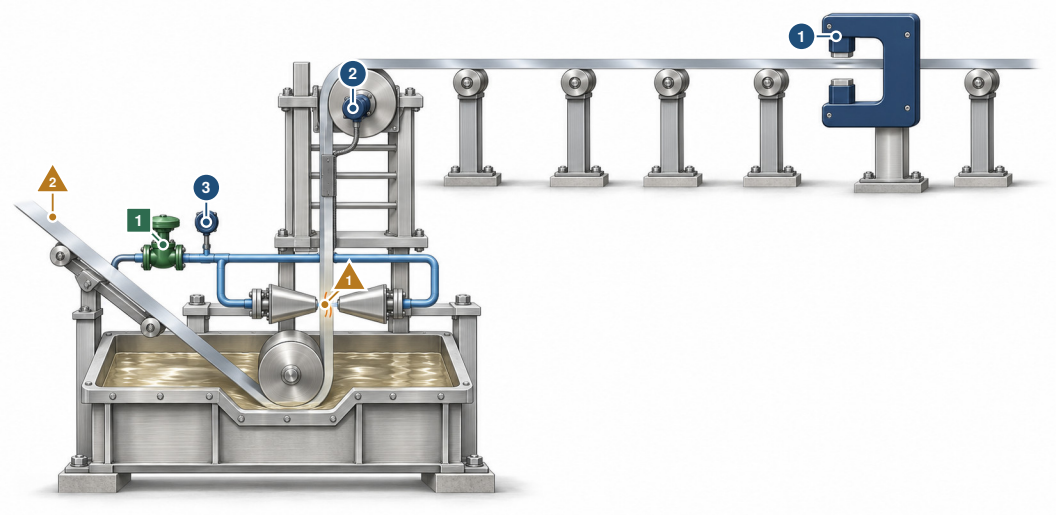
The app for this chapter
https://apmonitor.com/apps/galvanizing/

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

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

CHAPTER 10

Galvanizing Line Simulator for Coating Weight Control With a Delay Fixed in Length



The controller moves 1 air pressure to the knives

The controller reads 1 coating gauge, far down the strip 2 line speed 3 air pressure at the knives

What disturbs it 1 coil shape and flutter at the knives 2 a change of line speed

Notice that the knives act just above the zinc, while the gauge sees the coating only after a long run of strip, so every correction is late. Illustration, not to scale.

IN PLAIN TERMS **Holding a Zinc Coating With a Late Gauge**

The main idea. A galvanizing line pulls steel strip out of melted zinc, and two flat jets of air blow the extra zinc back. Harder air leaves a thinner coat, so the controller sets the air pressure to hold the coat on target.

The challenge. The coat can be measured only after it cools, far down the line, so the gauge always reports old news. When the line slows, that wait grows, and a controller that counts it in seconds instead of in length of strip starts to swing.

Why it matters. Every gram of zinc kept as a safety margin costs money, so a steadier coat saves zinc.

THREE WAYS TO READ THIS CHAPTER

Explore	12 min	The first page, the history and fact boxes of the Introduction, the history box of Section 10.2, the try-it and fact boxes of Section 10.6, and the key points
Study	35 min	The whole chapter with its pause points and the twelve exercises
Practice	15 min	The first page, Section 10.6, Section 10.7, Section 10.8 and its open challenges

The control question The coating gauge of a galvanizing line sits a fixed length of strip after the air knives, so its delay changes with every line speed, and the coating's sensitivity to knife pressure changes with every grade. Which loop holds the coating through both, and what is each gram of spread worth in zinc?

Learning objectives

- Compute the knife pressure, process gain and gauge delay of the wiping line at a given speed and coating.
- Explain why more gain on a loop with a long delay widens the coating spread.
- Show why a Smith predictor must keep a transport delay in meters, and find the speeds at which one kept in seconds fails.
- Use speed feedforward to keep the loop gain constant through a slowdown.
- Price coating variance as target margin and zinc, and say why a heavy lower tail needs more margin than the normal formula.

Prerequisites Sections 1.4.5 and 1.4.6 (delay as a length, wiping law and coating gain), 1.6.2 and 1.6.4 (the PI algorithm, scenarios and scorecard).

The app https://apmonitor.com/apps/galvanizing/

See also Volume 5, Chapter 5 for a Smith predictor on a moving ribbon, Volume 5, Chapter 4 for acting ahead of a late analyzer and Volume 5, Chapter 3 for tuning through a slow measurement. Volume 3, Chapter 6 for feedforward from a measured disturbance with a late trim, and Volume 3, Chapter 7 for a ratio paced by throughput.

10.1 Introduction

Many measurements in the process industries are taken downstream of the point where the product is made. Basis weight on a paper machine, film thickness after a die, the composition of a blend at the end of a pipeline and the gauge of rolled strip are all read after the material has traveled some distance, and in each case the dead time the controller faces is that distance divided by a throughput, as Section 1.4.5 sets out. Dead-time compensation began with the Smith predictor and has grown into a family of methods, and the stability of these compensators when their model of the delay is wrong has been analyzed since 1980 [1, 2]. A dead time also sets a floor under the variance that any feedback controller can reach, because nothing the controller does reaches the measurement before it, which is the basis of minimum-variance performance assessment [3, 4].

> ⧖ HISTORY **A Predictor for Dead Time (1957)**
>
> The Smith predictor was proposed in 1957. It splits the model of a process into a fast part without delay and a pure delay, and from them it builds a signal that anticipates the process output. The controller acts on that signal, and the delay drops out of the characteristic equation of the closed loop, but only while the model matches the plant. The original design was meant for plants that are stable in open loop, and later work extended it to integrating and unstable processes. A review of 2022 calls it the best known and most widely used method of dead-time compensation for a single loop, in teaching and in industry [1].

Hot-dip galvanizing is a clear example of this class. Zinc coatings protect steel sheet against corrosion, and the metallurgy of zinc-coated steel is reviewed in [5]. On a continuous line the strip leaves a bath of molten zinc vertically, and two gas jets, the air knives, blow the excess back into the bath [6, 7]. The coating is measured only after it has frozen, here 100 m of strip later. At 150 m/min that is 40 s, and at 90 m/min it is 67 s. A transfer function with a fixed dead time hides this, and it also hides that the sensitivity of the coating to knife pressure moves with the operating point. The simulator is built to correct two beliefs. The first is that a dead time is a number of seconds, which leads to a Smith predictor that works at one speed and oscillates at another. The second is that a tighter loop gives a narrower product, which a long delay makes false.

The economic side is equally direct. The coating target must sit above the specification minimum by a margin that covers the spread, so a reduction in variance is worth the zinc saved by moving the target down, and economic assessments of advanced control value variance reduction in this way [8]. On the 1.5 m strip of the simulator one gram per square meter on both sides at 150 m/min is 27 kg of zinc, or about $81, every hour.

> ✦ DID YOU KNOW? **Most Zinc Ends Up on Steel**
>
> World mine production of zinc was about 12 million tonnes in 2024, an estimate. In the United States galvanized steel was the leading use of refined zinc, and aluminum and plastics are the substitutes for galvanized sheet in automobiles [9].

This chapter works through a galvanizing line simulator in which feedback on the late gauge, speed feedforward and a Smith predictor with its delay kept in seconds or in meters are compared on the same coils, speed changes and grade changes, with every result priced in zinc. All of these strategies are established ones. Section 10.2 reviews coating weight control. Section 10.3 gives a model complete enough to rebuild, including the strip-coordinate history that makes the gauge late. Section 10.4 describes the loop and the two forms of the predictor, and Section 10.5 analyzes the linearized loop at each operating point. Section 10.6 reports eleven scenarios that separate what a delay fixed in length does to feedback, to a predictor and to the economics of the coating target. Section 10.7 describes the hardware of real lines, Section 10.8 states the limitations of the model and Section 10.9 sums up.

10.2 Coating Weight Control on Galvanizing Lines

The history of zinc-coated steel, and the metallurgy of the coatings made on continuous strip lines, are reviewed in [5, 10]. On a continuous line the coating is no longer set by the dipping of each piece but by what is removed from a moving strip as it leaves the bath, and gas jet finishing became the means of removal. An analytical description of jet finishing for hot-dip coatings on strip appeared in 1976 [6], and an analysis of jet stripping of liquid coatings followed in 1984 [11]. The strip speed, the nozzle pressure and the standoff distance from the nozzle to the strip are the process parameters of jet wiping [7], and they are the three quantities of the correlation used here.

Closed-loop control of coating weight arrived in the same decade, and its first steps are told in the history box that follows. The coating gauge of such lines is necessarily downstream of the knives, because the coating must freeze before it can be measured, so every one of these loops faces a transport delay.

> ⧖ HISTORY **Coating Weight Under Computer Control (1975 to 1992)**
>
> Computer control of zinc coating weight on a galvanizing line was reported in 1975 [12], and the regulation of coating thickness in continuous galvanizing of sheet was described in 1977 [13]. A review of 2006 dates the idea of coating weight control on continuous lines to 1976. The analytical wiping model of the time was too complicated for an on-line computer, so a simple exponential model in line speed, knife distance and pressure took its place in the controllers. In 1992 an adaptive controller was built on a linearized form of that model [14]. The wiping law of Equation (10.1) has the same exponential form.

Physical modeling of the wiping has continued since. An engineering model of jet wiping predicts the final coating thickness in agreement with experiments, and identifies splashing of the film as the limit on how hard the jets can wipe [7]. A coating weight model for the continuous hot-dip process was published in the same year [15]. Automatic coating weight control systems for continuous galvanizing lines, and coating weight control on an industrial line at the Florange works, were reported in 2008 and 2010 [16, 17].

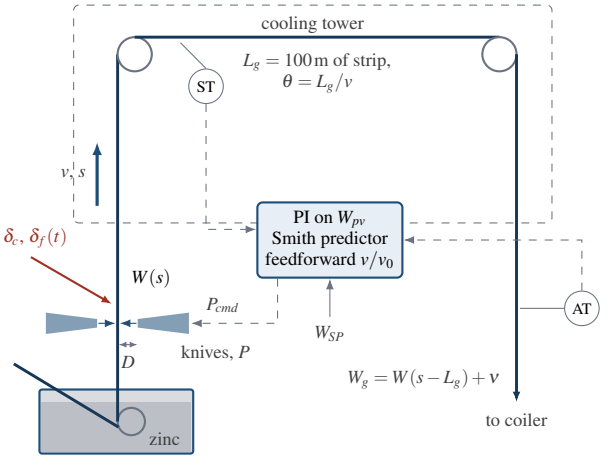


Figure 10.1 Continuous galvanizing line as modeled. Strip enters the zinc pot, turns under the sink roll and rises at speed v between two air knives at pressure P and distance D from the strip, which leave a coating $W(s)$ at strip coordinate s. The coil shape offset δ_c and the flutter δ_f move D unmeasured. The strip travels $L_g = 100$ m over the cooling tower to the gauge AT, which reads the coating made at $s - L_g$ plus noise ν. The controller receives the gauge, the speed from ST and the target W_{SP}, and sends the pressure command P_{cmd} to the knives.

Table 10.1 Variables of the Coating Loop

Symbol	Role	Nominal	Range
W_g	CV, gauge coating per side	76 g/m^2	60–160
W	coating at the knives	76 g/m^2	—
P_{cmd}	MV, knife pressure command	26.2 kPa	2–70
P	state, knife pressure	26.2 kPa	2–70
v	state, line speed	2.5 m/s	1–3.5
s	state, strip coordinate	0	≥ 0
A, A_b	states, area coated, below W_{min}	0	m^2
Z, Z_x	states, zinc used, above W_{min}	0	kg
v_{SP}	DV, line speed setting	150 m/min	60–210
δ_c	DV, coil shape offset (1σ)	4 %	0–8
δ_f	DV, flutter (1σ)	1 %	0–6
L_c	DV, coil length	2000 m	500–5000
ν	DV, gauge noise (1σ)	1 g/m^2	0–5
W_{min}	DV, specification minimum	70 g/m^2	30–150

10.3 Model of the Wiping Line

The line runs a strip 1.5 m wide at a reference speed of 150 m/min. The specification requires at least 70 g/m^2 of zinc on each side, and the default target is 76 g/m^2. The strip, the knives, the cooling tower and the gauge are shown in Figure 10.1.

Table 10.1 classifies the variables. The controlled variable is the gauge reading, not the coating at the knives, which no instrument sees. The model carries seven states, the knife pressure P, the line speed v, the strip coordinate s and four accumulators for the scorecard. The knife distance D and the gauge are algebraic, updated once per controller step.

10.3.1 Wiping, Knives and Line

The coating left on each side follows a power law in speed, distance and pressure, Equation (10.1), which is the wiping law of Equation (1.11) with this line's constants.

$$W = W_0 \left(\frac{v}{v_0}\right)^{0.6} \left(\frac{D}{D_0}\right) \left(\frac{P}{P_0}\right)^{-0.6} \tag{10.1}$$

where $W_0 = 70\,\mathrm{g/m^2}$ is the coating at the reference speed $v_0 = 2.5\,\mathrm{m/s}$, distance $D_0 = 10\,\mathrm{mm}$ and pressure $P_0 = 30\,\mathrm{kPa}$. The code bounds v below at 0.05 m/s and P at 0.5 kPa inside the power law. A faster strip drags out more zinc, knives farther away wipe less, and a higher pressure wipes more, which are the directions of the jet stripping analyses [6, 11]. The pressure that gives a coating W follows by inversion, and it gives $P = 26.16\,\mathrm{kPa}$ for $76\,\mathrm{g/m^2}$ at the reference condition.

The sensitivity of the coating to pressure is the process gain of the loop, Equation (10.2).

$$K_p = \frac{\partial W}{\partial P} = -0.6\,\frac{W}{P} \tag{10.2}$$

It is $-1.74\,\mathrm{g/m^2}$ per kPa at the reference condition. Holding $76\,\mathrm{g/m^2}$ at 90 m/min needs $P = 15.7\,\mathrm{kPa}$ and gives $K_p = -2.91$, and a heavy grade of $145\,\mathrm{g/m^2}$ at 150 m/min needs 8.91 kPa and gives $K_p = -9.76$, which is 5.6 times the nominal value. The gain rises whenever the coating rises or the pressure falls, and both happen on a heavy grade.

The knife pressure follows its command through a first-order lag with the command clipped to the actuator range, and the line speed ramps toward the speed the line sets, Equation (10.3).

$$\begin{aligned} \tau_P \frac{\mathrm{d}P}{\mathrm{d}t} &= \mathrm{clamp}(P_{cmd}, P_{min}, P_{max}) - P \\ \frac{\mathrm{d}v}{\mathrm{d}t} &= \mathrm{clamp}\left(\frac{v_{SP} - v}{2\,\mathrm{s}},\ -a_v,\ a_v\right), \qquad \frac{\mathrm{d}s}{\mathrm{d}t} = v \end{aligned} \tag{10.3}$$

with $\tau_P = 1\,\mathrm{s}$, $P_{min} = 2\,\mathrm{kPa}$, $P_{max} = 70\,\mathrm{kPa}$ and $a_v = 0.02\,\mathrm{m/s^2}$. A slowdown from 150 to 90 m/min therefore takes about 50 s.

10.3.2 Strip Shape

The knife distance is not measured and not manipulated. It carries a step at every coil and a flutter, Equation (10.4).

$$D = D_0\left(1 + \delta_c + \delta_f(t)\right) \tag{10.4}$$

Each time a further coil length L_c of strip has passed the knives, δ_c is redrawn as $\sigma_c \xi$ with ξ standard normal and $\sigma_c = 4\,\%$. The first change comes after $0.6L_c$, 8 min into a run at 150 m/min, and later coils last 13.3 min at that speed. The flutter is a first-order filtered Gaussian sequence with a correlation time of 60 s, scaled to a stationary standard deviation $\sigma_f = 1\,\%$, Equation (10.5).

$$x_k = a x_{k-1} + \sqrt{1 - a^2}\,\xi_k, \quad a = \mathrm{e}^{-\Delta t/60\,\mathrm{s}}, \quad \delta_f = \sigma_f x_k \tag{10.5}$$

started from $x_0 = 0$. Because W is proportional to D, a coil step of one standard deviation moves the coating by $3.0\,\mathrm{g/m^2}$ and the flutter by $0.76\,\mathrm{g/m^2}$ at the target.

10.3.3 The Gauge and the Strip Coordinate

Every controller step records the coating made at the knives against both the time t and the strip coordinate s in a ring buffer of 6000 entries, 20 min at the 0.2 s step. The gauge reads the entry for the strip that is now L_g past the knives, Equation (10.6).

$$W_g(t) = W\left(s(t) - L_g\right) + \nu(t), \qquad \theta(t) = \frac{L_g}{v(t)} \tag{10.6}$$

Table 10.2 Parameters of the Line, Adjustable in the Parameter Panel Unless Marked Fixed

Symbol	Meaning	Value	Class
W_0	reference coating	70 g/m^2	adjustable
v_0	reference speed	150 m/min	adjustable
P_0	reference pressure	30 kPa	adjustable
D_0	reference knife distance	10 mm	adjustable
—	speed, pressure exponents	0.6, −0.6	fixed
τ_P	knife pressure lag	1.0 s	adjustable
P_{min}, P_{max}	pressure limits	2, 70 kPa	adjustable
L_g	knives to gauge	100 m	adjustable
a_v	line speed ramp	0.02 m/s^2	adjustable
—	speed approach time	2 s	fixed
b	strip width	1.5 m	adjustable
p_{Zn}	zinc price	\$3.0/kg	adjustable
W_{min}	minimum per side	70 g/m^2	adjustable
—	flutter, noise correlation	60, 2 s	fixed

The lookup interpolates linearly between the two entries that bracket $s - L_g$. Equation (10.6) is the delay defined by a length of Equation (1.7), and the X-ray gauge of Chapter 8 is keyed to strip length in the same way. The noise ν is a second filtered Gaussian sequence of the form of Equation (10.5) with a 2 s correlation time and a standard deviation of 1 g/m^2. The delay θ is 40 s at 150 m/min, 67 s at 90 m/min and 33 s at 180 m/min. During a speed change the gauge reads strip that was coated at whatever speed it had when it passed the knives, which Equation (10.6) handles without a separate delay calculation.

10.3.4 Zinc and the Scorecard

The accumulators integrate the strip area, the zinc used and the zinc above the minimum on both sides, Equation (10.7).

$$\begin{aligned} \frac{\mathrm{d}A}{\mathrm{d}t} &= vb, \qquad \frac{\mathrm{d}Z}{\mathrm{d}t} = \frac{2Wvb}{1000} \\ \frac{\mathrm{d}Z_x}{\mathrm{d}t} &= \frac{2\max(0, W - W_{min})\,vb}{1000} \end{aligned} \tag{10.7}$$

in m^2/s and kg/s, with $b = 1.5$ m, and A_b grows at vb while $W < W_{min}$. The scorecard reports the mean $\bar{W}$ and standard deviation σ_W of the coating at the knives weighted by strip length since the start of the run, the strip below the minimum $100A_b/A$, the zinc cost $p_{Zn}Z_x/t$ at $p_{Zn} = \$3$ per kilogram averaged over the run, and a target this spread allows, $W_{min} + 2.33\sigma_W$, which would put 1 % of the strip below the minimum if the coating were normally distributed. One gram per square meter on both sides costs $2vbp_{Zn}$, which is \$81 an hour at 150 m/min and \$49 at 90 m/min.

The parameters are collected in Table 10.2, and the disturbance sliders and their ranges in Table 10.1, where the gauge noise is shown only in the advanced view. The line starts at $v = v_{SP}$ with the pressure that puts the coating on target in automatic, or at the manual pressure, and $s = 0$. The history is filled for 120 s behind the knives with that initial coating, so the gauge agrees with the knives at the start. The states advance in 0.2 s steps of the fourth-order Runge-Kutta engine of Section 1.6.1, with D and P_{cmd} held over each step, and a scenario replays exactly at its seed.

10.3.5 Assumptions

The wiping correlation has fixed exponents and a coating proportional to the knife distance. Jet wiping models also depend on the nozzle geometry, the tilt of the jet and the film properties, and the jets cannot wipe beyond the onset of splashing [7].

Both sides and the whole width are one coating. Real lines coat the two sides with separate knives and see crossbow and edge effects across the width, so the distance and the coating vary across the strip.

The knife distance is disturbed but never moved. On real lines the knife position is a second actuator used for large changes, and electromagnetic or mechanical stabilization reduces the flutter at its source.

The gauge is a point measurement with light filtered noise. A traversing gauge scans the width and mixes time with position, which lengthens the effective delay.

The strip does not stretch between the knives and the gauge, so the delay is exactly L_g/v, and the zinc price, the width and the minimum are constant apart from the grade change of two scenarios.

10.4 The Coating Loop as Implemented

One loop moves the knife pressure to hold the coating. It runs every 0.2 s, the same step as the physics, before each integration step. In the convention of Section 1.6.2 the loop is direct acting, $a = -1$, because more pressure leaves less zinc, so a coating above target raises the pressure. The app's Show Model dialog uses the same name.

10.4.1 Discrete PI Algorithm

The loop uses the velocity-form controller of Section 1.6.2 with the derivative term at zero. Equations (10.8) to (10.10) are Equations (1.13) to (1.16) written for this loop, with the output limits divided by the feedforward gain.

$$e_k = a\,(W_{SP} - W_{pv,k}), \quad a = -1, \quad p_k = K_c e_k \tag{10.8}$$

$$S_k = S_{k-1} + p_k - p_{k-1}, \quad c_k = S_k + I_{k-1} \tag{10.9}$$

$$u_k = \text{clamp}\left(c_k, \frac{P_{min}}{g_k}, \frac{P_{max}}{g_k}\right)$$

$$I_k = I_{k-1} + \left[\frac{K_c}{\tau_I} e_k + \frac{u_k - c_k}{\tau_I}\right] \Delta t \tag{10.10}$$

where $\Delta t = 0.2\,\text{s}$, W_{pv} is the loop measurement of the next subsection, and g_k is the feedforward gain of Section 10.4.3, equal to one when feedforward is off. The integral is unwound by back-calculation with a tracking time equal to τ_I. There is no derivative term, no rate limit and no output filter, so the pressure lag τ_P is the only smoothing between the controller and the knives. The command is $P_{cmd,k} = \text{clamp}(g_k u_k, P_{min}, P_{max})$. At reset the integral is seeded with the initial pressure divided by g, and in manual the same expression is solved for I_k against the operator's pressure, which makes both transfers bumpless. The default tuning is $K_c = 0.12\,\text{kPa per g/m}^2$ and $\tau_I = 15\,\text{s}$, and the predictor scenarios use $K_c = 0.3$ and $\tau_I = 4\,\text{s}$.

10.4.2 Three Loop Measurements

The measurement W_{pv} is chosen in the app. The first choice is the gauge itself, $W_{pv} = W_g$. The other two are Smith predictors built on the nominal wiping model, Equation (10.11).

$$\hat{W}(t) = W_0 \left(\frac{v}{v_0}\right)^{0.6} \left(\frac{P}{P_0}\right)^{-0.6} \tag{10.11}$$

which is Equation (10.1) with the distance at D_0, evaluated with the measured speed and the actual knife pressure, and recorded in the same history as W. The predictor in meters compares the gauge with the model's value for the strip the gauge is reading, and the predictor in seconds looks the model up a fixed time $\theta_0 = L_g/v_d$ back, Equation (10.12). The design speed $v_d = 150\,\text{m/min}$ is fixed in the

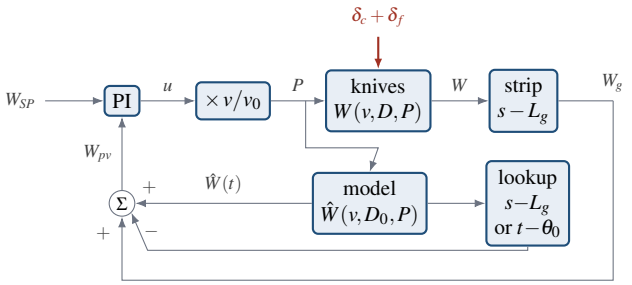


Figure 10.2 Coating loop with the Smith predictor. The model runs on the measured speed and pressure at the nominal distance, its value now enters the loop measurement with the gauge, and its value for the strip the gauge reads is subtracted. The lookup is by strip coordinate in one option and by a fixed 40 s in the other. With the gauge option the model path is absent, and with feedforward off the multiplier is one.

code and does not follow the reference speed v_0 of the wiping correlation, so $\theta_0 = 40\,\text{s}$ for the 100 m gauge distance.

$$\begin{aligned} W_{pv}^m &= \hat{W}(t) + W_g(t) - \hat{W}\left(s - L_g\right) \\ W_{pv}^s &= \hat{W}(t) + W_g(t) - \hat{W}\left(t - \theta_0\right) \end{aligned} \tag{10.12}$$

The loop structure is drawn in Figure 10.2. Substituting Equation (10.6) into the predictor in seconds separates what it sees, Equation (10.13).

$$\begin{aligned} W_{pv}^s = \hat{W}(t) &+ \left[W - \hat{W}\right]_{s-L_g} + \nu \\ &+ \left[\hat{W}(s - L_g) - \hat{W}(t - \theta_0)\right] \end{aligned} \tag{10.13}$$

The first bracket is the model error for one piece of strip, which is the correction a Smith predictor is meant to apply. The second bracket is zero only while $\theta = \theta_0$. At 90 m/min it carries the negative of every change in the model coating for the 27 s between the time the predictor assumes the change reached the gauge and the time it really does, and the loop reads that as a disturbance. The predictor in meters has no second bracket at any speed. Volume 5, Chapter 5 applies a Smith predictor on another moving product, a glass ribbon whose gauge reads glass formed long before.

10.4.3 Speed Feedforward

Equation (10.1) holds the coating constant when P is proportional to v. With feedforward on, the controller output u is the pressure the loop would use at the reference speed, and the knives receive $g_k u_k$ with $g_k = v_k/v_0$, which gives the coating of Equation (10.14).

$$W = W_0 \left(\frac{D}{D_0}\right) \left(\frac{u}{P_0}\right)^{-0.6} \tag{10.14}$$

The speed has disappeared, so a speed change no longer moves the coating, and the gain seen by the controller, $\partial W/\partial u = -0.6W/u$, is the nominal $-1.74\,\text{g/m}^2$ per kPa at every speed for the $76\,\text{g/m}^2$ target. The limits on u are divided by g_k so that the anti-windup acts on the pressure actually available. The feedforward still sees the ramp through the 1 s pressure lag. Scaling a manipulated variable by a measured throughput is the speed feedforward of Section 1.4.6. Chapter 9 feeds the line speed forward to the firing of its furnace, and Volume 3, Chapter 6 computes a coagulant dose from a measured raw-water disturbance and leaves the late measurement a slow trim.

Table 10.3 Linearized PI Loop at the Scenario Operating Points. GM Is the Gain Margin and T_u the Period at Phase Crossover

Operating point	K_c, τ_I	P	$\lvert K_p \rvert$	θ (s)	GM	T_u (s)
150 m/min, case 2	0.12, 15	26.2	1.74	40	3.05	114
150 m/min, case 3	0.25, 10	26.2	1.74	40	1.02	127
90 m/min, case 4	0.12, 15	15.7	2.91	67	1.16	214
90 m/min, feedforward	0.12, 15	26.2[a]	1.74	67	1.93	214
145 g/m^2, case 10	0.12, 15	8.9	9.76	40	0.55	114

[a] controller output u, the pressure at the reference speed

10.4.4 Three Models Kept Apart

The simulation model is Equations (10.1) to (10.7) with the true, disturbed distance. The controller's internal model is Equation (10.11), which shares the plant's parameters and exponents and differs only in the distance, so its error is exactly the coating change caused by $\delta_c + \delta_f$. No model is identified in the application. The tunings are supplied with the scenarios, and the control model used to analyze them in Section 10.5 is the linearization of the plant at each operating point.

10.5 The Linearized Loop

Around an operating point the coating responds to the pressure command as a first-order lag with dead time, $G(s) = K_p \mathrm{e}^{-\theta s}/(\tau_P s + 1)$, where K_p comes from Equation (10.2), θ from Equation (10.6) and $\tau_P = 1$ s. With the PI controller the open loop is $K_c(1 + 1/\tau_I s)G(s)$, and its gain margin and the period at its phase crossover were computed for the operating points of the scenarios. The results are listed in Table 10.3.

The default tuning has a gain margin of 3.05 at 150 m/min. Doubling the gain and shortening the integral time leaves a margin of 1.02, so case 3 sits on its stability limit. A slowdown to 90 m/min multiplies the gain by 1.67 and the delay by 1.67 at once and leaves 1.16. Feedforward restores the nominal gain, and the margin rises to 1.93 against the longer delay alone. On the heavy grade the margin is 0.55 and the linear loop is unstable. The simulated pressure rings near these periods at the 0.2 s step. It crosses its mean upward every 121 to 155 s in the bursts of case 3, its peaks come 191 to 236 s apart after the slowdown of case 4, and case 10 settles into a saturated limit cycle of 144 s.

The delay also bounds what the loop can remove, because the best achievable variance is set by the disturbance content faster than θ [3, 4]. At 150 m/min the coil shape steps last twenty delay times and the flutter correlation time is 1.5 delay times, so the loop can remove the first and barely touch the second. A Smith predictor removes the delay from the stability of the loop, but the gauge still reports each piece of strip θ late, so the predictor cannot lower this floor.

10.6 Scenarios and Measured Results

Table 10.4 gives the eleven scenarios, which follow the pattern of Section 1.6.4. All values below come from runs of the engine at its default seed, and the statistics are those of the coating at the knives weighted by strip length. Unless stated otherwise, the standard deviation, range and strip below the minimum exclude the first 5 min, and the zinc cost is the scorecard's average over the whole run. The scorecard itself accumulates from $t = 0$, so it reads slightly different values, for example 1.39 g/m^2 in case 2 and 4.58 g/m^2 with 9.4 % below the minimum in case 4. The coil shapes are random, so a result that rests on a few coil changes moves with the seed, and the spread over ten other seeds is given where a comparison depends on it. Cases 9 and 10 compare two loop measurements on the same coils by pressing Reset, which restarts the case from $t = 0$ and keeps every setting the learner changed. Figure 10.4 draws the rows of Table 10.5 as a chart.



Figure 10.3 Case 7, the Smith predictor with its delay in meters, at $t = 15$ min after a slowdown from 150 to 90 m/min at 5 min, cropped to the badges, the plant view and the trends. The badges (1) read 76.9 g/m^2 at the gauge, 75.1 g/m^2 at the knives, 14.9 kPa, 90 m/min and a 67 s delay. The air knives (2) sit just above the zinc, the strip path over the cooling tower (3) carries the 100 m that make the delay, and the gauge (4) reads the coating at its end. The histogram of the last 10 min sits well above the 70 g/m^2 minimum and around the 76 g/m^2 target, and the line under it reads 1.17 g/m^2 and $336 an hour of zinc above the minimum. The trends over 15 min show the coating at the knives between 73.4 and 79.0 g/m^2 through the slowdown, the knife pressure falling from about 26 to 15 kPa as the speed ramps down and the delay rises from 40 to 67 s, the unmeasured knife distance stepping 2.9 % closer at a coil change at 9.7 min (5), after which the coating at the knives dips to 71.6 g/m^2, and the average zinc cost.

Table 10.4 Scenarios and What Each Teaches

#	Scenario	Lesson
1	A fixed knife pressure (manual)	open loop, coils and speed
2	A PI loop on the gauge	feedback on a late signal
3	More gain, more variance	the best gain is finite
4	The line slows down	delay and gain change together
5	Speed feedforward	a measured disturbance
6	Smith predictor, delay in s	a delay in the wrong unit
7	Smith predictor, delay in m	a delay in the right unit
8	What a gram is worth	the target and the spread
9	A strip that flutters	variance control cannot remove
10	A heavier coating grade	a gain that grows 5.6 times
11	Challenge, one hour	grade and two speed changes

Table 10.5 Measured Scenarios. σ_W and Range in g/m^2, Below Is the Strip Under W_{min}, Cost Is Zinc Above W_{min} in \$/h

#	Configuration and window	σ_W	Range	Below	Cost
1	manual, 1–20 min	3.45	73.6–83.9	0.0 %	785
1	manual, 1–30 min	10.32	53.7–83.9	22.8 %	531
2	PI, 5–60 min	1.43	70.6–82.8	0.0 %	482
3	PI, K_c 0.25, τ_I 10	4.49	62.9–94.4	7.8 %	504
4	PI, slowdown, 5–40 min	5.08	56.9–89.0	11.6 %	311
5	PI and feedforward	1.25	72.5–78.9	0.0 %	310
6	Smith (s), 1–5 min	0.89	73.2–77.5	0.0 %	—
6	Smith (s), 5–40 min	53.35	28.7–218.2	40.1 %	1312
7	Smith (m), 5–40 min	1.24	71.5–79.2	0.0 %	308
8	PI, W_{SP} 73.5	1.41	68.3–80.1	2.0 %	281
9	PI, flutter 3 %	2.66	67.6–83.3	1.9 %	484
9	Smith (m), flutter 3 %	2.61	67.7–83.5	1.8 %	485
10	PI, heavy, 10–30 min	114.87	43.6–389.5	64.6 %	2915
10	Smith (m), heavy	2.94	75.6[a]–150.1	0.65 %	572

[a] the minimum is the coating at the instant of the grade change

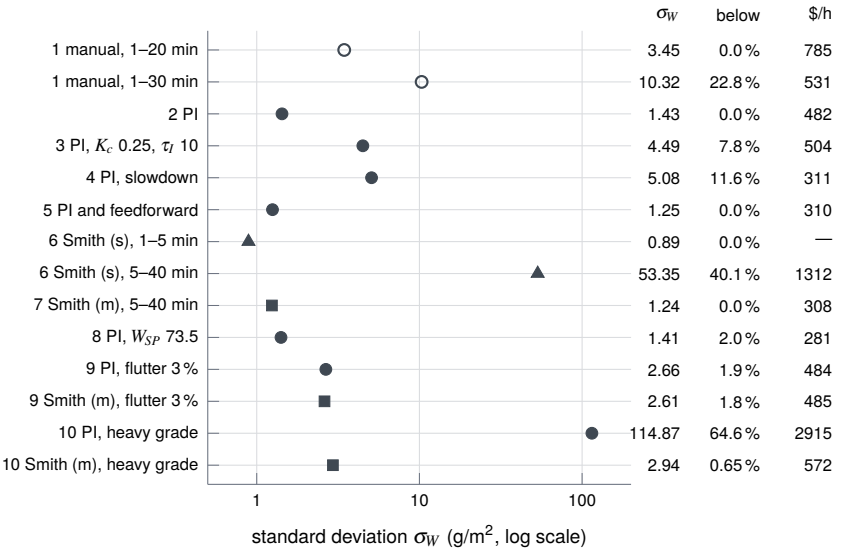


Figure 10.4 The fourteen rows of Table 10.5 as a chart: the standard deviation of the coating on a log scale, with its value, the strip below the minimum and the zinc cost above the minimum in the three columns at the right. Open circles mark the manual runs, filled circles the PI loop on the gauge, triangles the predictor with its delay in seconds and squares the predictor with its delay in meters. Notice that the two points near the right edge, the predictor in seconds after the slowdown and the PI loop on the heavy grade, sit more than a decade to the right of the predictor in meters facing the same slowdown and the same grade.

10.6.1 A Fixed Pressure and a Loop on the Gauge

Held at 26.2 kPa, the coating wanders between 73.6 and 83.9 g/m^2 with the coil shapes over the first 20 min, averaging 79.9 g/m^2 because the first coil change, at 8 min, moves the strip 0.93 mm farther from the knives, and zinc above the minimum costs \$785 an hour. When the line slows to 90 m/min at 20 min, less zinc is dragged out, the coating falls to between 54 and 57 g/m^2, and 22.8 % of the strip

in the half hour is below the minimum. The PI loop on the gauge holds 76.0 g/m^2 with a standard deviation of 1.43 g/m^2 and none of the strip below the minimum, at $482 an hour.

> **TRY IT IN THE APP** **1 · A Fixed Knife Pressure (MANUAL)**
>
> Load scenario 1 and press Run. Watch the coating at the knives on the top trend. For the first 20 min it wanders between 73.6 and 83.9 g/m^2 as the coils change. When the line slows to 90 m/min at 20 min, less zinc is dragged out, and the coating falls to between 54 and 57 g/m^2, below the 70 g/m^2 minimum.

The residual is not mainly flutter. With the flutter set to zero the standard deviation falls only to 1.25 g/m^2, while with the coil shape changes set to zero it falls to 0.87 g/m^2, and removing the gauge noise gives 1.40 g/m^2. Most of the variance left by the loop is the transient after each coil step, which the loop cannot see for 40 s and then removes over a couple of minutes. The default gain is near the best one for this delay. At $\tau_I = 15$ s the standard deviation is 1.68, 1.52, 1.43, 1.40, 1.40 and 1.42 g/m^2 for gains of 0.06, 0.09, 0.12, 0.15, 0.18 and 0.20. In case 3 the loop is at the edge of stability, as Table 10.3 predicts, and rings in bursts after some coil changes. The standard deviation rises to 4.49 g/m^2 and 7.8 % of the strip falls below the minimum. With the coil changes switched off the same tuning still gives 3.30 g/m^2, so the ringing is sustained by the flutter and the noise as well.

Pause. Try Exercise 10.1 before reading on.

10.6.2 A Slowdown, Feedforward and Two Predictors

At 5 min in cases 4 to 7 the line slows to 90 m/min. With the PI loop alone the coating first drops to about 57 g/m^2, because the speed step reaches the knives at once and the gauge 67 s later, and the loop then rings with every coil change. Over 40 min the standard deviation is 5.08 g/m^2 and 11.6 % of the strip is below the minimum. With the speed set to 90 m/min before Run in case 2, so that the line runs at 90 m/min from the start, the same loop still gives 3.36 g/m^2 and 3.9 % below from 5 to 60 min, more than twice the 1.43 g/m^2 at 150 m/min, so the ringing belongs to the operating point and not to the transition. Speed feedforward keeps the coating above 74.6 g/m^2 through the slowdown and lowers the standard deviation to 1.25 g/m^2 with none of the strip below the minimum, no more than the 1.44 g/m^2 of case 2 at 150 m/min over the same window. At 90 m/min from the start feedforward lowers the standard deviation from 3.36 to 2.29 g/m^2, because it restores the nominal loop gain. Over ten other seeds case 4 gives 4.50 to 7.85 g/m^2 and case 5 gives 0.93 to 1.62 g/m^2. At 90 m/min from the start the ten seeds give 1.57 to 2.70 g/m^2 without feedforward, below the 3.36 of the shipped seed, and 1.12 to 1.76 g/m^2 with it, lower on every seed, so the benefit of feedforward does not depend on the coil shapes drawn.

The Smith predictor with its delay in seconds and the tighter tuning holds 0.89 g/m^2 over the first 5 min at 150 m/min. After the slowdown the second bracket of Equation (10.13) is no longer zero, and the loop oscillates with a period of about 45 s and a growing amplitude until the pressure command reaches its upper limit at 1065 s. The coating swings between 28.7 and 218 g/m^2 and 40.1 % of the strip is below the minimum. The same predictor keeping its delay as 100 m of strip stays stable with the same tuning, at 1.24 g/m^2 with none of the strip below the minimum. The wiping model includes the speed, so this predictor sees the slowdown at once, as the feedforward does, and adding feedforward to it changes the result by little (1.20 g/m^2).

> **TRY IT IN THE APP** **6 · A Smith Predictor With Its Delay in Seconds**
>
> Load scenario 6 and run past the slowdown at 5 min. The loop oscillates with a period of about 45 s and a growing amplitude until the pressure command reaches its upper limit at 1065 s, and the coating swings between 28.7 and 218 g/m^2. Then load scenario 7, the same predictor with its delay kept as 100 m of strip. It stays stable at 1.24 g/m^2.

The time-based predictor was also run at constant speeds, as case 2 with the predictor and tuning of case 6 and the line speed set before Run, and measured from 10 min to 60 min. It is stable at 122.5, 125, 130, 135, 140, 160, 180, 200 and 210 m/min, with standard deviations of 1.08 to 1.79 g/m^2. At 120 m/min, where the true delay is 50 s against the assumed 40 s, the standard deviation grows from 1.91 g/m^2 between 10 and 30 min to 10.22 g/m^2 between 40 and 60 min, and at 105 m/min it reaches 68.2 g/m^2. A delay 9 s too short is tolerated and one 10 s too short is not, which the linearized loop of the predictor confirms, while a delay 11 s too long, at 210 m/min, is still stable. Robustness studies of dead-time compensators address exactly this sensitivity [1, 2].

Case 7 ten minutes after the slowdown is shown in Figure 10.3. No strip falls below the minimum over the 40 min, although the coil change at 9.7 min brings the strip 2.9 % closer to the knives and the coating dips to 71.6 g/m^2 before the loop sees it.

10.6.3 The Price of the Spread

Lowering the target of case 2 from 76 to 73.5 g/m^2 cuts the zinc cost from \$482 to \$281 an hour, about \$200 an hour or \$1.75 million over a year of continuous running. The strip below the minimum rises from none to 2.0 %, although the scorecard's normal-theory target for this spread is 73.2 g/m^2. At that target 3.0 % of the strip is below the minimum. A sweep of the target gives 1.39, 1.18, 1.04, 0.92 and 0.45 % below at 74.0, 74.5, 74.7, 74.8 and 75.0 g/m^2, so 1 % below needs a margin of about 4.75 g/m^2, 3.3 standard deviations rather than 2.33. The coil steps give the coating a heavy lower tail, and a target set from the standard deviation alone understates the margin. Each gram per square meter of standard deviation still costs at least $2.33 \times \$81 = \$189$ an hour in target margin at 150 m/min.

> **DID YOU KNOW?** **Twice the Gain, Three Times the Spread**
>
> From the simulator: raising the loop gain from 0.12 to 0.25 and shortening the integral time from 15 to 10 s widens the spread of the coating from 1.43 to 4.49 g/m^2, and the strip below the minimum rises from none to 7.8 % (Table 10.5). On a late gauge a tighter loop gives a wider product.

Tripling the flutter to 3 % raises the standard deviation to 2.66 g/m^2 with the PI loop and 2.61 g/m^2 with the tightly tuned predictor in meters, and the target this spread allows to about 76 g/m^2. The predictor does no better, and over ten other seeds the two differ by at most 0.3 g/m^2, because the gauge still reports the flutter 40 s late. Variance of this kind has to be removed at the strip, not in the controller.

Pause. Try Exercise 10.2 before reading on.

10.6.4 A Heavier Grade and the Challenge

At 10 min in case 10 the minimum rises to 137.5 g/m^2 and the target to 145 g/m^2. The knives drop toward 9 kPa, the gain rises 5.6 times, and the PI loop tuned for the light grade goes into a limit cycle between the pressure limits, with the coating between 43.6 and 389 g/m^2 and 64.6 % of the strip below the new minimum. The predictor in meters carries the gain in its model, and its inner loop sees the model rather than the delayed gauge, so the same tuning settles on the new target at 2.94 g/m^2 with 0.65 % of the strip below the minimum, all of it in the first 8 s after the change, while the coating climbs past the new minimum.

Pause. Try Exercise 10.3 before reading on.

Table 10.6 Challenge: Heavy Grade at 10 min, 90 m/min at 30 min, 180 m/min at 45 min. Strip Below W_{min} Over the Hour

Configuration	Below
as loaded, PI 0.12, 15 s, no feedforward	50.5 %
PI 0.03, 30 s, no feedforward	14.9 %
PI 0.03, 30 s, feedforward	12.2 %
Smith (s) 0.3, 4 s, feedforward	24.2 %
Smith (m) 0.3, 4 s, feedforward (suggested)	0.22 %

The challenge runs one hour with the grade change of case 10, a slowdown to 90 m/min at 30 min and a speed-up to 180 m/min at 45 min. It is met with at most 2 % of the strip below the minimum over the hour and the gauge within $\pm 20\,\mathrm{g/m^2}$ of the target over the last 20 min. The measured configurations are listed in Table 10.6. As loaded the loop goes unstable at the grade change and 50.5 % of the strip is below the minimum. The app's verdict reads 50.45 % and a worst gauge error over the last 20 min of 260.9 g/m^2, so both clauses fail. Detuning the PI loop four times keeps it stable on the heavy grade, but the coating then takes minutes to climb to the new minimum, and 34.7 % of the strip between 10 and 30 min is below it. Without feedforward the slowdown adds 16.4 % of the strip between 30 and 45 min, and with it none. The predictor with its delay in seconds fails at the two speed changes. The predictor in meters with feedforward meets the target with 0.22 % below, all of it in the first 8 s after the grade change, and over the last 20 min, which include the speed-up, the coating at the knives stays between 140.4 and 152.0 g/m^2 with a standard deviation of 1.83 g/m^2, while the gauge stays within 8.0 g/m^2 of the target. The verdict therefore passes 1.78 points inside the limit on the strip below the minimum and 12.0 g/m^2 inside the band, and the app gives this one run whether the suggested settings are revealed before Run or carried through Reset. The same predictor without feedforward gives the same result, because its model already carries the speed. The grade change needs the model's gain, and the speed changes need the delay in meters. Over ten other seeds the suggested settings leave 0.25 to 2.6 % of the strip below the minimum, and two of the ten exceed 2 %, so the margin of the graded run depends on the coil shapes it draws.

10.7 Technology in Practice

The loop of Section 10.4 runs on a coating gauge, a speed measurement and the pressure of two air knives. Figure 10.5 draws it with that hardware in place of the schematic's symbols.

What Measures It

On a working line the coating gauge stands where the chapter puts it, behind the cooling tower, and the air knife settings are adjusted from the deviation between the preset and the measured coating [18]. The gauge measures the top and the bottom coating, and the controller averages each side over one scan of the head [14]. The chapter has one fixed gauge, AT-1001, and treats the two sides as one coating. The predictor reads two more signals, the line speed ST-1002 and the knife pressure PT-1003, because its wiping model needs both.

> **TECHNOLOGY** **Sensor: An X-Ray Gauge Past the Cooling Tower**
>
> An X-ray fluorescence gauge measures a coating by the fluorescent X-rays that its atoms emit when irradiated, or by how much the coating absorbs the X-rays coming back from the steel beneath. The method serves coatings up to about 20 µm thick, galvanized strip among them [19]. The coating must be solid first, so the gauge stands behind the cooling tower, 100 to 200 m of strip past the knives [18]. Its head often traverses the width, so the points it measures trace a dashed diagonal line on the moving strip [14]. The chapter's gauge AT is a fixed point 100 m past the knives with 1 g/m^2 of filtered noise.

line controller: AC-1001 with AY-1001 and SY-1002 executes every 0.2 s, the step of the physics, against a gauge delay of 40 s at 150 m/min and 67 s at 90 m/min

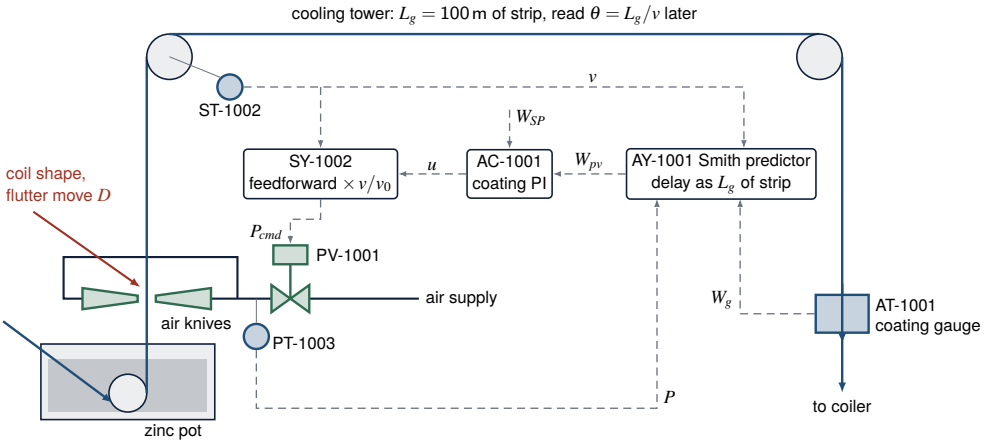


Figure 10.5 The coating loop drawn with its hardware. The air knives blow both sides of the strip just above the zinc pot, fed through the knife pressure PV-1001. The strip travels 100 m over the cooling tower to the coating gauge AT-1001, an X-ray gauge in practice. The Smith predictor AY-1001 combines the gauge with its wiping model, which reads the line speed ST-1002 and the knife pressure PT-1003, and looks the model up by strip length. The PI block AC-1001 acts on that signal, and the feedforward SY-1002 scales its output by the speed. Notice how far apart the knives that act and the gauge that sees are along the strip.

What Moves It

Real lines blow each side of the strip with its own knife. The controller computes a pressure for the top and the bottom knife, and it corrects a difference between the two sides by moving the knife distances [14]. The chapter has one pressure for both knives, PV-1001, and a knife distance that only the disturbances move.

> **TECHNOLOGY** **Final Element: The Air Knife**
>
> An air knife is a slot nozzle that blows a flat gas jet at the strip just above the bath. The pressure gradient and the shear stress of the jet on the liquid zinc set the final film thickness, and most of the zinc runs back into the pot [20]. The handle is the pressure in the knife chamber [14]. It cannot be raised without end. A harder jet brings tonal noise, splashing and an uneven coating, which is why lines hold their speed down instead [20], and splash leaves first from the strip edges [21]. A pressure difference between the two knives also shakes the strip [22]. The chapter's knives follow the pressure command through a 1 s lag between 2 and 70 kPa.

Where the Controller Runs

A line of this kind is run from a layered control system. A supervisor sends set points to distributed controllers, which drive the actuators from sensor feedback, and loops at both levels run with cycle times from milliseconds to minutes [23]. On a galvanizing line the coating control has a preset part that acts when a weld passes the knives and a feedback part that trims along the coil [14]. The chapter keeps only the feedback part. Its PI block AC-1001 executes every 0.2 s, the step of the physics and far shorter than the 40 s delay at 150 m/min. The predictor AY-1001 and the speed feedforward SY-1002 belong to the same block, and the coating target is the set point that a supervisor sends to it.

IN PRACTICE **How a Coating Loop Is Run on a Line**

On a working line the coating is set in two ways. When the weld between two coils passes the knives, a new pressure and knife distance are preset from a model that learns from the coating measured on earlier coils, because no feedback can reach the head of a coil in time. Feedback then trims through a Smith predictor whose delay is taken from the line speed and the distance from the knives to the gauge. It ignores errors within $\pm 1\,\mathrm{g/m^2}$, a dead band, and a feedforward meets the changes of line speed that the annealing furnace ahead of the pot calls for [14].

What Is Modernizing It

Three developments touch the loop of this chapter. Electromagnetic actuators now act on the strip without touching it. One system of five pairs of actuators and controllers, designed on a reduced finite element model of the strip, was shown effective for vibration and flatness control in simulation and in experiments under various tensions [22]. Data-driven models predict the coating from production data. Models trained for three bath chemistries found the knife pressure, the strip speed and the knife distance to be the main factors, and the sensitivity of the coating to each of them differed with the coating type [18], which is the moving gain of Equation (10.2) seen from the data. The knives themselves are changing too. A multi-slot knife with a weak auxiliary jet beside the main jet wiped a thinner coating than a single slot under similar main jet conditions in laboratory tests [20].

10.8 Limitations

The wiping correlation is a power law with fixed exponents, and it ignores the nozzle geometry, the jet tilt, the zinc temperature and the splashing limit [7].

OPEN CHALLENGE **Thinner Coatings Without Splashing**

Zinc resources are limited, and saving zinc means thinner coatings. Gas wiping meets two limits: the liquid film splashes as the jet pressure rises, and the wiping capacity of the equipment runs out. Hot-dip sheet with less than $30\,\mathrm{g/m^2}$ is difficult to make as a standard product, and zinc that begins to solidify under the cooling jet appears to resist wiping [21]. For now lines keep the strip speed low to stay clear of these troubles [20]. Progress would be a knife that wipes thinner at full speed, and a loop that holds its target closer to the minimum.

The two sides and the full width are one coating, so crossbow, edge build-up and side-to-side control are absent. The controller's model shares every parameter of the plant except the knife distance. A model with the wrong reference coating or exponent would add a steady model error that the gauge corrects slowly, which none of the scenarios tests.

The knife distance is only a disturbance. Real lines also move the knives, and they stabilize the strip mechanically or electromagnetically, which is the remedy case 9 points to.

OPEN CHALLENGE A Strip That Will Not Hold Still

Curvature and vibration of the strip at the knives make the coating uneven. More tension calms the strip but wears out the bushes and sleeves of the sink roll, and rolls that squeeze the strip can scratch it [22]. Active control is hard. The strip vibrates in many modes at once, only a non-contact position sensor can be used on a moving strip, and a controller built on the full state risks spill-over into modes it does not model [22]. Vibration also grows harder to reduce as the strip speeds up [24]. Progress would be a strip held still at full speed, which would remove the variance that case 9 shows the coating loop cannot.

The gauge is a fixed point sensor, without the scan period of a traversing gauge.

10.9 Conclusion

The galvanizing line combines a dead time fixed in meters, a process gain that moves with the coating and the speed, and a product specification that turns variance into zinc. A PI loop on the gauge holds 1.43 g/m^2, most of it the transient after each coil step, and doubling its gain triples the spread. A slowdown to 90 m/min raises the gain and the delay together and leaves 11.6 % of the strip below the minimum, and speed feedforward removes that loss by restoring the nominal gain. A Smith predictor that keeps its delay in seconds oscillates once the true delay exceeds its own by about 10 s, while the same predictor keeping the delay as 100 m of strip stays at 1.24 g/m^2 and also carries the 5.6-fold gain of a heavy grade. Lowering the target by 2.5 g/m^2 saves about $200 an hour, but the lopsided spread needs a margin of 3.3 standard deviations for 1 % of the strip below the minimum. The model treats both sides and the width as one coating with fixed wiping exponents. A fuller model would add two-sided knives, knife position as a second actuator and a traversing gauge. The simulator offers an accessible way to practice the control of a process whose dead time is set by a distance.

KEY POINTS

- The coating can be measured only after it has frozen, a fixed length of strip past the knives, so the gauge reads it late, and later still when the line slows.
- A PI loop on the gauge holds 1.43 g/m^2, and doubling its gain triples the spread.
- A slowdown raises the gain and the delay together, and speed feedforward removes the loss by restoring the nominal gain.
- A Smith predictor that keeps its delay in seconds oscillates once the true delay exceeds its own by about 10 s, while the same predictor keeping the delay as 100 m of strip stays stable.

Where else this applies. A delay set by a distance and a throughput also governs the glass ribbon of Volume 5, Chapter 5 and the strip gauge of Chapter 8.

References

[1] J. E. Normey-Rico, T. L. M. Santos, R. C. C. Flesch, and B. C. Torrico, “Control of dead-time process: From the Smith predictor to general multi-input multi-output dead-time compensators,” *Frontiers in Control Engineering*, vol. 3, p. 953768, 2022.

[2] Z. Palmor, “Stability properties of Smith dead-time compensator controllers,” *International Journal of Control*, vol. 32, no. 6, pp. 937–949, 1980.

[3] T. J. Harris, “Assessment of control loop performance,” *The Canadian Journal of Chemical Engineering*, vol. 67, no. 5, pp. 856–861, 1989.

[4] L. Desborough and T. Harris, “Performance assessment measures for univariate feedforward/feedback control,” *The Canadian Journal of Chemical Engineering*, vol. 71, no. 4, pp. 605–616, 1993.

[5] A. Marder, “The metallurgy of zinc-coated steel,” *Progress in Materials Science*, vol. 45, no. 3, pp. 191–271, 2000.

[6] J. A. Thornton and H. F. Graff, “An analytical description of the jet finishing process for hot-dip metallic coatings on strip,” *Metallurgical Transactions B*, vol. 7, no. 4, pp. 607–618, 1976.

[7] A. Gosset and J.-M. Buchlin, "Jet Wiping in Hot-Dip Galvanization," *Journal of Fluids Engineering*, vol. 129, no. 4, pp. 466–475, 2007.

[8] M. Bauer and I. K. Craig, "Economic assessment of advanced process control – A survey and framework," *Journal of Process Control*, vol. 18, no. 1, pp. 2–18, 2008.

[9] U.S. Geological Survey, "Mineral commodity summaries 2025: Zinc," U.S. Geological Survey, Reston, VA, Tech. Rep., 2025.

[10] A. R. Marder and F. E. Goodwin, "History of zinc-coated steel," in *The Metallurgy of Zinc Coated Steels*. Elsevier, 2023, pp. 1–14.

[11] C. H. Ellen and C. V. Tu, "An Analysis of Jet Stripping of Liquid Coatings," *Journal of Fluids Engineering*, vol. 106, no. 4, pp. 399–404, 1984.

[12] T. Noe and T. Schuerger, "Computer Control of Zinc-Coating Weight on the Fairfield No. 4 Galvanizing Line," *IFAC Proceedings Volumes*, vol. 8, no. 1, pp. 404–411, 1975.

[13] D. L. Grinberg, "Regulating coating thickness in continuous galvanizing of steel sheet," *Metallurgist*, vol. 21, no. 3, pp. 185–188, 1977.

[14] K.-T. Shin, H.-D. Park, and W.-K. Chung, "Synthesis method for the modelling and robust control of coating weight at galvanizing process," *ISIJ International*, vol. 46, no. 10, pp. 1442–1451, 2006.

[15] E. Elsaadawy, G. Hanumanth, A. Balthazaar, J. McDermid, A. Hrymak, and J. Forbes, "Coating Weight Model for the Continuous Hot-Dip Galvanizing Process," *Metallurgical and Materials Transactions B*, vol. 38, no. 3, pp. 413–424, 2007.

[16] Y.-T. Kim, "An automatic coating weight control for continuous galvanizing line," in *2008 International Conference on Control, Automation and Systems*, 2008, pp. 1715–1720.

[17] N. Guelton and A. Lerouge, "Coating weight control on ArcelorMittal's galvanizing line at Florange Works," *Control Engineering Practice*, vol. 18, no. 10, pp. 1220–1229, 2010.

[18] H. Wang, K. Zha, K. Yang, W. Zhang, S. Zhu, F. Xu, and Q. Xie, "Data-driven prediction and process parameter analysis of zinc coating thickness in hot-dip galvanizing," *ISIJ International*, vol. 66, no. 7, pp. 857–866, 2026.

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

[20] A. Yahyaee Soufiani, J. R. McDermid, A. N. Hrymak, and F. E. Goodwin, "A parametric study of a multi-slot air knife for coating thickness reduction," *ISIJ International*, vol. 60, no. 5, pp. 1040–1051, 2020.

[21] H. Kobayashi, G. Takeda, K. Katoh, and T. Wakimoto, "Coating weight reduction technology in gas wiping of hot-dip galvanizing on steel strip," *ISIJ International*, vol. 64, no. 11, pp. 1716–1722, 2024.

[22] W. Jung, Y. J. Jang, S. Lim, and S. Won, "Active vibration and flatness control of a strip in a continuous galvanizing line using positive position feedback control," *ISIJ International*, vol. 53, no. 5, pp. 854–865, 2013.

[23] 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.

[24] J. Li, Y.-H. Yan, X.-H. Guo, and Y.-Q. Wang, "Research on vibration control method of steel strip for a continuous hot-dip galvanizing line," *ISIJ International*, vol. 52, no. 6, pp. 1072–1079, 2012.

10.10 Exercises

10.1 ⌕ Concept check · 2 min

In scenario 3, More gain, more variance, the PI loop on the gauge runs with $K_c = 0.25$ and $\tau_I = 10\,\text{s}$ instead of 0.12 and 15 s. Measured from 5 to 60 min, the standard deviation of the coating rises from 1.43 to 4.49 g/m^2. Which explanation is right?

(a) The knife pressure now reaches its limits after each coil change, so the clipped loop cannot finish its corrections

(b) The coating becomes less sensitive to pressure as the loop raises it, so the loop is detuned just when it is needed

(c) The gain margin falls from 3.05 to 1.02, so a loop that sees each correction 40 s late rings near its stability limit

(d) The loop now follows the gauge noise with its 2 s correlation time, and the knives copy that noise onto the strip

10.2 Concept check · 1 min

In scenario 9, A strip that flutters, the flutter is tripled to 3 %. The PI loop leaves 2.66 g/m^2 and the tightly tuned Smith predictor in meters 2.61 g/m^2. Why does the predictor do no better?

(a) Its tighter tuning makes the loop ring, which cancels what it gains on the flutter

(b) The flutter is too small against the coil steps to change the spread of either loop

(c) The flutter changes the line speed, which the predictor handles no better than the gauge

(d) Its model runs at the nominal distance, so the flutter reaches it only through the late gauge

10.3 Concept check · 2 min

In scenario 10, A heavier coating grade, the target rises from 76 to 145 g/m^2 at 10 min. The PI loop tuned for the light grade goes into a limit cycle, while the Smith predictor in meters, with a higher gain of 0.3, settles on the new target. What explains the difference?

(a) The process gain grows 5.6 times at the low pressure, and the predictor's inner loop sees it without the delay

(b) The heavy grade lengthens the gauge delay, which the predictor in meters tracks and the gauge loop does not

(c) The higher minimum leaves less room below the target, so the same spread puts more strip below the minimum

(d) The heavy grade needs a pressure below the knives' lower limit, which only a predictor can work around

10.4 Theory and calculation · 6 min

Use the wiping correlation of Equation (10.1) with $W_0 = 70\,\mathrm{g/m^2}$, $v_0 = 150\,\mathrm{m/min}$, $P_0 = 30\,\mathrm{kPa}$, $D = D_0$ and the gauge 100 m after the knives.

(a) Find the pressure that gives 76 g/m^2 at 150 m/min, and the process gain K_p there. *Checkpoint:* 26.16 kPa and −1.74 g/m^2 per kPa.

(b) Repeat at 90 m/min, and give the gauge delay at both speeds. *Checkpoint:* 15.7 kPa, −2.91 g/m^2 per kPa, 40 s and 66.7 s.

(c) Repeat for 145 g/m^2 at 150 m/min, and show that the ratio of the two gains at 150 m/min is $(145/76)^{1+1/0.6}$. *Checkpoint:* 8.91 kPa, −9.76 g/m^2 per kPa, ratio 5.6.

10.5 Theory and calculation · 6 min

At 90 m/min the linearized loop of Section 10.5 has $G(s) = K_p \mathrm{e}^{-\theta s}/(\tau_P s + 1)$ with $\theta = 66.7\,\mathrm{s}$ and $\tau_P = 1\,\mathrm{s}$, under the default PI tuning $K_c = 0.12$ and $\tau_I = 15\,\mathrm{s}$. Table 10.3 gives the period at phase crossover as 214 s.

(a) Check that the open-loop phase is -180° at $\omega = 2\pi/214$, adding the phases of the integral term, the delay and the lag. *Checkpoint:* $66.2^\circ + 112.1^\circ + 1.7^\circ = 180.0^\circ$.

(b) Find the gain margin with $|K_p| = 2.91$, the gauge loop of case 4. *Checkpoint:* 1.16.

(c) Find the gain margin with speed feedforward, which makes the gain seen by the controller 1.74 g/m^2 per kPa, and say why the period does not change. *Checkpoint:* 1.93.

10.6 Theory and calculation · 6 min

The Smith predictor in seconds looks its model up a fixed $\theta_0 = 40$ s back, while the gauge reads strip $L_g = 100$ m back.

(a) Find the true delay and the mismatch $\theta - \theta_0$ at 120, 130 and 210 m/min. *Checkpoint:* 50 s and +10 s, 46.2 s and +6.2 s, 28.6 s and −11.4 s.

(b) The chapter reports that a delay 9 s too short is tolerated and one 10 s too short is not. Find the speed at which the true delay is $\theta_0 + 9$ s. *Checkpoint:* 122.4 m/min.

(c) At a constant 120 m/min the model coating $\hat{W}$ steps by $\Delta\hat{W}$ at $t = 0$ and the model is otherwise exact. Using Equation (10.13), give the second bracket as a function of time. *Checkpoint:* $-\Delta\hat{W}$ from 40 to 50 s, zero otherwise.

10.7 App challenge · 8 min

Use the default seed and speed 120, and change no setting. Near the end of each run lower the speed to 5 and press Pause when the clock reads 60 min.

(a) Load scenario 2, A PI loop on the gauge, run it, and read the Coating, mean $\pm$ σ, Below the minimum and Zinc above the minimum tiles. *Checkpoint:* 76.0 ± 1.39 g/m^2, 0.00 %, \$482 an hour.

(b) Load scenario 3, More gain, more variance, and read the same tiles. *Checkpoint:* 76.0 ± 4.3 g/m^2, about 7.1 %, about \$500 an hour.

Read the standard deviation within 0.2 g/m^2 and the strip below within 0.2 points. Explain why the zinc cost hardly changes while the spread triples, and what the spread would cost if the target were set from it.

10.8 App challenge · 8 min

Use the default seed and speed 120, change no setting, and pause each run when the clock reads 40 min. Read Coating, mean $\pm$ σ, Below the minimum and Gauge delay.

(a) Scenario 4, The line slows down. *Checkpoint:* 75.4 ± 4.58 g/m^2, 9.3 % to 9.4 %, 67 s.

(b) Scenario 5, Speed feedforward. *Checkpoint:* 75.9 ± 1.18 g/m^2, 0.00 %, 67 s.

(c) Scenario 7, A Smith predictor with its delay in meters. *Checkpoint:* 75.8 ± 1.18 g/m^2, 0.00 %, 67 s.

Read the standard deviations within 0.05 g/m^2. Explain why two different remedies give the same result after the slowdown.

10.9 App challenge · 8 min

Load scenario 2, A PI loop on the gauge. Before Run, set Loop measurement to Smith, delay in s, Coating loop gain to 0.30, Coating loop integral time to 4 s and Line speed to 120 m/min. Run at speed 120 with the default seed and pause when the clock reads 60 min. Then press Reset, set Line speed to 130 m/min, keep the other settings and run the same way.

(a) Read Coating, mean $\pm$ σ, Below the minimum and Gauge delay at 120 m/min, the standard deviation within 0.2 g/m^2. *Checkpoint:* 76.0 ± 6.2 g/m^2, about 12.5 %, 50 s.

(b) Read the same tiles at 130 m/min, the standard deviation within 0.05 g/m^2. *Checkpoint:* 76.0 ± 1.12 g/m^2, 0.00 %, 46 s.

Using the mismatch of each run, explain why 10 m/min decides between a stable and an unstable loop.

10.10 Application area · 4 min

A line holds the coating with the PI loop of case 2 at 150 m/min and 1.43 g/m^2. The chapter measures 1.25 g/m^2 with the flutter set to zero and 0.87 g/m^2 with the coil shape changes set to zero, and 2.66 g/m^2 when the flutter is tripled to 3 %. Each gram per square meter of standard deviation costs at least $189 an hour in target margin.

(a) An electromagnetic stabilizer would remove the flutter, and better coil shape or knife position control would remove the coil steps. Estimate the margin each saves on the line as it is. *Checkpoint:* About $34 and $106 an hour.

(b) Estimate what the stabilizer saves on a line whose strip flutters at 3 %, if it brings the flutter back to 1 %. *Checkpoint:* About $232 an hour.

Say why these estimates are lower bounds for the coil steps in particular.

10.11 Application area · 3 min

The coating gauge of the simulated line is 100 m of strip after the knives, beyond the cooling tower.

(a) Say why the gauge cannot simply be placed at the knives. *Checkpoint:* No numeric answer.

(b) A redesigned cooling section would allow a gauge 60 m after the knives. Find its delay at 150 and 90 m/min. *Checkpoint:* 24 s and 40 s.

Say what the shorter delay would buy and what it would not change.

10.12 Review the agent · 5 min

An AI assistant is asked to pass the challenge, A heavy grade at three speeds, and replies as follows. “Keep the gauge as the loop measurement, because a predictor adds a model that can be wrong. Table 10.3 gives the default PI loop a gain margin of 0.55 on the heavy grade, so cut the gain four times, to 0.03 kPa per g/m^2, with an integral time of 30 s. That restores a margin of about 3. Turn speed feedforward on, so the slowdown and the speed-up do not move the loop gain. Both clauses will pass.” Find the flaw in the proposal and say how it could be checked.