INTERACTIVE CONTROL ENGINEERING
Principles and Applications

Volume 7 Nonlinear Dynamics, Oscillations, and Stabilization

CHAPTER 4

Anti-Lock Braking at the Friction Peak Where the Process Gain Changes Sign

Pages 77–106

Edited by John D. Hedengren

How to cite this chapter. "Anti-Lock Braking at the Friction Peak Where the Process Gain Changes Sign," in *Nonlinear Dynamics, Oscillations, and Stabilization*, John D. Hedengren, Ed., vol. 7 of *Interactive Control Engineering: Principles and Applications*. Advanced Process Solutions, LLC, 2026, ch. 4, pp. 77–106.

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

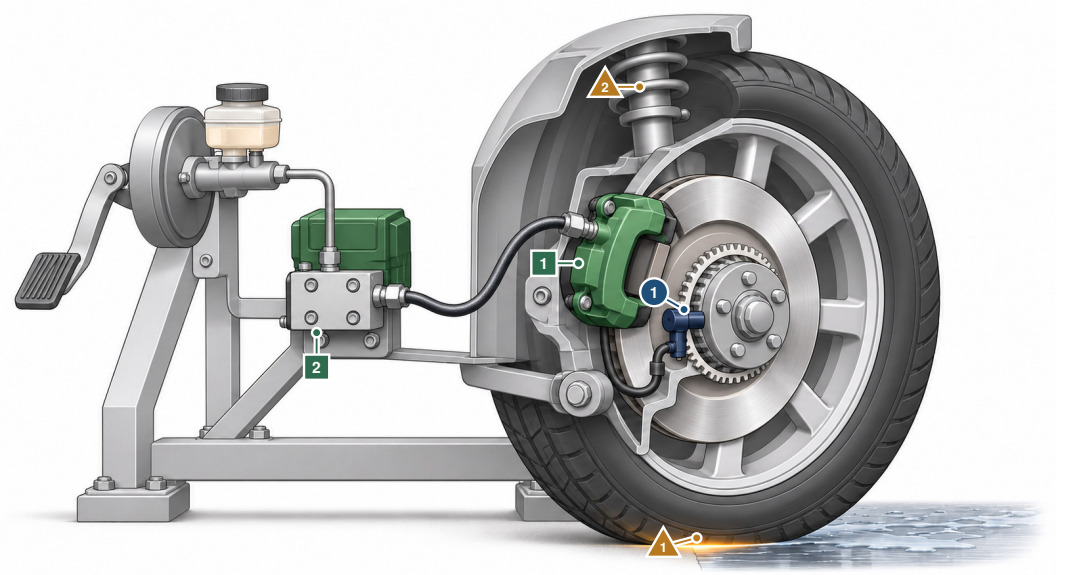
The app for this chapter
https://apmonitor.com/apps/abs-braking/

The volume online: contents, other chapters and citation records
https://apmonitor.com/apps/books/v07-nonlinear-stabilization/#abs-braking

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

CHAPTER 4

Anti-Lock Braking at the Friction Peak Where the Process Gain Changes Sign



The controller moves 1 the brake, squeezing the disc on the wheel
2 the hydraulic modulator, which sets the brake pressure the controller asks for

The controller reads 1 the wheel-speed sensor at the toothed ring, how fast the wheel turns

What disturbs it 1 a change in the road surface under the tire 2 a change in the load the spring puts on the wheel

Notice that only the wheel's speed is read, while the road under the tire, which sets how much braking the wheel can take, is never measured. Illustration, not to scale.

IN PLAIN TERMS Braking Hard Without Locking the Wheel

The main idea. A braked tire grips the road best while it turns a little slower than the car is moving. An anti-lock system keeps each wheel turning just that much slower, at the point where the grip is greatest.

The challenge. Past that point, more brake gives less grip, and the wheel runs away into a skid within a fraction of a second. The best point moves with the road, and the system cannot tell dry pavement from ice before the tire does.

Why it matters. Many machines work best at the top of a curve that falls away beyond it.

THREE WAYS TO READ THIS CHAPTER

Explore	12 min	the first page, the history and the fact on heavy trucks in Section 4.2, the two try-its and the fact on the ideal distance in Section 4.6, and the key points
Study	60 min	the whole chapter with its three pause points and the twelve exercises at its end
Practice	25 min	the first page, the results (Section 4.6), Technology in Practice (Section 4.7), Limitations (Section 4.8) and the open challenge

The control question A braked tire delivers its largest force at an intermediate slip, and past that slip more brake torque gives less force, so the process gain changes sign at the setpoint an anti-lock loop wants. How well can a fixed-gain PI loop hold slip at that peak when the road surface, and with it the peak, changes without warning?

Learning objectives

- Derive the hold torque and the linearized slip pole from the two momentum balances, and locate the stability boundary of a constant pedal.
- Compute the phase margin of the slip loop at a given speed and explain why it falls as the car slows.
- Explain why a sudden loss of friction is a windup problem for the PI loop although its output is far from the upper limit.
- Compare detuning the gain with lowering the slip setpoint as remedies for a slow brake actuator, and justify the choice from the phase margin.
- Judge a sliding-mode law against the PI loop from its boundary-layer gain and from measured runs.

Prerequisites Sections 1.4.2 (the fold of the holding input), 1.4.3 (sign from the immediate response), 1.4.7 (a lag ahead of an unstable pole), 1.4.9 (relay and describing function) and 1.6.2 (velocity form and back-calculation). Chapter 2 meets the same fold on an exothermic reactor.

The app https://apmonitor.com/apps/abs-braking/

See also Volume 4, Chapter 2 for a protective loop held off the peak of a compressor characteristic, and Volume 4, Chapter 12 for a slip loop on the same Burckhardt tire. Volume 20, Chapter 9 for speed limited by stopping distance and wheel slip, and Volume 20, Chapter 2 for a vehicle loop whose behavior moves with speed.

4.1 Introduction

Almost every loop in an introductory control course regulates a variable whose plant responds monotonically to the actuator. Opening a valve raises a level and adding heat raises a temperature, and the design questions concern the magnitude and the speed of the response [1]. A braked wheel is different. Pressing the brake pedal harder increases the retarding force up to a point, and past that point pressing harder reduces it, because the friction available at the contact patch is a function of how fast the tire is sliding over the road and that function has a maximum. The operating point that a braking controller wants is that maximum, which means the setpoint is deliberately placed at the one slip value where the derivative of the plant characteristic is zero.

The consequence dominates the control design. Below the maximum the tire is part of the negative feedback, because extra slip raises the friction coefficient, raises the force opposing the brake at the rim, and limits the growth of slip. Above the maximum the same mechanism has changed sign, so a disturbance that pushes slip up keeps pushing until the wheel stops turning, and a stopped wheel has no lateral capacity left either, so the vehicle cannot be steered while it slides. The friction peak is therefore not a flat optimum but a stability boundary that has to be held from the stable side inside a transient lasting two or three seconds [2]. A centrifugal compressor meets a boundary of the same kind at the peak of its pressure-rise characteristic, as Volume 4, Chapter 2 shows.

Two further properties complicate the design. The position of the maximum and the friction available there both depend on the road surface, which the controller cannot measure, so one fixed setpoint has to serve dry asphalt, standing water, snow and ice. The linearized plant gain also grows without bound as the vehicle slows, so a fixed-gain loop tuned with adequate margin at the start of a stop operates near its stability boundary at the end of one.

The physical content of the model is two conservation statements, a momentum balance on the translating quarter car and an angular momentum balance on the wheel, coupled through a single friction force that is the only driving force in the problem [3]. The slip dynamics are one to two orders of magnitude faster than the speed dynamics, and this time-scale separation is what permits a frozen-speed linearization of the loop [4]. The model brings together concepts a process dynamics and control course already carries, including a process gain that changes sign inside the operating

range, an open-loop unstable operating point, integral windup against an actuator limit, actuator lag as a phase budget, and a measurement whose noise grows as the signal shrinks [1, 2].

Textbook treatments of anti-lock braking usually present the friction curve as a static figure and then move directly to a sliding-mode or rule-based controller [5]. However, a static figure does not show the closed-loop consequence of operating on the falling branch, and it does not expose the interaction between the integral term, the actuator limit and a surface change that arrives without warning. This chapter specifies a browser simulation of quarter-car braking with a Burckhardt tire and measures what a fixed-gain slip loop can and cannot do when its setpoint sits on a maximum whose location is unmeasured.

Section 4.2 traces the problem from aircraft antiskid systems to electronic stability control and sets out the tire models and control approaches in use. Section 4.3 gives the four-state simulator, its equations, coefficients, limits, instrument chain and sample periods, at the level of detail a reimplementation needs. Section 4.4 derives the linearized slip dynamics including the vehicle-deceleration term, shows that the constant-torque equilibrium loses stability at the maximum of the hold-torque curve rather than at the maximum of the friction curve, and establishes that this stability boundary is independent of vehicle speed while the loop crossover is not. Section 4.5 describes the three automatic algorithms: the PI slip loop with its low-speed exit, a rule-based pressure cycle and a sliding-mode slip law. Section 4.6 runs thirteen configurations in which one property of the plant or the loop is changed at a time, reads each outcome from the application's own scorecard, and in Section 4.6.6 sets the sliding-mode law against the fixed-gain loop on an ice patch, with counterfactual runs that remove each part of the law in turn. The tire model and all three control laws are standard ones. Section 4.7 describes the sensors, valves and control units of real antilock systems, Section 4.8 states the limitations and Section 4.9 summarizes the measured results.

4.2 Background

Locked-wheel skidding was recognized as a control problem on aircraft long before it was addressed on cars, because an aircraft touching down carries far more kinetic energy per tire than a car and loses directional control completely when a main wheel stops turning. Aircraft antiskid systems modulate brake pressure against a measured wheel deceleration or a slip estimate, and the design of those systems is still organized around a tire-runway friction characteristic with an interior maximum [6]. Rail vehicles adopted the same idea for wheel-slide protection, where a slid wheelset both lengthens the stop and produces flat spots on the tread.

> ⧖ HISTORY **From Wheel Deceleration to Slip on Airliners, 1929 to 1976**
>
> Anti-lock braking was used on aircraft as early as 1929 [2]. A NASA analysis published in 1976 traced how airliner systems grew. In the earlier ones a speed sensor on each braked wheel watched for a deceleration past a preset threshold, and an on-off valve then cut the brake pressure to zero and reapplied it in full. Later systems modulated the pressure by the depth of the previous skid. The latest regulated slip velocity or slip ratio, which needs the airplane's ground speed as well as the wheel speed. Investigators had found flat spots from locked wheels on tires whose brakes and skid-control units all worked within tolerance, which suggested that wrong information was reaching the logic on slippery runways. Normal antiskid action was assured with high wheel spin-up accelerations or a continuous, accurate ground speed reference [7].

Road vehicles adopted the technology when the sensing and the actuation became inexpensive enough. Early automotive systems used analog electronics and mechanical modulators and were confined to expensive cars and to the rear axle. The step that made anti-lock braking a mass-market item was the replacement of analog logic by a digital control unit, which allowed the pressure-modulation rules to be revised in software rather than in hardware and made per-wheel regulation affordable [8]. Electronic brake control then became the platform for traction control and for electronic

stability control, which reuses the same wheel-speed sensors and the same hydraulic modulator to apply an individual brake for yaw regulation.

The safety record has been measured directly. Fleet studies of fatal crashes before and after anti-lock brakes were fitted found the effect mixed, with reductions in some crash types offset by increases in others, and concluded that the benefit is not a simple reduction in stopping distance [9]. Electronic stability control, built on the same actuator, has a much clearer record, with reviews of the empirical evidence reporting large reductions in single-vehicle and rollover crashes [10]. The distinction matters here, because on a low-friction surface anti-lock braking yields little reduction in stopping distance and instead preserves steering.

> **+ DID YOU KNOW? Antilock by Rule on Heavy Trucks**
>
> In the United States, Federal Motor Vehicle Safety Standard No. 121 requires an antilock brake system on each air-braked truck tractor manufactured on or after March 1, 1997, and on each single-unit vehicle from March 1, 1998. It must directly control the wheels of at least one front axle and at least one rear axle [11].

The modeling side developed alongside the hardware. The Magic Formula is a curve-fitting form that reproduces measured longitudinal and lateral tire characteristics from a small number of coefficients and is the standard representation inside vehicle-dynamics simulation [12]. The Burckhardt form used here is simpler, with three coefficients per surface, and is the representation normally taught alongside quarter-car braking models [3]. Both are static maps from slip to friction coefficient. Dynamic alternatives treat the contact patch as a distributed bristle with its own state, which is the LuGre construction originally proposed for friction in mechanisms [13] and later adapted to the tire-road interface, where the lumped and distributed forms are shown to reproduce the static curve at steady sliding while adding a relaxation length during transients [14].

Control of the resulting plant has been approached from several directions. Sliding-mode designs exploit the fact that the controller does not need the friction curve if it only needs to drive a switching surface to zero [15, 16], and a sliding-mode observer can search for the extremum of the friction characteristic directly rather than regulating to an assumed slip value [17]. Gain-scheduled linear designs address the same nonlinearity by scheduling on the operating point and by formulating the wheel-slip loop against a measured or estimated surface [18]. Extremum-seeking control removes the fixed setpoint entirely and dithers the brake pressure to climb the characteristic, which additionally improves the lateral force available during a stop [19]. Regulating a different variable is also an option, and blending wheel slip with vehicle deceleration keeps the loop from being left holding a signal that becomes meaningless at low speed [5]. All of these require some knowledge of where the peak is, and estimating the tire-road friction on line from the slip and the driving force is a recognized identification problem in its own right [20].

4.3 Vehicle Model

One corner of a passenger car is simulated during a straight-line stop. The vertical dynamics of such a corner, a body and a wheel on a spring and a strut, are the plant of Chapter 11. Here the corner moves only along the road. The controlled variable is the brake slip reconstructed from the wheel speed and the vehicle speed, and the manipulated variable is the brake torque commanded to the hydraulic modulator. Road surface, wheel-speed sensor noise and normal-load change remain operator inputs in both modes, and the controller receives none of them, which makes all three unmeasured loads. Figure 4.1 places every symbol in the equations at the point on the vehicle where it applies.

4.3.1 Variables and Signal Classification

Table 4.1 lists what each signal is and where it may be moved from. The single controlled variable is dimensionless, the single manipulated variable is a torque, and the vehicle speed is a state that the controller is allowed to read without error.

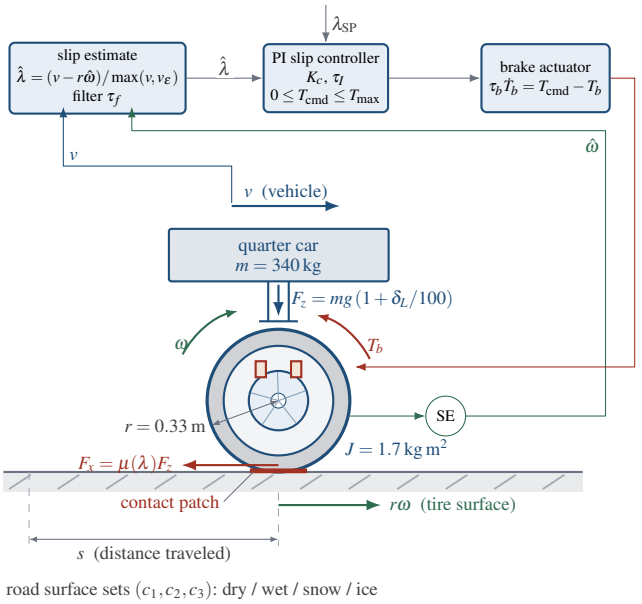


Figure 4.1 Quarter-car braking model with the symbols used in the state equations. Only the wheel speed is instrumented, and the vehicle speed v reaches the slip estimator without error.

Table 4.1 Roles and Ranges of the Signals

Symbol	Role	Unit	Range or default
λ	CV, brake slip	–	$[-1, 1]$
λ_{SP}	setpoint	–	0.02 to 0.60, def. 0.170
T_b	MV, applied torque	N m	0 to 2200
T_{cmd}	controller output	N m	0 to T_{max}
v	state, vehicle speed	m/s	$v_0 = 27$ at reset
ω	state, wheel speed	rad/s	v_0/r at reset
s	state, distance	m	0 at reset
surface	DV, road	–	dry / wet / snow / ice
σ_ω	DV, sensor noise	rad/s	0 to 4, def. 0
δ_L	DV, normal-load change	%	-20 to $+40$, def. 0
$\hat{\omega}$	measurement	rad/s	≥ 0
$\hat{\lambda}$	filtered CV estimate	–	what the loop sees
μ	output, friction coef.	–	0 to 1.17
F_x	output, tire force	N	μF_z

4.3.2 Translation, Rotation and the Slip Definition

Four states are integrated. The vehicle carries a quarter of the mass and is retarded only by the longitudinal tire force, the wheel is accelerated by that force at the rim and retarded by the brake, the brake torque follows its command through a first-order lag, and distance is the integral of speed, as shown in Equations (4.1) to (4.4).

$$m\dot{v} = -F_x \tag{4.1}$$
$$J\dot{\omega} = rF_x - T_b \tag{4.2}$$
$$\tau_b \dot{T}_b = T_{\text{cmd}} - T_b \tag{4.3}$$
$$\dot{s} = v \tag{4.4}$$

where m is the quarter-car mass, v the vehicle speed, F_x the longitudinal tire force, J the wheel rotational inertia, ω the wheel angular speed, r the loaded wheel radius, T_b the applied brake torque, τ_b the hydraulic time constant, T_{cmd} the commanded torque and s the distance traveled. Equation (4.3) runs whether or not the vehicle is moving, because brake pressure continues to build after the car has stopped. The remaining three derivatives are set to zero once $v \le 0$.

The tire force is the product of a friction coefficient and the normal load on the corner, as shown in Equation (4.5).

$$F_x = \mu(\lambda) F_z, \qquad F_z = mg\left(1 + \tfrac{\delta_L}{100}\right) \tag{4.5}$$

where δ_L is the normal-load change in percent and g is gravitational acceleration. With $\delta_L = 0$ the normal load is the static mg, so the deceleration in Equation (4.1) is exactly μg and the ideal stopping distance from speed v is $v^2/(2\mu_{\text{peak}} g)$.

Brake slip is the fraction of the contact-patch velocity that is sliding rather than rolling, and it is regularized in the denominator and clamped, as shown in Equation (4.6).

$$\lambda = \text{clamp}\left(\frac{v - r\omega}{\max(v, v_\varepsilon)}, -1, 1\right) \tag{4.6}$$

where $v_\varepsilon = 0.5$ m/s is the regularization speed. A free-rolling wheel has $r\omega = v$ and $\lambda = 0$, and a locked wheel has $\omega = 0$ and $\lambda = 1$. The exact ratio is singular at the end of a stop, because the numerator and the denominator both approach zero, and the $\max(v, v_\varepsilon)$ guard bounds both the signal and its derivative at the cost of understating the true slip below v_ε.

4.3.3 Burckhardt Tire Model and Surface Coefficients

The friction coefficient is the three-coefficient Burckhardt form, made odd in λ so that a wheel briefly turning faster than the vehicle is pulled back toward free rolling, as shown in Equation (4.7).

$$\mu(\lambda) = \text{sgn}(\lambda)\left[c_1\left(1 - \mathrm{e}^{-c_2|\lambda|_1}\right) - c_3|\lambda|_1\right] \tag{4.7}$$
$$\mu'(\lambda) = c_1 c_2 \mathrm{e}^{-c_2|\lambda|_1} - c_3 \tag{4.8}$$
$$\lambda^* = \min\left(1,\ \tfrac{1}{c_2}\ln\tfrac{c_1 c_2}{c_3}\right), \quad c_3 > 0,\ c_1 c_2 > c_3 \tag{4.9}$$

where $|\lambda|_1 = \min(|\lambda|, 1)$ clips the slip magnitude so the correlation is never extrapolated past locked wheel, c_1 sets the asymptotic level, c_2 the sharpness of the rise out of the origin and c_3 the slope of the falling branch. Every surface gives $\mu(0) = 0$, so a free-rolling wheel transmits no longitudinal force at all. Table 4.2 reports the coefficients that ship with the application together with the quantities they imply, computed from Equations (4.7) to (4.9) rather than quoted from a textbook, and Figure 4.2 draws the four curves. The same curve, with the same dry-asphalt coefficients, drives the wheel of Volume 4, Chapter 12, where wheelspin rather than braking carries the slip past the peak.

Three properties of Table 4.2 shape the control problem. The peak moves by more than a factor of eleven across the four surfaces, from $\lambda^* = 0.1700$ on dry asphalt to 0.0150 on ice, while the friction

Table 4.2 Shipped Burckhardt Coefficients and Derived Peak Quantities

Surface	c_1	c_2	c_3	λ^*	μ^*	$\mu(1)$	$\mu(0.170)$
Dry asphalt	1.2801	23.990	0.5200	0.1700	1.1700	0.7601	1.1700
Wet asphalt	0.8570	33.822	0.3470	0.1308	0.8013	0.5100	0.7953
Snow	0.1946	94.129	0.0646	0.0600	0.1900	0.1300	0.1836
Ice	0.0500	306.39	0	0.0150[a]	0.0500[a]	0.0500	0.0500

[a] Ice has no interior maximum. λ^* is the 99% point of the asymptote and μ^* is the asymptote c_1.

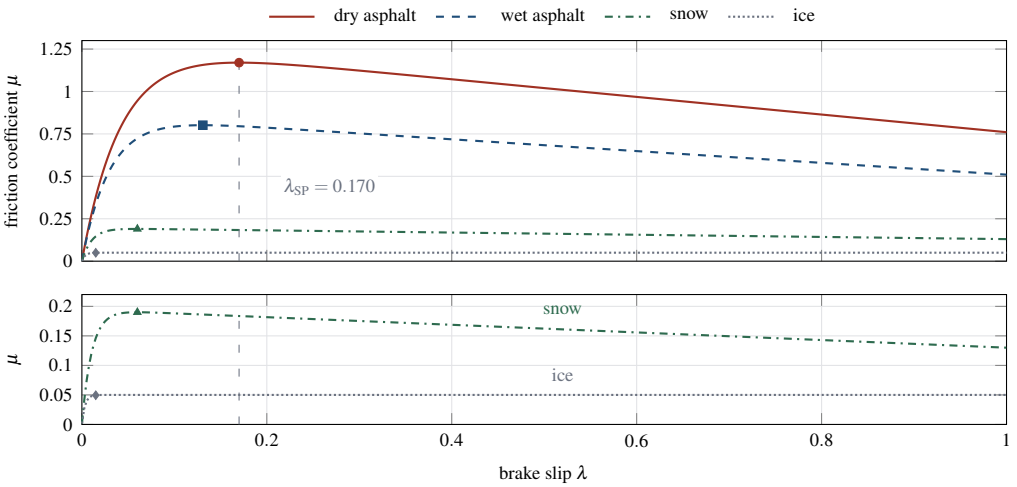


Figure 4.2 Burckhardt friction curves for the four shipped surfaces: dry asphalt solid with a circle at its peak, wet asphalt dashed with a square, snow dash-dotted with a triangle and ice dotted with a diamond (the 99% point of the asymptote for ice). The default slip setpoint of 0.170, the vertical line, is the dry peak and sits on the falling branch of the other three curves. The lower panel repeats snow and ice on an axis that resolves them.

available at the peak moves by a factor of 23.4, from 1.1700 to 0.0500. A locked wheel on dry asphalt gives $\mu(1) = 0.7601$, which is 35.0% below the 1.1700 available at the peak, and the corresponding losses on wet asphalt and snow are 36.4% and 31.6%. Ice has $c_3 = 0$, so its curve never turns over and no interior maximum exists. In that case the application places the displayed peak slip at the 99% point of the asymptote, $\lambda^* = \ln(100)/c_2 = 0.0150$, where $\mu = 0.0495$, which is a convention rather than a maximum. The available friction μ^* is taken as the asymptote $c_1 = 0.0500$ itself, so the ideal stopping distance on ice is the true minimum and cannot be beaten by the closed loop.

4.3.4 Brake Actuator, Wheel-Speed Sensor and the Slip Estimate

The brake modulator is a first-order lag with $\tau_b = 0.020$ s on the commanded torque, with no separate slew limit, and the command is clamped to $[0, T_{\max}]$ with $T_{\max} = 2200$ N m by default. Only the wheel speed is instrumented. The sensor adds Gaussian noise passed through a first-order filter of correlation time 0.01 s and rescaled to a stationary standard deviation of σ_ω, and the sum is clamped at zero, as shown in Equation (4.10).

$$\hat{\omega} = \max(\omega + n(t), 0), \qquad \operatorname{std}[n] = \sigma_\omega \tag{4.10}$$

The vehicle reference speed is handed to the estimator exactly. Slip is then reconstructed with Equation (4.6) using $\hat{\omega}$ in place of ω, and a first-order filter of time constant τ_f smooths the result, as shown in Equation (4.11).

Table 4.3 Constants of the Plant and of the Loop

Symbol	Meaning	Default	Unit	Range
m	quarter-car mass	340	kg	220 to 520
r	wheel radius	0.33	m	0.26 to 0.42
J	wheel inertia	1.7	kg m^2	0.8 to 3.5
g	gravity (hidden)	9.81	m/s^2	9.7 to 9.9
v_0	initial speed	27	m/s	5 to 40
$T_{\max}$	brake torque limit	2200	N m	800 to 2200
τ_b	brake actuator lag	0.020	s	0.005 to 0.200
R^+	rule-based apply rate	12000	N m/s	500 to 30000
R^-	rule-based release rate	15000	N m/s	2000 to 60000
$\Delta\lambda$	rule-based deadband	0.080	–	0.005 to 0.150
v_ε	slip regularization	0.5	m/s	0.2 to 2.0
v_{exit}	slip-control exit	3.0	m/s	1.0 to 6.0
v_{stop}	stop threshold (hidden)	0.05	m/s	0.01 to 0.20
K_c	proportional gain	6000	N m	0 to 30000
τ_I	integral time	0.12	s	0.02 to 2
τ_D	derivative time (adv.)	0	s	0 to 0.020
τ_f	slip filter (adv.)	0.004	s	0 to 0.050
λ_{SP}	slip setpoint	0.170	–	0.02 to 0.60
ρ	boundary-layer scale (adv.)	1	–	0 to 3

$$\hat{\lambda}_k = \hat{\lambda}_{k-1}\, a + \lambda_{\text{raw},k}\,(1-a), \qquad a = \mathrm{e}^{-\Delta t_c/\tau_f} \tag{4.11}$$

where Δt_c is the controller sample period and $\tau_f = 0.004$ s by default, adjustable from 0 to 0.050 s. The exact zero-order-hold form in Equation (4.11) is used rather than an Euler difference, so the filter is stable for any τ_f. There is no transport delay and no sensor bias anywhere in the application. Inferring the vehicle reference speed from the wheel speeds is the hardest part of a production anti-lock system and it is absent here by design.

🕮 **Pause.** Try Exercise 4.1 before reading on.

4.3.5 Parameters, Hard Events and Numerics

Table 4.3 lists every constant. All of them except g and v_{stop} are operator-adjustable inputs, and the two exceptions are fixed because they are not quantities a calibration engineer would adjust.

The initial condition is a free-rolling wheel with the brake released, namely $v = v_0$, $\omega = v_0/r$, $T_b = 0$ and $s = 0$, which gives $\lambda = 0$, $\mu = 0$ and $\dot{v} = 0$ exactly. Three hard events are handled in a post-step routine, outside the right-hand side, so the integrator never meets a discontinuity. The wheel angular speed is clamped at zero so the wheel cannot turn backward. The vehicle is declared stopped once $v \le v_{\text{stop}} = 0.05$ m/s, at which point both v and ω are set to zero and frozen. A wheel-lock flag is raised while $v > \max(1.0, v_\varepsilon)$ and $\lambda > 0.92$, and the time spent in that condition is accumulated and reported. The theoretical best distance $v^2/(2\mu^* g)$ is evaluated once, at the instant the applied torque first exceeds 5 N m, and then held. Recomputing it during a stop would make it a moving target, so on a surface transition it continues to quote the figure for the surface the car started on.

The engine of Section 1.6.1 advances the states with the classical fourth-order Runge-Kutta scheme at $\Delta t = 0.002$ s. The controller executes on the same 2 ms grid, which is a 500 Hz sample rate, and the historian logs at 50 Hz. The fastest genuine dynamics in the plant are the 20 ms brake lag and the runaway to lock, which takes of order 100 ms, so the 2 ms step resolves both. The noise generator is seeded with 20260824, so noisy runs replay identically. Time is in seconds throughout, and every logged signal exports at 50 Hz for offline analysis.

4.3.6 Assumptions

The model is one wheel and one quarter of the mass with a single translational degree of freedom. Pitch, roll, yaw, the other three wheels and the suspension are absent, so nothing in the application can show the loss of steering that locking a front wheel actually causes, and that loss has to be asserted in the text rather than demonstrated. The coupling between longitudinal and lateral tire force that causes it is the central constraint of Chapter 10.

The normal load is static unless the normal-load input is changed. A real 1 g stop shifts roughly 30% of the load onto the front axle, which is the reason production systems regulate each wheel separately, and the input is a manual stand-in for that transfer rather than a model of it.

The Burckhardt curve is a steady-state fit. A real tire reaches its force through a relaxation length of a few decimeters, which at 27 m/s is a further 10 to 20 ms of lag that this model does not carry, and that missing lag is part of what makes a production anti-lock system harder to tune than this one. A dynamic friction model would supply it [14].

The friction coefficient depends on slip alone, while speed, temperature, tread wear, water film thickness and vertical load all move the real curve, so the coefficients in Table 4.2 represent one measurement of one tire on one surface. Slip is regularized below $v_\varepsilon = 0.5$ m/s and slip control is abandoned below $v_{\text{exit}} = 3.0$ m/s, which are numerical necessities that together make the last half meter optimistic by about 1 cm of a 30 m stop.

The stop is a hard event with no brake creep, no driveline windup and no rollback on a slope. Aerodynamic drag and rolling resistance are omitted entirely. That is a small error on dry asphalt, where the tire supplies $\mu^* g = 11.5$ m/s^2, but on ice the tire supplies only 0.491 m/s^2 and the omitted resistances are of the same order, so the ice stopping distances in Section 4.6 should be read as upper bounds. It also means the coasting configuration coasts forever.

4.4 Stability Analysis of the Slip Dynamics

4.4.1 Linearized Slip Dynamics

Differentiating Equation (4.6) above the regularization speed and substituting Equations (4.1) and (4.2) gives the slip rate exactly, as shown in Equation (4.12).

$$\dot{\lambda} = \frac{r}{Jv}\left(T_b - r\,\mu(\lambda) F_z\right) - \frac{(1-\lambda)\,\mu(\lambda)\,F_z}{m\,v} \tag{4.12}$$

The first term is the torque imbalance at the wheel and the second is the contribution of the vehicle's own deceleration, which is small but not negligible and is the term a wheel-only linearization omits. Setting $\dot{\lambda} = 0$ and using $F_z = mg$ gives the brake torque that holds a constant slip, as shown in Equation (4.13).

$$T_{\text{hold}}(\lambda) = r\,\mu(\lambda) F_z + \frac{J(1-\lambda)\,\mu(\lambda)\,g}{r} \tag{4.13}$$

Equation (4.13) carries no speed at all, which is why a well-regulated stop shows an almost flat manipulated-variable trace. On dry asphalt at the peak the two terms are 1287.8 N m and 49.1 N m, giving 1336.9 N m in total, and the wheel term is therefore 3.7% of the total.

Linearizing Equation (4.12) about a constant-slip operating point, holding v fixed because the speed dynamics are two orders of magnitude slower, gives an input gain and a pole, as shown in Equations (4.14) and (4.15).

$$k = \frac{\partial \dot{\lambda}}{\partial T_b} = \frac{r}{Jv} \tag{4.14}$$

$$p = \frac{\partial \dot{\lambda}}{\partial \lambda} = -\mu'(\lambda)\left[\frac{r^2 F_z}{Jv} + \frac{(1-\lambda) g}{v}\right] + \frac{\mu(\lambda)\,g}{v} \tag{4.15}$$

where μ' is the tire slope from Equation (4.8). The sign of μ' sets the character of the plant. On dry asphalt at $\lambda = 0.05$ the slope is 8.73 and the pole is $-71.8\ \mathrm{s}^{-1}$, so slip is strongly self-regulating. At $\lambda = 0.12$ the slope has fallen to 1.21 and the pole to $-9.51\ \mathrm{s}^{-1}$. At $\lambda = 0.25$ the slope is -0.444 and the pole is $+4.05\ \mathrm{s}^{-1}$, an unstable mode with a 0.25 s time constant that takes the wheel to lock in a few tenths of a second. The plant is first-order and stable on one side of the setpoint and first-order and unstable on the other, and the same three parameters describe both.

4.4.2 Stability Boundary of the Constant-Torque Equilibrium

Multiplying Equation (4.15) by v removes the speed entirely, so the slip at which the pole crosses zero is a property of the tire, the wheel and the load and not of how fast the car is going. That crossing satisfies Equation (4.16).

$$\mu'(\lambda_c)\left[\frac{r^2 F_z}{J} + (1-\lambda_c)g\right] = \mu(\lambda_c)\, g \tag{4.16}$$

Differentiating Equation (4.13) and setting the result to zero produces the same condition, so the constant-torque equilibrium loses stability exactly where the hold-torque curve is maximized rather than where the friction curve is maximized. This is the picture of Section 1.4.2: the input that holds a steady state has a maximum, the equilibrium loses stability at that maximum, and a larger input has no equilibrium at all. On dry asphalt that is $\lambda_c = 0.1661$ against a friction peak at $\lambda^* = 0.1700$, with $T_{\mathrm{hold}}(\lambda_c) = 1337.0$ N m against 1336.9 N m at the friction peak. The corresponding values are 0.1280 and 917.3 N m on wet asphalt, 0.0587 and 218.2 N m on snow, and 0.0289 and 57.5 N m on ice, where an interior maximum of the hold torque exists even though the friction curve has none.

This boundary is directly checkable in manual mode. A constant pedal of 1337 N m on dry asphalt stops the car in 32.7 m with the wheel never locking, while 1338 N m locks the wheel for 0.27 s and 1340 N m locks it for 0.87 s and costs 1.4 m. The equilibrium does not degrade as the torque is raised past 1337.0 N m but ceases to exist.

4.4.3 Speed Dependence of the Loop Transfer Function

Both the input gain in Equation (4.14) and the pole in Equation (4.15) are inversely proportional to v, so as the car slows the plant keeps its shape and compresses its time scale. What does not compress is the brake actuator, whose 20 ms lag is fixed by the hydraulics. A lag in series with an unstable pole limits the loop in the way Section 1.4.7 shows for the arm. The loop transfer function at the setpoint is given by Equation (4.17).

$$L(s) = K_c\,\frac{\tau_I s + 1}{\tau_I s}\cdot\frac{k}{(s-p)(\tau_b s + 1)} \tag{4.17}$$

Crossover therefore rises roughly as $v^{-1/2}$ while the actuator corner stays where it is, and the phase margin falls monotonically through the stop. Table 4.4 evaluates Equation (4.17) on dry asphalt at the default tuning and at two of the scenario tunings.

Three readings follow from Table 4.4. The default tuning gives a comfortable 40.6° at the start of a stop and 15.2° at the speed where slip control is handed over, which quantifies the need for a low-speed exit. The high-gain tuning is already at 8.4° at 27 m/s and reaches 3.6° at 5 m/s, so the cycling it produces gets faster and less damped as the car slows. The slow-actuator case is not marginal but unstable at the setpoint, at -1.3°, and moving the setpoint down to $\lambda_{\mathrm{SP}} = 0.12$ restores $+28.1^\circ$ by moving the plant pole from $+0.42$ to $-9.51\ \mathrm{s}^{-1}$. Where the actuator prevents an adequate margin at the peak, lowering the setpoint restores the margin.

Table 4.4 Linearized Loop at $\lambda = \lambda_{SP}$ on Dry Asphalt

Condition	k [(N m s)$^{-1}$]	p [s^{-1}]	ω_c [rad/s]	PM [deg]
$v = 27$, default	0.00719	+0.42	36.0	+40.6
$v = 20$, default	0.00971	+0.57	44.3	+37.1
$v = 10$, default	0.01941	+1.15	68.9	+28.1
$v = 5$, default	0.03882	+2.29	102.5	+20.1
$v = 3$, default	0.06471	+3.82	135.0	+15.2
$v = 27$, K_c=25000, τ_I=0.03	0.00719	+0.42	91.6	+8.4
$v = 5$, K_c=25000, τ_I=0.03	0.03882	+2.29	218.8	+3.6
$v = 27$, $\tau_b = 0.12$	0.00719	+0.42	19.0	−1.3
$v = 27$, τ_b=0.12, λ_{SP}=0.12	0.00719	−9.51	17.8	+28.1

4.5 Controller Design

4.5.1 Discrete Form, Direction and Anti-Windup

The primary controller is the velocity-form PID of Section 1.6.2, with the derivative taken on the measurement, executed every $\Delta t_c = 0.002$ s. Written in increments for this loop, it is given by Equations (4.18) to (4.22).

$$e_k = a\left(\lambda_{SP} - \hat{\lambda}_k\right), \qquad a = +1 \tag{4.18}$$

$$\Delta u_k^P = K_c\left(e_k - e_{k-1}\right) \tag{4.19}$$

$$\Delta u_k^I = \frac{K_c}{\tau_I} e_k \Delta t_c \tag{4.20}$$

$$\Delta u_k^D = -K_c \tau_D a \frac{\hat{\lambda}_k^f - 2\hat{\lambda}_{k-1}^f + \hat{\lambda}_{k-2}^f}{\Delta t_c} \tag{4.21}$$

$$\begin{aligned} s_k &= s_{k-1} + \Delta u_k^P + \Delta u_k^D \\ c_k &= s_k + I_{k-1} \\ T_{\mathrm{cmd},k} &= \mathrm{clamp}(c_k, 0, T_{\max}) \\ I_k &= I_{k-1} + \Delta u_k^I + \frac{\Delta t_c}{\tau_t}\left(T_{\mathrm{cmd},k} - c_k\right) \end{aligned} \tag{4.22}$$

where a is the controller direction, I is the integral state, $\hat{\lambda}^f$ is the derivative-path estimate filtered with time constant τ_f, and $\tau_t = \tau_I$ is the tracking time of the back-calculation. These are Equations (1.12) to (1.15) with s_k for the running sum S_k, c_k for the demanded torque v_k, $T_{\mathrm{cmd},k}$ for the clipped output u_k and Δt_c for Δt. The bias term is zero, so the loop starts from no brake pressure.

The controller direction follows from the sign of the process gain. More brake torque slows the wheel relative to the vehicle, which raises slip, so the process gain from T_b to λ is positive, a slip below setpoint produces a positive error and calls for more torque, and $a = +1$. The sign is read from the immediate response, the input gain k of Equation (4.14), which is positive on both branches, and not from the steady-state gain, which reverses at the maximum of the hold-torque curve. Section 1.4.3 makes the same distinction on the pendulum. A shipped numerical test checks that the output is larger below setpoint than on setpoint and smaller above it, so an edit that inverts the loop fails the test.

Anti-windup is the back-calculation term of Equation (4.22). While the command is clipped, $(\Delta t_c/\tau_t)(T_{\mathrm{cmd},k} - c_k)$ drives I toward the value that reproduces the applied torque with a tracking time of τ_I, so the integral alone is corrected and the running sum s_k is never truncated [21]. At the default $\tau_I = 0.12$ s each scan removes $1/60$ of the clipped excess. The tracking time affects the reported results. A recursion that absorbed the whole excess in one scan would release the brake command 0.12 s sooner after the ice patch of Section 4.6. Anti-windup matters more here than in most loops of

this size, because the danger is not an integral that has grown too large against a limit but an integral that is suddenly wrong. When the friction available collapses mid-stop, most of the commanded torque is stored in I, and the proportional term can react within one sample while the integral can only unwind at a rate of order K_c/τ_I. The mechanism is integral windup even though the output is nowhere near its upper limit.

Manual mode and the transfer between modes work as Section 1.6.2 describes, so the pedal and the loop hand over without a bump in both directions. The idle algorithms always track the applied torque, so switching among the three automatic algorithms is bumpless as well.

4.5.2 The Low-Speed Exit

Below $v_{\text{exit}} = 3.0$ m/s the controller stops regulating slip. On the first sample below that speed it latches the torque that holds the present slip against the road force read from the car's own deceleration, as shown in Equation (4.23).

$$T_{\text{exit}} = \text{clamp}\left(m\hat{a}\left(r + \frac{J(1-\hat{\lambda})}{mr} \right), 0, T_{\max} \right) \tag{4.23}$$

where $\hat{a}$ is the fall rate of the vehicle-speed signal filtered with a 0.01 s time constant. The value is commanded for the rest of the stop while the PID continues to track it. Whenever $\hat{a}$ falls below 0.049 m/s^2, a tenth of the peak deceleration on ice, the held torque rises at the modulator's 12000 N m/s, and once that ramp has acted a measured slip above 0.8 with the car above 1 m/s resets the hold to the value of Equation (4.23). The controllers are given the vehicle's mass, wheel radius and inertia, as a production unit is coded with the car and its tire size, and never the road surface or the load. An earlier version held $\max(T_{\text{cmd}}, T_{\text{hold}}(\lambda_{\text{SP}}))$ with T_{hold} from Equation (4.13) on the true surface and load, which a controller cannot know, and the larger of the two also froze a command caught mid-swing. The change moves every run that reaches the exit, and the results below are measured with Equation (4.23). Across 1014 runs (thirteen configurations, 26 settings each and three algorithms, 40 s apiece) every automatic stop either finishes or is still decelerating at what the road carries when the 40 s run ends, with one exception that predates the change, the rule-based cycle at $\lambda_{\text{SP}} = 0.02$ on the wet-to-dry configuration, which coasts at 3.12 m/s with the brake released. Table 4.4 gives the engineering reason for the handover, because the phase margin at 3 m/s is 15.2° and falling, and Section 4.6 gives the measurement reason, because the slip signal itself becomes uninformative at the same time. A production system does the same below roughly 5 km/h.

4.5.3 The Rule-Based Alternative

The second automatic algorithm is the classical increase, hold and decrease pressure cycle, requiring no model and no gains. Its logic is given by Equation (4.24).

$$u_k = \begin{cases} u_{k-1} - R^{-}\Delta t_c, & e_k < 0 \text{ and } \dot{\hat{\lambda}}_k > 0 \\ u_{k-1} + R^{+}\Delta t_c, & e_k > \min(\Delta\lambda, \lambda_{\text{SP}}/2) \\ u_{k-1}, & \text{otherwise} \end{cases} \tag{4.24}$$

followed by a clamp to $[0, T_{\max}]$, where $\dot{\hat{\lambda}}$ is a filtered slip rate with a 0.01 s time constant. Three asymmetries are deliberate. The deadband $\Delta\lambda = 0.080$ sits entirely below the setpoint, so pressure comes off as soon as slip crosses λ_{SP} and goes back on only after slip has fallen a full deadband below it. The deadband is capped at half the setpoint, so a low target such as the 0.060 snow peak can still call for pressure. Release stops as soon as the wheel starts recovering rather than waiting for slip to return, which prevents the cycle from dumping far more pressure than it needs. Apply and release run at fixed rates rather than in proportion to error, exactly as a modulator valve does. All three exist because the plant is asymmetric, since operating slightly below the peak costs a few percent of force while operating past the peak drives the wheel toward lock.

4.5.4 Sliding-Mode Slip Control

The third automatic algorithm is a sliding-mode law, the classical design for a plant whose gain is uncertain but bounded [22]. It predicts nothing and solves no optimization. Equation (4.12) rearranges into the input-affine form of Equation (4.25), in which the brake torque enters with a positive gain that depends only on speed.

$$\dot{\lambda} = b(v)\Big(T_b - F_x\,\kappa(\lambda)\Big), \quad b = \frac{r}{Jv}, \quad \kappa = r + \frac{J(1-\lambda)}{mr} \tag{4.25}$$

The torque that holds any slip is therefore $F_x\kappa$, whatever the speed, which is Equation (4.13) written for an unknown force. The switching surface is the slip error itself, $\sigma = \hat{\lambda} - \lambda_{\mathrm{SP}}$, read from the same filtered estimate the PI loop uses, and the law is given by Equations (4.26) to (4.28).

$$T_{\mathrm{cmd}} = \mathrm{clamp}\Big(\hat{F}_x\,\kappa(\hat{\lambda}) - K\,\mathrm{sat}(\sigma/\phi),\, 0,\, T_{\max}\Big) \tag{4.26}$$

$$K = \bar{F}_x\,\kappa(\hat{\lambda}) + \eta/b(v) \tag{4.27}$$

$$\dot{\hat{F}}_x = -\beta^2\,\sigma/(b\,\kappa), \qquad 0 \le \hat{F}_x \le \bar{F}_x \tag{4.28}$$

where $\hat{F}_x\kappa$ is the equivalent control (the torque that would hold the slip if the friction force were $\hat{F}_x$), sat is the unit saturation, ϕ is the boundary-layer width, $\eta = 1\ \mathrm{s}^{-1}$ is the reaching margin and β is the closed-loop pole defined below. The friction bound $\bar{F}_x = \bar{\mu}\,mg\,(1+0.40) = 5463$ N is the only use of the tire table, with $\bar{\mu} = 1.170$ the largest coefficient any shipped surface reaches (the dry peak) and 1.40 the top of the normal-load input. The estimate $\hat{F}_x$ is an integral of the slip error, projected onto $[0, \bar{F}_x]$, frozen while the command sits on a limit in the direction the error pushes, and updated after the output with a forward Euler step of Δt_c. It is a force rather than a slip, so a constant road needs a constant estimate while the car slows. The law executes every 2 ms with the PI loop tracking its applied torque, it shares the low-speed exit, and a command that is not finite or not inside $[0, T_{\max}]$ hands the brake to the tracking PI loop until the next valid one. Its mean cost is below 1.5 microseconds per execution in Node on every case of Section 4.6.6.

Outside the layer, $|\sigma| \ge \phi$, the gain of Equation (4.27) covers any force and any estimate inside the bound, so the unclamped command gives $\sigma\dot{\sigma} \le -\eta|\sigma|$, the reaching condition [23]. A check over 8.96×10^6 states outside the layer (five vehicles, two torque limits, ten speeds, five setpoints, six layers, three filters, four surfaces, three loads, 51 slips and five estimates) finds no violation for the unclamped command. The clamp breaks it on one side. A fully released brake leaves only the tire to spin the wheel back, at $bF_x\kappa$, which exceeds η only if $F_x \ge \eta Jv/(r\kappa)$, or 412 N at 27 m/s, while ice carries at most 233 N. On ice above about 14 m/s, and on snow near the top of the speed range, the slip error therefore shrinks more slowly than the condition requires.

The brake lag and the sample delay are not in Equation (4.25), and a relay acting through them chatters [24]. The boundary layer replaces the relay by a straight line inside $|\sigma| < \phi$ [23]. With the balance condition $K/\phi = 2\beta/b$ the law inside the layer is a proportional gain $2\beta Jv/r$ on slip plus the integral of Equation (4.28), and the slip error obeys $(d/dt + \beta)^2\sigma = 0$ at every speed. The pole is placed at a third of the bandwidth of the lags between the command and the measured slip, $\beta = 1/(3(\tau_b + \tau_f)) = 13.9$ rad/s with the declared 20 ms brake and the 4 ms filter, and the width follows as $\phi = \rho\,bK/(2\beta)$ with ρ the boundary-layer scale of Table 4.3. At $\rho = 1$ the layer is 0.52 of slip at 27 m/s, 0.97 at 14 m/s and 4.4 at 3 m/s, the proportional gain is 3864, 2003 and 429 N m per unit slip, and the integral time is $2/\beta = 0.144$ s. The loop inside the layer, from the command to the measured slip, is given by Equation (4.29).

$$L_\sigma(s) = \frac{2\beta s + \beta^2}{s^2} \cdot \frac{\mathrm{e}^{-s\Delta t_c/2}}{(\tau_b s + 1)(\tau_f s + 1)} \tag{4.29}$$

Table 4.5 Shipped Configurations

#	Name	Mode	What is changed
1	Rolling wheel	MAN	dry, $T_b = 0$, then constant-pedal trials
2	Dry pavement	AUTO	reference stop, $v_0 = 27$
3	Wet pavement	AUTO	wet, $v_0 = 25$
4	Snow	AUTO	snow, $v_0 = 16$
5	Ice	AUTO	ice, $v_0 = 8$
6	Dry to ice	AUTO	surface to ice at $t = 1.2$ s
7	Wet to dry	AUTO	surface to dry at $t = 1.5$ s, $v_0 = 25$
8	Slow actuator	AUTO	$\tau_b = 0.12$ s
9	Excessive gain	AUTO	$K_c = 25000$, $\tau_I = 0.03$ s
10	Sensor noise	AUTO	$\sigma_\omega = 2.0$ rad/s
11	Manual maximum	MAN	$T_b = 2200$ N m held
12	Sliding mode	AUTO	configuration 6 under the sliding-mode law, $\rho = 1$
13	Challenge	AUTO	wet to ice at 1.5 s, $v_0 = 20$, $K_c = 1200$, $\tau_I = 0.8$ s

It carries no speed, because the law's gain falls as the plant gain $r/(Jv)$ rises. Evaluated exactly in discrete time (zero-order hold at 2 ms, the filter of Equation (4.11), the integral after the output) it crosses over at 25.4 rad/s with 41.8° of phase margin and 18.9 dB of gain margin. A brake lag of 0.12 s, that of configuration 8, leaves 0.6° and 1.7 dB, and the closed-loop spectral radius rises from 0.981 to 0.9998. At $\rho = 0$ the law is a relay between the clamps, 0 and 2200 N m, and a describing function (Section 1.4.9) with the half-sample delay predicts a limit cycle at 15.6 Hz.

Three models are kept apart. The plant is Equations (4.1) to (4.4) with the Burckhardt curve and the run's brake lag and load. The instrument is Equations (4.10) and (4.11). The law's model is Equation (4.25) with the vehicle's coded m, r and J, the declared 20 ms brake and g, and the friction bound. It never reads the road, the load, the true brake lag or the friction curve, so a brake-lag change is model error that the law has to absorb.

4.6 Results

The shipped configurations appear in Table 4.5, and the scorecard at the end of each run appears in Table 4.6, and Figure 4.3 draws three of its columns. The application lists them as numbered scenarios in the same order, with the challenge last and unnumbered, and credits each one by the rules of Section 1.6.4. Every value comes from the running application, with each configuration loaded as shipped, the one change a row names made before the run, and the run taken to the stop. The application gives the same run whether that change is made before the run or carried through Reset. Distance is measured from the instant the applied torque first exceeds 5 N m, IAE is $\int |\lambda_{\mathrm{SP}} - \lambda|\, dt$ over the whole run, MV travel is $\sum_k |T_{b,k} - T_{b,k-1}|$ over the 50 Hz log, and saturation counts only the time the brake sits at $T_{\max}$, because a released brake is resting rather than out of authority. Peak slip, event times and oscillation frequencies are read from the simulated trajectory at the 2 ms step, because the 50 Hz log undersamples the fastest chatter, whose frequency approaches the 157 rad/s Nyquist limit of the log.

4.6.1 Reference Stop and Locked-Wheel Comparison

Configuration 2 is the case every other run is compared against. The loop brings slip onto 0.170 and holds it, the manipulated variable settles at 1336.9 N m and stays there, and the car stops in 33.0 m against a theoretical 31.8 m. The 1.2 m gap is bought in the first third of a second, while the brake pressure is still building through its 20 ms lag and the friction coefficient is still climbing from zero. The friction coefficient reaches 99% of its peak 0.24 s after the brake is applied, and from 0.33 s to

Table 4.6 Measured Outcome of Each Configuration, Read From the Application Scorecard

Configuration	Distance [m]	Ideal [m]	Locked [s]	Peak λ	IAE	MV travel [N m]	Saturated [%]	Note
2 Dry, PI	33.0	31.8	0.00	0.170	0.025	730	0.0	reference stop
2 Dry, rule-based	36.7	31.8	0.00	0.552	0.256	17910	0.0	3.7 m longer, 25× the MV travel
3 Wet	40.5	39.8	0.00	0.170	0.022	448	0.0	λ_{SP} 0.039 past the wet peak
4 Snow	71.1	68.7	0.00	0.236	0.037	1144	0.0	λ_{SP} nearly 3× the snow peak
4 Snow, rule-based	77.7	68.7	1.42	1.000	3.221	38313	0.0	fixed rates do not transfer
5 Ice	65.2	65.2	0.00	0.370	0.100	955	0.0	ideal from the μ asymptote
6 Dry to ice	218.1	31.8	0.00	0.655	0.280	2267	0.0	ideal frozen at the dry figure
7 Wet to dry	36.9	39.8	0.00	0.170	0.025	1235	0.0	no excursion above the setpoint
8 Slow actuator	35.0	31.8	0.00	0.264	0.153	6368	0.0	unstable at the setpoint
9 Excessive gain	33.3	31.8	0.00	0.344	0.161	68023	0.0	93× the MV travel, 0.3 m worse
10 Sensor noise	34.0	31.8	0.00	0.906	0.205	27477	0.0	no lock at the exit
11 Manual maximum	47.9	31.8	3.30	1.000	2.854	809	92.7	14.9 m lost to a locked wheel
12 Sliding mode	239.3	31.8	0.00	0.849	0.488	2334	0.0	ice patch under the law
13 Challenge, as loaded	269.7	25.4	0.00	0.410	0.503	913	0.0	never reaches the wet peak
13 Challenge, retuned	94.3	25.4	0.00	0.629	0.197	1704	0.0	$K_c = 6000$, $\tau_I = 0.12$ s

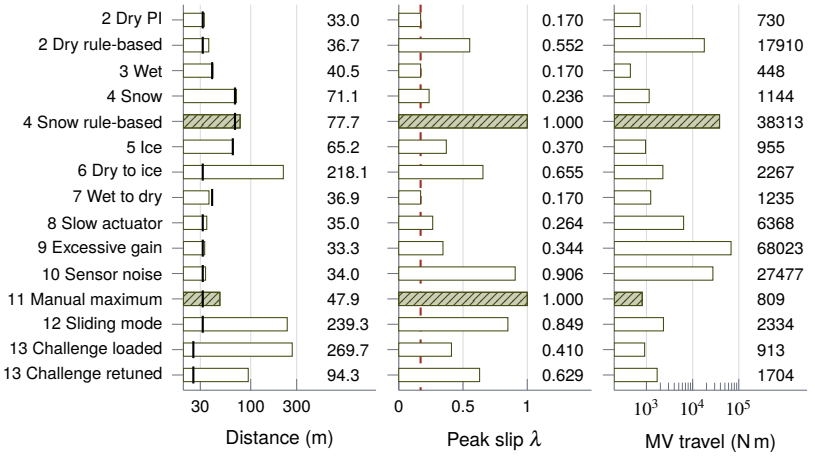


Figure 4.3 The fifteen runs of Table 4.6: stopping distance with the ideal as a black tick (left, logarithmic), peak slip with the 0.170 setpoint dashed (middle) and the travel of the brake torque (right, logarithmic). Hatched bars are the two runs that lock the wheel. Notice that configurations 2 and 9 stop within 0.3 m of each other while their torque travel differs by a factor of 93, and that both runs that lock the wheel carry a peak slip of 1.000.

the end of the stop the car is 0.55 m/s faster than an ideal stop, so no tuning removes the gap and tunings differ only in how quickly they stop paying it.

> **TRY IT IN THE APP** **2 · Dry Pavement · ABS Holding the Peak**
>
> Load the scenario and press Run. Watch the slip trace climb onto the 0.170 setpoint and stay there while the brake torque settles at 1336.9 N m and stays flat. The car stops in 33.0 m against a theoretical 31.8 m (Table 4.6).

Configuration 11 quantifies the loss that anti-lock braking exists to prevent. Holding the pedal at the 2200 N m limit locks the wheel 0.15 s after the brake is applied and keeps it locked for 3.30 s of a 3.57 s stop, at which point the tire is delivering $\mu(1) = 0.7601$ instead of the 1.1700 available at the peak. The stop takes 47.9 m against 33.0 m for the regulated case, so braking harder lengthens the stop by 14.9 m. Handing the sliding wheel to the PI loop at $t = 1.5$ s, at 15.5 m/s, unlocks it 0.07 s later, brings slip back within ± 0.03 of the setpoint by $t = 1.83$ s and finishes the stop in 42.8 m. In the all-manual stop the manipulated variable is saturated for 92.7% of the run while its total travel is 809 N m, barely more than the well-behaved reference stop. A loop can be quiet, saturated and badly wrong at the same time, and no single scorecard number detects the failure.

Pause. Try Exercise 4.2 before reading on.

4.6.2 Fixed Setpoint Across Four Surfaces

Configurations 3, 4 and 5 keep the setpoint at 0.170 and change the road. A fixed setpoint costs little braking force and a great deal of stability margin. On wet asphalt $\mu(0.170) = 0.7953$ against a peak of 0.8013, which is 99.2% of what is available, and the stop takes 40.5 m against an ideal 39.8 m. Aiming at the wet peak itself, $\lambda_{\mathrm{SP}} = 0.131$, gives 40.7 m, and aiming far past it at 0.30 still stops in 42.5 m without a lock. On snow the setpoint is nearly three times past the peak of 0.0600, yet $\mu(0.170) = 0.1836$ against 0.1900, which is 96.6%, and the stop is 71.1 m against an ideal 68.7 m, or 68.9 m with the setpoint moved to the snow peak. The falling branch is gentle on both surfaces because the $c_3\lambda$ term is small, which is why production systems can use a single target slip.

What the fixed setpoint costs is the stabilizing term. On wet asphalt at $\lambda = 0.170$ the tire slope is -0.255 and the open-loop pole is $+2.38\ \mathrm{s}^{-1}$ at 27 m/s, and on snow it is -0.0646 and $+0.60\ \mathrm{s}^{-1}$, so the loop runs on the runaway branch of both surfaces and only feedback keeps the wheel turning. Trading 0.8% of the available friction for the whole restoring term is why a calibration engineer aims below a peak rather than at it.

On ice the regulated stop takes 65.2 m against an ideal of 65.2 m, because the controller operates at $\lambda = 0.170$ on the flat asymptote, where $\mu = 0.0500$ is exactly the μ^* the ideal is computed from. A pedal held at 2200 N m keeps the wheel locked for 14.2 s of the 16.2 s stop and also stops in 65.2 m, so the whole benefit of anti-lock braking on ice is the lateral force that a rolling tire still has.

> **DID YOU KNOW?** **A Stop That Beats Its Own Ideal**
>
> From the simulator: configuration 7 stops in 36.9 m against an ideal of 39.8 m (Table 4.6). The ideal is computed once, for the wet asphalt the stop starts on, and the dry patch at 1.5 s then offers more grip than that figure assumes.

4.6.3 Surface Transition From Dry Asphalt to Ice

Configuration 6 is the central demonstration of the application, and it is reproduced in Figure 4.4. The loop has settled at 1336.9 N m holding $\lambda = 0.170$ on dry asphalt when the road becomes ice at $t = 1.2$ s, where the torque that holds the same slip is 57.1 N m. That is a factor of 23.4 applied in one step, and Equation (4.13) makes it a step in the plant and not in the setpoint.

Slip reaches 0.655 at $t = 1.27$ s, which is far out on what would be the runaway branch on any surface with $c_3 > 0$, and the wheel does not lock. The proportional term reacts within one sample but the stored integral has to be unwound, the command reaches zero at $t = 1.22$ s and stays clamped there until $t = 1.90$ s, and slip returns to within ± 0.03 of setpoint at $t = 1.86$ s, which is 0.66 s after

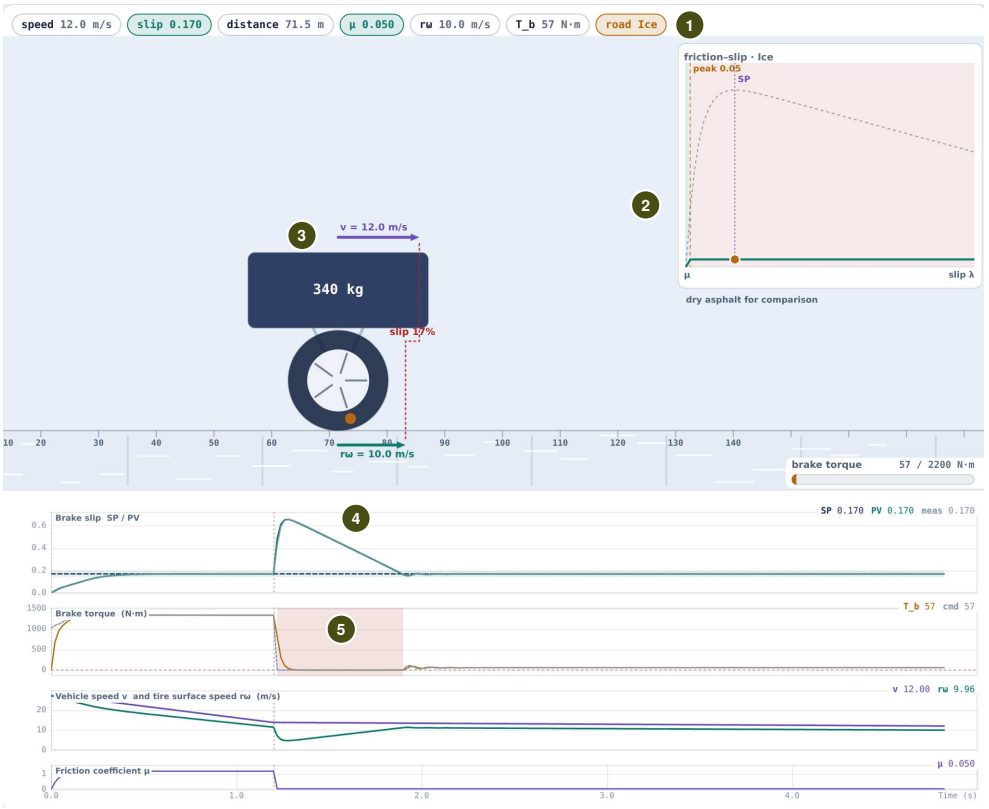


Figure 4.4 Dry asphalt turning to ice 1.2 s into a stop from 27 m/s, paused at $t = 4.83$ s, where the trends of the 5 s window end. The picture is a composite of the plant view (top) and the four trends (bottom) cut from one screenshot. (1) The badges: at the instant shown the car has covered 71.5 m and still travels at 12.0 m/s, with slip on its 0.170 setpoint and μ at 0.050 on ice. (2) The friction-slip view of ice, its operating point on the flat curve. (3) The vehicle speed above the car and the tire surface speed below the wheel. (4) Slip spikes from the setpoint to 0.655 (Table 4.6). (5) The brake command is held against its lower limit (shaded), from 1.22 s to 1.90 s, times the screen does not print, while the friction coefficient collapses from 1.170 to 0.050 (Table 4.2) and the vehicle and tire surface speeds separate and rejoin.

the event. Raising τ_I to 1 s slows the unwind without locking the wheel, because the slower integral has built only 1254.5 N m by the time the ice arrives. Slip then peaks at 0.498, is still 0.016 above setpoint at $t = 3.2$ s and does not return within ± 0.03 until $t = 2.57$ s. No integral time in the 0.02 to 2 s range locks the wheel on this transition, while a gain between 600 and 1500 N m does, for 0.21 s at $K_c = 1000$ N m. This is the clearest demonstration in the application that the sign of a disturbance matters more than its size for this plant, because configuration 7 applies the mirror-image change and produces no excursion above the setpoint, with a peak slip of 0.170 and an IAE of 0.025 that are indistinguishable from an undisturbed stop. Its only excursion is toward free rolling, a dip to 0.074 at $t = 1.53$ s that is back within ± 0.03 by $t = 1.57$ s.

Pause. Try Exercise 4.3 before reading on.

Reducing the torque limit makes the recovery better rather than worse, and the integral state explains why. With $T_{\max}$ lowered to 900 N m the loop spends the entire dry phase asking for more torque than the brake can deliver, back-calculation holds the integral near the value that reproduces the applied output instead of letting it grow, and there is therefore less stored torque to unwind when

the ice arrives. Peak slip after the transition falls from 0.655 to 0.430, and slip is back within ± 0.03 of setpoint 0.47 s after the event instead of 0.66 s. The dry phase costs distance, because the car meets the ice at 18.0 m/s instead of 13.8 m/s, and the whole stop grows from 218.1 m to 356.7 m, while the ice transition costs less.

> **TRY IT IN THE APP** **6 · Dry → Ice, One Second Into the Stop**
>
> Load the scenario and run it to 3.5 s: slip spikes to 0.655 and is back within ± 0.03 of the setpoint 0.66 s after the ice. Reload the scenario, set Brake torque limit to 900 N m under Parameters in Show Model, and run again. Slip now peaks at 0.430 and recovers in 0.47 s.

4.6.4 Actuator Lag, Excessive Gain and Sensor Noise

Configurations 8, 9 and 10 each attack the loop in a different way and produce three different failures. The slow actuator at $\tau_b = 0.12$ s adds phase without adding gain, and Table 4.4 puts the loop at -1.3°, which is unstable at the setpoint. The simulated trajectory agrees. Slip never settles and instead runs a sustained oscillation between 0.084 and 0.264 for the whole regulated part of the stop. Its period shortens from 0.32 s in the first cycle to 0.16 s in the eighth as the car slows, a frequency rising from 20 to 41 rad/s with a mean of 26.7 rad/s, against a predicted crossover of 19.0 rad/s at 27 m/s rising to 31.1 rad/s at 10 m/s. The wheel never locks. At the low-speed exit at $t = 2.24$ s the hold is the 1268 N m the road is carrying, not the 1710 N m the oscillation was demanding at that instant. The stop takes 35.0 m instead of 33.0 m.

Two repairs behave very differently. Halving the gain to $K_c = 3000$ lengthens the stop to 37.1 m and leaves the slip IAE at 0.153, because the oscillation is driven by pressure that was already commanded arriving late rather than by loop gain as such. Moving the setpoint to $\lambda_{\mathrm{SP}} = 0.120$ gives 36.1 m and drops the slip IAE from 0.153 to 0.031, a factor of five, and Table 4.4 says why, because the plant pole moves from $+0.42$ to -9.51 s^{-1} and the margin returns to $+28.1^\circ$. Lowering the setpoint restores the stability margin that the actuator lag consumed.

The excessive-gain case demonstrates that stopping distance alone does not discriminate between tunings. At $K_c = 25000$ and $\tau_I = 0.03$ s the car stops in 33.3 m against 33.0 m, a difference the driver would never feel, while the total variation of the brake torque rises from 730 to 68023 N m, a factor of 93, and the slip IAE gets six times worse at 0.161 against 0.025. Every one of those cycles is a real modulator valve opening and closing. The cycling is present from the start, with 22% of the travel spent in the first half second, and what changes through the stop is its frequency. Slip cycles at 75 rad/s over the first half second and at 123 rad/s over the last half second before the low-speed exit, rising with the predicted crossover of 91.6 rad/s at 27 m/s and 218.8 rad/s at 5 m/s while the phase margin falls from 8.4° to 3.6°.

Sensor noise of $\sigma_\omega = 2.0$ rad/s is 0.66 m/s at the tire surface, which is 0.024 of slip at 27 m/s and 0.13 of slip at 5 m/s, because Equation (4.6) divides by the speed. The measurement therefore degrades in exactly the same proportion as the plant gain grows, and the actuator travel records it, because 10% of the total 27477 N m is spent in the first half second and 68% in the last 1.3 s. The stop grows by 1.1 m. At the low-speed exit at $t = 2.26$ s the noisy command is 2200 N m, the actuator limit, but the exit holds the 895 N m the road is carrying and the wheel does not lock. Raising the slip filter to $\tau_f = 0.02$ s removes only 16% of the actuator travel, because the noise already has a 0.01 s correlation time. Halving K_c to 3000 N m removes 48% of the travel but lengthens the stop by a further 1.8 m. The absent lock is partly a property of the seed. Over noise seeds 1 to 10 the default gain locks the wheel in four runs, three of them after the exit latch, and the halved gain in five, two of them after the latch. Neither repair makes the slip signal informative at 3 m/s, which is the argument for the exit.

4.6.5 Rule-Based Comparison and the Graded Challenge

The rule-based cycle reaches the peak faster than the PI loop because it ramps pressure on at a fixed 12000 N m/s regardless of the error, and it then hunts. On dry asphalt it stops in 36.7 m against 33.0 m and spends 17910 N m of actuator travel against 730, and that pressure cycling is the pulsation a driver feels at an anti-lock pedal. On snow the same fixed rates fail, because 12000 N m/s is an enormous step next to the 218 N m the surface can hold, and the cycle overshoots into lock for 1.42 s and stops in 77.7 m against 71.1 m. Moving its setpoint to the 0.060 snow peak does not help, giving 80.6 m, and although the wheel no longer locks, slip still peaks at 0.32. Fixed modulation rates do not transfer between surfaces, which parallels the fixed-setpoint result of configurations 3 to 5.

The graded challenge starts on wet asphalt at 20 m/s with $K_c = 1200$ and $\tau_I = 0.8$ s and turns the road to black ice at $t = 1.5$ s. The run is scored at $t = 3.6$ s, and it passes when slip stays within ± 0.03 of setpoint over the final third, from $t = 2.4$ s, and the wheel spends no more than 0.02 s locked in total, with the model parameters and the road as given. As loaded, the loop is so slow that slip has reached only 0.023 by the time the ice arrives, far short of the wet peak. The ice excursion then peaks at 0.410 and does not return within ± 0.03 until $t = 3.17$ s, so the worst error over the scored window is 0.086, nearly three times the band, although the wheel never locks. Restoring $K_c = 6000$ and $\tau_I = 0.12$ s brings slip back within the band 0.41 s after the ice and leaves a worst error of 0.0007 over the scored window, inside the band with 0.029 to spare. Run on to the stop, the retuned loop cuts the distance from 269.7 m to 94.3 m and the slip IAE from 0.503 to 0.197, mostly by braking properly during the 1.5 s of wet asphalt rather than by doing anything clever on the ice, because the car meets the ice at 8.5 m/s instead of 15.4 m/s. Raising K_c helps twice, because it enlarges the proportional response and it also speeds the integral unwind for a given τ_I, while shortening τ_I helps once and spends stability margin.

Configuration 1 has no controller and precedes the closed-loop configurations. The wheel rolls freely at 27 m/s with $\lambda = 0$ and $\mu = 0$, so the car does not decelerate at all. Separate stops at constant pedal torques, each restarted from 27 m/s, move the operating point along the stable left branch of Figure 4.2, from a slip of 0.092 at 1250 N m to 0.155 at 1337 N m one second into the stop, until the equilibrium disappears at 1337.0 N m. Walking the pedal up within a single run does not reproduce this, because with 100 N m steps every 0.5 s the car stops after 5.4 s with the pedal at only 1100 N m. The scorecard is not meaningful for this configuration, accumulating the full 0.170 setpoint error as IAE (0.51 over 3 s of coasting) with no loop closed, while it correctly reports 0.0% saturation for a released brake.

4.6.6 Sliding-Mode Law Against the PI Loop

Configuration 12 is configuration 6 with the sliding-mode law of Section 4.5.4 selected at $\rho = 1$. The comparison in Table 4.7 runs each case to the stop under four structures with the same seed, instruments and noise: the shipped PI loop (and, on the challenge, its suggested tuning), the rule-based cycle, the law, and the PI loop retuned once by the law's own rule, with $K_c = 2\beta J v / r$ set every execution from the speed it reads and $\tau_I = 2/\beta = 0.144$ s. Nothing is retuned per case. Torque chatter is the root-mean-square difference between the 2 ms command and its 42 ms centered mean over the regulated part of the stop (automatic, $t \geq 0.1$ s, $v \geq v_{\text{exit}}$). Peak slip error is the largest $|\lambda - \lambda_{\text{SP}}|$ while regulating after 0.5 s, and recovery is the time after a surface change until slip stays within ± 0.03 of the setpoint for 0.2 s.

The law stops later than the PI loop on the two high-friction roads, 34.58 m against 32.97 m on dry asphalt and 41.37 m against 40.54 m on wet asphalt, and ties it on snow (71.20 m against 71.13 m) and ice (65.28 m against 65.24 m), as listed in Table 4.7. At 27 m/s its proportional gain is 3864 N m per unit slip against the PI loop's fixed 6000, so slip reaches the setpoint later. The same falling gain lowers its peaks on snow and ice, 0.170 and 0.193 against 0.235 and 0.370. On the ice patch of configurations 6 and 12 the law's slip peaks at 0.849 against 0.655 and recovers in 1.00 s against 0.66 s, because at the peak (14.4 m/s, $t = 1.29$ s) its gain is 2066 N m per unit slip, a third of the PI

Table 4.7 Sliding-Mode Law Against the Shipped Structures, Each Case Run to the Stop

Case	Structure	Distance [m]	Locked [s]	Peak λ	Peak error	Recovery [s]	IAE	MV travel [N m]	Chatter [N m]
2 Dry	PI	32.97	0.000	0.170	0.003	–	0.025	730	0.4
	rule-based	36.66	0.000	0.552	0.356	–	0.256	17910	30.7
	sliding mode	34.58	0.000	0.170	0.032	–	0.052	937	0.2
	scheduled PI	34.58	0.000	0.170	0.031	–	0.052	938	0.2
3 Wet	PI	40.54	0.000	0.170	0.001	–	0.022	448	0.2
	rule-based	47.61	0.414	1.000	0.830	–	0.952	27748	29.2
	sliding mode	41.37	0.000	0.170	0.006	–	0.039	626	0.2
	scheduled PI	41.37	0.000	0.170	0.006	–	0.039	627	0.2
4 Snow	PI	71.13	0.000	0.235	0.000	–	0.037	1144	3.2
	rule-based	77.66	1.424	1.000	0.631	–	3.221	38313	29.7
	sliding mode	71.20	0.000	0.170	0.001	–	0.040	297	0.3
	scheduled PI	71.20	0.000	0.170	0.001	–	0.040	297	0.3
5 Ice	PI	65.24	0.000	0.370	0.003	–	0.100	955	2.7
	rule-based	65.31	0.000	0.851	0.422	–	4.311	20538	26.4
	sliding mode	65.28	0.000	0.193	0.000	–	0.078	189	0.2
	scheduled PI	65.28	0.000	0.194	0.000	–	0.078	189	0.2
6, 12 Dry to ice	PI	218.09	0.000	0.655	0.485	0.66	0.280	2267	11.0
	rule-based	257.70	0.214	1.000	0.830	–	3.914	55154	23.8
	sliding mode	239.29	0.000	0.849	0.679	1.00	0.488	2334	4.6
	scheduled PI	239.23	0.000	0.857	0.687	0.97	0.495	2373	4.5
7 Wet to dry	PI	36.88	0.000	0.170	0.096	0.13	0.025	1235	13.2
	rule-based	41.57	0.230	1.000	0.231	–	0.570	21178	29.4
	sliding mode	37.90	0.000	0.170	0.114	0.50	0.077	1057	3.8
	scheduled PI	37.90	0.000	0.170	0.115	0.49	0.076	1059	3.8
13 Challenge	PI as loaded	269.74	0.000	0.410	0.240	1.67	0.503	913	0.7
	PI suggested	94.28	0.000	0.629	0.459	0.41	0.197	1704	11.9
	rule-based	128.83	0.176	1.000	0.830	–	3.648	40919	25.4
	sliding mode	103.02	0.000	0.895	0.725	0.73	0.394	1642	4.1
	scheduled PI	103.01	0.000	0.905	0.735	0.80	0.398	1680	4.0

loop's, to unwind the 1337 N m stored for dry asphalt. The law releases the brake fully and never locks the wheel there, as shown in Figure 4.5. On the challenge it meets the target with a worst error of 0.006 over the scored window, where the PI loop as loaded leaves 0.086 and its suggested tuning 0.001. It beats the rule-based cycle on every case in distance, lock, actuator travel and chatter.

The law and the scheduled PI loop agree within 0.06 m on every case in Table 4.7. At $\rho = 1$ the switching term saturates for 0.00% of the regulated time in every configuration, and replacing $\mathrm{sat}(\sigma/\phi)$ by σ/ϕ gives a bit-identical command and slip on all eight comparison cases. The layer sized to the brake lag is wider than any slip error these roads produce, so at the shipped setting the law is a proportional-integral loop whose gain is scheduled on speed and whose integral is bounded by the friction bound. Its robustness argument, the reaching condition, is never called on.

Three counterfactual runs remove one part of the law at a time, as listed in Table 4.8. Without the boundary layer ($\rho = 0$) the switching term acts every 2 ms as a relay between 0 and 2200 N m. The full switching gain catches the dry peak at once and the dry stop shortens to 32.77 m, but the chatter rises from 0.2 to 892.6 N m and the actuator travel from 937 to 69778 N m. In configuration 12 run to $t = 3.5$ s the ice-patch peak falls from 0.849 to 0.655 while the chatter rises from 12.3 to 538.8 N m, and run to the stop the wheel locks for 0.27 s after the low-speed exit. The engine's limit cycle runs at

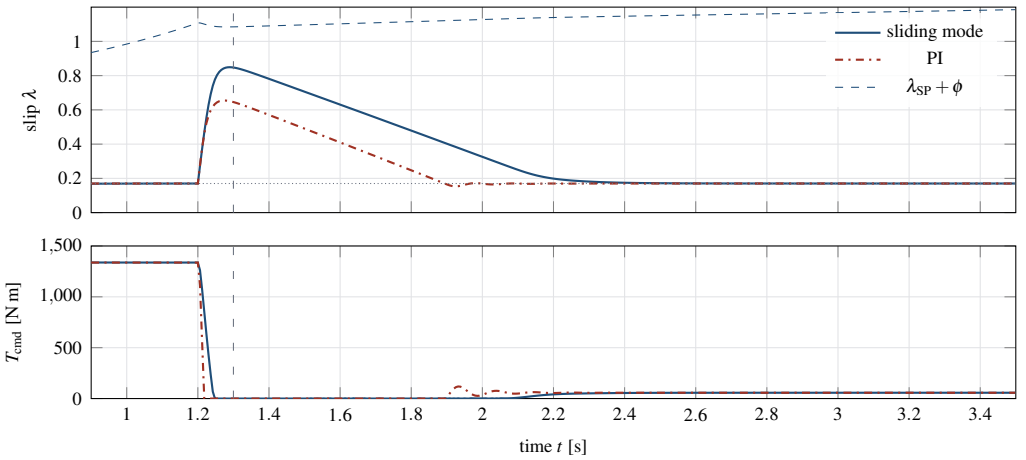


Figure 4.5 Configuration 12, dry asphalt turning to ice at $t = 1.2$ s, under the sliding-mode law at $\rho = 1$ (solid) and under the PI loop (dash-dotted), with the upper edge of the law's boundary layer dashed and the setpoint dotted. At $t = 1.3$ s (vertical line) the law's slip is 0.847 with $\sigma = 0.678$ inside a layer of $\phi = 0.915$ and its command is clamped at zero, while the PI loop's slip is 0.646. The upper edge of the layer stays above every slip reached.

Table 4.8 Counterfactual Variants of the Sliding-Mode Law, Run to the Stop

Case	Variant	Dist. [m]	Locked [s]	Peak λ	Chatter [N m]
Dry	$\rho = 0$	32.77	0.000	0.309	892.6
	no switching	38.15	0.108	1.000	1.2
	no integral	36.89	0.000	0.445	0.2
Wet	$\rho = 0$	41.05	0.000	0.535	792.0
	no switching	49.57	0.080	1.000	2.0
	no integral	40.91	0.282	1.000	0.9
Snow	$\rho = 0$	73.73	0.000	0.685	548.1
	no switching	84.30	0.924	1.000	2.0
	no integral	87.57	4.692	1.000	4.9
Ice	$\rho = 0$	65.23	0.000	0.713	308.5
	no switching	67.12	0.000	0.436	1.2
	no integral	65.25	14.128	1.000	5.6
Dry to ice	$\rho = 0$	212.15	0.272	1.000	372.5
	no switching	284.23	0.236	1.000	1.8
	no integral	249.66	14.068	1.000	4.8
Challenge	$\rho = 0$	95.99	0.000	0.503	403.4
	no switching	142.05	0.196	1.000	1.7
	no integral	94.91	14.180	1.000	6.1

16.7, 18.3 and 20.0 Hz at 24.6, 15.6 and 8.6 m/s against the 15.6 Hz the describing function predicts, with slip half-swings of 0.046, 0.072 and 0.101 against a predicted 0.087, 0.133 and 0.238. The relay is lopsided between the clamps, which halves the amplitude, and the tire's slope below the peak reduces it further, while the growth as $1/v$ survives. A thin layer sits between the two settings, with $\rho = 0.25$ giving a 0.610 peak and 31.4 N m of chatter in configuration 12 (saturated 9.5% of the time) and $\rho = 0.5$ giving 0.706 and 22.4 N m. Doubling the layer to $\rho = 2$ locks the wheel for 0.18 s on the ice patch. Without the switching term the equivalent control alone is an integral-only law, and most cases lock (snow 0.92 s, dry to ice 0.24 s). Without the integral, with $\hat{F}_x$ frozen at the middle of the bound, the wide layer leaves an offset that locks the wheel for 14.1 s on ice.

Safety is judged over every existing configuration as shipped and with each single extreme a learner can set (setpoint, slip filter, every vehicle and loop constant at its minimum and maximum,

Table 4.9 Worst Outcome per Structure Over Every Configuration and Single Extreme

Structure	Runs	With a lock	Longest lock [s]
PI	336	80	5.04 (6, noise 4 rad/s)
Rule-based	348	201	9.11 (6, $\tau_b = 0.2$ s)
Sliding mode	300	60	4.97 (11, $v_0 = 40$, manual)

the load at -20 and $+40\%$, and 4 rad/s of sensor noise), 276 cases shared by the three structures, plus each structure's own settings at their extremes. Table 4.9 gives the worst outcome per structure. The law locks the wheel in fewer runs, 60 of 300 against 80 of 336 for the PI loop and 201 of 348 for the rule-based cycle. Over the shared cases it locks where the PI loop does not, or for more than 0.05 s longer, in 19 cases, and the PI loop does so against it in 36. The longest of the 19 is 1.41 s on the challenge with a 0.2 s brake, where the PI loop locks for 0.50 s, and in 24 further cases the law's slip overshoots the PI loop's by more than 0.1. One of the 19 is a shipped configuration. Under the law configuration 8 locks the wheel for 0.126 s after the low-speed exit, because the law reaches the latch at $t = 2.27$ s with slip at 0.194, just past the dry peak and rising, and the lagged brake applying 1340 N m against the 1334 N m hold, so the falling branch carries the wheel to lock. The PI loop reaches its latch with slip at 0.090. The cause is the margin of Equation (4.29) at a 0.12 s brake, 0.6° against 41.8° at the declared 20 ms, since the law keeps its gain from the declared lag.

The friction bound shapes the load trend. Inside the bound the law's loss to the PI loop on dry asphalt grows with the load, from 1.61 m at 0% (34.58 against 32.97 m) to 2.57 m at $+40\%$ (27.20 against 24.63 m), and the scheduled PI loop loses the same (27.19 m), so the cause is the lower gain and not the projection. A load of $+60\%$, past the bound and past what the input allows, pins $\hat{F}_x$ at $\bar{F}_x$. The law then under-brakes at a mean slip of 0.091 and stops in 25.81 m against 22.18 m, without locking.

4.6.7 The Application Display

The application runs on the framework of Section 1.6, with Manual and Auto modes and a full export of the logged trajectory. The vehicle speed and the tire surface speed are displayed together because their divergence is the definition of the controlled variable, and Figure 4.4 shows configuration 6 paused 4.83 s into the stop.

4.7 Technology in Practice

What Measures It

Slip is not measured. It is computed from two speeds, and only one of them comes from an instrument at the wheel. Airplane antiskid systems first used the braked-wheel sensor only to detect a fast deceleration, and the later ones that regulate slip also need a reference for the ground speed [7]. Tests on a transport airplane's main wheel found the largest braking friction at a slip of about 0.1, while the side force fell so fast that above about 0.2 little cornering capability was left, so about 0.1 was suggested as the best target for an antiskid design [25]. The chapter's sensor model is the wheel speed with optional noise of 0.01 s correlation time, the vehicle speed reaches the estimator exactly, and its setpoint of 0.170 sits at the dry peak itself.

> **TECHNOLOGY** **Sensor: Wheel Speed and a Ground-Speed Reference**
>
> Each braked wheel of an airliner carries a sensor whose signal is proportional to the wheel's angular velocity. Most antiskid systems derive the ground speed from these sensors too, and some read it from an unbraked nose wheel or a small auxiliary wheel [7]. A slip-ratio system that holds the wheel about 10 % below the ground speed never lets it spin up, so it needs a reference from elsewhere [25]. When the brakes went on before the wheels had spun up on a wet runway, the logic took a ground speed well below the true one, and the traces of the front wheels showed deep skids for about 30 s [7].

What Moves It

The brake is moved by fluid pressure. In a car's antilock system the controlling device sends its signals to modulators, which adjust the brake actuating forces [26]. In the NASA tests the antiskid control valve sat between the pilot's metering valve and the brake, and a control box fed by the wheel-speed sensor drove it [25]. The chapter's modulator is a 20 ms first-order lag on torque with a clamp at 2200 N m, and its rule-based cycle applies and releases at fixed rates, as the early on-off valves did [7].

> **TECHNOLOGY** **Final Element: A Servo Valve and a Brake With Hysteresis**
>
> An airliner's antiskid logic modulates the hydraulic pressure at each wheel brake through a servo control valve. Early valves were on-off and applied or released pressure at a single rate, while later ones are constant-gain or variable-gain servo valves [7]. Behind the valve, the brake does not turn pressure into torque along a single curve. Tests during antiskid cycling showed large hysteresis loops, so one pressure gave a wide range of torque. Neither a linear spring with viscous damping nor an undamped nonlinear spring reproduced them, while a nonlinear spring with friction damping cut the torque error significantly [25]. The chapter's lag on torque carries no hysteresis.

Where the Controller Runs

In a car, control is executed by electronic control units, microprocessors with input and output interfaces, and a modern car carries more than a hundred of them. Once computers were introduced they took on anti-lock braking, electronic braking and stability control, braking the wheels one by one [2]. No open, vendor-neutral source read for this chapter gives the execution period of a production antilock unit. The chapter's laws run every 2 ms, and one function holds its structure: the velocity-form PI with back-calculation and the low-speed latch, or in its place the rule-based cycle or the sliding-mode law. Figure 4.6 draws the loop on vehicle hardware.

What Is Modernizing It

The first development replaces hydraulics at the brake. In aircraft braking, electromechanically actuated brakes have found wide acceptance as an alternative to hydraulic ones, and advanced antiskid controllers take advantage of them [27]. The second is electric braking under the antilock function. The United States light-vehicle brake standard covers electrically actuated service brakes and electric transmission of the brake signal, and an electric vehicle whose regenerative braking is part of the service brake must have that braking controlled by its antilock system [26]. The third is learning the peak on line. One design estimates the vehicle's states and the tire curve's parameters with a particle filter while it brakes, choosing each action for what it will both earn and reveal, and in simulation it stopped up to 8.5 % shorter than extremum seeking [28]. Another drives automated vehicles to high slip on purpose during empty trips, to estimate the peak friction of the road [29].

> **IN PRACTICE** **What a Brake Test Counts as a Locked Wheel**
>
> In the United States, Standard No. 135 (49 CFR 571.135) governs light-vehicle brakes. It defines an antilock system as the part of the service brake that controls wheel slip by sensing wheel rotation, passing those signals to controlling devices and sending their outputs to modulators. Wheel lockup means 100 % slip, and in many of its stops no wheel may lock for longer than 0.1 s above 15 km/h. Any electrical failure of the antilock system must light the brake warning indicator, and the test sequence includes stops with the antilock function failed [26]. The chapter's lock flag counts the time spent above a slip of 0.92.

4.8 Limitations

The controller is given the vehicle reference speed exactly. Estimating that speed from four wheel speeds when all four may be slipping is the hardest part of a production anti-lock system, and none of it is present here, so the application should not be read as a model of anti-lock braking end to end.

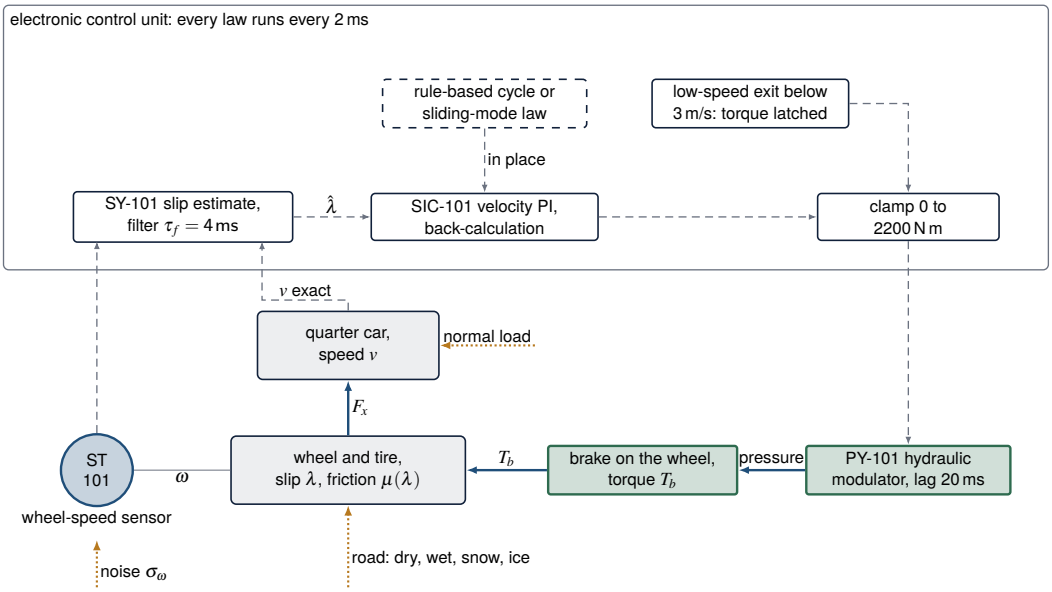


Figure 4.6 The slip loop on vehicle hardware. The wheel-speed sensor ST-101 reads the wheel, and the slip estimate SY-101 divides by the vehicle speed, which the chapter hands over exactly. The velocity-form PI SIC-101 (or, in its place, the rule-based cycle or the sliding-mode law, dashed) sends a torque command, clamped to the brake limit, to the hydraulic modulator PY-101, and below 3 m/s the exit latches a hold torque. Notice that the controller sees neither the road nor the load (dotted arrows), so a change in either reaches it only through the slip it causes.

The friction map is static and single valued. Relaxation length, temperature, wear and water film thickness are absent, and a dynamic friction model would change both the achievable bandwidth and the shape of the transition in configuration 6 [14]. On ice the curve has no peak at all, and the displayed peak slip of 0.0150 is a convention rather than a maximum.

⚠ OPEN CHALLENGE Finding a Peak Without Falling Off It

Estimates of the slip and of the road are often made apart from the antilock controller, which leads to sub-optimal braking. The peak is harder to identify on snow or ice, where it is less defined than on dry or wet roads [28]. Ordinary driving gives little help, because it stays at slips below about 0.05, too low to show the region near the peak [29]. For aircraft braking on a runway that can change fast, the estimation is called an extremely challenging and still open task [27]. Progress would be a loop that learns the peak while braking on it, without the lock it exists to prevent.

The scorecard was designed for regulation against a setpoint, and two of its figures need care here. Overshoot and settling are not reported, because the slip setpoint never steps, and the peak slip error reported instead includes the initial approach from $\lambda = 0$, so it reads 0.151 even for the reference stop. The theoretical best distance is frozen when the brake is applied and is therefore wrong by construction after a surface change, which is deliberate but has to be stated whenever configuration 6 is quoted.

The linearization in Section 4.4 holds v fixed, which is a good approximation because the slip dynamics are one to two orders of magnitude faster than the speed dynamics at every speed in Table 4.4. The phase margins quoted there are still properties of a frozen plant rather than guarantees over a whole stop.

The sliding-mode law's model leaves out the brake lag, the slip filter and the sample delay, which enter only through the pole β and the layer width, and it keeps β from the declared 20 ms brake, so a slower brake in the plant erodes its margin without the law knowing, to 0.6° at the 0.12 s of configuration 8. Its friction bound covers the shipped roads up to dry asphalt's peak at $+40\%$ load. A force beyond it leaves the estimate pinned and the wheel under-braked. Its reaching condition holds only while the brake can deliver the switching torque, and at the shipped layer it is never exercised, so the application demonstrates the boundary-layer trade rather than the robustness of an ideal sliding mode. The vehicle's mass, radius and inertia are treated as coded specifications for every controller, so a learner's change to them reaches the controllers as a recalibrated vehicle rather than as model error.

4.9 Conclusion

This chapter has used a browser simulation of quarter-car braking with a Burckhardt tire to study regulation to the maximum of a plant characteristic. Four states are integrated with fourth-order Runge-Kutta on a 2 ms step, the friction coefficient comes from three coefficients per surface, and the four shipped surfaces place the friction peak between $\lambda^* = 0.0150$ on ice and 0.1700 on dry asphalt with peak friction between 0.0500 and 1.1700. The constant-torque equilibrium is shown to lose stability at the maximum of the hold-torque curve rather than at the maximum of the friction curve, at $\lambda_c = 0.1661$ and 1337.0 N m on dry asphalt, and that boundary carries no speed dependence while the loop crossover does. A fixed-gain proportional-integral loop on slip stops the car in 33.0 m against a theoretical 31.8 m, where a locked wheel needs 47.9 m, and the same tuning falls from 40.6° of phase margin at 27 m/s to 15.2° at the 3 m/s handover. A mid-stop change from dry asphalt to ice steps the required torque by a factor of 23.4 and produces a slip excursion to 0.655 that recovers in 0.66 s, and the mirror-image change produces no excursion above the setpoint. A sliding-mode law with its boundary layer sized to the brake lag never saturates its switching term, acts as a speed-scheduled PI loop that stops 1.6 m later on dry asphalt and peaks at 0.849 on the ice patch, and locks the wheel in 60 of 300 stress runs against 80 of 336 for the fixed-gain loop. The natural extensions would estimate the surface on line, add a relaxation length to the tire, model the reference-speed estimator, and carry the lateral force that makes anti-lock braking worth fitting on ice. A moving horizon estimator or a Kalman filter could treat the Burckhardt coefficients as slowly varying parameters and reconstruct the friction peak from the measured wheel speed, which is the identification problem of Section 4.2 [20]. Nonlinear model predictive control could enforce the torque limit and a slip ceiling

as explicit constraints instead of relying on anti-windup and a low-speed exit [30]. The application gives a learner a control loop whose setpoint sits exactly where its process gain vanishes.

KEY POINTS

- A braked tire gives its largest force at a small slip, and past that slip more brake torque gives less force.
- The constant-torque equilibrium loses stability at the maximum of the hold-torque curve, 1337.0 N m on dry asphalt, a boundary with no speed in it.
- The same PI tuning falls from 40.6° of phase margin at 27 m/s to 15.2° at the 3 m/s handover.
- An ice patch is a windup problem with the output far from its limit: slip runs to 0.655 and recovers in 0.66 s.

Where else this applies. Volume 4, Chapter 2 holds a compressor off the peak of its characteristic, and Chapter 2 meets the same fold of the holding input.

References

[1] D. E. Seborg, T. F. Edgar, D. A. Mellichamp, and F. J. Doyle, III, *Process Dynamics and Control*, 4th ed. Hoboken, NJ: John Wiley & Sons, 2016.

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

[3] U. Kiencke and L. Nielsen, *Automotive Control Systems: For Engine, Driveline, and Vehicle*, 2nd ed. Berlin: Springer Berlin Heidelberg, 2005.

[4] B. W. Bequette, *Process Control: Modeling, Design, and Simulation*. Upper Saddle River, NJ: Prentice Hall, 2003.

[5] S. M. Savaresi and M. Tanelli, *Active Braking Control Systems Design for Vehicles*. London: Springer London, 2010.

[6] Z. Jiao, X. Liu, F. Li, and Y. Shang, "Aircraft antiskid braking control method based on tire–runway friction model," *Journal of Aircraft*, vol. 54, no. 1, pp. 75–84, 2017.

[7] W. B. Horne, J. L. McCarty, and J. A. Tanner, "Some effects of adverse weather conditions on performance of airplane antiskid braking systems," NASA Langley Research Center, Hampton, VA, Technical Note NASA TN D-8202, 1976, nASA Technical Reports Server 19760019083.

[8] H. Leiber and A. Czinczel, "Antiskid system for passenger cars with a digital electronic control unit," SAE International, SAE Technical Paper 790458, 1979.

[9] C. M. Farmer, "New evidence concerning fatal crashes of passenger vehicles before and after adding antilock braking systems," *Accident Analysis & Prevention*, vol. 33, no. 3, pp. 361–369, 2001.

[10] A. Erke, "Effects of electronic stability control (ESC) on accidents: A review of empirical evidence," *Accident Analysis & Prevention*, vol. 40, no. 1, pp. 167–173, 2008.

[11] National Highway Traffic Safety Administration, "Standard no. 121; air brake systems, 49 CFR 571.121," Code of Federal Regulations, Title 49, Chapter V, edition of October 1, 2024, U.S. Government Publishing Office, 2024. [Online]. Available: https://www.govinfo.gov/content/pkg/CFR-2024-title49-vol6/pdf/CFR-2024-title49-vol6-sec571-121.pdf

[12] H. B. Pacejka and E. Bakker, "The magic formula tyre model," *Vehicle System Dynamics*, vol. 21, no. sup001, pp. 1–18, 1992.

[13] C. Canudas de Wit, H. Olsson, K. J. Åström, and P. Lischinsky, "A new model for control of systems with friction," *IEEE Transactions on Automatic Control*, vol. 40, no. 3, pp. 419–425, 1995.

[14] C. Canudas-de Wit, P. Tsiotras, E. Velenis, M. Basset, and G. Gissinger, "Dynamic friction models for road/tire longitudinal interaction," *Vehicle System Dynamics*, vol. 39, no. 3, pp. 189–226, 2003.

[15] C. Unsal and P. Kachroo, "Sliding mode measurement feedback control for antilock braking systems," *IEEE Transactions on Control Systems Technology*, vol. 7, no. 2, pp. 271–281, 1999.

[16] A. Harifi, A. Aghagolzadeh, G. Alizadeh, and M. Sadeghi, "Designing a sliding mode controller for slip control of antilock brake systems," *Transportation Research Part C: Emerging Technologies*, vol. 16, no. 6, pp. 731–741, 2008.

[17] S. Drakunov, U. Ozguner, P. Dix, and B. Ashrafi, "ABS control using optimum search via sliding modes," *IEEE Transactions on Control Systems Technology*, vol. 3, no. 1, pp. 79–85, 1995.

[18] T. A. Johansen, I. Petersen, J. Kalkkuhl, and J. Ludemann, "Gain-scheduled wheel slip control in automotive brake systems," *IEEE Transactions on Control Systems Technology*, vol. 11, no. 6, pp. 799–811, 2003.

[19] E. Dincmen, B. A. Guvenc, and T. Acarman, "Extremum-seeking control of ABS braking in road vehicles with lateral force improvement," *IEEE Transactions on Control Systems Technology*, vol. 22, no. 1, pp. 230–237, 2014.

[20] F. Gustafsson, "Slip-based tire-road friction estimation," *Automatica*, vol. 33, no. 6, pp. 1087–1099, 1997.

[21] K. J. Åström and T. Hägglund, *Advanced PID Control*. Research Triangle Park, NC: ISA – The Instrumentation, Systems, and Automation Society, 2006.

[22] V. Utkin, "Variable structure systems with sliding modes," *IEEE Transactions on Automatic Control*, vol. 22, no. 2, pp. 212–222, 1977.

[23] J.-J. E. Slotine and W. Li, *Applied Nonlinear Control*. Prentice Hall, 1991.

[24] K. D. Young, V. I. Utkin, and Ü. Özgüner, "A control engineer's guide to sliding mode control," *IEEE Transactions on Control Systems Technology*, vol. 7, no. 3, pp. 328–342, 1999.

[25] S. M. Stubbs and J. A. Tanner, "Review of antiskid and brake dynamics research," in *1980 Aircraft Safety and Operating Problems, Part 2*. Hampton, VA: NASA Langley Research Center, 1981, pp. 555–562, nASA Technical Reports Server 19810010540.

[26] National Highway Traffic Safety Administration, "Standard no. 135; light vehicle brake systems, 49 CFR 571.135," Code of Federal Regulations, Title 49, Chapter V, edition of October 1, 2024, U.S. Government Publishing Office, 2024. [Online]. Available: https://www.govinfo.gov/content/pkg/CFR-2024-title49-vol6/pdf/CFR-2024-title49-vol6-sec571-135.pdf

[27] F. Crocetti, G. Costante, M. L. Fravolini, and P. Valigi, "Tire-road friction estimation and uncertainty assessment to improve electric aircraft braking system," in *29th Mediterranean Conference on Control and Automation (MED)*, 2021, pp. 330–335.

[28] B. Sullivan, J. Jiang, G. Mavros, and W.-H. Chen, "An exploration-exploitation approach to anti-lock brake systems," arXiv preprint 2306.14730 [eess.SY], 2023. [Online]. Available: https://arxiv.org/abs/2306.14730

[29] Z. Liang, H. Zhou, H. Huang, and X. Li, "High-slip-ratio control for peak tire–road friction estimation using automated vehicles," in *2026 IEEE Intelligent Vehicles Symposium (IV)*, 2026, pp. 1901–1906.

[30] J. D. Hedengren, R. Asgharzadeh Shishavan, K. M. Powell, and T. F. Edgar, "Nonlinear modeling, estimation and predictive control in APMonitor," *Computers & Chemical Engineering*, vol. 70, pp. 133–148, 2014.

4.10 Exercises

4.1 ⌕ Concept check · 1 min

Only one signal of the quarter-car model in Figure 4.1 reaches the slip estimator through an instrument, and the app can add noise to it. Which signal is it?

(a) The vehicle speed v, read by a ground-speed sensor on the body
(b) The friction coefficient μ, inferred from the brake pressure
(c) The wheel angular speed ω, read by the wheel-speed sensor
(d) The applied brake torque T_b, read by a load cell at the caliper

4.2 ⌕ Concept check · 2 min

In configuration 11 the pedal is held at the 2200 N m limit on dry asphalt, and the car stops in 47.9 m against 33.0 m under the PI loop. Why does braking harder lengthen the stop?

(a) The brake command sits on its limit, so the wheel receives less torque than the pedal asks
(b) The low-speed exit releases the brake below 3 m/s, and the car coasts to rest
(c) The 20 ms hydraulic lag withholds the full torque for most of the stop
(d) The wheel locks, and a sliding tire delivers $\mu(1) = 0.7601$ instead of the peak 1.1700

4.3 ⌕ Concept check · 2 min

A PI slip loop holds $\lambda = 0.170$ when the road changes without warning. Which change threatens the wheel, and why?

(a) A fall in friction, because the integral holds a torque the new road cannot carry and pushes slip onto the falling branch

(b) A rise in friction, because the wheel spins up toward free rolling and the integral then winds up against the upper limit

(c) Both equally, because the loop is linear about its setpoint and answers either sign of change in the same way

(d) Neither, because back-calculation holds the integral at the applied torque through any change of the road

4.4 Theory and calculation · 6 min

Wet asphalt has the Burckhardt coefficients $c_1 = 0.8570$, $c_2 = 33.822$ and $c_3 = 0.3470$. Use Equations (4.7) to (4.9) and $g = 9.81\,\mathrm{m/s^2}$.

(a) Find the peak slip λ^* and the peak friction μ^*. *Checkpoint:* 0.1308 and 0.8013.

(b) Find the friction of a locked wheel, $\mu(1)$, and the percentage of the peak force it loses. *Checkpoint:* 0.5100, a loss of 36.4 %.

(c) Find μ at the default setpoint 0.170 as a fraction of μ^*, and the ideal stopping distance from 25 m/s. *Checkpoint:* 0.7953, 99.2 % of the peak, and 39.8 m.

4.5 Theory and calculation · 8 min

The default vehicle has $m = 340\,\mathrm{kg}$, $r = 0.33\,\mathrm{m}$, $J = 1.7\,\mathrm{kg\,m^2}$ and $g = 9.81\,\mathrm{m/s^2}$, with no load change, on dry asphalt ($c_1 = 1.2801$, $c_2 = 23.990$, $c_3 = 0.5200$, $\mu(0.170) = 1.1700$).

(a) Use Equation (4.13) to find the brake torque that holds $\lambda = 0.170$, and the share of it that decelerates the wheel itself. *Checkpoint:* 1336.9 N m, of which 49.1 N m or 3.7 %.

(b) Find the tire slope μ' and the pole p of Equation (4.15) at $\lambda = 0.25$ and $v = 27\,\mathrm{m/s}$, and the time constant of the mode. *Checkpoint:* $\mu' = -0.444$, $p = +4.05/\mathrm{s}$, about 0.25 s.

(c) Find the same pole at 5 m/s without repeating the calculation, and say what the result means for the slip loop. *Checkpoint:* 21.9/s, a time constant of 46 ms.

4.6 Theory and calculation · 8 min

Take the loop of Equation (4.17) on dry asphalt at $\lambda_{\mathrm{SP}} = 0.170$ with the default tuning $K_c = 6000\,\mathrm{N\,m}$ and $\tau_I = 0.12\,\mathrm{s}$.

(a) Show that for an unstable plant pole $p > 0$ the phase margin is $\arctan(\tau_I \omega_c) - 90^\circ + \arctan(\omega_c/p) - \arctan(\tau_b \omega_c)$, and evaluate it at 27 m/s, where $p = 0.42/\mathrm{s}$, $\tau_b = 0.020\,\mathrm{s}$ and the crossover is 36.0 rad/s. *Checkpoint:* 40.6°.

(b) Evaluate it at 3 m/s, where $p = 3.82/\mathrm{s}$ and the crossover is 135.0 rad/s. *Checkpoint:* 15.2°.

(c) Evaluate it for the slow actuator of configuration 8, $\tau_b = 0.12\,\mathrm{s}$, at 27 m/s, where the crossover is 19.0 rad/s, and say what the run will show. *Checkpoint:* −1.3°.

4.7 App challenge · 8 min

Load scenario 1, Rolling wheel, brake released. It runs in Manual from 27 m/s on dry asphalt at the default seed. Make four stops, each with Brake torque set before pressing Run and with Reset between them (Reset keeps the pedal): 1250, 1337, 1338 and 1340 N m. Run each to the stop, which takes under 3 s, at a playback speed of 0.25, and read the Stopping distance and Wheel locked tiles. Export the file of the 1250 and 1337 N m runs.

(a) Read both tiles for 1337 and 1338 N m. *Checkpoint:* 32.7 m with 0.00 s, and 32.9 m with 0.27 s.

(b) Read both tiles for 1250 and 1340 N m. *Checkpoint:* 34.8 m with 0.00 s, and 34.1 m with 0.87 s.

(c) Read column `slip` at 1.0 s in the two exported files, and compare the values with the slip λ_c at which the hold torque is largest. *Checkpoint:* 0.092 and 0.155, within 0.002, against $\lambda_c = 0.1661$.

Accept distances within 0.1 m and lock times within 0.01 s.

4.8 App challenge · 8 min

Load scenario 6, Dry → ice, one second into the stop, which runs the PI loop in Auto with the road turning to ice at 1.2 s, at the default seed. Make three runs to 3.5 s at a playback speed of 0.25, and export the file of each. For each run read the Peak slip error tile at 3.5 s, the first logged time after 1.2 s from which column `slip` stays within ± 0.03 of the setpoint 0.170, and column `vehicle_speed_mps` at 1.2 s.

(a) Run the scenario as loaded. *Checkpoint:* 0.484, 1.86 s and 13.77 m/s.

(b) Press Reset, set Integral time to 1 s before pressing Run, and run again. *Checkpoint:* 0.328, 2.58 s and 16.16 m/s.

(c) Reload the scenario, set Brake torque limit to 900 N m in the Parameters of Show Model before pressing Run, and run again. Explain why this run recovers fastest. *Checkpoint:* 0.259, 1.68 s and 17.98 m/s.

Accept the tile within 0.002, the time within 0.02 s and the speed within 0.02 m/s.

4.9 App challenge · 7 min

Load the graded challenge, listed last as wet asphalt that turns to ice. It stops from 20 m/s with $K_c = 1200\,\mathrm{N\,m}$ and $\tau_I = 0.8\,\mathrm{s}$, the road turns to ice at 1.5 s, and the target is slip within ± 0.03 of the setpoint from 2.4 s to 3.6 s with no more than 0.02 s of lock. Use the default seed and a playback speed of 0.25, run to 3.6 s, and export the file of each run. The worst error is the largest $|\texttt{slip} - \texttt{slip_sp}|$ over the rows from 2.40 to 3.60 s, as the app's verdict reports it.

(a) Run the challenge as loaded. Find the worst error, the largest slip before 1.5 s and column `vehicle_speed_mps` at 1.5 s. *Checkpoint:* 0.086, 0.023 and 15.43 m/s, no lock.

(b) Press Reset, set Proportional gain to 6000 N m and Integral time to 0.12 s before pressing Run, and run again. Find the worst error and the speed at 1.5 s. *Checkpoint:* 0.0007, within 0.001, and 8.45 m/s, no lock.

(c) The retuned loop lets slip peak higher on the ice than the loaded one. Explain why it still passes.

4.10 Application area · 4 min

A wheel-speed sensor on a wheel of radius 0.33 m carries noise of standard deviation $\sigma_\omega = 2.0\,\mathrm{rad/s}$, as in configuration 10, while the vehicle speed is known exactly.

(a) Find the standard deviation of the slip noise at 27, 5 and 3 m/s. *Checkpoint:* 0.024, 0.13 and 0.22.

(b) Give two reasons, one from the measurement and one from the plant, why the controller stops regulating slip at 3 m/s.

4.11 Application area · 3 min

The model omits rolling resistance and aerodynamic drag. Take a rolling resistance coefficient of 0.01, which adds a deceleration of $0.01g$.

(a) Compare that deceleration with what the tire supplies at its peak on ice and on dry asphalt. *Checkpoint:* 0.0981 m/s^2, 20 % of 0.491 on ice and under 1 % of 11.5 on dry asphalt.

(b) Find the ideal stopping distance on ice from 8 m/s with and without it. *Checkpoint:* 54.4 m against 65.2 m.

4.12 Review the agent · 5 min

In scenario 8, Slow brake actuator ($\tau_b = 0.12\,\mathrm{s}$), slip oscillates between 0.084 and 0.264 at the default tuning. An AI assistant replies as follows. “The loop oscillates because its gain is too high for the slower actuator. Halve the proportional gain to 3000 N m and keep $\tau_I = 0.12\,\mathrm{s}$. A lower gain moves the crossover down to where the 0.12 s lag costs less phase, so the phase margin rises and the oscillation dies out.” Find the flaw in the proposal and say how it could be checked.