INTERACTIVE CONTROL ENGINEERING
Principles and Applications

Volume 9 Model Predictive Control and Economic Optimization

CHAPTER 12

Automatic Train Operation Simulator Where the Timetable Decides Where to Coast

Pages 287–310

Edited by John D. Hedengren

How to cite this chapter. "Automatic Train Operation Simulator Where the Timetable Decides Where to Coast," in *Model Predictive Control and Economic Optimization*, John D. Hedengren, Ed., vol. 9 of *Interactive Control Engineering: Principles and Applications*. Advanced Process Solutions, LLC, 2026, ch. 12, pp. 287–310.

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

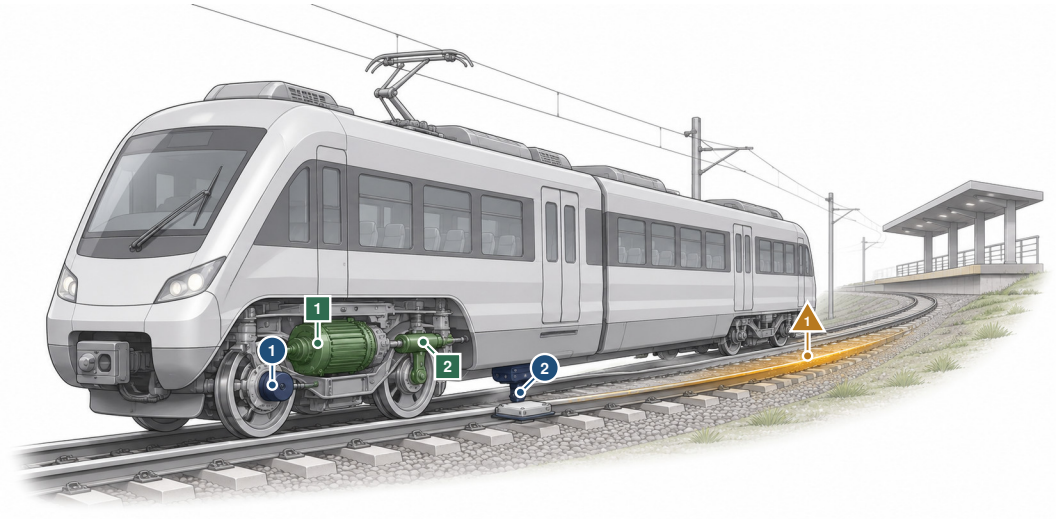
The app for this chapter
https://apmonitor.com/apps/train-ato/

The volume online: contents, other chapters and citation records
https://apmonitor.com/apps/books/v09-mpc-optimization/#train-ato

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

CHAPTER 12

Automatic Train Operation Simulator Where the Timetable Decides Where to Coast



The controller moves 1 traction motor and drive 2 brake

The controller reads 1 wheel speed sensor 2 balise reader under the car

What disturbs it 1 wet rail ahead

Notice where the force acts: the motors push the train, the brakes stop it at the platform, and on the wet rail ahead the wheels can slip. Illustration, not to scale.

IN PLAIN TERMS Where a Train Should Stop Pushing

The main idea. An electric train must reach the next platform at a set time, and it usually has a few seconds to spare. It can spend them by running slower. It can also switch off its motors early and roll, which saves electric power.

The challenge. Speed limits on curves and hills change where rolling pays, so a planner works out the whole run before the train leaves. Wet rails, a heavier load and a late start then push the train off its plan, and the loops that steer it must bring it back.

Why it matters. On the chapter's line, one well-placed switch from power to rolling cuts the energy of a run by 31 percent.

THREE WAYS TO READ THIS CHAPTER

Explore	12 min	the first page, the history in Section 12.2, the fact and try-it segments in the Introduction and Section 12.6, and the key points
Study	45 min	the whole chapter with its three pause points and the twelve exercises at its end
Practice	15 min	the first page, the results (Section 12.6), Technology in Practice (Section 12.7), Limitations (Section 12.8) and the open challenges

The control question An electric train due at the next platform at a set time can spend its spare seconds on speed or on coasting. Where should the force change along the run so that the train arrives on time with the least energy, and what does a plan need from the loop that tracks it?

Learning objectives

- Derive from a Hamiltonian linear in the force why an energy-optimal run uses only full power, hold, coast and brake.
- Interpret the price of time λ as the energy one more second of budget saves, and relate it to the holding speed.
- Compare one coast point with a dynamic program on a line where speed limits interleave, and explain the reversal of the classical rule.
- Explain how wheel slip, a heavier train and a late departure each defeat a plan, and which part of the stack recovers each.
- Recognize a timetable that no plan can meet from the planner's earliest arrival.

Prerequisites Sections 1.4.9, 1.4.7, 1.4.6, 1.6.2 and 1.6.4, and Volume 2, Chapter 2 for a vehicle speed loop in velocity form.

The app https://apmonitor.com/apps/train-ato/

See also Volume 20, Chapter 9 for a speed limited by the distance needed to stop, Volume 20, Chapter 2 for a vehicle whose loops depend on its speed, and Volume 7, Chapter 4 for slip regulated at the peak of a friction curve.

12.1 Introduction

A suburban train leaving station A has 5.0 km and 275 s to reach the platform mark at station B. In the simulator of this chapter the fastest drive that the train and the speed limits allow takes 250.8 s on its plan, and driven flat out the train arrives 21.0 s early with 28.2 kWh on the energy meter at the pantograph. Leaving that drive at one coast point 2.45 km from A, and coasting from there to the final braking curve, the train arrives 3.1 s early with 19.5 kWh, which is 31 % less. A dynamic program that places its own switches stops powering 0.35 km from A at 81 km/h and arrives 1.7 s early with 16.4 kWh. The three runs differ only in where the force changes, and the 24.2 s between the flat-out time and the timetable is what the second and the third spend.

The misconception this plant corrects is that a timetable's slack has one best use whatever the line, a belief that the textbook rule, coast rather than slow down, can leave behind. On a level 5 km test line in the same model, a plan barred from coasting needs 7.2 to 8.8 % more energy than the plan with one coast point, at budgets from 5 to 30 % over the flat-out time, which is the classical result. On the A to B run, where a 60 km/h curve begins 1.5 km out, the order reverses at the 271 s budget. The one-coast-point plan, which holds 98 km/h and brakes into the curve before it coasts, plans 18.76 kWh. The same dynamic program barred from coasting holds about 78 km/h and plans 16.32 kWh, and the unrestricted program, which combines a lower top speed with coasting, plans 15.45 kWh. Where the limits fall decides which rule wins, and that is the argument for planning the whole run.

The problem belongs to a general class, the least-energy transfer between two states of rest in a fixed time, under state constraints that change with position. Energy-efficient driving of road vehicles has been posed as the same optimal control problem [1], and the recursion that solves such a problem on a grid is the one Bellman published in 1957 [2]. A batch brought to a hold temperature by a deadline, or a tank filled by a pump before a tariff changes, has the same shape, with a slack in time that can be spent on rate or on letting the system run. What the train adds is a structure known in advance. Pontryagin's maximum principle limits the energy-optimal drive between stops to full power, speed holding, coasting and full braking [3, 4], so a numerical plan can be checked against the regimes it ought to contain.

DID YOU KNOW? Energy, a Large Line in a Transit Budget

A review of how urban rail systems recover braking energy reports that payments to electric utilities for energy and peak power demand reach up to 35 percent of a rail transit operator's operating expense. Studies that time braking trains to feed accelerating ones claimed savings of 4 to 34.5 percent [5].

A course usually meets the two halves of this problem apart. The theory of optimal train control is written for analysis [4, 6, 7], dynamic programming of running profiles is reported as a numerical method [8, 9], and automatic train operation appears as a survey of research and practice [10]. Speed control elsewhere in the series is a regulator on one setpoint, as in the cruise control of Volume 2, Chapter 2, while here a dynamic program plans the whole trajectory. This chapter works with a browser train simulator in which an analytic four-regime plan and a dynamic program plan the same runs, a tracking loop drives those plans through a lagged drive and a wet rail, and the plan, the run, the clock and the meter are logged side by side. Its measurements are made on one declared line, among them the reversal above and the price of a lost minute. The dynamic program has eleven decisions on a 50 m by 1 km/h grid and solves a plan in 156 to 239 ms in a browser. It is checked against brute force, a closed form and two energy balances, behind a measurement boundary that a test enforces. Pontryagin's structure, Davis resistance, the creep law and dynamic programming are all standard, and Section 1.4.9 sets out the same optimal control and recursion on the heated tank of Chapter 1. Each of the ten scenarios carries one completion rule, and the exercises at the end of the chapter are built on them.

Section 12.2 traces the domain from the driver's judgment to planned profiles. Section 12.3 gives the train, the line, the rail and the instruments. Section 12.4 derives the structure of an energy-optimal run and shows how its switching points move with the budget. Section 12.5 describes the planners, the tracking and stopping loops, the slip override and re-planning. Section 12.6 measures the nine cases and the challenge and verifies the code, Section 12.7 describes the instruments, drives and protection of real trains, Section 12.8 states the limitations, and Section 12.9 draws the results together.

12.2 From a Driver's Judgment to Planned Profiles

For most of the history of electric traction, how a train was driven between stations was left to the driver, who had a timetable and the feel of the line. A 1968 study that posed the run as an optimal control problem is told in the box below, and a solution of the energetically optimal control of a train by the maximum principle followed in 1985 [11].

HISTORY 1968: The Run Between Stations Becomes an Optimal Control Problem

A paper received in 1967 and published in 1968 asked for the train operation that uses the least energy to take a train from one station to the next, arriving at a specified instant and stopping there. Because the speed limit must be observed strictly, it treated the run as an optimization problem with a bounded state variable. Pontryagin and others had given a jump condition of the costate for such problems, but the paper judged it not easily applied to the synthesis of a controller. Its author used a simple method of his own instead and obtained the optimal control law. He believed the work to be the first scientific research on the economization of train operation [12]. The speed limit of that paper is the state constraint $v \le v_{lim}(s) - \Delta v_m$ of Equation (12.6).

The structure that this simulator teaches was settled over the following two decades. A monograph on energy-efficient train control appeared in 1995 [6]. In 2000 the optimal control of a train was characterized as a sequence of a few control modes [3], and a numerical algorithm for a variable grade under arbitrary speed restrictions was built on the analytical properties of the maximum-principle solution [13], with further work on energy-efficient operation of rail vehicles in 2003 [14]. A systematic account of the key principles, from strategies of optimal type and the location of optimal

Table 12.1 Roles of the Variables, and Their Values at $t = 470$ s When the Challenge Runs on Its Suggested Settings

Variable	Role	At 470 s	Range
$\hat{v}$	CV, measured speed	67.4 km/h	0 to v_{lim}
Δt	CV, projected arrival at C	+3.6 s	±15 s
u_{cmd}	MV, lever	51 %	−100 to 100
F_{cap}	MV of the override, motor cap	71.0 kN	5 kN to F_t
$\mu(s)$	DV, adhesion (hidden)	0.08	0.25, 0.08
κ	DV, line receptivity	100 %	0 to 100 %
—	DV, restriction, door hold, load	on, spent, 0	events
s, $\hat{s}$	state, position, odometer	6.822, 6.822 km	0 to 12
v	state, true speed	67.3 km/h	≥ 0
w, $\hat{w}$	state, slip, measured	0.41, 0.42 m/s	—
F_d, F_m	state, demand, motor force	74.7, 71.0 kN	—
E	state, meter	53.0 kWh	—
E_t, E_r	states, drawn, returned	69.7, 16.7 kWh	—
θ	grade at the train	+2.0 %	±2.0 %

switching points to existence, uniqueness, a local energy minimization principle and computational techniques, was published in two parts in 2016 [4, 7].

Dynamic programming entered the problem as the numerical alternative that accepts whatever the line contains. A discrete dynamic programming algorithm that minimized traction energy on a nonlinear point-mass model, subject to the travel time, the speed limits and the available power, was reported in 2000 with energy savings between 10 and 30 % in simulation and in measurements, and was proposed as the basis of a model predictive controller running in hard real time [8]. A 2005 study applied Bellman's dynamic programming to the running profile with complicated motor and brake characteristics, speed limits, nonlinear resistance and variable grades, found computation time to be its obstacle, and reduced it with a confined state space and an irregular lattice [9]. The 50 m by 1 km/h grid used here is a small case of the same approach, sized so that a plan is solved inside one control step of a browser. A 2017 review covers energy-efficient train control together with its integration into the timetable [15], the connection on which cases 8 and 9 of Section 12.6 turn.

Automatic train operation, the onboard function that drives a train from one platform to the next, has its own research and development literature, which a 2017 survey collects [10]. Two physical questions sit beside the optimization. Regenerated braking energy is worth something only if the line can take it, and strategies and technologies for managing it in urban rail have been reviewed [16]. A traction vehicle at the limit of adhesion needs a creep-force law that holds beyond its maximum [17], which is where a wet rail puts a train that its plan expected to find dry.

12.3 The Train, the Line and the Rail as Modeled

The simulated train is a generic four-car electric unit with two motored and two trailer cars, 190 t loaded, on a declared 12 km line, as shown in Figure 12.1. Stations A, B and C stand at 0, 5.0 and 12.0 km. The line speed is 100 km/h, with permanent limits of 60, 80, 80 and 70 km/h over 1.5 to 2.3, 3.1 to 3.8, 7.6 to 8.3 and 10.6 to 11.2 km, a climb of +2.0 % from 5.2 to 7.4 km, a summit 44 m above A, and a descent of −2.0 % from 8.2 to 9.8 km. A temporary restriction of 40 km/h can be imposed from 5.9 to 6.3 km, and the rail can be wet from 6.4 to 7.4 km, the upper part of the climb, or along the whole line. Twelve balises reset the odometer. All of these numbers were chosen for the lesson, and none comes from a railway.

Table 12.1 gives each variable its role. Eight states are integrated, the true position s and speed v, the slip velocity w of the motored wheelsets, the drive's lagged force demand F_d, and four energy integrals, the meter E, the traction energy drawn E_t, the regenerated energy returned E_r, and the braking work done at the rail E_b. The base unit is the second, and the trend axis reads minutes past one minute.

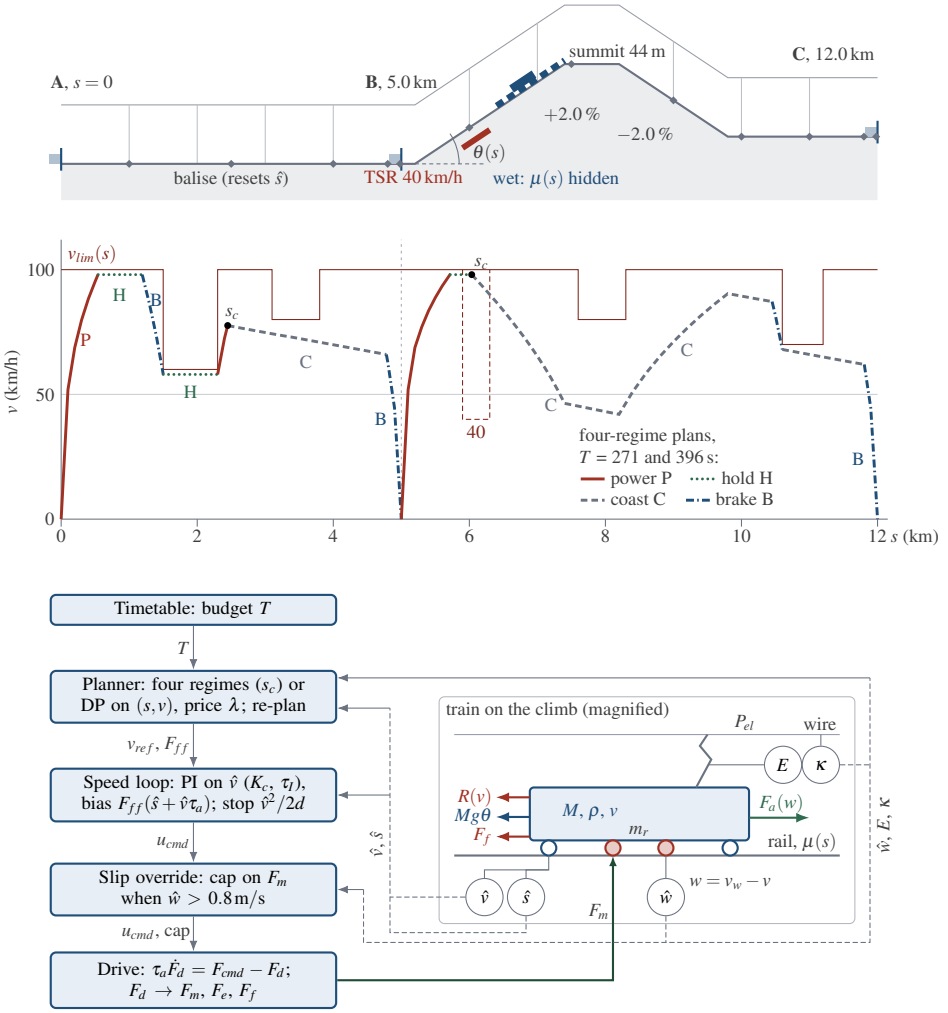


Figure 12.1 The declared line, the train and its automatic train operation, after the simulator's Line view. Top, the profile with heights exaggerated, the stations, the overhead wire, the balises that reset the odometer $\hat{s}$, the 40 km/h restriction, the grade $\theta(s)$ and the wet stretch where the hidden adhesion $\mu(s)$ falls, with the train of Figure 12.4 on it. Middle, speed against distance with the posted limits $v_{lim}(s)$ and the simulator's four-regime plans for the timetable less the 4 s margin, drawn by regime (power solid, hold dotted, coast dashed, brake dash-dotted), with their coast points s_c. Bottom right, the train magnified, with the forces of Equations (12.1) to (12.3), the instruments $\hat{v}$, $\hat{s}$, $\hat{w}$, E and κ. Bottom left, the control layers from the timetable's budget T down to the drive, with the measurements returning on dashed lines. Not to scale.

12.3.1 Motion in Time

The train is one point mass on the route profile, with the motored wheelsets split off so that they can slip, as shown in Equation (12.1).

$$\begin{aligned} \frac{\mathrm{d}s}{\mathrm{d}t} &= v, \qquad R(v) = aM + bMv + cv^2 \\ (\rho M - m_r)\frac{\mathrm{d}v}{\mathrm{d}t} &= F_a(w) - F_f - R(v) - Mg\,\theta(s) \end{aligned} \tag{12.1}$$

where M is the mass aboard, 190 t declared plus any added load, $\rho = 1.08$ is the rotating-mass factor, $m_r = 0.05M$ is the equivalent mass of the motored wheelsets, F_a is the force the rail transmits to them, F_f is the friction brake on the trailer axles, $\theta(s)$ is the grade as a fraction, and $R(v)$ is Davis resistance with $a = 13\,\mathrm{N/t}$, $b = 0.20\,\mathrm{N/(t\,m/s)}$ and $c = 6.0\,\mathrm{N/(m/s)^2}$, the Davis equation whose coefficients the train-resistance literature estimates for suburban and high-speed trains [18]. The resistance is 4.77 kN at 60 km/h and 8.16 kN at 100 km/h, where a coasting train on the level slows at $0.040\,\mathrm{m/s^2}$. Speed is clipped at zero, resistance and friction act only while the train moves, and a train at rest stays there unless the rail pushes it forward.

12.3.2 The Motored Wheelset and the Creep Curve

The motored wheelsets turn at $v_w = v + w$, driven by the motor force F_m and held back by the rail, with a grip that rises with the slip velocity to a peak and falls beyond it, as shown in Equation (12.2).

$$\begin{aligned} m_r \frac{\mathrm{d}v_w}{\mathrm{d}t} &= F_m - F_a(w), \qquad w = v_w - v \\ F_a(w) &= \mu(s)\,\varphi M g\,\frac{2\,w/w_{pk}}{1 + (w/w_{pk})^2} \end{aligned} \tag{12.2}$$

where $\varphi = 0.5$ is the share of the mass on motored axles and $w_{pk} = 0.6\,\mathrm{m/s}$ is the slip of peak grip. The rail gives at most 233.0 kN dry at $\mu = 0.25$ and 74.6 kN wet at $\mu = 0.08$. Past the peak more slip means less grip, so a wheelset asked for more than the peak accelerates itself. A braking tire has the same peak, and the anti-lock brake of Volume 7, Chapter 4 regulates the slip at it. The curve keeps only the rising and falling branches that the creep-force literature describes [17], and its shape is a declared teaching choice. Its time constant, $m_r/(2\mu\varphi M g/w_{pk})$, is 12.2 ms on a dry rail and 38.2 ms on a wet one, and it sets the integration step.

12.3.3 Drive, Brakes and the Energy Meter

The drive's demand follows the command with a first-order lag, and the force that reaches the wheels is the demand cut by the tractive-effort curve, the motors' top speed and the slip override's cap, as shown in Equation (12.3).

$$\begin{aligned} \tau_a \frac{\mathrm{d}F_d}{\mathrm{d}t} &= F_{cmd} - F_d \\ F_m &= \min\big(F_d,\ \chi(v_w)\,F_{t,\max}(v_w),\ F_{cap}\big) \quad (F_d \ge 0) \\ F_B &= \min(-F_d,\ b_{\max}\rho M) \quad (F_d < 0) \\ F_e &= \min\big(F_B,\ F_{r,\max}(v),\ F_{cap,b}\big), \qquad F_m = -F_e \\ F_f &= \min\big(F_B - F_e,\ \mu(s)(1-\varphi)Mg\big) \\ F_{t,\max}(v) &= \min(F_t,\ P_t/v) \\ F_{r,\max}(v) &= \min(F_r,\ P_r/v)\min(1,\ v/v_f) \end{aligned} \tag{12.3}$$

where $\tau_a = 1.5\,\mathrm{s}$, $F_t = 220\,\mathrm{kN}$ and $P_t = 2.8\,\mathrm{MW}$ give a base speed of 45.8 km/h, the electric brake of $F_r = 200\,\mathrm{kN}$ and $P_r = 3.2\,\mathrm{MW}$ fades below $v_f = 2\,\mathrm{m/s}$, and $b_{\max} = 1.0\,\mathrm{m/s^2}$ is the full service brake. The taper $\chi = \min\big(1, \max(0, (V_w - v_w)/\Delta V_w)\big)$ with $V_w = 125\,\mathrm{km/h}$ and $\Delta V_w = 10\,\mathrm{km/h}$ makes an unchecked spin end at the motors' top speed rather than run away. Braking blends, with the electric brake taking what it can through the motored wheelsets and the friction brake on the trailers taking the rest under ideal wheel-slide protection. The meter integrates the electrical power at the pantograph, as shown in Equation (12.4).

$$\frac{\mathrm{d}E}{\mathrm{d}t} = P_{el} = \frac{F_m^{+}\,v_w}{\eta_t} - \kappa\,\eta_r\,F_e\,v_w \tag{12.4}$$

where $\eta_t = 0.85$ and $\eta_r = 0.80$, so a kilowatt-hour that goes out as traction and comes back through the brake keeps 0.68 of itself, and the receptivity κ is the share of regenerated power the line accepts. The line voltage $V_0 - R_f P_{el}/V_0$, with $V_0 = 1500\,\mathrm{V}$ and $R_f = 0.06\,\Omega$, is displayed and drives nothing.

12.3.4 Timetable and Instruments

The timetable allows 275 s from A to B, a dwell of 30 s, and 400 s from B to C, with the first departure 10 s after a case starts. The flat-out plans take 250.8 and 311.7 s, so the runs carry 24.2 and 88.3 s of slack, the second large enough to recover a minute. The doors close at the later of the dwell's end and the timetable's departure, plus any door hold. The controllers read only the instruments of Equation (12.5).

$$\begin{aligned} \hat{v} &= q\big[v(1+\varepsilon) + \sigma_v\, n(t)\min(1, v/v_1)\big], \quad \hat{v} = 0 \text{ at rest} \\ \frac{\mathrm{d}\hat{s}}{\mathrm{d}t} &= v(1+\varepsilon), \qquad \hat{s} \leftarrow s_b + (s - s_b)(1+\varepsilon) \text{ at a balise } s_b \\ \hat{w} &= w(1+\varepsilon) \end{aligned} \tag{12.5}$$

where the tachometer on a trailer axle reads high by the hidden scale error $\varepsilon = 0.3\,\%$, $\sigma_v = 0.03\,\mathrm{m/s}$, $n(t)$ is colored unit noise with a correlation time of 0.5 s drawn from the case's seed, $v_1 = 1\,\mathrm{m/s}$, and q rounds to 0.01 km/h. Balises stand at 1.0, 2.5, 4.0, 6.0, 7.5, 9.0, 10.0 and 11.0 km and 200 and 20 m before the marks at B and C. The meter reads to 0.01 kWh, and the drive's force, the line's receptivity, the signaled restriction and a door hold arrive as signals. The adhesion, the true position and speed and the mass aboard belong to the simulator alone. They are written to log columns whose names begin with `truth_` and drawn dashed as teaching truth, and a test that scans the source fails if the controller or the planner names one of them.

12.3.5 Constants, Starting State and Integration

Table 12.2 collects the constants. A case starts with the train at rest at its first station's mark, with the doors open, $w = F_d = 0$ and every energy integral at zero, and the controller holds the brake at –40 % of the lever while the doors are open. Fixed-step fourth-order Runge-Kutta integration advances the states every 10 ms, the controllers run every 0.1 s, and the log samples at 5 Hz, the three clocks of Section 1.6.1. Hard events, a balise passed, the wet stretch entered, an overspeed or a slip begun and a platform stop, are each handled once, after the integration step in which they occur, and marked on the trends.

12.3.6 Assumptions

The four cars are one rigid mass, so coupler forces and the train's length on a change of grade are absent, and a train 80 m long that meets a grade change over 3.6 s at 80 km/h feels it at once. Resistance has constant coefficients, with no curving term, tunnel factor or wind, and curves appear only as speed limits. One lumped wheelset with a declared creep curve stands for eight motored axles, so neither the cleaning of the rail by the leading axles nor the variation of adhesion along the rail is represented, and a long spin is where the model should be trusted least. The line is a receptivity and a feeder resistance, with no substations, no voltage limit on regeneration and no other trains except through κ. The drive's lag is a first-order teaching value, and the traction and regeneration efficiencies do not change with the operating point. Every speed limit, the 2 km/h overspeed margin and the ± 0.5 m stopping accuracy are teaching values. On a railway they are set by the infrastructure manager's line data, the train protection system and the operator's platform specification.

12.4 The Structure of an Energy-Optimal Run

Written in distance, because the route is, the planning problem asks for the applied force $F(s)$ that spends the least electrical energy between a start and the next stop within a time budget, as shown in

Table 12.2 Parameters. All Are Teaching Values, Adjustable in the Advanced Panel Unless Marked Fixed

Symbol	Meaning	Value	Unit
M, ρ	declared mass, rotating-mass factor	190, 1.08	t, —
φ, m_r/M	motored share, wheelset inertia	0.5, 0.05	—
F_t, P_t	starting effort, traction power	220, 2.8	kN, MW
F_r, P_r, v_f	electric brake, power, fade	200, 3.2, 2	kN, MW, m/s
b_{plan}, $b_{\max}$	planning, full service brake	0.75, 1.0	m/s^2
τ_a	drive lag	1.5	s
η_t, η_r	traction, regeneration efficiency	0.85, 0.80	—
a, b, c	Davis coefficients	13, 0.20, 6.0	see text
μ_{dry}, μ_{wet}	adhesion (hidden)	0.25, 0.08	—
w_{pk}	slip of peak grip (hidden)	0.6	m/s
V_0, R_f	line voltage, feeder resistance	1500, 0.06	V, Ω
ε, σ_v	tachometer error (hidden), noise	0.3, 0.03	%, m/s
Δv_m	plan margin under each limit, fixed	2	km/h
K_c, τ_I	speed loop (shipped default)	10, 20	%/(km/h), s
Δt_m	punctuality margin	4	s
Δt_{rp}	re-plan threshold	5	s
Δs, Δv	DP grid	50, 1	m, km/h
Δt, Δt_c	integration step, control period, fixed	0.01, 0.1	s

Equation (12.6).

$$\begin{aligned} &\min_{F(s)} \int_{s_0}^{s_N} \left(\frac{F^+}{\eta_t} - \hat{\kappa}\,\eta_r F_e \right) ds \\ &\text{s.t. } \tfrac{1}{2}\rho M \frac{\mathrm{d}(v^2)}{\mathrm{d}s} = F - R(v) - Mg\,\theta(s) \\ &\qquad -F_b \le F \le \min\!\left(F_{t,\max}(v),\ \hat{\mu}(s)\varphi Mg\right) \\ &\qquad v \le v_{lim}(s) - \Delta v_m, \qquad \int_{s_0}^{s_N} \frac{ds}{v} \le T \end{aligned} \tag{12.6}$$

where $F_b = b_{plan}\rho M = 153.9\,\mathrm{kN}$ is the planning brake, $\hat{\mu}$ is the adhesion the planner is told, dry unless the wet-rail report says otherwise, $\hat{\kappa}$ is the receptivity it assumes, 1 unless the forecast is on, and T is the timetable's running time less the punctuality margin, or, with re-planning, the time left to the scheduled arrival. The planner never reads the rail's true adhesion, the mass aboard or the model's state.

Attaching a price λ to time and a co-state p to the speed gives the Hamiltonian of Equation (12.7), which is linear in the force as the tank's is in Equation (1.10).

$$H = -\left(\frac{F^+}{\eta_t} - \hat{\kappa}\eta_r F_e \right) - \frac{\lambda}{v} + \frac{p}{\rho M v}\left(F - R(v) - Mg\theta\right) \tag{12.7}$$

The force that maximizes H therefore sits at a bound or at zero, full power when $p/(\rho M v) > 1/\eta_t$, coasting between $\hat{\kappa}\eta_r$ and $1/\eta_t$, and full braking below $\hat{\kappa}\eta_r$, and speed holding is the singular arc on which $p/(\rho M v)$ stays at $1/\eta_t$. On a level line the arc requires $\lambda\eta_t = V^2 R'(V)$, so the price of time sets the holding speed, 77.8 kW for 60 km/h, 177.0 kW for 80 km/h and 317.9 kW for 98 km/h with the parameters of Table 12.2. This is the four-regime structure of the literature [3, 4]. It says what kind of drive to look for and leaves the switching points to be found.

🕮 **Pause.** Try Exercise 12.1 before reading on.

The switching points move with the budget, as shown in Figure 12.2. On the A to B run the one-coast-point plan coasts from 3.99 km at a budget of 252 s, from 2.45 km at 271 s and from 1.91 km at 330 s, and its planned energy falls from 26.81 to 15.46 kWh over that range, against 28.49 kWh planned flat out. The dynamic program plans 25.22, 15.45 and 9.60 kWh at the same three budgets, so its advantage grows from 6 to 38 % as the slack grows. Between 268 and 275 s its energy falls

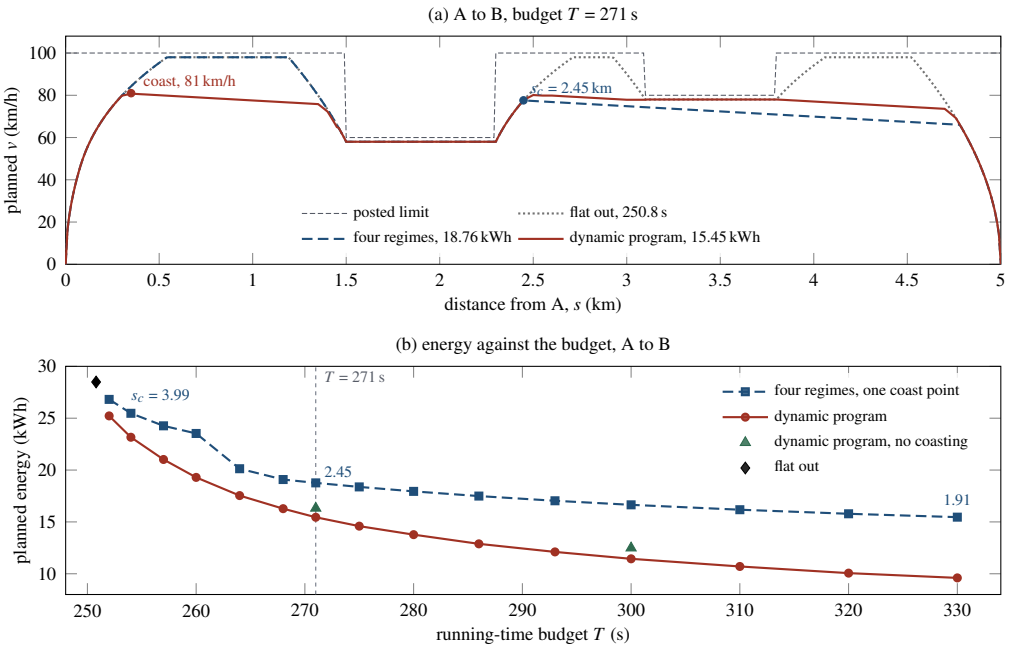


Figure 12.2 Planned drives on the A to B run, from the simulator's own planners. (a) Speed against distance at the 271 s budget. The four-regime plan (dashed) holds 98 km/h, brakes into the 60 km/h curve and coasts from s_c, while the dynamic program (solid) coasts from 0.35 km into the curve. The flat-out plan is dotted and the posted limit thin dashed. (b) Planned energy against the budget, the four-regime plan dashed with squares and its coast point written beside three of its points, the dynamic program solid with circles, the same dynamic program barred from coasting at 271 and 300 s as triangles, and the flat-out plan as a diamond.

by 0.241 kWh per second of budget, 867 kW, against a converged price of 925 kW at 271 s, which is the meaning of λ as the marginal value of time. The gap between the two planners is not a defect of either. The analytic plan holds the line limit and coasts once, while the true optimum lowers its hold speed as λ falls and coasts wherever a limit or the platform would otherwise take the speed away. The no-coasting plans of Figure 12.2(b) were computed by calling the simulator's planner outside its interface, which offers the three planners and not a program barred from coasting. The flat 5 km line shows the other side. There the two planners agree within 0.21 % at 5, 10 and 15 % over the flat-out time, and a plan barred from coasting needs 7.2 to 8.8 % more energy than the four-regime plan.

🕮 **Pause.** Try Exercise 12.2 before reading on.

12.5 Planners and Tracking as Implemented

A plan is made at each departure for the run to the next station, and again whenever the supervisor re-plans. Each plan is a profile on a 5 m grid of speed, force, regime, time and energy that the tracking loop follows.

12.5.1 Flat Out and Four Regimes

The fastest drive applies full traction until the speed meets the braking envelope $v_b(s)$, the highest speed from which the planning brake still meets every lower limit ahead and the stop, and then

follows it, as shown in Equation (12.8).

$$\begin{aligned} \tfrac{1}{2}\rho M \frac{\mathrm{d}(v_b^2)}{\mathrm{d}s} &= -\big(F_b + R(v_b) + Mg\theta\big) \\ v_b(s_N) &= 0, \qquad v_b \le v_{lim} - \Delta v_m \end{aligned} \tag{12.8}$$

integrated backward from the mark by the fourth-order Runge-Kutta method in v^2 with 5 m steps. Following the envelope means holding a limit, or braking along the curve into a lower limit and into the platform. A speed bounded by the distance needed to stop is also how the warehouse robot of Volume 20, Chapter 9 limits its speed. The four-regime plan is the same drive left at a coast point s_c, after which the force is zero except where the envelope caps the speed, and it never powers again. A later coast point always arrives sooner, so s_c is found by bisection until the arrival is within 1 ms of T or after 60 halvings. On both runs of the line at their timetables the worst miss is 0.99 ms.

12.5.2 The Dynamic Program

The run is cut into stages of $\Delta s = 50\,\mathrm{m}$ and the speed into nodes $\Delta v = 1\,\mathrm{km/h}$ apart, 100 stages by 99 nodes from A to B and 140 by 99 from B to C. The cost-to-go is computed backward from the platform, as shown in Equation (12.9).

$$\begin{aligned} V_N(0) &= 0 \\ V_i(v_j) &= \min_{d \in \mathscr{D}} \Big[e_i(v_j, d) + \lambda\, t_i(v_j, d) + V_{i+1}\big(v_i^{+}(v_j, d)\big) \Big] \end{aligned} \tag{12.9}$$

where e_i and t_i are the energy and the time of stage i under decision d, v_i^{+} is the speed at its end, and V_{i+1} is interpolated linearly between nodes and infinite above a limit or above the braking envelope. It is the recursion of Equation (1.11) written in distance instead of time, with time priced by λ, and the same backward recursion made stochastic gives the water value of Chapter 7. The decision set $\mathscr{D}$ holds eleven moves, full traction, hold, coast and the planning brake, traction at three quarters, a half and a quarter, a half brake, power and coast each capped by the braking envelope, and a terminal stop. Every move is integrated over its stage in ten 5 m substeps of the fourth-order Runge-Kutta method in v^2, energy per substep is Fh/η_t for $F \ge 0$ and $-\hat{\kappa}\eta_r \min(-F, F_{r,\max})h$ for braking, and a capped move that would have to brake harder than the planning brake to follow the envelope is refused. The terminal stop is the constant braking force $F_{end} = \tfrac{1}{2}\rho M v^2/d - aM - \tfrac{2}{3}bMv - \tfrac{1}{2}cv^2 - \bar{G}$ that stops the train at the mark, where $\bar{G}$ is the mean grade force over the distance d, allowed between $0.3F_b$ and F_b. The part forces let a switch fall inside a stage, so the energy is not quantized by the stage length. The four regimes are still the only regimes the plans contain, but the program is told neither their order nor how many switches to make.

The budget enters through the price. A forward pass from the measured start takes at each stage the decision that minimizes the bracket of Equation (12.9) at the exact speed reached, and $\ln\lambda$ is bisected between 10 W and 1 GW until the interval is below 10^{-4}, which takes 18 bisections, keeping the lowest price whose plan meets T. Because the decisions are discrete, the plan at the bracketing higher price arrives early and the one at the lower price late, so at the first stage where the two differ the force is blended as $(1-\alpha)F_{hi} + \alpha F_{lo}$, with α bisected 30 times so that the arrival lands on the budget. If even the highest price misses T, the run cannot be made in time. The controller then drives the flat-out profile and reports the timetable impossible only if flat out also misses it. Transitions do not depend on λ and are computed once per run and condition, 81 008 feasible ones from A to B and 119 789 from B to C, so a re-plan reuses them.

12.5.3 Tracking, Stopping and the Slip Override

The speed loop is the velocity-form proportional-integral law of Section 1.6.2 with the plan's force fed forward as its bias, run every $\Delta t_c = 0.1\,\mathrm{s}$ on the measured speed in km/h, as shown in Equation (12.10).

$$\begin{aligned} e_k &= v^*_k - \hat{v}_k, \qquad P_k = K_c e_k \\ \sigma_k &= \sigma_{k-1} + P_k - P_{k-1}, \qquad \hat{u}_k = L_{ff,k} + \sigma_k + I_{k-1} \\ u_{cmd,k} &= \min\bigl(100, \max(-100, \hat{u}_k)\bigr) \\ I_k &= I_{k-1} + \Bigl[\frac{K_c}{\tau_I} e_k + \frac{u_{cmd,k} - \hat{u}_k}{\tau_I}\Bigr]\Delta t_c \end{aligned} \tag{12.10}$$

where the reference v^* is the plan's speed $v_{ref}(\hat{s})$, capped by the braking envelope of the limits in force now, and on a coasting stretch it equals $\hat{v}$ while the error is within 2 km/h, so that a coast stays a coast. The bias L_{ff} is the plan's force at $\hat{s} + \hat{v}\tau_a$, read one drive lag ahead and expressed as a lever, and the lever maps to the command force as $F_{cmd} = (u_{cmd}/100)F_{t,\max}(\hat{v})$ in power and $(u_{cmd}/100)\,b_{\max}\rho M$ in braking, with the declared mass. The loop is reverse acting, so a measured speed below the reference raises the lever. The running sum and the back-calculation are those of Equations (1.15) to (1.17), with the tracking time τ_I, and $\tau_I = 0$ removes integral action.

The stopping controller takes over where the plan's final braking begins, or earlier once $\hat{v}^2/(2d)$ reaches 85 % of b_{plan} within 1500 m of the mark. It feeds forward the braking force $F_S = \rho M\hat{v}^2/(2d) - R(\hat{v}) - Mg\theta$ through the declared mass and runs the same loop against the curve $\sqrt{2b_{plan}d}$, where $d = s_{mark} - \hat{s}$. A train at rest short of the mark by more than 10 m creeps in at 15 % of the lever, and the doors open after 0.4 s at rest within ± 10 m.

The slip override acts on the drive's cap every control period, as shown in Equation (12.11).

$$F_{cap,k} = \begin{cases} \max\bigl(F_{lo},\ 0.6\min(|F_m|, F_{cap,k-1})\bigr) & \hat{w} > w_{on} \\ \min\bigl(F_t,\ F_{cap,k-1} + r\Delta t_c\bigr) & \hat{w} < w_{off} \end{cases} \tag{12.11}$$

where $w_{on} = 0.8\,\mathrm{m/s}$, $w_{off} = 0.4\,\mathrm{m/s}$, $F_{lo} = 5\,\mathrm{kN}$, $r = 10\,\mathrm{kN/s}$, and the cap is removed once it recovers to F_t, with the mirror rule on $F_{cap,b}$ for negative slip in electric braking. The cap acts on the motor force directly, not through the lag, and while it binds the speed loop tracks the capped lever so that its integral does not wind up. The force therefore hunts about the rail's peak.

12.5.4 Re-Planning and the Three Models

Every control period the supervisor projects the arrival as the time now plus the plan's time still to run from $\hat{s}$, and compares it with the timetable's arrival less the margin. When the difference Δt exceeds the threshold Δt_{rp}, or the restriction changes, it plans again from the next 50 m node with the measured speed and the time left, at most once every 20 s, not within 400 m of the mark and not when the plan in force is already flagged impossible and late. The note it writes states the energy the new plan spends beyond the old plan's remainder. Three models stay apart. The plant is everything in Section 12.3. The planner's model is Equation (12.6) with the declared train, a dry rail unless told and no drive lag. The stopping controller uses the declared mass, so a heavier train than declared is corrected only by feedback.

12.6 Nine Cases and a Challenge, Measured

Table 12.3 lists the cases with the rule that completes each. Every result below is a headless replay of the train engine on its default seed, in the 10 ms steps a browser tab would take. A setting that a sentence names is applied before the first step, which primes the run exactly as Reset does (Section 1.6.1). A case is credited for the reader's actions together with the state the simulated train reaches, a scripted event earns no credit, and a parameter, wet-rail or receptivity change during the challenge voids the attempt. The three ways of reaching a case (loading it first, picking it after another

Table 12.3 The Nine Cases and the Challenge, With What Completes Each

#	Case, settings and events	Teaching point	Required to complete
1	By hand, A to B, Manual, lever in full power	the clock and the meter	Leave A in Manual, stop within 20 m of the B mark, and confirm reading the arrival and the energy.
2	Flat out, A to B	the fastest run and its cost	Run flat out to B and confirm reading drawn, returned and net energy.
3	Four regimes, A to B	one coast point	Run the four-regime plan to B, choose Flat out and Reset, and run it to B.
4	Four regimes, B to C on a 330 s express path	hold uphill, coast over the top	Watch the plan hold up the climb, switch to coasting between 7.0 and 8.2 km, and run to C.
5	Dynamic program, A to B, Grid view	where limits interleave	Run the program to B, choose Four regimes and Reset, and run it to B.
6	Dynamic program, $K_c = 1$, no integral, 40 t unplanned load	tracking and the platform mark	Raise K_c to 8 or more with integral action and stop within ± 0.5 m of the B mark.
7	B to C, override off, wet top of the climb, restriction	slip and what the plan knew	Spin the wheels past 5 m/s with the override off, turn it on, report the wet stretch, Reset, and reach C within ± 15 s with the wet slip under 2 m/s.
8	B to C, doors held 60 s, re-planning off	the price of a minute	Arrive late by more than 15 s, turn re-planning on, Reset, and arrive within ± 15 s.
9	B to C, restriction, doors held 90 s	a timetable no plan can meet	Watch the planner report the timetable impossible while running, run to C, and confirm the late arrival.
10	Challenge, A to C, wet top of the climb, doors held 25 s at B, restriction announced at 330 s	the whole stack	Both arrivals within ± 15 s, both stops within ± 0.5 m, no overspeed, and the meter at C below 75 % of the flat-out run on this case, 55.1 kWh of 73.4 kWh, with the case's train, line and events.[a]

[a] Shipped with the four-regime plan, $K_c = 1$, no integral, no re-planning and no slip override, it stalls on the wet climb. The suggested settings are the dynamic program, $K_c = 10$, $\tau_I = 20$ s, re-planning, the override and the wet-rail report for the top of the climb.

case, and Reset part-way through) give logs with identical hashes for all ten, the application's suite passes 43 of 43 checks, and a run paced at 1 × and at 30 × by wall-clock frames ends bit-identical.

12.6.1 Coasting, Limits and a Grade

Driven by hand with the lever left in full power, the train of case 1 enters the 60 km/h curve 63.0 km/h over its limit. A cautious drive, full power to 25 s, the lever at 3 % to 320 s and then −60 %, never exceeds 58.6 km/h and arrives 62.9 s late and 7.7 m short with 10.5 kWh. It saves energy by spending the passengers' time. Flat out in case 2 draws 58.9 kWh and returns 30.7 kWh, 75 % of the braking work at the rail, for a net 28.2 kWh, and the 21.0 s it gains buy nothing because the train waits at B for the timetable. The coast point of case 3 and the dynamic program of case 5 reproduce the comparison of the Introduction, 19.5 kWh against 16.4 kWh when driven, 16 % less, as listed in Table 12.4.

Table 12.4 Measured Case Results. Arrival Against the Timetable (+ Late), True Stopping Error (+ Past the Mark) and Energy at the Mark

#	Configuration	Arrival (s)	Stop (m)	Energy (kWh)
1	cautious hand drive	B +62.9	−7.7	10.5
2	flat out	B −21.0	−0.05	28.2
3	four regimes, s_c 2.45 km	B −3.1	−0.02	19.5
4	four regimes, s_c 6.96 km	C +0.6	−0.02	32.7
	flat out	C −14.7	—	43.6
	margin 8 s, s_c 8.35 km	C −3.3	—	35.9
5	dynamic program	B −1.7	−0.01	16.4
6	$K_c = 1$, no integral	B +6.6	+6.86	17.6
	$K_c = 10$, $\tau_I = 20$ s	B −0.5	+0.11	19.6
	$K_c = 10$, no integral	—	+0.56	—
7	override off	stalls at 7.10 km, slip 34.6 m/s		
	override on	C −0.6	0.00	30.8
	override and wet report	C −1.2	+0.02	29.2
8	no re-planning	C +60.8	−0.06	26.1
	re-planning	C −0.6	−0.00	30.5
9	any planner, margin or gain	C +42.4	−0.03	47.2

TRY IT IN THE APP **3 · The Coast Point: One Decision, and What It Saves**

Load the scenario and press Run. The train leaves A under full power on the four-regime plan. Watch the regime change to coasting 2.45 km from A, after which the train never powers again before B. It arrives 3.1 s early with 19.5 kWh on the meter.

Case 4 runs B to C on a 330 s express path. Holding 98 km/h up the 2.0 % grade takes 45.2 kN of the 102.9 kN the motors give at that speed, and the plan coasts from 6.96 km, 443 m below the summit, so the train crosses it at 84 km/h and the descent restores the speed with no power. Coasting at 80 km/h the train loses 0.212 m/s^2 up the grade and gains 0.151 m/s^2 down it. It arrives 0.6 s late with 32.7 kWh against 43.6 kWh flat out. An 8 s margin shrinks the budget by 4 s and moves the coast point past the summit to 8.35 km, for 35.9 kWh, while on the ordinary 400 s timetable the coast point falls at 6.04 km, near the foot of the climb. Figure 12.3 draws the energies and arrivals of Table 12.4 side by side.

12.6.2 Tracking a Heavier Train Than Planned

In case 6, 40 t more than the declared mass boards at the first second. With $K_c = 1$ % per km/h and no integral action the train stops 6.86 m past the B mark and 6.6 s late, because the stopping feedforward brakes for 190 t with 230 t aboard. With $K_c = 10$ and $\tau_I = 20$ s it stops 0.11 m past and 0.5 s early, and with $K_c = 10$ and no integral action 0.56 m past, outside the declared accuracy. Proportional action leaves an error in proportion to the missing force, and the integral removes it.

Pause. Try Exercise 12.3 before reading on.

12.6.3 A Wet Climb

Case 7 leaves B with the restriction in force and the rail wet over the upper climb, which the plan does not know. Accelerating out of the restriction with no override, the motors can deliver 168 kN at 60 km/h where the rail gives 74.6 kN, the wheels spin up to 34.6 m/s of slip at the motors' top speed, and the train stalls at 7.10 km. With the override on, the largest slip on the wet stretch is 1.89 m/s, the train climbs at the rail's limit, falls behind its dry-rail plan, re-plans three times and arrives 0.6 s early with 30.8 kWh. With the wet-rail report as well the plan expects 74.6 kN there, the slip peaks at 1.57 m/s, one re-plan suffices and the arrival is 1.2 s early with 29.2 kWh. Over six seeds the override alone holds the wet slip between 1.38 and 2.37 m/s, above the declared 2 m/s at one seed, and with the report at or below 1.72 m/s. The peak comes within one 0.1 s control period of the wet entry.

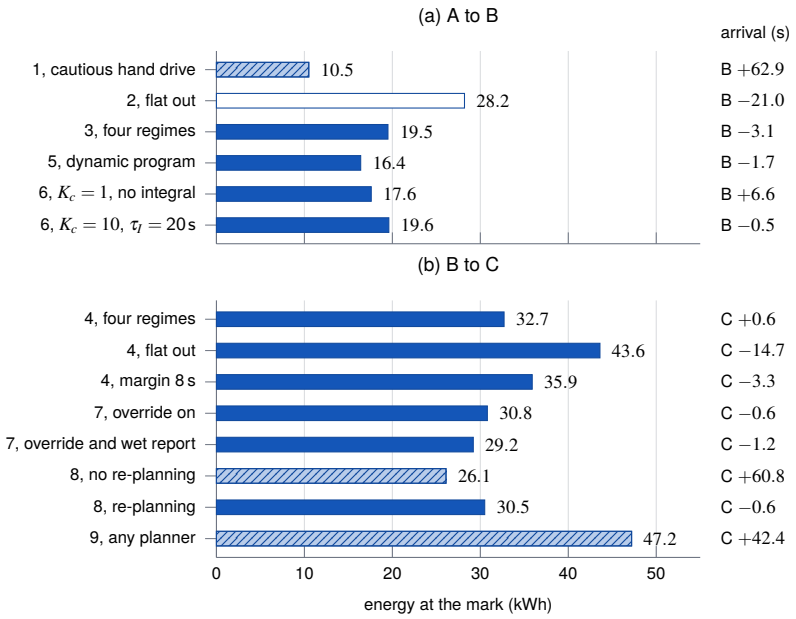


Figure 12.3 The measured cases of Table 12.4: energy at the mark (bars, with values) and arrival against the timetable (right, + late) for (a) the runs from A to B and (b) the runs from B to C. Dark bars arrive within 15 s of the timetable, the open bar more than 15 s early and the hatched bars more than 15 s late. Notice that on each leg the cheapest run is a late one, and that among the runs on time from A to B the dynamic program spends the least.

> **TRY IT IN THE APP** **7 · The Wet Rail: Slip, the Override, and What the Plan Knew**
>
> Load the scenario and press Run. The slip override is off, and the top of the climb is wet, which the plan does not know. Watch the wheel slip as the train accelerates out of the restriction: the wheels spin up to 34.6 m/s of slip, and the train stalls at 7.10 km.

12.6.4 A Minute Late, and a Timetable No Plan Can Meet

In case 8 the doors at B are held 60 s. Without re-planning the train runs its on-time plan and arrives 60.8 s late with 26.1 kWh. Re-planned at departure with a 336 s budget it arrives 0.6 s early with 30.5 kWh, against 26.1 kWh for the same run leaving on time, so the minute costs 4.4 kWh, 17 % more. Case 9 adds the restriction to a 90 s hold. The fastest plan from B to C with the restriction takes 348.5 s, the timetable leaves 310 s, and the planner reports the timetable impossible and hands the loop the flat-out profile. The train arrives 42.4 s late with 47.2 kWh, and a logged sweep over the three planners, margins of 0 and 10 s and gains of 10 and 20 finds no arrival earlier than 42.4 s late. The flag switches between door holds of 47 and 48 s, where flat out stops meeting the budget.

12.6.5 The Challenge

The challenge runs A to C with the wet top of the climb, a 25 s door hold at B and the restriction announced at 330 s, while the train stands at B. Of the 88.3 s of slack on the B to C timetable, the hold takes 25 s, the restriction 36.8 s and the wet rail, once the plan knows it, 8.4 s more. As shipped the train reaches B 0.6 s early and stalls on the wet climb at 7.10 km on six of six seeds. The suggested strategy passes on six of six, as listed in Table 12.5, with C reached between 2.8 and 1.4 s early, 51.7 to 52.0 kWh and both stops within 0.08 m. Removing any one element fails the run. The four-regime plan meets the clock and spends 58.8 kWh, the run without the wet-rail report arrives 15.1 s late with

Table 12.5 The Challenge, Default Seed. Arrival (s) and Stopping Error (m) at Each Station. Energy at C, or at 760 s for Stalled Runs

Configuration	B	C	kWh	Verdict
shipped	−0.6, −0.07	stall 7.10 km	61.3	fails
suggested	−1.7, −0.01	−1.4, −0.04	51.8	passes
flat out, rest suggested	−21.0, −0.05	−10.6, −0.06	73.4	fails
four regimes, rest suggested	−3.1, −0.02	+0.9, −0.03	58.8	fails
suggested, no wet report	−1.7, −0.01	+15.1, −0.00	57.9	fails
suggested, no re-planning	−1.7, −0.01	+28.7, −0.01	44.6	fails
suggested, no override	−1.7, −0.01	stall 7.08 km	55.4	fails

Table 12.6 Code Verification

Check	Result
meter against $\int P_{el}\,dt$, every 10 ms, A to B	1.7×10^{-7}
work balance, challenge suggested, wet, hold, restriction	1.3×10^{-5}
work balance, case 7 spin and stall	2.6×10^{-5}
free rolling against $R/\rho M$, 40 to 100 km/h	1.6×10^{-9}
coasting times against the closed form, three bands	$\leq 1.3 \times 10^{-4}$
halving Δt, A to C energy	0.007 %
grid 50 m by 1 km/h against 25 m by 0.5 km/h	0.05 %
priced DP against enumeration, three prices	$\leq$ 0.316 %
DP without blend against constrained enumeration	0 to 5.42 %
DP against four regimes, flat 5 km line	$\leq$ 0.21 %
DP fastest plan against flat out	$\leq$ 0.19 s
four-regime arrival against T	$\leq$ 0.99 ms

57.9 kWh, the run without re-planning arrives 28.7 s late, and the run without the override stalls at 7.08 km. The suggested run on the wet climb is shown in Figure 12.4.

> ✦ DID YOU KNOW? **The Cheapest Challenge Run Fails**
>
> From the simulator: in the challenge, the suggested settings without re-planning reach C with 44.6 kWh, the least energy of every run in Table 12.5, and still fail, because the train arrives 28.7 s late. The suggested run spends 51.8 kWh and arrives 1.4 s early. The challenge accepts an arrival only within ±15 s (Table 12.3).

12.6.6 Code Verification

The checks of Table 12.6 test the implementation against itself and against closed forms, and none of them compares the model with a railway. The meter equals the integral of the logged power to 1.7×10^{-7}. Over the suggested challenge the work the motors do at the wheels, 27.6 kWh, equals the change of kinetic energy plus the slip loss of 1.85 kWh, the work against resistance and friction, and the gravity term of 6.213 kWh, which matches $Mg\Delta h$, to a residual of 1.3×10^{-5}. Halving the step to 5 ms moves a whole A to C trip's energy by 0.007 % and keeps every challenge verdict, the largest case change being an arrival of case 7 that moves from 0.6 to 1.6 s early. Refining the grid from 100 m by 2 km/h to 50 m by 1 km/h and to 25 m by 0.5 km/h gives 15.589, 15.449 and 15.441 kWh for the A to B plan at 271 s, changes of 0.91 and 0.05 %. On a 600 m instance with a grade and a limit, the priced policy matches an enumeration of all 44 244 sequences of the four regimes and the stop exactly at two prices and within 0.316 % at a third. On a 300 m instance with all eleven decisions and the budget as a hard constraint, the program without its blend lies 0.00 to 5.42 % above the enumerated optimum over ten budgets, and with the blend at or below it, because the blend uses a force the enumeration lacks. A plan with its transitions takes 163 to 224 ms in Node.js, and two timed runs of the challenge at 30 × in a desktop browser took 24.2 s of wall time each, with plans of 156 to 239 ms and re-plans of 65 to 84 ms.

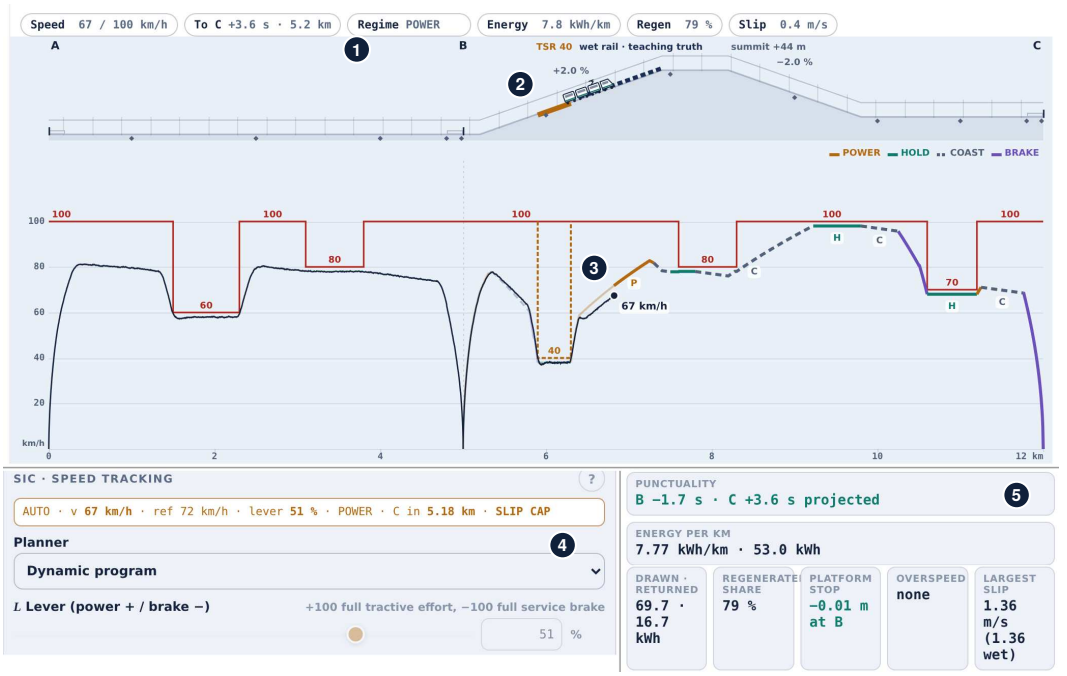


Figure 12.4 Suggested settings in the challenge, $t = 470$ s, on the Line view. The figure joins three parts of the screen, the Line view with its badges (top), the speed-tracking card (bottom left) and four tiles of the Learn panel (bottom right), and leaves out the header, the rest of the control panel, the trends and the other tiles. In this run the train left B 26.1 s after the timetable. (1) The badges read 67 km/h, C projected 3.6 s late, the regime and the regenerated share. (2) The train ran through the 40 km/h restriction from 5.9 to 6.3 km and climbs the wet stretch. (3) The dynamic program's plan to C lies ahead under the limits, lettered by regime, and the run lies behind. (4) The speed-tracking card reads 67 km/h against a reference of 72 km/h, the lever at 51 % and SLIP CAP. At 6.82 km the slip override holds the motor force at 71.0 kN under the wet rail's 74.6 kN, with a measured slip of 0.42 m/s (Table 12.1). (5) The tiles read B 1.7 s early and 0.01 m short, C projected 3.6 s late, 53.0 kWh on the meter with 69.7 drawn and 16.7 kWh returned, 79 % of the braking work returned, and a largest slip of 1.36 m/s. A re-plan follows at 492.7 s, and the train reaches C 1.4 s early with 51.8 kWh (Table 12.5).

12.6.7 Agreement With the Literature

The agreement claimed with published work is structural only. Every plan the dynamic program returns is made of the four regimes that the maximum principle predicts [3, 4], and on a flat 5 km line a direct search over power, hold at 50 to 98 km/h, coast and a final brake of 0.3 to 1.0 times the planning brake finds plans within 0.26 % of both planners at 5, 10 and 15 % over the flat-out time. No number in the model was fitted to measured trains, and the Davis coefficients are teaching values in the quadratic form [18], not a measured curve. The energy savings of 10 to 30 % reported for planned driving in [8] are against a baseline that is not the flat-out run used here, so the 31 % of case 3 is not compared with them.

12.7 Technology in Practice

What Measures It

Odometry, the onboard function that estimates a train's speed and travelled distance by dead reckoning, is the base of automatic train protection and control. Besides wheel angular speed sensors it can use

Doppler radar, accelerometers and gyroscopes, and a single sensor is reliable only over a limited range of conditions. Satellite positioning is held back on railways by lines that are not georeferenced and by signal outages, in tunnels for example. European train control requirements set the speed error and the distance error that an odometry algorithm may make [19]. In this chapter the controllers read only the tachometer, the odometer with its balise resets, the measured slip and the signals of the line (Equation (12.5)), and the energy meter is logged.

> **TECHNOLOGY** **Sensor: Tachometers and Balises for Speed and Position**
>
> Onboard odometry estimates a train's speed and travelled distance by dead reckoning, usually from wheel angular speed sensors that count pulses per revolution. In pure rolling the speed is the wheel radius times the angular speed, but wheels wear and are periodically returned, so the radius may not be reliable. When adhesion is degraded and the train accelerates or brakes, the wheels slide: a driven wheel runs faster than the train and a braked one slower, and the error may become very high. The estimate of distance is reset each time the train passes a balise [19]. The chapter's tachometer sits on a trailer axle and reads 0.3 % high, and twelve balises reset its odometer (Equation (12.5)).

What Moves It

Electric train operation divides into four modes, acceleration, cruising, coasting and braking. On urban lines, where stations are close together, cruising is typically omitted: trains accelerate rapidly to their top speed and brake shortly afterwards for the next stop [5]. The same wheel angular speed sensors serve the wheel slide protection and anti-skid systems that act at the wheels when adhesion is poor [19]. In this chapter the drive answers the lever with a 1.5 s lag, the electric brake works through the motored wheelsets, and the friction brake on the trailers takes the rest (Equation (12.3)).

> **TECHNOLOGY** **Final Element: Traction Motors That Also Brake**
>
> In regenerative braking, common in electric rail systems, a train slows by reversing the operation of its motors, which then act as generators. The energy they produce supplies the train's own auxiliary loads, and the surplus is fed back to the line that supplies the trains, where a nearby train that is accelerating can use it. With no train near, the line voltage rises toward an overvoltage limit, the braking train may be unable to return its energy, and the excess is burned as heat in onboard or wayside resistors [5]. The receptivity κ of Equation (12.4) is the share the line accepts, and the friction brake on the trailers takes what the electric brake cannot.

Where the Controller Runs

An automatic train operation system first generates a target speed curve from the condition of the train and of the line, and then sends speed control commands so that the train tracks it [20]. Under it a train control system enforces the speeds that protection allows [21]. No open source read for this chapter gives the scan period of an onboard computer. In this chapter the planner solves a plan in 156 to 239 ms in a browser, and the supervisor, the speed loop, the stopping controller and the slip override run every 0.1 s (Figure 12.5), where a real drive's slip control runs in milliseconds (Section 12.8).

> **IN PRACTICE** **Protection Sits Under the Plan**
>
> On railroads in the United States, a federal rule requires each positive train control system to reliably prevent overspeed derailments, including those related to civil engineering speed restrictions and slow orders. An automatic train control system must apply the brakes automatically when the train's speed exceeds the rate its speed control mechanism sets [21]. Protection depends on odometry, and an error in position may overestimate the distance available for braking. One national protection system is forced to overestimate speed for safety, so its emergency braking occurs more often than required [19]. The plans of this chapter keep 2 km/h below each limit (Table 12.2).

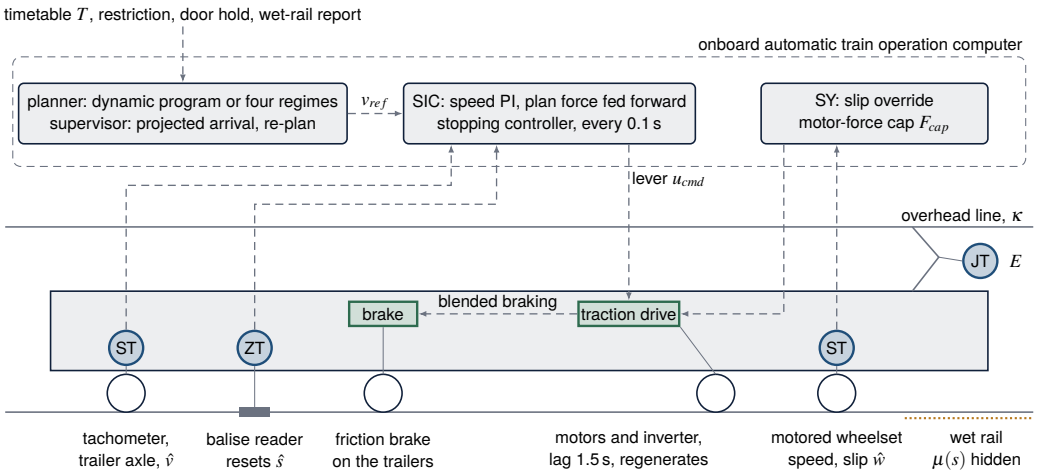


Figure 12.5 The chapter's control layers drawn with real devices. The onboard computer plans the run from the timetable and the signals it is given, the speed loop (SIC) tracks the plan with the lever of the traction drive, and the slip override (SY) caps the motor force when the motored wheelsets slip. The tachometer (ST) on a trailer axle gives $\hat{v}$, the balise reader (ZT) resets the odometer $\hat{s}$, a speed sensor on the motored wheelsets gives the slip $\hat{w}$, and the energy meter (JT) at the pantograph integrates the power drawn and returned. Notice that no instrument reads the rail's adhesion: the plan learns of a wet rail only from a report.

What Is Modernizing It

Three developments bear on this chapter's method. The first is odometry that fuses tachometers with inertial measurement units through Kalman filters, aimed at accuracy in critical adhesion conditions. Its preliminary results estimated position and speed markedly better than an algorithm in service [19]. The second is the recovery of braking energy that the line cannot take. Wayside energy storage has reported savings of up to 30 percent in service, and reversible substations feed regenerated energy back to the AC grid [5]. The third is learning. Algorithms built from the records of experienced subway drivers and refined by reinforcement learning used less energy than manual driving and existing automatic operation in simulations with field data from one line [20]. In this chapter the receptivity forecast and the dynamic program stand where such storage and such learning would enter.

12.8 Limitations

The train, the line and the timetable are teaching values that make a consistent run, not data from any railway, and a simulated pass of the challenge is not evidence about any real train's punctuality or energy. The point mass, the single lumped wheelset and the constant-coefficient resistance leave out coupler forces, the train's length on a grade, curving resistance and the way adhesion recovers along a train, and the line has no substations and no other trains except through the receptivity.

> ⚠ **OPEN CHALLENGE** **Automatic Operation Still Hands Back to the Driver**
>
> A study of automatic train operation on subways listed problems that remain unsolved. Most work weighs one or two objectives, although safety, punctuality, energy and riding comfort all count. The automatic system tracks a speed profile optimized before the run, which limits its flexibility, and when the trip time or the track's resistance changes, the subway transfers to manual driving. The study's answer adapts to such changes without an offline speed profile [20]. The supervisor of this chapter re-plans for a late departure (Section 12.5).

The slip override runs at the controller's 10 Hz, where a real drive's slip control runs in milliseconds, and on case 7 without the wet-rail report its peak slip reaches 2.37 m/s at one of six seeds, above the declared 2 m/s. The dynamic program is exact on its grid only. Its switches fall on 50 m stages, its cost-to-go is linear between nodes, and a budget that falls between two prices is met by blending one stage rather than by an exact constrained optimum, a duality gap that the blend closes only approximately. The four-regime planner holds each line limit and coasts once, which is a restricted family rather than the full Pontryagin solution with a hold speed set by λ, and its distance from the dynamic program on the A to B run measures that restriction as much as the line. In the challenge the restriction is announced while the train stands at B, so the plan to C is solved twice, at the arrival and at the departure, and each solve costs a visible frame of about 0.2 s. The receptivity forecast and the punctuality margin are the only economics, with no price of energy that varies in time.

> **OPEN CHALLENGE** **Knowing the Rail Before the Wheels Reach It**
>
> Most slip control on trains acts after a wheel has begun to slide, so damage to wheels and track cannot be avoided. Predictive slip controls have assumed constant traction over the prediction horizon and fail to avoid sliding when traction changes fast. Rain mixing with dust into a muddy film on the rails is one cause of poor traction. One proposal records traction conditions in a map along the line and lets a predictive controller read them ahead, and it names the online adaptation of that map from measurements as future work [22]. In case 7 the plan learns of the wet rail only from a report.

12.9 Conclusion

This chapter studied a four-car electric train on a declared 12 km line, a plant on which the energy of a run is decided by where the regime changes. Flat out between A and B costs 28.2 kWh for an arrival 21.0 s earlier than the timetable needs. One coast point brings it to 19.5 kWh on time, and a dynamic program on a 50 m by 1 km/h grid, which coasts from 0.35 km into a 60 km/h curve, to 16.4 kWh. The program's advantage over the one-coast-point plan grows from 6 to 38 % of planned energy as the budget grows from 252 to 330 s, and on that run a drive barred from coasting beats the one-coast-point plan by 13 %, while on a level line it loses by 7.2 to 8.8 %. A minute's delay recovered by re-planning costs 4.4 kWh, a slip override at 10 Hz keeps a train climbing a wet rail that otherwise stalls it at 7.10 km, and a timetable left 38.5 s short of the fastest possible run is reported impossible rather than planned. The model is a point mass with a lumped wheelset and every constant is a teaching value. A fuller model would add a train of several masses, a substation network that decides the receptivity, an energy price that varies in time, a drive-level slip controller, and a plan over several stations. The simulator shows that a timetable's slack is an energy budget whose best use depends on the line.

KEY POINTS

- A timetable's spare seconds are an energy budget, and the best way to spend them depends on where the limits of the line fall.
- From A to B, flat out costs 28.2 kWh, one coast point 19.5 kWh and a dynamic program 16.4 kWh, and the last two arrive a few seconds early.
- Pontryagin's principle limits an energy-optimal run to full power, speed holding, coasting and full braking, and the price of time sets the holding speed.
- Feedback finishes the plan: integral action stops a heavier train at the mark, a slip override keeps it climbing a wet rail, and re-planning recovers a minute for 4.4 kWh.

Where else this applies. The anti-lock brake of Volume 7, Chapter 4 regulates the slip of a braking tire at the same kind of peak that the slip override here hunts about.

References

[1] A. Sciarretta, G. De Nunzio, and L. L. Ojeda, "Optimal ecodriving control: Energy-efficient driving of road vehicles as an optimal control problem," *IEEE Control Systems*, vol. 35, no. 5, pp. 71–90, 2015.

[2] R. E. Bellman, *Dynamic Programming*. Princeton, NJ: Princeton University Press, 1957, reprinted 2010.

[3] P. Howlett, "The optimal control of a train," *Annals of Operations Research*, vol. 98, no. 1–4, pp. 65–87, 2000.

[4] A. Albrecht, P. Howlett, P. Pudney, X. Vu, and P. Zhou, "The key principles of optimal train control—Part 1: Formulation of the model, strategies of optimal type, evolutionary lines, location of optimal switching points," *Transportation Research Part B: Methodological*, vol. 94, pp. 482–508, 2016.

[5] M. Khodaparastan, A. A. Mohamed, and W. Brandauer, "Recuperation of regenerative braking energy in electric rail transit systems," *IEEE Transactions on Intelligent Transportation Systems*, vol. 20, no. 8, pp. 2831–2847, 2019.

[6] P. G. Howlett and P. J. Pudney, *Energy-Efficient Train Control*, ser. Advances in Industrial Control. London: Springer, 1995.

[7] A. Albrecht, P. Howlett, P. Pudney, X. Vu, and P. Zhou, "The key principles of optimal train control—Part 2: Existence of an optimal strategy, the local energy minimization principle, uniqueness, computational techniques," *Transportation Research Part B: Methodological*, vol. 94, pp. 509–538, 2016.

[8] R. Franke, P. Terwiesch, and M. Meyer, "An algorithm for the optimal control of the driving of trains," in *Proceedings of the 39th IEEE Conference on Decision and Control*, vol. 3, Sydney, NSW, Australia, 2000, pp. 2123–2128.

[9] H. Ko, T. Koseki, and M. Miyatake, "Numerical study on dynamic programming applied to optimization of running profile of a train," *IEEJ Transactions on Industry Applications*, vol. 125, no. 12, pp. 1084–1092, 2005.

[10] J. Yin, T. Tang, L. Yang, J. Xun, Y. Huang, and Z. Gao, "Research and development of automatic train operation for railway transportation systems: A survey," *Transportation Research Part C: Emerging Technologies*, vol. 85, pp. 548–572, 2017.

[11] I. A. Asnis, A. V. Dmitruk, and N. P. Osmolovskii, "Solution of the problem of the energetically optimal control of the motion of a train by the maximum principle," *USSR Computational Mathematics and Mathematical Physics*, vol. 25, no. 6, pp. 37–44, 1985.

[12] K. Ichikawa, "Application of optimization theory for bounded state variable problems to the operation of train," *Bulletin of JSME*, vol. 11, no. 47, pp. 857–865, 1968.

[13] E. Khmelnitsky, "On an optimal control problem of train operation," *IEEE Transactions on Automatic Control*, vol. 45, no. 7, pp. 1257–1266, 2000.

[14] R. Liu and I. M. Golovitcher, "Energy-efficient operation of rail vehicles," *Transportation Research Part A: Policy and Practice*, vol. 37, no. 10, pp. 917–932, 2003.

[15] G. M. Scheepmaker, R. M. P. Goverde, and L. G. Kroon, "Review of energy-efficient train control and timetabling," *European Journal of Operational Research*, vol. 257, no. 2, pp. 355–376, 2017.

[16] A. González-Gil, R. Palacin, and P. Batty, "Sustainable urban rail systems: Strategies and technologies for optimal management of regenerative braking energy," *Energy Conversion and Management*, vol. 75, pp. 374–388, 2013.

[17] O. Polach, "Creep forces in simulations of traction vehicles running on adhesion limit," *Wear*, vol. 258, no. 7–8, pp. 992–1000, 2005.

[18] B. P. Rochard and F. Schmid, "A review of methods to measure and calculate train resistances," *Proceedings of the Institution of Mechanical Engineers, Part F: Journal of Rail and Rapid Transit*, vol. 214, no. 4, pp. 185–199, 2000.

[19] M. Malvezzi, G. Vettori, B. Allotta, L. Pugi, A. Ridolfi, and A. Rindi, "A localization algorithm for railway vehicles based on sensor fusion between tachometers and inertial measurement units," *Proceedings of the Institution of Mechanical Engineers, Part F: Journal of Rail and Rapid Transit*, vol. 228, no. 4, pp. 431–448, 2014.

[20] K. Zhou, S. Song, A. Xue, K. You, and H. Wu, "Smart train operation algorithms based on expert knowledge and reinforcement learning," *IEEE Transactions on Systems, Man, and Cybernetics: Systems*, vol. 52, no. 2, pp. 716–727, 2022.

[21] U.S. Federal Railroad Administration, "Rules, standards, and instructions governing the installation, inspection, maintenance, and repair of signal and train control systems, devices, and appliances," Code of Federal Regulations, Title 49, Part 236 (October 1, 2024 edition), U.S. Government Publishing Office, 2024. [Online]. Available: https://www.govinfo.gov/content/pkg/CFR-2024-title49-vol4/pdf/CFR-2024-title49-vol4-part236.pdf

[22] M. Hauck, P. Schmidt, A. Kobelski, and S. Streif, "A map-based model predictive control approach for train operation," in *2023 European Control Conference (ECC)*, Bucharest, Romania, 2023, pp. 1–6.

12.10 Exercises

12.1 Concept check · 2 min

The dynamic program of this chapter may choose among eleven moves at every stage, including traction at three quarters, a half and a quarter. Yet every plan it returns is made of full power, speed holding, coasting and braking. Why?

(a) The Hamiltonian is linear in the force, so the best force sits at a bound or at zero except on the arc that holds a speed

(b) The grid of 1 km/h is too coarse to resolve a part force, so the program rounds each one to a full regime

(c) The tractive-effort curve is flat below the base speed, so a part force costs the same energy as full power there

(d) The drive's lag of 1.5 s filters out a part force, so the tracking loop can only follow the four regimes

12.2 Concept check · 1 min

At the 271 s budget from A to B, the one-coast-point plan holds 98 km/h and plans 18.76 kWh, while the dynamic program barred from coasting holds about 78 km/h and plans 16.32 kWh. On the level 5 km test line the no-coasting plan loses by 7.2 to 8.8 %. What reverses the order on the A to B run?

(a) The dynamic program's grid is finer than the bisection that places the single coast point

(b) The fast plan must brake into the 60 km/h curve at 1.5 km and throws away speed it paid for

(c) Regenerative braking returns a larger share of the energy at low speed than at high speed

(d) Davis resistance is lower on the A to B run because that run has no grade

12.3 Concept check · 2 min

In Scenario 6, 40 t more than the declared 190 t boards at the first second. With $K_c = 10$ and no integral action the train stops 0.56 m past the B mark, outside the 0.5 m accuracy. With $\tau_I = 20$ s added it stops 0.11 m past. Why does the gain alone not suffice?

(a) The tachometer reads 0.3 % high between balises, and only an integral can correct a scale error

(b) The drive's lag of 1.5 s delays the brake, and only an integral can anticipate a lag

(c) The stopping feedforward brakes for 190 t, and a proportional loop needs a standing error to supply the missing force

(d) The gain of 10 makes the loop oscillate about the braking curve, and the integral adds the damping that stops it

12.4 Theory and calculation · 6 min

Use Table 12.2 with $M = 190\,\text{t}$ and $g = 9.81\,\text{m/s}^2$. Davis resistance is $R(v) = aM + bMv + cv^2$ with v in m/s, and the motors give at most $\min(F_t, P_t/v)$.

(a) Compute R at 60 and 100 km/h, and the deceleration of a train coasting on the level at 100 km/h, with the inertia ρM. *Checkpoint:* 4.77 kN, 8.16 kN, $0.040\,\text{m/s}^2$.

(b) Compute the force that holds 98 km/h up the +2.0 % climb and the most the motors give at that speed. *Checkpoint:* 45.2 kN of 102.9 kN.

(c) Compute the deceleration of a train coasting at 80 km/h up the climb and its acceleration coasting down the −2.0 % descent. Why does case 4 coast before the summit rather than after it? *Checkpoint:* $0.212\,\text{m/s}^2$ up, $0.151\,\text{m/s}^2$ down.

12.5 Theory and calculation · 6 min

The creep curve of Equation (12.2) peaks at $\mu\varphi Mg$ with $\varphi = 0.5$, $\mu_{dry} = 0.25$, $\mu_{wet} = 0.08$, $w_{pk} = 0.6\,\text{m/s}$ and $m_r = 0.05M$.

(a) Compute the largest force the rail transmits to the motored wheelsets, dry and wet. *Checkpoint:* 233.0 kN and 74.6 kN.

(b) Compute the base speed P_t/F_t and the motor force available at 60 km/h, where case 7 accelerates out of the restriction. What happens on the wet rail if nothing limits the motors? *Checkpoint:* 45.8 km/h; 168 kN, over twice the wet peak.

(c) Compute the creep time constant $m_r/(2\mu\varphi Mg/w_{pk})$ on each rail, and how many of the wet ones pass in one 0.1 s control period. *Checkpoint:* 12.2 ms and 38.2 ms; about 2.6.

12.6 Theory and calculation · 7 min

On the singular arc of a level line the price of time sets the holding speed through $\lambda\eta_t = V^2R'(V)$, with $R'(V) = bM + 2cV$ and $\eta_t = 0.85$.

(a) Compute λ for holding 60, 80 and 98 km/h. *Checkpoint:* 77.8, 177.0 and 317.9 kW.

(b) Between budgets of 268 and 275 s from A to B the dynamic program's planned energy falls by 0.241 kWh per second of budget. Express that in kilowatts and compare it with the converged price of 925 kW at 271 s. *Checkpoint:* 867 kW.

(c) In case 8 the lost minute cost 4.4 kWh driven. The plans log $\lambda = 165\,\text{kW}$ at the 396 s budget and 444 kW at 336 s. What is the average price of the minute, and why does it lie between the two? *Checkpoint:* About 264 kW.

12.7 App challenge · 8 min

Load Scenario 3, *3 · The coast point: one decision, and what it saves*, at the default seed and the shipped speed of 10 (each run to B takes about half a minute).

(a) Before Run, read the coast point on the *Timetable and plan* card. Run to B and read *Punctuality*, *Energy per km* (its kWh), *Drawn · returned* and *Regenerated share*. *Checkpoint:* 2.45 km; B −3.1 s, 19.5 kWh, 37.1 and 17.6 kWh, 76 %.

(b) Choose *Flat out* as the Planner, press Reset and run to B. Read the same tiles, and compute the saving of the coast point in net energy and in energy drawn. *Checkpoint:* B −21.0 s, 28.2 kWh, 58.9 and 30.7 kWh, 75 %. 31 % net, 37 % drawn.

(c) Load Scenario 5, *5 · The dynamic program against the analytic plan*, and run it to B. Read *Punctuality* and the energy, and say where on the run the program saves over the coast point. *Checkpoint:* B −1.7 s, 16.4 kWh, about 16 % less, before the 60 km/h curve.

12.8 App challenge · 8 min

Load Scenario 8, *8 · Running late: re-planning, and the energy a minute costs*, at the default seed. Choose a speed of 30 (the run to C takes about 16 s). The doors at B are held 60 s and *Re-plan when late* is Off.

(a) Run to C. Read *Punctuality* and the energy, and export the CSV. In any row with `phase` RUN, read `plan_T_s` and `plan_lambda_kW`. *Checkpoint:* C +60.8 s, 26.1 kWh; 396 s and 165.3 kW.

(b) Turn *Re-plan when late* On, press Reset and repeat. *Checkpoint:* C −0.6 s, 30.5 kWh; 336 s and 443.9 kW.

(c) What did the minute cost, and why does the price λ of the new plan rise? *Checkpoint:* 4.4 kWh, 17 %.

12.9 App challenge · 9 min

Load the challenge, *Challenge · the full route: wet rail, a restriction, a late departure*, at the default seed and choose a speed of 30 (the route takes about 24 s).

(a) Run the shipped settings. The run pauses itself when the attempt cannot finish. Read the *Challenge target* tile and *Largest slip*. *Checkpoint:* Not met, B −0.6 s −0.07 m, stalled at 7.10 km; slip 34.63 m/s.

(b) Press *Reveal suggested tuning*, then Reset, and run to C. Read the tile, *Largest slip* and *Re-plans*. *Checkpoint:* Met, B −1.7 s −0.01 m, C −1.4 s −0.04 m, 51.8 of 55.1 kWh; 1.36 m/s; 1.

(c) Set *Wet-rail report* back to None, press Reset and run to C. Explain why one setting fails both the clock and the energy. *Checkpoint:* C +15.1 s, 57.9 of 55.1 kWh, slip 2.64 m/s, *CANNOT MAKE IT* on the plan card.

12.10 Application area · 4 min

The meter of Equation (12.4) credits regenerated power only in the share κ that the line accepts, and the line has no substations or other trains in the model except through κ.

(a) Flat out from A to B draws 58.9 kWh and returns 30.7 kWh at $\kappa = 1$, and the four-regime plan draws 37.1 kWh. What does each meter read with *Line receptivity* at 0 %, and what is the coast point's saving then? Check one in the app. *Checkpoint:* 58.9 kWh and 37.1 kWh; 37 %.

(b) What on a real line decides κ, and what does that mean for the value of coasting? *Checkpoint:* Another train drawing nearby, or equipment that takes the power. Coasting is worth more where the line cannot receive.

12.11 Application area · 4 min

The slip override runs at the controller's 10 Hz. In case 7 with no wet-rail report its peak slip on the wet stretch ranges from 1.38 to 2.37 m/s over six seeds, above the declared 2 m/s at one, and at most 1.72 m/s with the report.

(a) Why does the peak vary from seed to seed, and why does the report lower it? *Checkpoint:* The wet entry falls at a random point in the control period; the report lowers the force demanded on entry.

(b) As the engineer specifying the traction drive for this line, what would you require, and how would you check it in service? *Checkpoint:* Slip control inside the drive at millisecond rates, and adhesion reports passed to the planner.

12.12 Review the agent · 5 min

An AI assistant reviews the challenge and writes: *"Pontryagin's principle proves that the energy-optimal drive is full power, hold, coast and brake. The Four regimes planner builds exactly that drive, so it is the optimum. Keep the other suggested settings, choose Four regimes, and the run will meet*

the clock and the energy target with room to spare.” Find the flaw and say how you would check it in the app.