INTERACTIVE CONTROL ENGINEERING
Principles and Applications

Volume 16 Control in Iron, Steel, and Metals Production

CHAPTER 1

Introduction: From Slow Furnaces to Fast Strip

Pages 1–22

Edited by John D. Hedengren

How to cite this chapter. "Introduction: From Slow Furnaces to Fast Strip," 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. 1, pp. 1–22.

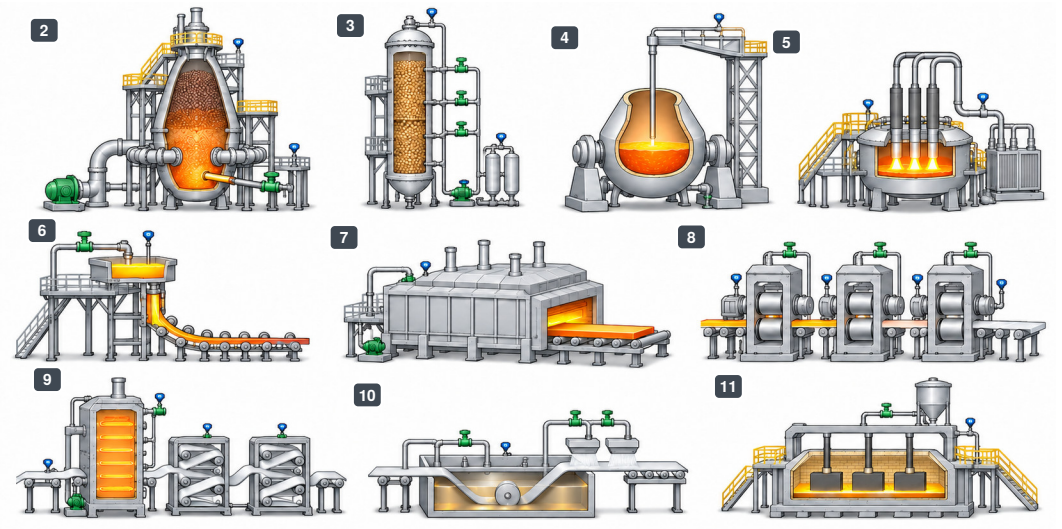
The page numbers in this file are those of the printed volume.

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

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

CHAPTER 1

Introduction: From Slow Furnaces to Fast Strip



The chapters 2 blast furnace 3 hydrogen shaft furnace 4 oxygen converter 5 arc furnace 6 caster mold 7 reheat furnace 8 hot strip mill 9 annealing line 10 galvanizing line 11 aluminum smelting cell

Notice that the top row makes iron and steel, the middle row casts, reheats and rolls it, and the bottom row anneals and coats the strip and smelts aluminum. Illustration, not to scale.

IN PLAIN TERMS **Slow Furnaces Feeding Fast Strip**

The main idea. A steel works turns ore into iron, refines the iron into steel and rolls the steel into thin sheet. Some of its furnaces take hours to answer a change of fuel, while the rolling and coating lines move sheet past each machine in seconds.

The challenge. The slow furnaces must still deliver what the fast lines need. The things that matter most, such as the heat inside a slab or the thickness of the sheet, are seldom measured where the work is done.

Why it matters. Good control makes metal of the right heat, size and coating without wasting fuel, zinc or the metal itself.

THREE WAYS TO READ THIS CHAPTER

Explore	7 min	This page, Section 1.2 with its timeline and two History boxes, Figure 1.2, the Did you know? boxes in Sections 1.4.1, 1.4.3 and 1.4.6, and Key points in Section 1.8. No equations.
Study	45 min	The whole chapter in order, working each number of the example W-1, then Exercises 1 to 10.
Practice	5 min	This page, Section 1.4.7 with Table 1.1, Section 1.5 with Figures 1.7 and 1.10, and Section 1.8.

1.1 The Question This Volume Answers

A steel works turns ore into coated strip through a chain of units that run on very different clocks. A blast furnace holds its burden for hours and answers a change of coke five hours later. A converter refines a heat in about twenty minutes and must stop it once, at the right carbon and temperature. A reheat furnace keeps a slab for two and a half hours, and the mill that receives it corrects its gauge every few milliseconds. Annealing and coating lines move strip at a few meters a second through furnaces that need minutes to change their temperature. This volume answers one question: how are

slow furnaces and fast strip controlled so that each unit delivers what the next one needs, and which techniques recur from the furnace to the coil?

Seven control problems define the industry, and every chapter meets several of them.

- *Furnaces are slow.* A hearth, a slab or a smelting cell's frozen ledge stores so much heat that a change of firing shows in the product minutes to days later.
- *The product is seldom measured where it matters.* A hearth is dipped once a cast, a slab's mean temperature never, a strip's gauge at the stand only by inference and its coating a hundred meters downstream.
- *Steelmaking runs in batches judged at the end.* A heat is blown once, and carbon removed or iron burned cannot be put back.
- *Material moves, so delay is a distance.* A gauge reads strip that left the actuator meters ago, and the time that takes changes with every change of speed.
- *A fast, limited actuator works beside a slow, wide one.* Coal and coke, line speed and firing, and the electrode and the transformer tap each share one load.
- *Continuous lines are fed in pieces.* Heats, slabs and coils arrive one at a time, and the stops between them must be absorbed by an inventory or felt by the product.
- *Gains move and loops share nodes.* A heavier coating multiplies the gain of the coating loop, and three arc furnace electrodes share one floating neutral.

Ten chapters answer the question on ten plants, each with an app that the reader runs in a browser. Chapters 2 and 3 make iron, in a blast furnace and in a hydrogen shaft furnace. Chapters 4 and 5 make steel, by oxygen and by electric arcs. Chapter 6 casts it, Chapters 7 and 8 reheat and roll it, and Chapters 9 and 10 anneal and coat the strip. Chapter 11 adds a second metal, aluminum, made in an electrolytic cell whose heat balance is as slow as any furnace's.

A reader who has worked through the volume can do the following on a metals plant that is new to the reader.

- Estimate how fast a furnace and the product inside it answer, and say why a loop on the furnace does not hold the product.
- Divide a load between a fast, limited actuator and a slow, wide one, and check that the slow loop does not fight the fast one.
- Plan a batch from its measured charge, set the aim that a one-way late correction needs, and learn a drifting constant between batches.
- Write a transport delay as a length, track a piece of material along a line, and size an inventory against the stops it must carry.
- Say where an inferred gauge or coating is exact and where it is wrong, and which slower measurement must correct it.

This chapter develops each of these on one small example and names the chapter that uses it. PID forms and tuning are taken as known from Volume 2, feedforward, ratio and cascade from Volume 3, selectors, override and mid-ranging in general from Volume 4, dead time, Smith predictors and estimation from Volume 5, interaction and the relative gain array from Volume 6, batch and run-to-run control from Volume 8 and predictive control from Volume 9.

1.2 A Short History

The milestones of Figure 1.1 follow the plants of this volume, from the first fuel blown into a blast furnace to iron reduced by hydrogen made without fossil fuel.

Iron and steel came first. Fuel injection through the tuyeres of a blast furnace dates back to experiments with coal in France in 1840 [1], the practice whose rate Chapter 2 controls. Steel by blowing followed sixteen years later.

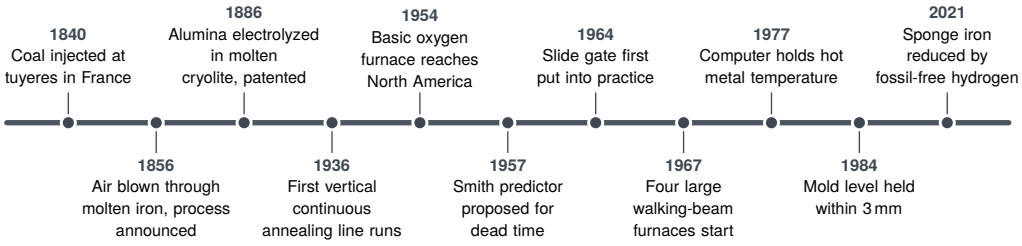


Figure 1.1 Milestones of iron, steel and aluminum making and their control. The dates are from [1] (1840), [2] (1856), [3] (1886), [4] (1936), [5] (1954), [6] (1957), [7] (1964), [8] (1967), [9] (1977), [10] (1984) and [11] (2021). The mold level of 1984 is dated by its report. Notice that each process came first and the loops that control its product decades later. Events are evenly spaced, not to a time scale.

> ⧖ HISTORY **1856 to 1969: From Air to Oxygen**
>
> In 1856 a process was announced at a meeting of the British Association that decarbonized cast iron by forcing a blast of air through the molten metal. Many metallurgists doubted it until they saw melted cast iron converted to malleable iron in a perfectly fluid state. Early trials by ironmasters failed, and only after two more years of experiments did the process turn out a product as good as that of the older methods [2]. A century later oxygen took the place of air. The basic oxygen furnace came to North America in 1954 and runs on its own heat when 70 percent or more of its charge is hot metal. By 1969 it made as much steel there as the open hearth, 60 million tons a year each [5].

Aluminum came by another route. A French patent of April 1886 covered a process based on alumina dissolved in molten cryolite, and in July of the same year a comparable patent was applied for in the United States. Apart from scale, hooding and fume recovery, the prebake cell of today keeps the basic design of the first commercial cells [3], the cell of Chapter 11.

The strip lines and their furnaces grew in the twentieth century. The first vertical continuous annealing line went into operation in 1936 at 25 m/min [4]. The sliding gate under a ladle, patented in the United States in 1885, was first put into practice in West Germany in 1964 [7], and four large walking-beam furnaces began to run at one Japanese works in June 1967 [8]. Control theory had supplied a tool for the delays of these lines in 1957, when the Smith predictor was proposed for processes with dead time [6].

> ⧖ HISTORY **Before 1960 to 1984: Measuring What the Loop Needs**
>
> Before 1960 automatic gauge control by load cell brought a gauge meter formula. The exit thickness at the head of each strip was computed from gap and force and stored, a step called lock-on, and the gap was then moved to hold the rest of the strip to that value [12]. In 1977 a blast furnace put its hot metal temperature under a computer that moved the injected oil, and 92 percent of the temperatures fell within 15 °C of the aim, against about 70 percent by hand [9]. In 1982 a level sensor for the caster mold was still called one of the most important and difficult tasks of the process [13], and in 1984 one slab caster held its level within 3 mm [10].

Iron is now being made in a new way. A pilot plant in Sweden, commissioned in August 2020, made the first sponge iron reduced with fossil-free hydrogen in June 2021 [11], the route of Chapter 3.

1.3 Notation and Conventions

Temperatures are in degrees Celsius, strip speeds in meters per second or per minute, thicknesses in millimeters, roll forces in meganewtons and coating weights in grams per square meter on each

side. A furnace or zone has a heat capacity C and a conductance G. A fast actuator u_f and a slow actuator u_s have gains K_f and K_s. In a converter C with a subscript is a carbon content in percent, λ the weight of a run-to-run update and δ the drift it follows. Strip runs at speed v, a gauge sits a length L downstream and θ is a delay. In a rolling stand d is the strip thickness, S the unloaded roll gap, F the roll force, M the mill modulus and Q the plastic modulus of the pass. The framework's h is its integration step and its a the controller action. Neither is used for anything else in this chapter, and the framework's Q in a predictive controller's objective is unrelated to the plastic modulus. Loop action follows Section 1.6.2: $a = +1$ is reverse acting, for a process whose measurement rises with the manipulated variable, and $a = -1$ is direct acting. A zone or strip temperature loop on the firing is reverse acting. A coating loop on the knife pressure is direct acting, because more pressure leaves less zinc.

The chapters keep their own symbols, and several collide. φ is the radiant exchange factor in Chapter 7, the heat-loss factor in Chapter 4 and the back-up roll angle in Chapter 8, while Chapter 6 writes the bulging phase ϕ. λ is the learning weight in Chapter 4, the laboratory bias weight in Chapter 3 and the roll pitch in Chapter 6, while Chapter 2 writes its bias weight α. Chapter 9 writes the controller action as α and Chapter 3 as σ, which is also a strip tension in Chapter 8 and the Stefan-Boltzmann constant in Chapter 7. h is the exit thickness in Chapter 8 and the slab thickness in Chapter 7, where this chapter writes d and 2ℓ. g is the pellet grade in Chapter 3, the exit factor of the strip in Chapter 9, the correction gain in Chapter 7, the feedforward multiplier in Chapter 10 and gravity in Chapter 6. M is the bath mass in Chapter 4 and the metallization in Chapter 3. K_I is the integrating gain of the mold in Chapter 6 and the current regulator's gain in Chapter 5. τ_c is the coke lag in Chapter 2 and a closed-loop time constant elsewhere, and T is the slab thickness in Chapter 6. In Chapter 11 φ is the share of a feeder shot that sinks as sludge, λ the factor on the metal pad's roll during an instability, g the slowdown of dissolution in a cold bath, h the bath-to-ledge heat transfer coefficient and M_b the bath mass.

1.4 From Furnace to Strip

1.4.1 The Flow Scheme

Figure 1.2 places each chapter's plant on the path from ore to coated strip. A blast furnace reduces sintered ore and pellets with coke and injected coal and taps liquid hot metal, which a basic oxygen converter refines to steel with scrap as coolant [14, 15]. This route and the one described next are the two main routes from ore to steel [16]. A shaft furnace can instead reduce pellets with hydrogen to solid direct reduced iron, which an electric arc furnace melts with scrap [17]. Either steel is cast continuously into slabs [18]. A reheat furnace brings the slabs back to rolling temperature, a hot strip mill rolls them to a few millimeters, and after cold rolling, which this volume does not cover, a continuous line anneals the strip and coats it with zinc. Aluminum has a chain of its own, in which electrolytic cells reduce alumina dissolved in a molten salt [19].

The techniques the chapters share are developed below on one running example, a works in miniature called W-1. Its hearth furnace makes hot metal on two fuels, its converter blows heats to a carbon and temperature aim, its reheat furnace heats 200 mm slabs for a hot mill, and one strip line anneals and galvanizes the rolled product. Every number of W-1 is declared here or computed from these declarations by the script kept with this chapter's record. The numbers are chosen for clear arithmetic and are not taken from any one chapter.

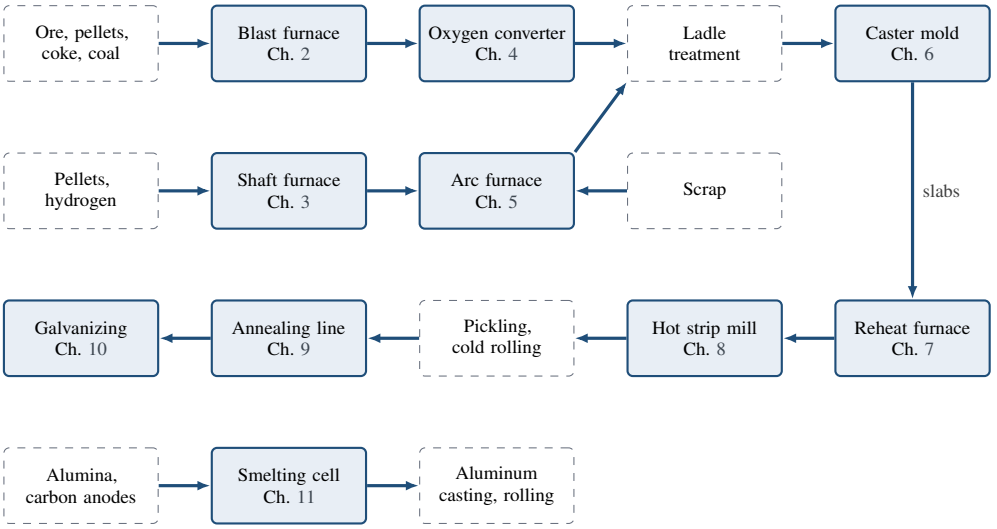


Figure 1.2 The Plants of the Volume on the Path From Ore to Coated Strip. Filled Blocks Are the Plants of Chapters 2 to 11, Dashed Blocks Lie Outside the Volume. Iron Is Made in the Top Two Rows, Steel Is Cast at the Right and Rolled, Annealed and Coated Along the Third Row, and Aluminum Has a Chain of Its Own. The Arrows Carry Hot Metal, Direct Reduced Iron (DRI), Steel, Slabs and Coils. Scrap Is Also Charged to the Converter as a Coolant (Not Drawn).

✦ **DID YOU KNOW? Two Billion Tonnes a Year**

The world made an estimated 1.9 billion metric tons of raw steel and 1.3 billion tons of pig iron in 2024 [20]. Its smelters made about 72 million tons of primary aluminum the same year [21], each ton of it in cells like the one of Chapter 11.

1.4.2 Thermal Inertia of Furnaces

A furnace zone stores heat in its walls, its gas and its load. If a zone of heat capacity C is fired at a rate Q_f and loses heat in proportion to its excess over the surroundings through a conductance G that lumps the losses and the linearized load, its temperature follows Equation (1.1).

$$C\frac{\mathrm{d}T_z}{\mathrm{d}t} = Q_f - G(T_z - T_0), \qquad \tau_z = \frac{C}{G} \tag{1.1}$$

A W-1 zone with $C = 90\,\mathrm{MJ/K}$ and $G = 0.15\,\mathrm{MW/K}$ answers a firing step with $\tau_z = 600\,\mathrm{s}$, ten minutes. A zone loop on that lag is an ordinary PI loop of Volume 2.

The slab inside the zone is slower. Heated from both faces, a slab of thickness 2ℓ conducts heat to its mid-plane on the time scale t_k of Equation (1.2), and its surface runs ahead of its center by an amount that the Biot number Bi measures.

$$t_k = \frac{\ell^2 \rho c_p}{k}, \qquad \mathrm{Bi} = \frac{h_r \ell}{k}, \qquad h_r = 4\varphi\sigma_{SB}T^3 \tag{1.2}$$

Here k is the conductivity, ρc_p the volumetric heat capacity, φ an exchange factor and h_r the radiant heat transfer coefficient at the absolute temperature T. W-1's 200 mm slab has $\ell = 0.1\,\mathrm{m}$, $k = 30\,\mathrm{W/(m\,K)}$ and $\rho c_p = 5.4\,\mathrm{MJ/(m^3\,K)}$, so $t_k = 30\,\mathrm{min}$. At 1250 °C with $\varphi = 0.6$, $h_r = 481\,\mathrm{W/(m^2\,K)}$ and $\mathrm{Bi} = 1.6$. The slab spends 120 min in three zones held at 1150, 1320 and 1235 °C. A five-cell conduction model of half the slab, radiated from its zones, discharges it with a mean of 1200 °C, a surface of 1212 °C and a core of 1194 °C, as Figure 1.3a shows. Three clocks are now in view: the zone answers in ten minutes, the slab's center in half an hour, and the product leaves after two hours.

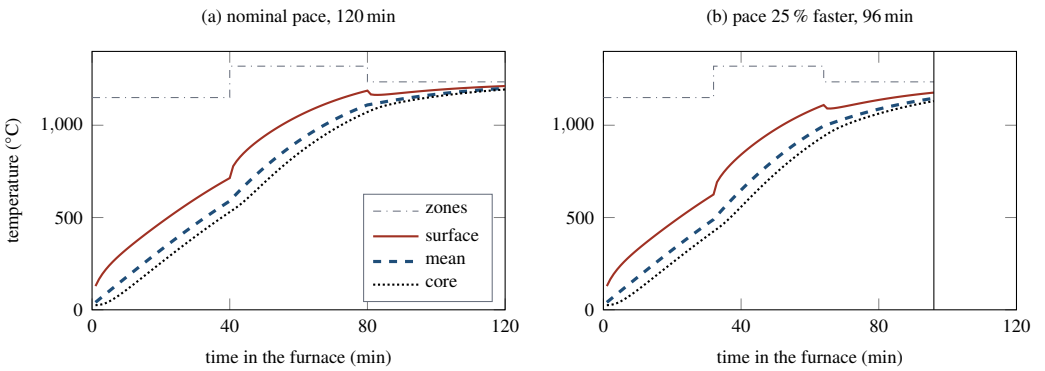


Figure 1.3 A W-1 Slab, 200 mm Thick and Charged Cold, Through Three Zones Held at Their Setpoints. The Surface (Solid), Mean (Dashed) and Core (Dotted) Against Time, With the Zone Temperatures (Dash-Dotted). (a) At the Nominal Pace the Slab Leaves With a Mean of 1200 °C. (b) At a Pace 25 % Faster, With the Zones Still on Setpoint, It Leaves at 96 min (Thin Vertical Line) With a Mean of 1146 °C.

The mill then asks for slabs 25 % faster. Each slab spends 96 min in the furnace, and with every zone loop holding its setpoint exactly the slab leaves with a mean of 1146 °C, 54 °C cold, as Figure 1.3b shows. Its surface is at 1177 °C, and the gap between surface and mean has grown from 12 to 31 °C. Nothing in the zone loops can see this, because they hold the furnace and not the product. A loop that raised every zone until the surface read its old 1212 °C would need 28 °C and would still discharge a mean of 1184 °C, because the wider gap breaks the fixed relation between surface and mean. A model that tracks the slab finds that a uniform offset of 40 °C restores the mean of 1200 °C, with a surface of 1228 °C.

Chapter 7 measures exactly this on a walking-beam furnace, whose 250 mm slabs have a conduction time of about 44 min and a Biot number of about 1.9 [22]. When the pace rises from one walk every 300 to every 240 s, zone loops that hold their setpoints deliver slabs 55 °C cold, a pyrometer trim on the surface settles the mean between 1174 and 1177 °C for every gain tried, and a slab-tracking model puts all 117 slabs on target, the approach that model-based furnace control takes [23]. The other furnaces of the volume add their own clocks. The coke of Chapter 2 reaches the hearth after five hours of burden descent and a three-hour lag. The product of Chapter 3 turns over a residence of six hours, while its top gas answers within three minutes. The furnace of Chapter 9 follows its firing with a lag of 240 s while strip crosses it in 80 s. The cell of Chapter 11 has a bath that answers in 1.30 h and a frozen ledge that integrates a heat imbalance for days, checked by one dip a day. Chapter 5 carries its bath temperature between dips forward by the metered energy.

1.4.3 Two Actuators on Two Clocks: Mid-Ranging

W-1's hearth is held at 1500 °C by two fuels. Injected fuel u_f acts with a gain $K_f = 1.0\,^{\circ}\mathrm{C}$ per kg/t of hot metal after a delay of 1 h and a lag of 1.5 h, but its injectors allow only 100 to 150 kg/t. Coke u_s acts with $K_s = 2.0\,^{\circ}\mathrm{C}$ per kg/t after a delay of 4 h and a lag of 2 h, over a wide range of 300 to 420 kg/t. A PI loop on the injected fuel, tuned by the SIMC rule of Volume 2 with a closed-loop time constant equal to the delay, has $K_c = 0.75\,\mathrm{kg/t}$ per °C and $\tau_I = 1.5\,\mathrm{h}$, and it is reverse acting. At 2 h a harder burden starts to demand 40 °C more heat through a lag of 2 h.

At steady state the two fuels must together supply the demand d, as Equation (1.3) states.

$$K_f\,\Delta u_f + K_s\,\Delta u_s = d, \qquad d \le K_f\,r_f \text{ (fast fuel alone)}, \qquad \Delta u_s \to \frac{d}{K_s} \text{ (mid-ranging)} \tag{1.3}$$

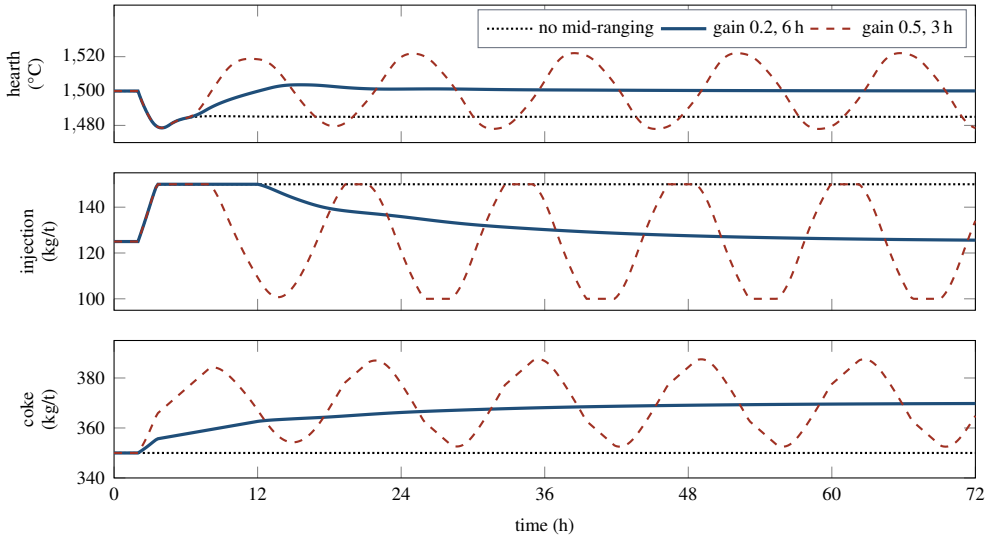


Figure 1.4 W-1's Hearth After a Burden Change That Demands 40 °C More Heat From 2 h. Hearth Temperature, Injected Fuel and Coke With the Injection Loop Alone (Dotted), With a Slow Mid-Ranging Loop on Coke, Gain 0.2 and Integral Time 6 h (Solid), and With a Faster One, Gain 0.5 and 3 h (Dashed). The Injection Saturates at 150 kg/t Without Mid-Ranging.

where r_f is the fast actuator's room above its operating point. W-1 runs its injection at 125 kg/t, so $r_f = 25\,\mathrm{kg/t}$ is worth only 25 °C. Alone, the injection loop reaches its limit at 3.7 h and stays there, and the hearth settles at 1485 °C, 15 °C cold, as the dotted curve of Figure 1.4 shows. Integral action cannot remove that offset, because the actuator cannot move.

Mid-ranging [24, 25] adds a second, slow PI loop whose measurement is the fast actuator's output and whose setpoint is the operating point the fast actuator should rest at, here 125 kg/t. Its output is the coke rate. More coke lets the injection fall, so the loop raises its output when its measurement rises, and it is direct acting. With a gain of 0.2 kg of coke per kg of injection and an integral time of 6 h, the coke rises to 369.8 kg/t, close to the 370 kg/t that Equation (1.3) requires, the injection leaves its limit at 12.0 h, the hearth is within 2 °C of setpoint from 19.5 h with a peak of 1503.6 °C, and the injection is back to 125.6 kg/t at 72 h. Neither structure prevents the first dip to 1478.5 °C at 4 h, because the coke arrives four hours after it is charged. A mid-ranging loop five times faster in integral action, with a gain of 0.5 and an integral time of 3 h, fights the injection loop through the coke's delay. The hearth then cycles between 1478 and 1522 °C with a period of about 13.6 h and never settles.

Three rules follow. The fast loop must be the faster of the two. The mid-ranging loop must be slow against the slow actuator's own delay and lag, here six hours, in gain as well as in integral time: in Chapter 2 a gain of 1.0 spends 7.0 h below the alarm even with a 12 h integral time. The resting point of the fast actuator is an economic choice. Resting it near its limit saves money when the fast fuel is the cheaper one, as coal is beside coke, but leaves less room for the next disturbance, which can then reach the limit before the slow fuel arrives. Anti-windup between the two loops is a refinement of its own [26], and Volume 4 treats mid-ranging in general.

Chapter 2 applies the structure with coal as the fast fuel and coke as the slow one, on a hearth dipped once every two hours. Two burden changes there hold the coal on its limit for about 26 h and leave the last half day at about 1477 °C. Mid-ranging raises the coke to 382 kg/t, keeps the coal off its limit and holds the last half day at about 1504 °C, for $1.73 against $1.68 million of fuel a day. A mid-ranging gain of 1.0 makes the two loops fight and spends 7.0 h below the cold-hearth alarm.

The chapter's predictive controller plans the coke five hours ahead instead. Chapter 9 uses the same shape without a mid-ranging loop. Line speed, borrowed from the entry looper, is its fast handle on strip temperature, the firing is its slow one, and a feedforward from speed to firing joins them. In Chapter 5 the electrode regulators correct each arc continuously while the transformer tap moves the operating point in coarse 50 V steps.

> **DID YOU KNOW? Two Handles on One Beam**
>
> Mid-ranging is not only a furnace technique. The storage ring of Volume 6, Chapter 6 steers its electron beam with a fast, small set of correctors and a slow, large one. It offloads the fast set to the slow one by the same mid-ranging [24] that lets coke take over from coal in Chapter 2.

1.4.4 Batch Endpoints in Steelmaking

A converter heat is a batch with one chance at its end. W-1 charges hot metal with 4.4 % carbon and blows oxygen until the bath reaches 0.05 % carbon within ±0.01 % and 1650 °C within ±10 °C. A charge model plans the oxygen and the coolant from the measured charge by mass and heat balances, which is feedforward from the initial state. Its plan is only as good as what it is told, and the scatter of what it is not told, the true carbon, the oxygen efficiency and the vessel's heat loss, leaves W-1's planned end temperature in error by a standard deviation of 18 °C.

Carbon has a measurement the temperature lacks. While carbon is plentiful, decarburization is limited by the oxygen supply. Below a critical carbon C^* the rate is limited by carbon transport and falls in proportion to the carbon left above an equilibrium value C_{eq}, as Equation (1.4) states.

$$r_C = r_{max}\,\frac{C - C_{eq}}{C^* - C_{eq}} \quad (C < C^*), \qquad \hat{C} = C_{eq} + (C^* - C_{eq})\,\frac{r_{C,m}}{r_{max}} \tag{1.4}$$

The off-gas flow and analysis give the measured rate $r_{C,m}$. Integrated from the charge, it estimates the carbon removed, and a 3 % meter error on W-1's 4.35 % of carbon removed is 0.13 %, thirteen times the half-width of the window. Read below C^* as an instantaneous rate, the same signal estimates the carbon by the second form of Equation (1.4), where the same relative error multiplies only $C - C_{eq}$. With $C^* = 0.40\,\%$ and $C_{eq} = 0.02\,\%$ that is 0.0009 % at the aim, 145 times smaller. Chapter 4 measures three hits and nine reblows in ten heats when it stops on the integral, against nine hits and no reblow on the rate with a 25 °C hot aim.

Temperature has one late measurement and a correction that works in one direction only. A sublance reading near the end of the blow lets the model re-plan, and coolant can then take heat out. Nothing puts heat back without blowing carbon below its aim. If e is the static plan's error, m a margin added to the aim and r the error left after the late re-plan, the heat ends at the temperature of Equation (1.5).

$$T_{end} - T_{aim} = \begin{cases} r, & m + e > 0 \text{ (too hot, cooled)} \\ m + e + r, & m + e \le 0 \text{ (too cold, left)} \end{cases} \tag{1.5}$$

With e and r normal, of standard deviations 18 and 4 °C, Figure 1.5 gives the share of W-1's heats inside the 10 °C window. Without the late reading, 42 % hit. With it and no margin, half the heats are cold and cannot be helped, so 70 % hit at 13.3 °C rms. A margin of one standard deviation raises the hits to 92.5 %, and two standard deviations, 36 °C, to 98.2 % at 4.2 °C rms, close to the 98.8 % that a two-way correction would reach. The margin is paid in coolant and in heats that must be cooled. In Chapter 4 all three misses of the sublance campaign without a margin are cold, and a hot aim of 25 °C cuts the rms error from 18.2 to 7.2 °C.

Between heats the model can learn. If one constant, such as the vessel's heat-loss factor, drifts by δ per heat as the lining wears, and each heat moves the model's estimate a fraction λ of the way toward the value that would have explained it, the estimate settles behind the truth by the lag of

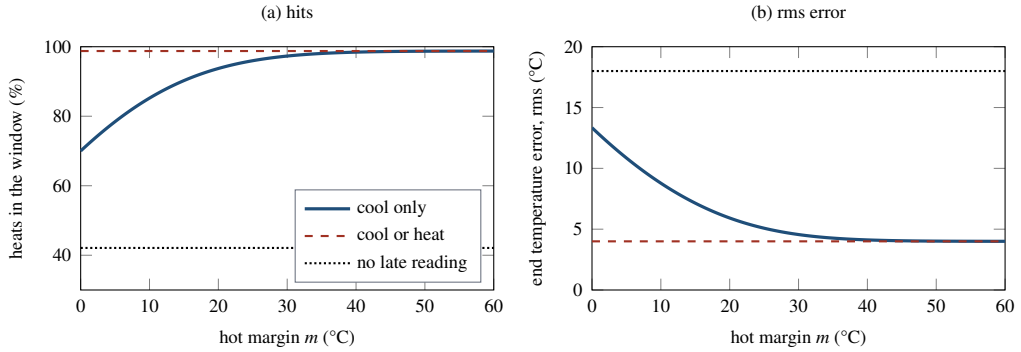


Figure 1.5 W-1's Converter Heats Against the Margin Added to the Temperature Aim. Static Error 18 °C, Error After the Late Re-Plan 4 °C, Window ±10 °C, All Normal. (a) The Share of Heats in the Window and (b) the rms End Temperature Error When the Late Correction Can Only Cool (Solid), When It Could Also Heat (Dashed) and Without the Late Reading (Dotted).

Equation (1.6) [27].

$$\hat{\varphi}_{k+1} = \hat{\varphi}_k + \lambda\,(\varphi_k - \hat{\varphi}_k) \quad \Rightarrow \quad \varphi_{k+1} - \hat{\varphi}_{k+1} \to \frac{\delta}{\lambda} \tag{1.6}$$

A W-1 lining whose loss factor grows by 0.02 a heat, learned at $\lambda = 0.5$, leaves a lag of 0.04, about 8 °C when each 0.01 of the factor is worth 2 °C. A larger λ shortens the lag and copies more of each heat's random error into the next plan. Chapter 4 derives a lag of 0.06, about 13 °C, and learning raises its hits on a wearing lining from five to seven in ten. Run-to-run control is the subject of Volume 8. The other batches of the volume have the same shape. Chapter 5 steps its power program on the metered energy per tonne and calls the tap when its estimate of the bath reaches 1630 °C. Chapter 11 ends each underfeed when a slope of the cell's resistance passes a threshold, a derivative read where a level says too little, as in the rate endpoint.

1.4.5 Tracking Material Along a Line

On a strip line the delay between an actuator and a gauge is a length. A gauge L downstream reads, at time t, the strip that passed the actuator at the earlier time $t - \theta$ defined by Equation (1.7).

$$\int_{t-\theta}^{t} v(t')\,\mathrm{d}t' = L, \qquad \theta = \frac{L}{v} \text{ at constant speed} \tag{1.7}$$

W-1's coating gauge sits 120 m after the air knives, a delay of 60 s at 2.0 m/s and 80 s at 1.5 m/s. A simulator or a Smith predictor [28] that stores the delay as 60 s is 20 s short after the slowdown, while one that records each piece of strip against its position along the line is right at every speed and through every ramp. Chapter 10 measures the difference. Its predictor with a delay stored in seconds tolerates a true delay 9 s longer than its own but not one 10 s longer, while the same predictor keeping its delay as 100 m of strip stays stable after a slowdown, at a spread of 1.24 g/m^2. Chapter 8 keys its X-ray gauge and the span between its stands to strip length in the same way. A disturbance can carry the speed as well. The strand of Chapter 6 bulges between support rolls of pitch λ, so the mold level sees a period λ / v, 12.5 s at 1.2 m/min on 0.25 m rolls, and a canceller locked to that computable frequency removes it.

Tracking is the same bookkeeping applied to identity. The model of Chapter 7 shifts its copy of every slab with each walk of the beams, so it knows each slab's position, charge temperature and elapsed time. Chapter 9 starts each coil's record at the strip coordinate it will reach after the strip

stored ahead of it in the looper, and the pellets of Chapter 3 charged at a step of the gas reach the bustle one residence later, at hour 7 of a step made at hour 1.

An inventory decouples a line fed in pieces from a furnace that wants a steady speed. W-1's entry stops for 60 s to weld each coil to the next while the process section runs at $v = 2.0\,\mathrm{m/s}$, and the entry decelerates and accelerates at $\dot{v}_E = 0.1\,\mathrm{m/s^2}$ toward a top speed of $v_E = 2.5\,\mathrm{m/s}$. One weld draws from the looper the length ΔL_w of Equation (1.8), and coils of length ℓ_c supply on average the speed $\bar{v}_E$.

$$\Delta L_w = v\left(t_w + \frac{v}{\dot{v}_E}\right), \qquad \bar{v}_E = \frac{\ell_c}{\ell_c/v_E + t_w + v_E/\dot{v}_E} \tag{1.8}$$

A weld draws 160 m, and a 500 m looper carries a stop of up to 230 s at full speed [29]. Coils of 2000 m supply 2.26 m/s on average, more than the furnace takes, but coils shorter than 850 m supply less, and coils of 800 m supply 1.975 m/s. No tuning of the looper loop can raise that average. Chapter 9 shows both halves. A looper loop tuned to hold its level passes every weld into the process speed and puts 71.3 % of the strip outside its temperature window, while an averaging loop that lets the level swing puts none outside. When short coils lower the supply to 277 m/min, the averaging loop lets the looper run empty seven times, and a cap on the speed at 97 % of that supply removes every stop.

1.4.6 Inferred Gauge and Coating

A rolling stand stretches under its own force, so the strip leaves thicker than the unloaded roll gap S by the stretch F/M, where M is the mill modulus [30]. The force in turn falls as the strip thickens, with a plastic modulus $Q = -\partial F/\partial d$, and an offset c such as thermal crown or roll wear moves the roll surface without moving the gap transducer. The thickness and its sensitivity to the gap are given by Equation (1.9).

$$d = S + c + \frac{F(d)}{M}, \qquad \frac{\partial d}{\partial S} = \frac{M}{M+Q} \tag{1.9}$$

W-1's last stand has $M = 5\,\mathrm{MN/mm}$ and $Q = 20\,\mathrm{MN/mm}$ and rolls 3.00 mm from a gap of 1.00 mm at 10 MN. Only a fifth of a gap move reaches the strip, so a gap move of five times a gauge error cancels it in one step.

The gaugemeter turns Equation (1.9) into a thickness measured at the stand, $\hat{d} = S + F/\hat{M}$, from the gap transducer, the load cell and a calibrated modulus $\hat{M}$. Its error is Equation (1.10).

$$\hat{d} - d = -c + F\left(\frac{1}{\hat{M}} - \frac{1}{M}\right) \tag{1.10}$$

Nothing on the right depends on the strip. A harder or colder piece moves the force and the thickness together along the mill line, and the gaugemeter reads it exactly. An offset is the opposite case, as Figure 1.6 shows. Crown of $c = -0.10\,\mathrm{mm}$ thins the strip by 0.02 mm and raises the force to 10.4 MN, and the gaugemeter reads 3.08 mm. Its error is the whole offset with the opposite sign, five times the true change. A loop that holds the gaugemeter at 3.00 mm closes the gap to 0.60 mm and delivers 2.90 mm, five times worse than no loop at all. A modulus calibrated 5 % high adds a steady bias of −95 µm.

> **+ DID YOU KNOW? Most of the Stretch Is in the Rolls**
>
> Of the elastic stretch of a rolling mill under load, the rolls take 40 to 70 percent, the housing 10 to 16 percent and the screw-down device 4 to 20 percent [12]. The mill modulus M of Equation (1.9) lumps all three into one number.

The X-ray gauge measures the product itself, 6 m after the stand and 0.5 s late at 12 m/s. A proportional loop on a pure delay in series with a static gain of 0.2 reaches its stability limit at a

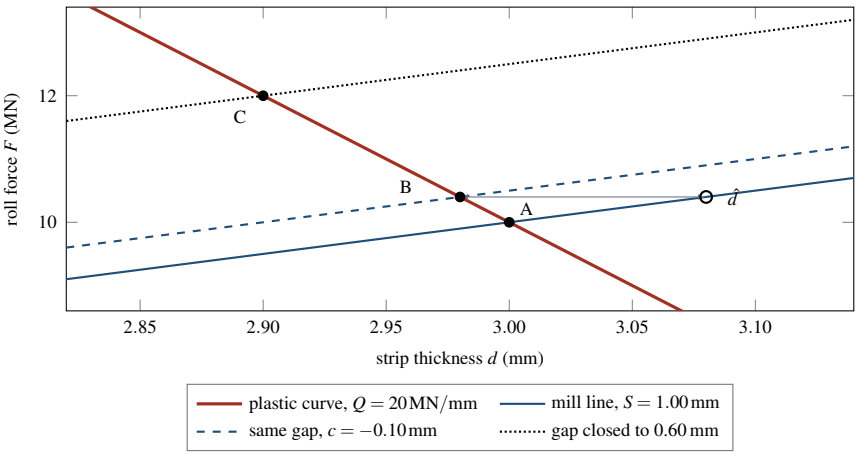


Figure 1.6 W-1's Last Stand. The Strip Leaves Where the Mill Line Crosses the Plastic Curve (Thick Solid), at A. Crown of −0.10 mm Moves the Mill Line (Dashed) and the Strip to B, 0.02 mm Thinner, While the Gaugemeter Reads the Higher Force Back Along the Undisplaced Line (Thin Solid) as $\hat{d} = 3.08$ mm (Open Circle). A Loop That Holds $\hat{d}$ at 3.00 mm Closes the Gap to 0.60 mm (Dotted) and Rolls 2.90 mm, at C.

gain of $1/0.2 = 5$ and a period of 1.0 s, so the delayed loop cannot use the gain the stand's own measurement allows. The two are combined by frequency: the gaugemeter closes the fast loop, and a slow trim from the X-ray removes its bias. Chapter 8 has a gap gain of 0.2104 and a deadbeat gain of 4.752. There, 200 µm of crown leaves the gaugemeter loop 198.6 µm thin against 41.9 µm with the loop open. On an entry wave the inferred loop holds 11.6 µm rms where the delayed loop at the same gains gives 173.8 µm, and a monitor trim at an eighth of the gauge loop's speed removes the crown.

A coating is inferred in the same spirit. Air knives wipe the zinc that a strip drags out of the bath, and the coating left on each side follows a power law in strip speed, knife distance D and knife pressure P [31], written in Equation (1.11) with the exponents Chapter 10 uses.

$$W = W_0 \left(\frac{v}{v_0}\right)^{0.6} \frac{D}{D_0} \left(\frac{P}{P_0}\right)^{-0.6}, \qquad \frac{\partial W}{\partial P} = -0.6\,\frac{W}{P} \tag{1.11}$$

W-1 coats 66 g/m^2 against a minimum of 60 g/m^2, with $W_0 = 60\,\mathrm{g/m^2}$ at $v_0 = 2.0\,\mathrm{m/s}$, $D = D_0$ and $P_0 = 25\,\mathrm{kPa}$. It needs 21.3 kPa, where the loop gain is −1.86 g/m^2 per kPa. Slowing to 1.5 m/s lowers the pressure needed to 16.0 kPa, which multiplies the gain by 1.33, and multiplies the delay by the same factor. Equation (1.11) holds the coating when P is proportional to v, so a speed feedforward removes the coating step and the gain change, but not the longer delay. A heavy grade of 140 g/m^2 at the same speed multiplies the gain by $(140/66)^{1+1/0.6} = 7.43$. Every gram per square meter on both sides of W-1's 1.25 m strip is 18 kg of zinc an hour, \$54 at \$3 per kilogram, so the spread of the coating sets the margin above the minimum and the zinc it costs. Chapter 10 finds a gain 5.6 times the nominal on its heavy grade, a spread cut from 5.08 to 1.25 g/m^2 by speed feedforward through a slowdown, and a coating tail heavy enough that 1 % below the minimum needs a margin of 3.3 standard deviations rather than 2.33.

The other chapters infer what they cannot measure. Chapter 2 reads a hearth soft sensor corrected at every cast [32], Chapter 3 closes an oxygen balance on the top gas and corrects it from a laboratory sample every two hours, Chapter 7 computes each slab's mean and corrects its exchange factor from a pyrometer, and Chapter 11 reads the alumina through a pseudo-resistance whose minimum hides which side of it the cell is on. The rule is the same in each [33]. An inferred measurement is exact for what its model contains, wrong and sometimes inverted for what the model leaves out, and it needs a

Table 1.1 For Each Chapter, the Slow Part, the Fast Part and What Joins Them

Ch.	Slow part	Fast part	What joins them
2	Coke, 5 h of descent	Coal, limited, hearth dipped every 2 h	Mid-ranging on coke, or a coke plan
3	Product over a 6 h residence	Top gas within minutes	Soft sensor with laboratory bias, buffer override
4	Lining wear from heat to heat	One blow of about 21 min	Charge model, late re-plan, run-to-run learning
5	Bath heat, taps in 50 V steps	Electrodes every 0.02 s	Power program on metered energy
6	Casting speed and strand	Mold level, sensor 0.3 s late	Canceller at the roll-passing frequency
7	Slab, 2.5 h in the furnace	Mill pace, zone loops	Slab-tracking setpoints every 120 s
8	X-ray 0.5 s late, crown drift	Gaugemeter, 4 ms scan	Monitor trim at an eighth the speed
9	Furnace, 240 s lag	Line speed, one 80 s residence	Looper, speed and firing feedforward
10	Gauge 100 m downstream	Knives, line speed ramps	Speed feedforward, predictor in meters
11	Heat and ledge, one dip a day	Resistance read every second	Target corrected from the daily dip

slower true measurement to correct its bias, with each sample compared to the inference at the instant the sample was taken, as Volume 5 develops.

1.4.7 Slow Furnaces and Fast Strip

A strip in a continuous furnace heats toward the furnace temperature T_f with a time constant proportional to its thickness, $\tau_1 d/d_1$, while it crosses a path of length L_f. Treated as plug flow, it leaves at the temperature of Equation (1.12).

$$T_{exit} = T_f - \left(T_f - T_{in}\right) \exp\left(-\frac{L_f}{v\,\tau_1 d/d_1}\right) \tag{1.12}$$

The exit temperature depends on speed and thickness only through the product vd. W-1's annealing section has $L_f = 120\,\text{m}$ and $\tau_1 = 30\,\text{s}$ for 1 mm strip, so at 2.0 m/s the strip stays 60 s and a furnace at 921.3 °C delivers 800 °C. A coil of 1.2 mm strip welded to it leaves at 752 °C, 48 °C cold. The furnace would need to rise by 59 °C, which takes minutes, while slowing the line to 1.67 m/s keeps vd and restores 800 °C within one residence. The looper absorbs the slower draw. Speed is the fast handle, firing the slow one, and a feedforward from speed to firing lets the furnace follow. Chapter 9 measures this. Strip of 1.3 mm with the speed held leaves 77.5 % of it outside its temperature window, against 2.5 % with the speed scaled by thickness.

Every chapter of the volume joins a slow part and a fast part in one of a few ways: a setpoint computed by a model of the slow part, a feedforward from the fast part's measured changes, an inventory that lets the two run at different rates, or a slow loop that returns a fast actuator to its range. Table 1.1 lists, for each chapter, what is slow, what is fast and what joins them. Chapter 12 adds the measured results and compares how slow furnaces and fast strip are coordinated.

1.5 Instruments and Final Elements in This Volume

Figure 1.7 lists what the controllers of each chapter read and what they move. A thermocouple, dipped into the metal or set in a furnace wall, is read in five chapters and a radiation pyrometer in two. Seven chapters control a value that only a model computes: a soft sensor, a slab mean, a gaugemeter, a carbon content from the off-gas rate, a bath temperature from the metered energy or a cell resistance. On the other side, solids are fed or added in five chapters, a hydraulic cylinder or a positioning drive moves the handle in four, and burners and valves in three each.

The hot metal of Chapter 2, the converter bath of Chapter 4 and the bath of Chapter 11 are read by a thermocouple dipped once a cast, once a heat or once a day. Thermocouples are typically less

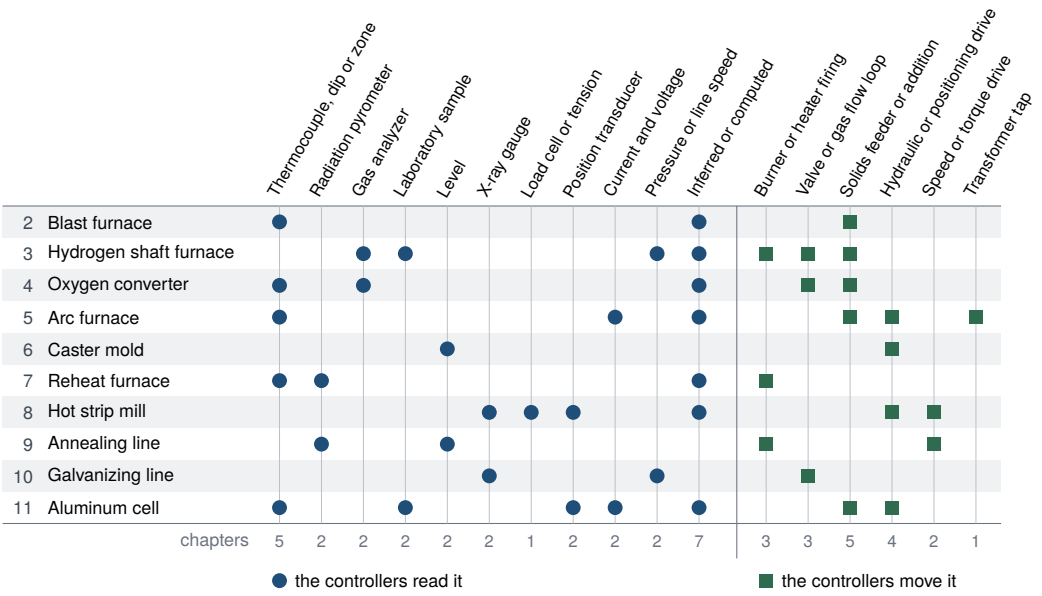


Figure 1.7 Sensors and final elements by chapter. A round mark means that the chapter's controllers read the device, a square mark that they move it. A mark is set only where the chapter's variables table, control structure or technology section names the device, and a flow station counts as a valve or flow loop. A blank says that the model leaves the device out. Notice that seven of the ten chapters control a computed value and that only the caster reads a single instrument.

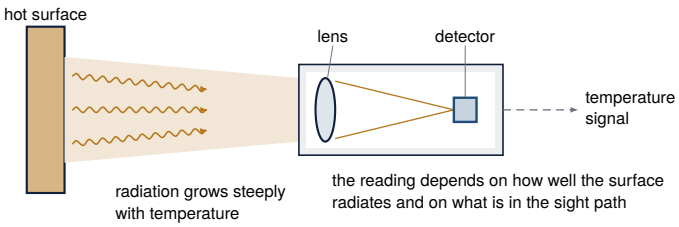


Figure 1.8 Principle of a radiation pyrometer. Notice that the instrument sees only the surface it points at and only through its emittance, so it reads the surface of a slab and never its mean, the quantity Chapter 7 must deliver [34].

accurate than resistance thermometers but more rugged, with greater spans [34], which suits a bath of molten metal. The reading is a sample, and each chapter holds it until the next one.

Hot product leaving a furnace is read without contact. A radiation pyrometer concentrates the radiation of a hot surface onto a detector, and the energy radiated grows with the fourth power of absolute temperature. Its main weakness is emittance, which depends on the material, its finish and its shape and is known for certain only by testing the object at a known temperature. Radiation reflected from other bodies adds to the reading [34]. Figure 1.8 draws the principle. Chapter 7 trims its slab model from a pyrometer, and Chapter 9 reads bright strip, the hard case.

The gauge of Chapter 8 is computed from two instruments at the stand, a transducer on the roll gap and load cells under the roll stack. A load cell is a steel element of precisely known modulus with strain gauges bonded to it, which compresses on a microscopic scale and changes their resistance [34]. Figure 1.9 draws one under the leg of a vessel. The gaugemeter of Equation (1.10) is only as good as this force and the modulus it is divided by.

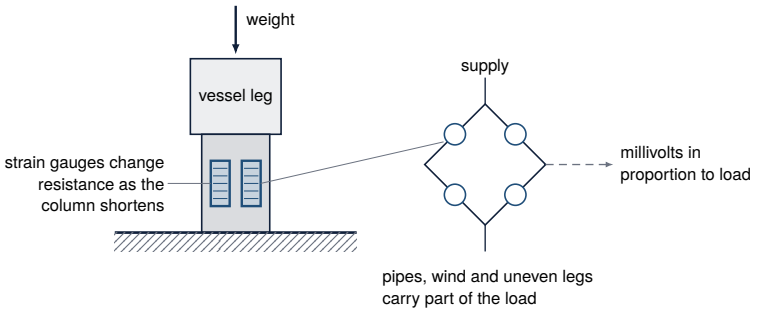


Figure 1.9 Principle of a strain-gauge load cell, drawn under the leg of a vessel. Notice that the cell reads only the load that passes through it. In a rolling stand the cells sit under the roll stack and carry the whole roll force [34].

Layer	What it does in this volume	Where it runs	How often
Setup, tracking and predictive control	plans a heat or a coil, tracks each slab, writes zone and fuel setpoints	**Level 2 process computer** that writes setpoints to the loops below	per heat or coil; 4 s to 30 min in the apps
Trims from late measurements	corrects a model or an inferred value from a dip, an assay or a gauge downstream		per sample: a cast, a heat, a laboratory report, a day
Regulatory control	PI loops on temperatures, flows, pressures, coating and the looper	**Distributed control system** or programmable controller	milliseconds to minutes; 0.2 s to 5 min in the apps
Fast servo loops	roll gap, electrode arms and the slide gate, inside the loops above	**Programmable controller** or a dedicated fast task	1 ms in one published mill; 4 ms to 0.05 s in the apps
Field devices	thermocouples, pyrometers, gauges, load cells, burners, cylinders and drives	on wires or a bus to the controller	continuously

Figure 1.10 The control layers of this volume, the platform that hosts each and how often each acts. The range of loop periods is that of [35], the 1 ms is the gap control of the published mill model that Chapter 8 cites, and the ranges in the apps are those of the chapters' technology sections. Notice that a trim acts once per cast, heat, report or day, while the servo loops below it act thousands of times a minute.

Figure 1.10 places the volume's loops on their platforms. Control loops in a plant run with cycle times from milliseconds to minutes, and a distributed control system uses a supervisory loop to mediate a group of localized controllers [35]. The apps run their fastest loops every 4 ms at the mill stand of Chapter 8 and every 0.02 s on the electrodes of Chapter 5, their regulatory loops from every 0.2 s in Chapter 10 to every 5 min in Chapter 2, and their setup and predictive layers from every 4 s in Chapter 9 to every 30 min in Chapter 3. The converter of Chapter 4 decides only at the events of each heat.

1.6 The Simulator Framework

Every chapter of this book has an app that runs in a web browser, collected with the others at https://apmonitor.com/apps/. About 240 apps exist across the series. All but a few are built on one shared framework, so the engine, the PID algorithm, the trends and the data export behave the same way in every chapter. The exceptions are standalone studios, such as the cascade studio and the two Temperature Control Lab apps, which follow the same ideas in their own code. This section describes the framework once and states where the standalone apps differ. Table 1.2 lists what a reader meets.

Table 1.2 What Every App Provides. Standalone Apps Are Described in Their Own Chapters.

Element	What it does
Run, Pause, Reset	Start, hold and restart the run. Reset keeps the reader's settings.
Speed	Simulated time per second of real time. It changes the pace and never the result.
Manual and Auto	Hand each actuator to the reader or to the controller, without a bump.
Scenario list	Numbered cases, each with a briefing, a completion rule and a reload button.
Control panel	Setpoints, manual outputs, tuning, options and disturbances, at a Beginner or an Advanced level.
Trends	SP and PV above the MV for each loop, with limits, event marks and a cursor.
Scorecard	IAE, overshoot, settling time, MV travel and time at a limit for each loop.
Learn card	Predict, Try It, two hints and Explain, opened by the reader in that order.
Show Model	Equations, live variable values, editable parameters and assumptions.
Export CSV	Every logged sample of the run as a text file with a header row.
Instructions	A separate page with background, an operating guide and the model's limitations.

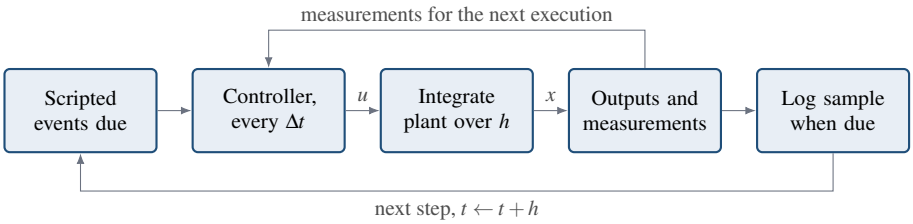


Figure 1.11 Work of the Engine in One Integration Step

1.6.1 The Simulation Engine

Each app carries a plant model written as ordinary differential equations $\frac{dx}{dt} = f(x,u,d,p)$ in the states x, the inputs u, the disturbances d and the parameters p. The engine integrates the model with the classical fourth-order Runge–Kutta method on a fixed step h set for the fastest dynamics of the app. Some apps supply their own step, such as explicit Euler for a discontinuous right-hand side or an exact update for a model driven by events.

Three clocks run at separate rates. The plant advances every step h. The controller executes every control period Δt, equal to h unless the app declares a slower controller, and holds its outputs between executions. The log takes one sample at a fixed interval of simulated time for the trends, the scorecard and the exported file. The order of work in one step is shown in Figure 1.11. The controller also executes once when a scenario is loaded, so every run starts from an initialized controller.

Drawing the screen is kept apart from these clocks. A speed of 5 means five units of the app's time base per second of real time. Speed, frame rate and pauses change how fast a run is shown and never what it computes, so a phone and a workstation produce the same trajectory and the same file.

Noise comes from a deterministic pseudorandom generator. Its seed is declared by the app, or by a scenario that needs its own noise record, and it is restored whenever a scenario is loaded or Reset is pressed. The framework gives the reader no seed setting, so *default seed* in this book means the seed the app ships with. A run repeats bit for bit when the same scenario is loaded and the same settings are made before Run. A change made by hand during a run lands on a different step each time, which is one reason numeric answers carry a tolerance. The standalone simulators draw noise from the browser's own generator, so a noisy run there repeats in its statistics and not sample for sample.

Run and Pause start and hold the clock. Reset restarts the run at $t = 0$, paused, and keeps the reader's settings: every setpoint, tuning value, option, model parameter, disturbance slider and mode the reader changed is put back before the plant and the controller are initialized. Scripted events replay, and a one-shot disturbance button pressed during the run is cleared. The reload button beside

the scenario list restores the scenario's own settings instead.

1.6.2 The Velocity-Form PID Algorithm

One PID implementation serves the PID loops of every framework app. It is the velocity (incremental) form: each execution adds the change in the proportional and derivative terms to a running sum, and the integral is a separate term, as shown in Equations (1.13) to (1.16) for execution k and control period Δt.

$$e_k = a\,(\mathrm{SP}_k - \mathrm{PV}_k), \qquad P_k = K_c\,e_k, \qquad D_k = -a\,K_c\,\tau_D\,\frac{\mathrm{PV}^f_k - \mathrm{PV}^f_{k-1}}{\Delta t} \tag{1.13}$$

$$S_k = S_{k-1} + (P_k - P_{k-1}) + (D_k - D_{k-1}) \tag{1.14}$$

$$v_k = u_b + S_k + I_{k-1}, \qquad u_k = \mathrm{sat}(v_k) \tag{1.15}$$

$$I_k = I_{k-1} + \left[\frac{K_c}{\tau_I}\,e_k + \frac{u_k - v_k}{T_t}\right]\Delta t \tag{1.16}$$

where $a = \pm 1$ is the controller action, PV^f is the measurement after an optional first-order filter, u_b is the output bias, S is the running sum of proportional and derivative increments, v is the output before limiting, u is the output sent to the actuator, I is the integral term and T_t is the tracking time constant, equal to τ_I unless a loop sets another value. The function sat clamps v_k to the actuator range and then applies a rate limit if the loop declares one.

The limit acts on the emitted output u_k only. The sum S_k is never clamped, so it equals $P_k + D_k$ and the output before limiting is the position form $u_b + K_c e_k + D_k + I_{k-1}$. Anti-windup touches the integral alone. While the output is limited, the last term of Equation (1.16) drives I toward the value that reproduces the output applied, which is back-calculation [36]. With no integral action, Equation (1.16) is skipped. Clamping the running sum would instead discard each increment that crosses a limit and keep each one that returns, so symmetric noise would push the output off the limit even in a proportional-only loop.

The derivative acts on the measurement, so a setpoint step gives a proportional kick and no derivative kick. A change of K_c acts on the whole error at the next execution, as in the position form. The action $a = +1$ suits a process whose PV rises with the MV, a reverse-acting controller in the ISA convention, and $a = -1$ gives the direct-acting one.

In Manual the controller keeps executing and sets $I_k = u_{\mathrm{man}} - u_b - S_k$, so the output it would send equals the manual value u_{man}. The first automatic output differs from it by one period's change in P and D. In Auto the engine copies each controller output to the manual station, so a switch to Manual is also bumpless. The standalone studios with PID loops use velocity-form code of their own, with the limit on the output and anti-windup on the integral alone, in which a gain change applies only to later increments.

1.6.3 The Predictive Control and Estimator Options

More than 100 apps offer model predictive control (MPC) as an option beside the structure they open with, which stays the default. Volume 9 develops the theory. In most of them the controller plans with a small linear model in discrete time as shown in Equation (1.17).

$$x_{k+1} = A\,x_k + B\,u_k + E\,w_k + c, \qquad y_k = C\,x_k \tag{1.17}$$

where w is a measured or scheduled disturbance and c is a constant left by linearization. The model comes from a step test at the design point, from the app's balances linearized there or again at every decision, or from the model the reader fits. It never reads the true states or parameters of the run, so a parameter the reader changes in Show Model is model mismatch, as on a real plant.

At each decision the controller minimizes the objective of Equation (1.18) over a horizon of N steps.

$$J = \sum_{k=1}^{N} (y_k - r_k)^{\top} Q (y_k - r_k) + \sum_{j} \Delta u_j^{\top} R \, \Delta u_j + \sum_{i} \rho_i s_i^2 \tag{1.18}$$

where r is the setpoint, Q weights the tracking error, R is the move suppression on each input change Δu_j and ρ_i is the price of a slack $s_i \geq 0$ that lets an output band be broken when it cannot be held. Actuator limits are hard constraints. The quadratic program is solved in the browser by a dual active-set method [37]. The first move is applied, and the plan is made again at the next decision.

The plan starts from an estimate of the model state. About half of these apps use a steady-state Kalman filter on the model with an added disturbance state, which absorbs what the model did not predict and removes steady offset without an integral term [38]. Others correct the model output with a filtered bias to the same end. About a dozen apps also offer moving horizon estimation (MHE), which solves the estimation problem over a window of recent samples and can hold bounds that a Kalman filter cannot, such as a concentration that is never negative [39].

With the MPC selected, the tuning fields of the loops it replaces leave the Control panel. In most of these apps the Advanced level shows up to three settings, each a multiple of the shipped value: the prediction horizon, the move suppression and the band price. The trend draws the accepted plan ahead of the present time, and a list of constraint prices shows which limits bind. A plan that is not optimal, not finite or outside its limits is rejected, and a fallback, in most apps the shipped controller, holds the actuator until a plan succeeds. The Model Fitting and Tuning Studio builds the same kind of controller from the reader's own data in its tab, MPC from the fit.

1.6.4 Scenarios, Progress and Data

An app teaches through scenarios, the numbered cases at the top of the Control panel, nine to thirteen in most apps. Loading one sets the model parameters, the controller settings, the disturbances, the mode, the initial state and any scripted events. The information button beside the list opens a briefing in four parts: what the case is, why it matters, what to do and one detail for the practicing engineer. The Learn card holds the lesson back, with Predict and Try It before the run, hints on request and Explain afterward.

Each scenario has a completion rule that the app judges from the run on the simulated clock: a run length, the scripted events passed, a setting made by the reader, a state held for a stated time or a target met. Choosing a case earns nothing. A completed case is marked in the list, and a badge row shows the count and a level, Explorer after one case, Practitioner at half and Completed at all. Progress is kept in the reader's browser with no account. A certificate of progress or of completion can be issued from it, as a record of self-paced work and not an assessed credential.

One case, usually the last, is the graded challenge. It states a target, such as an error band to hold through a sequence of upsets. Changing the model parameters or the disturbances voids the attempt in most apps, and a record made with Reveal suggested tuning says so.

Show Model opens the equations, live values of the variables, the editable parameters and the assumptions. The trends plot SP and PV above the MV for each loop, with actuator limits and marks at setpoint moves, disturbances, mode changes and scripted events. The scorecard reports for each loop the integral of the absolute error (IAE), overshoot, settling time, MV travel and share of time at a limit, with the app's own measures, and it does not combine them into one score. IAE and travel accumulate from Reset, so two runs are compared at the same clock reading.

Export CSV saves the log as a text file with a header row of column names that carry their units and one row per logged sample. In most apps the time column is in seconds, minutes or hours according to the run length. The log holds the most recent 60 000 samples unless the app sets another size. Apps with the MPC option begin the file with comment lines that start with #, which a script must skip.

An app exercise in this book names the scenario by its number and name as the app shows them, every setting that differs from the scenario's defaults, the run length in simulated time, the seed

(the default seed unless stated), what to read and where (a trend, a scorecard tile or a column of the exported file) and a tolerance on a numeric answer. Settings are made before Run unless the exercise gives a time.

1.7 Working With AI Agents While Studying

A reader of this book is likely to have an AI assistant within reach, and the apps make its work easy to test. This section gives the editor's view of how to use one while working through a chapter. In short, the assistant may do the algebra and the scripting, and the reader keeps the checking.

What to delegate. Tasks with a clear input and a result that can be tested are good candidates. An assistant can derive the linearization of a chapter's balance at an operating point and the transfer function that follows from it. It can write a short script that reads the exported CSV file and fits a first-order-plus-dead-time model to a step test. It can propose a first tuning from that fit with a named rule, convert units, explain a term in a model or outline a report. The assistant knows nothing about the app until it is told. Give it the equations and parameter values from Show Model, the column names of the exported file and the scenario that produced the data, as described in Section 1.6.4. Ask it to state its assumptions and the units of every result.

What to check, and how. An assistant returns a fluent answer whether or not the answer is right, so each result needs a test that it could fail. Four tests cover most of the errors a chapter will produce.

- *Units.* Every term of an equation carries the same units, and a process gain is in PV units per MV unit. A gain computed with the valve as a fraction is 100 times the gain of a loop whose valve is in percent.
- *Signs.* Decide from the physics which way the PV moves when the MV rises. Compare that with the sign of the fitted gain and with the controller action of Section 1.6.2.
- *Steady state.* Set the derivatives to zero. The model must return the operating point that Show Model reports, and a linearization must give zero deviation there.
- *A run on the app.* Load the named scenario at the default seed, enter the proposed numbers and run. The trend and the scorecard decide. A tuning that reads well on paper and holds the valve at a limit for half the run has failed.

A fitted model gets one more test: compare it with data it was not fitted to, such as a second step in the opposite direction. When a check fails, report the failure to the assistant with the evidence and ask for the cause, then check the correction the same way.

How the exercises train the habit. One exercise in each chapter is marked *Review the agent*. It presents a tuning, a model or a claim as an AI assistant might propose it, with a flaw, and asks the reader to find the flaw in about five minutes. The flaws are ordinary ones: a unit, a sign, a steady state that does not close, a result measured on the wrong scenario or at a different run length, or a limit that was ignored. Each is found with one of the tests above. The purpose is not distrust. An engineer who can say why an answer is right can hand an assistant much larger tasks and still stand behind the result. That larger picture, in which agents do most of the routine work and the engineer supervises and checks, is the subject of Chapter 12.

1.8 Chapter Roadmap

Table 1.3 lists the chapters with the lesson each adds to this one. The level says how much of Section 1.4 a chapter assumes. The time is an estimate for reading the chapter once and running each case of its app once, at about eight minutes a case.

Chapters 6, 9 and 2 form a first course: an integrating mold whose oscillation has three causes, an inventory that swings so the product does not, and two fuels on two clocks. They rest on Sections 1.4.3

Table 1.3 Roadmap of the Volume

Ch.	Plant	The lesson	Level	Time
2	Iron blast furnace	A slow, wide fuel returns a fast, limited one to its range.	intermediate	2.5 h
3	Hydrogen shaft furnace	The gas answers in minutes and the product over six hours.	advanced	2 h
4	Oxygen converter	One late reading and a one-way trim set the aim.	intermediate	2 h
5	Arc furnace	What a loop measures decides how its phases interact.	advanced	2 h
6	Caster mold	Diagnose an oscillation before treating it.	introductory	2 h
7	Reheat furnace	Holding the zones is not heating the slab.	intermediate	2.5 h
8	Hot strip mill	An inferred gauge is exact or inverted.	advanced	3 h
9	Annealing line	A looper is there to swing, not to hold its level.	introductory	2.5 h
10	Galvanizing line	A delay fixed in meters needs a model kept in meters.	intermediate	2 h
11	Aluminum smelting cell	A reading with a minimum cannot tell its side.	advanced	2.5 h

and 1.4.5. Chapters 4, 7 and 10 add a batch endpoint, a tracked product and a delay set by distance. Chapters 3, 5, 8 and 11 combine inferred measurements with interaction, inversion or an extremum and are best left until Section 1.4.6 is familiar. A reader who follows the metal can read the chapters in their order, which is the order of Figure 1.2.

KEY POINTS

- A steel works joins units on very different clocks, and each slow unit must deliver what the next one needs.
- Zone loops hold the furnace, not the product.
- A slow loop returns a fast, limited actuator to its range, and it must be slow against the slow actuator's delay.
- A one-way late correction needs a margin on the aim, paid in coolant.
- An inferred measurement is exact for what its model contains and wrong for what it leaves out, so a slower true measurement corrects it.

Where else this applies. Mid-ranging returns in Volume 4, delay and soft sensors in Volume 5 and run-to-run control in Volume 8.

References

[1] U.S. Department of Energy, "Comprehensive report to congress, clean coal technology program: Blast furnace granulated coal injection system demonstration project," U.S. Department of Energy, Washington, DC, Tech. Rep., 1990.

[2] "Bessemer, Sir Henry," in *Encyclopædia Britannica*, 11th ed. Cambridge: Cambridge University Press, 1910, vol. 3, p. 823, read in the Wikisource transcription (en.wikisource.org, 1911 Encyclopædia Britannica/Bessemer, Sir Henry), accessed November 19, 2026.

[3] T. R. Beck, "ECS classics: Hall and Héroult and the discovery of aluminum electrolysis," *The Electrochemical Society Interface*, vol. 23, no. 2, pp. 36–37, 2014.

[4] T. Fukushima, "Recent Technological Progress in High Speed Continuous Annealing," *Transactions of the Iron and Steel Institute of Japan*, vol. 25, no. 4, pp. 278–293, 1985.

[5] J. Stubbles, "Energy use in the U.S. steel industry: A historical perspective and future opportunities," U.S. Department of Energy, Office of Industrial Technologies, Tech. Rep., 2000.

[6] 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.

[7] M. Ohnishi, T. Nishio, and T. Terashima, "Utilization of sliding gate nozzle in Japan," *Transactions of the Iron and Steel Institute of Japan*, vol. 16, no. 3, pp. 164–167, 1976.

[8] K. Kono, M. Fujiura, K. Goto, and J. Yamasaki, "Walking beam slab reheating furnace," *Tetsu-to-Hagane*, vol. 55, no. 14, pp. 1392–1398, 1969, in Japanese.

[9] M. Hatano, Y. Matoba, K. Otsuka, M. Yoshiki, and T. Miyagi, "Automatic control of hot metal temperature of blast furnace," *Transactions of the Iron and Steel Institute of Japan*, vol. 22, no. 7, pp. 534–542, 1982.

[10] "Technology of automatic mold bath level control," *Transactions of the Iron and Steel Institute of Japan*, vol. 24, no. 9, p. 777, 1984, new Technology note.

[11] F. Superchi, A. Mati, C. Carcasci, and A. Bianchini, "Techno-economic analysis of wind-powered green hydrogen production to facilitate the decarbonization of hard-to-abate sectors: A case study on steelmaking," *Applied Energy*, vol. 342, p. 121198, 2023.

[12] 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.

[13] H. Imada and Y. Nimura, "A new electrode-type mold level sensor for continuous casting machine," *Transactions of the Iron and Steel Institute of Japan*, vol. 22, no. 12, pp. 991–996, 1982.

[14] S. Ueda, S. Natsui, H. Nogami, J.-i. Yagi, and T. Ariyama, "Recent progress and future perspective on mathematical modeling of blast furnace," *ISIJ International*, vol. 50, no. 7, pp. 914–923, 2010.

[15] N. Dogan, G. A. Brooks, and M. A. Rhamdhani, "Comprehensive model of oxygen steelmaking part 1: Model development and validation," *ISIJ International*, vol. 51, no. 7, pp. 1086–1092, 2011.

[16] L. Holappa, "A general vision for reduction of energy consumption and CO_2 emissions from the steel industry," *Metals*, vol. 10, no. 9, p. 1117, 2020.

[17] F. Patisson and O. Mirgaux, "Hydrogen ironmaking: How it works," *Metals*, vol. 10, no. 7, p. 922, 2020.

[18] B. G. Thomas, "Modeling of the continuous casting of steel—past, present, and future," *Metallurgical and Materials Transactions B*, vol. 33, no. 6, pp. 795–812, 2002.

[19] P. Homsi, J.-M. Peyneau, and M. Reverdy, "Overview of process control in reduction cells and potlines," in *Essential Readings in Light Metals: Volume 2, Aluminum Reduction Technology*. Cham: Springer, 2016, pp. 739–746.

[20] U.S. Geological Survey, "Mineral commodity summaries 2025: Iron and steel," U.S. Geological Survey, Reston, VA, Tech. Rep., 2025.

[21] ——, "Mineral commodity summaries 2025: Aluminum," U.S. Geological Survey, Reston, VA, Tech. Rep., 2025.

[22] A. Steinboeck, D. Wild, T. Kiefer, and A. Kugi, "A mathematical model of a slab reheating furnace with radiative heat transfer and non-participating gaseous media," *International Journal of Heat and Mass Transfer*, vol. 53, no. 25–26, pp. 5933–5946, 2010.

[23] A. Steinboeck, D. Wild, and A. Kugi, "Nonlinear model predictive control of a continuous slab reheating furnace," *Control Engineering Practice*, vol. 21, no. 4, pp. 495–508, 2013.

[24] B. J. Allison and A. J. Isaksson, "Design and performance of mid-ranging controllers," *Journal of Process Control*, vol. 8, no. 5–6, pp. 469–474, 1998.

[25] B. J. Allison and S. Ogawa, "Design and tuning of valve position controllers with industrial applications," *Transactions of the Institute of Measurement and Control*, vol. 25, no. 1, pp. 3–16, 2003.

[26] S. Haugwitz, M. Karlsson, S. Velut, and P. Hagander, "Anti-windup in mid-ranging control," in *Proceedings of the 44th IEEE Conference on Decision and Control*, Seville, Spain, 2005, pp. 7570–7575.

[27] A. Ingolfsson and E. Sachs, "Stability and sensitivity of an EWMA controller," *Journal of Quality Technology*, vol. 25, no. 4, pp. 271–287, 1993.

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

[29] P. R. Pagilla, S. S. Garimella, L. H. Dreinhoefer, and E. O. King, "Dynamics and control of accumulators in continuous strip processing lines," *IEEE Transactions on Industry Applications*, vol. 37, no. 3, pp. 934–940, 2001.

[30] 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.

[31] 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.

[32] H. Saxén, C. Gao, and Z. Gao, "Data-driven time discrete models for dynamic prediction of the hot metal silicon content in the blast furnace—a review," *IEEE Transactions on Industrial Informatics*, vol. 9, no. 4, pp. 2213–2225, 2013.

[33] 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.

[34] 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.

[35] 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.

[36] K. J. Åström and L. Rundqwist, "Integrator windup and how to avoid it," in *Proc. 1989 American Control Conference*, 1989, pp. 1693–1698.

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

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

[39] C. V. Rao, J. B. Rawlings, and D. Q. Mayne, "Constrained state estimation for nonlinear discrete-time systems: Stability and moving horizon approximations," *IEEE Transactions on Automatic Control*, vol. 48, no. 2, pp. 246–258, 2003.

1.9 Exercises

1.1 Theory and calculation · 3 min

A furnace zone has a heat capacity of 144 MJ/K, and its losses and load together act as a conductance of 0.2 MW/K. It heats slabs 250 mm thick with $k = 30\,\mathrm{W/(m\,K)}$ and $\rho c_p = 5.4\,\mathrm{MJ/(m^3\,K)}$.

(a) Find the zone's time constant. *Checkpoint:* 720 s, 12 min.

(b) Find the slab's conduction time and its Biot number for a radiant coefficient of $450\,\mathrm{W/(m^2\,K)}$. *Checkpoint:* 46.9 min and 1.88.

1.2 Concept check · 2 min

A reheat furnace runs with three zone loops that hold their setpoints exactly. The mill asks for slabs 25 % faster, and nobody changes the setpoints. What happens to the slabs that reach the mill?

(a) nothing, because the zone loops hold every zone on its setpoint

(b) they leave colder, with a larger gap between surface and mean

(c) they leave colder, with a smaller gap between surface and mean

(d) they leave hotter, because the zones fire harder for the larger load

1.3 Theory and calculation · 4 min

A hearth is held by a fast fuel with a gain of 1.5 °C per kg/t that rests at 20 kg/t below its upper limit, and a slow fuel with a gain of 3.0 °C per kg/t. A burden change demands 45 °C more heat.

(a) With the fast loop alone, how much of the demand can it cover, and what offset is left? *Checkpoint:* 30 °C, 15 °C cold.

(b) With a mid-ranging loop on the slow fuel, how far must the slow fuel move at steady state, and where does the fast fuel end? *Checkpoint:* 15 kg/t, back at its resting point.

1.4 Concept check · 2 min

A mid-ranging loop measures the output of a fast fuel loop and moves a slow fuel that answers after a 4 h delay and a 2 h lag. More slow fuel lets the fast fuel fall. Which statement about the mid-ranging loop is right?

(a) it is reverse acting and should be faster than the fast loop

(b) it is reverse acting and should be slower than the slow fuel's dynamics

(c) it is direct acting and should be faster than the fast loop

(d) it is direct acting and should be slower than the slow fuel's dynamics

1.5 Theory and calculation · 3 min

A converter charges hot metal with 4.2 % carbon and aims at 0.04 % within ±0.01 %. The equilibrium carbon is 0.02 %, and the off-gas meters read the decarburization rate 2 % high.

(a) What carbon error does the meter error give if the rate is integrated from the charge? *Checkpoint:* 0.083 %.

(b) What error does it give if the rate is read below the critical carbon as in Equation (1.4)? *Checkpoint:* 0.0004 %.

1.6 Theory and calculation · 4 min

A converter model learns its heat-loss factor after every heat with weight $\lambda = 0.3$, while the true factor grows by 0.03 a heat as the lining wears. Each 0.01 of the factor is worth 2.5 °C of end temperature.

(a) How far behind the truth does the estimate settle? *Checkpoint:* 0.10.

(b) What end temperature error does that lag leave, and in which direction? *Checkpoint:* About 25 °C, cold.

1.7 Theory and calculation · 4 min

A process section runs at 2.5 m/s from a looper of 600 m. The entry stops for 40 s at each weld and decelerates and accelerates at 0.125 m/s^2.

(a) How much strip does one weld draw from the looper, and what is the longest stop a full looper carries? *Checkpoint:* 150 m and 220 s.

(b) Coils are 1500 m long and the entry can run at 3.0 m/s. Can the entry keep up on average? *Checkpoint:* Yes, 2.66 m/s against 2.5 m/s.

1.8 Theory and calculation · 3 min

A coating gauge sits 90 m after the air knives.

(a) Find the delay at 1.8 and at 1.2 m/s, and the error of a Smith predictor that keeps its delay at the first value after the line slows to the second. *Checkpoint:* 50 and 75 s, 25 s short.

(b) A 1.5 mm coil follows a 1.0 mm coil through an annealing furnace run at 2.4 m/s. At what speed does the thicker strip leave at the same temperature, with the furnace unchanged? *Checkpoint:* 1.60 m/s.

1.9 Theory and calculation · 5 min

A stand has a mill modulus $M = 4\,\mathrm{MN/mm}$ and a plastic modulus $Q = 16\,\mathrm{MN/mm}$. It rolls with an unloaded gap of 1.2 mm at a force of 8 MN.

(a) Find the thickness and the fraction of a gap move that reaches it. *Checkpoint:* 3.20 mm and 0.20.

(b) Roll wear opens the gap by $c = +0.05\,\mathrm{mm}$ without moving the gap transducer. Find the true change and the gaugemeter's error. *Checkpoint:* $+10\,\mu\mathrm{m}$ and $-50\,\mu\mathrm{m}$.

(c) A loop holds the gaugemeter at 3.20 mm. What does it deliver? *Checkpoint:* 3.25 mm.

1.10 Concept check · 2 min

A coating loop on knife pressure is tuned for 60 g/m^2. The line switches to 120 g/m^2 at the same speed, and the coating follows Equation (1.11). How does the loop's gain change?

(a) it rises about sixfold, so the loop may go unstable

(b) it doubles, in proportion to the coating

(c) it halves, because the knives need less pressure

(d) it stays the same, because the exponents are fixed