INTERACTIVE CONTROL ENGINEERING
Principles and Applications

Volume 7 Nonlinear Dynamics, Oscillations, and Stabilization

CHAPTER 14

String Stability of Well-Tuned Followers in a Vehicle Platoon Simulator

Pages 338–365

Edited by John D. Hedengren

How to cite this chapter. "String Stability of Well-Tuned Followers in a Vehicle Platoon Simulator," 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. 14, pp. 338–365.

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

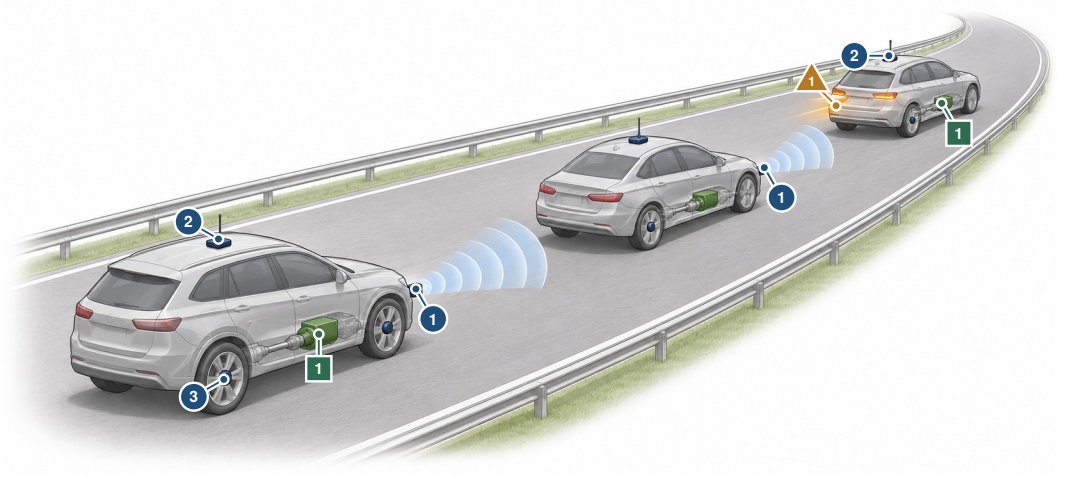
The app for this chapter
https://apmonitor.com/apps/vehicle-platoon/

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

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

CHAPTER 14

String Stability of Well-Tuned Followers in a Vehicle Platoon Simulator



The controller moves 1 each car's throttle and brakes, which follow its controller's command a little late

The controller reads 1 the radar that measures the gap to the car ahead and how fast it closes

2 the radio that hears the car ahead's acceleration 3 the car's own speed

What disturbs it 1 the lead car's brake tap

Notice that each car sees only the car directly ahead, by radar and by radio, so a change made by the lead reaches the last car only through every car between them. Illustration, not to scale.

IN PLAIN TERMS **When a Line of Cars Passes a Wave Back**

The main idea. A car that follows another on its own settles smoothly after the car ahead slows. In a line of ten such cars, a short tap on the first car's brake can grow into a much deeper slowdown at the back.

The challenge. Each car reacts only to the one ahead, and a little late. So a well-behaved car can still pass the wave back larger than it came. A longer gap or a radio link to the car ahead stops the growth, but one costs room on the road and the other can fail.

Why it matters. Growing waves like these are how traffic jams form with no crash or road work to cause them.

THREE WAYS TO READ THIS CHAPTER

Explore	12 min	the first page, the fact in the Introduction (Section 14.1), the two history segments in Section 14.2, the try-its and the fact in Section 14.6, and the key points
Study	60 min	the whole chapter with its three pause points and the twelve exercises at its end
Practice	20 min	the first page, the results (Section 14.6), Technology in Practice (Section 14.7), Limitations (Section 14.8) and the open challenges

The control question A car that follows another can settle smoothly after the car ahead changes speed and still pass the change back larger than it came. When does a line of well-damped followers let a brake tap die out, and what do the spacing policy, a radio link, human reaction delay and a predictive controller change?

Learning objectives

- Derive the car-to-car transfer function of a time-gap follower and the bound $h \geq \sqrt{2/k_p}$, and explain why constant spacing fails for every gain.
- Show that the damping of one follower does not decide string stability, using the roots and the $|\Gamma|$ peak of the same loop.
- Measure the amplification along the line in the app and compare it with the criterion, for radar ACC, cooperative ACC and drivers.
- Price string stability in lane flow, radio delay and packet loss, and explain what the fallback and the cut-in supervisor buy.
- Judge when a linearized criterion understates the growth of a large wave.

Prerequisites Sections 1.4.10 (string stability), 1.4.4 (eigenvalues), 1.4.7 (delay limits), 1.6.1 and 1.6.2 (engine and velocity form). Section 1.6.3 is needed only for the predictive controller. Volume 2, Chapter 2 for the cruise loop of one car.

The app https://apmonitor.com/apps/vehicle-platoon/

See also Volume 20, Chapter 2 for the speed dependence of an automated car's path-following loop, and Volume 20, Chapter 9 for a speed chosen from stopping distance. Volume 5, Chapter 8 for a loop closed through a seeded delay line with jitter and packet loss.

14.1 Introduction

An automated car that follows another regulates a gap whose setpoint moves with the car ahead, and its loop is first judged, like any other, by whether it settles. A line of such cars raises a question that no single loop answers, which is how much of the acceleration of the car ahead each car passes on to the car behind it. If that gain exceeds one at any frequency, a disturbance with energy there grows from car to car however well damped each loop is, and the requirement that it never does is called string stability [1, 2]. Section 1.4.10 states the criterion on a line of carts. Production adaptive cruise control (ACC) has made the question practical. Seven 2018 model-year systems identified from car-following experiments were all string unstable, and along a platoon of one model an initial disturbance of 6 mph grew to 25 mph [3].

This simulator makes the difference between a well-tuned loop and a well-behaved line measurable. At a time gap of 0.6 s each radar follower behind a car at constant speed has a closed-loop damping ratio of 0.50, yet the gain from car to car peaks at 1.306 at 1.12 rad/s, and a brake tap that takes the lead from 72 to 57.3 km/h takes car 10 down to 36.3 km/h. Stiffening the loop does not rescue the line. A spacing gain of 3 instead of 2 leaves the peak at 1.305, because the time gap must exceed $\sqrt{2/k_p}$, and over the whole range of the gain slider that bound never falls below 0.82 s. Holding a fixed distance instead of a time gap fails for every gain, as the vehicle-string literature showed for predecessor following on relative spacing alone [4], and the best of 90 gain pairs here still peaks at 1.050. A radio that delivers the predecessor's acceleration restores stability down to a 0.41 s gap at 50 ms of delay, but needs 0.74 s at 400 ms with one packet in five lost. None of this is visible in the poles of one follower, because the property belongs to the chain.

DID YOU KNOW? Production Time Gaps Span This Chapter's Cases

Models fitted to car-following data from seven production adaptive cruise controls gave effective time gaps from 0.544 s to 2.054 s across their shortest and longest settings, with lags, perception included, from 0.153 s to 0.994 s [5]. The 0.6 s and 1.5 s gaps of this chapter lie inside that range.

The structure is general. String stability was posed for countably infinite interconnections of nonlinear systems, with weak-coupling conditions that guarantee it [2], and $\mathscr{L}_p$ definitions extend it to nonlinear cascades [6]. Outside vehicles, the orders that each company in a supply chain places with its upstream neighbor vary more than its sales, and the distortion grows moving upstream, the bullwhip effect [7]. Volume 6, Chapter 11 measures that amplification on a dairy supply chain. A ring of such stages closes the chain on itself, and on a circuit of 22 cars a uniform flow breaks into a stop-and-go wave with no bottleneck at all [8]. In each case the stage-to-stage gain at every frequency, not the stability of any one stage, decides what the chain does.

However, the experimental record measures amplification on instrumented cars [3, 6, 9] and on test tracks [8, 10], so a student seldom has a plant on which the frequency-domain criterion and the measured response of the chain can be compared car by car while the policy, the gains, the radio or the drivers change one at a time. This chapter pairs the string-stability criterion, computed live from the linearization of the very law each car executes, with the amplification measured along the nonlinear chain of a browser platoon simulator. The spacing policies, the cooperative feedforward, the driver model and the criterion are established methods from the literature. In the simulator each controller sees the line only through a delayed radar, a lossy radio, its own wheel speed and accelerometer, or a driver's delayed eye, and every equation, constant, delay and limit appears below. Its ten cases and a challenge isolate the policy, the gains, the radio, a cut-in, a ring and a predictive controller in every follower.

Section 14.2 traces the problem from the first vehicle strings to cooperative platoons. Section 14.3 describes the line, its sensors, its radio and what is measured. Section 14.4 derives the criterion from the cars' own law, for radar ACC with a time gap, constant spacing, cooperation over a delayed radio and human drivers. Section 14.5 gives the lead's cruise control, the followers' laws, the supervisor and the distributed predictive controller. Section 14.6 reports the measured runs of the ten cases and the challenge, Section 14.7 sets the radar, the radio and the throttle and brakes beside their real counterparts, Section 14.8 states the limitations and Section 14.9 concludes.

14.2 From Vehicle Strings to Cooperative Platoons

Automatic following began as a problem of densely packed strings of high-speed vehicles, and the first history segment tells how string stability was named.

HISTORY Strings of Vehicles and Automated Highways, 1964 to 1980

Automatic following began as a problem of densely packed strings of high-speed vehicles. In 1966 the theory of optimal control gave a linear feedback that regulates the position and velocity of every vehicle in such a string, with analog computer simulations of a string of three [11]. A long-range program at The Ohio State University studied headway safety policy, longitudinal and lateral control and highway system operations from 1964 to 1980, and a 6.4 km automated vehicle control facility was almost complete when the work was terminated [12]. By 1974 string stability was named as an important performance criterion for longitudinal control. A vehicle that follows its own moving reference cell always has it, because no vehicle then reacts to another. Without such a reference, relative-motion systems could reach it only with both forward and rearward separation measurements [1].

The California PATH program was founded in 1986 as the first research program in North America on what is now called intelligent transportation systems [13]. In the decade that followed, a comparison of the two longitudinal policies gave a systematic design of constant-spacing controllers

that guarantee string stability and found that a constant headway, or time gap, assures it without communication between the vehicles, which suits it to traffic that mixes automated and manual vehicles and to autonomous intelligent cruise control [14]. String stability was then given a formal definition for countably infinite interconnections [2]. A team that included engineers of California PATH ran automated vehicles in platoons under an integrated longitudinal and lateral control system, string stable at very small inter-vehicle spacing from a stop to high-speed cruising and through split and join maneuvers under communication constraints, with experimental results from the NAHSC demonstration of automated highways in San Diego [15]. The limitation of constant spacing on relative information alone was later traced to a complementary sensitivity integral constraint [4].

Adaptive cruise control reached the market with the time gap as a setting the driver chooses. A microscopic simulation built on the time gaps that drivers from the general public used in a field experiment found that ACC was unlikely to change lane capacity by much, while cooperative ACC raised it greatly once its market penetration reached moderate to high percentages [16]. Whether production systems damp traffic waves was later measured directly on seven 2018 model-year vehicles widely available in the United States. Models identified from more than 1200 miles of car-following data showed all seven string unstable, and in a platoon of one model that offers ACC on all trim levels the growth of the disturbance ended with the last car's ACC disengaging [3].

Cooperative adaptive cruise control (CACC) adds the predecessor's acceleration received over a wireless link. A decentralized design with a link to the nearest preceding vehicle only was given a necessary and sufficient frequency-domain condition for string stability, which showed that the link permits small distances under a velocity-dependent spacing policy while a constant spacing cannot be guaranteed string stable, and it was validated on two equipped cars [9]. An $\mathscr{L}_p$ definition for nonlinear cascades was validated on a platoon of six passenger vehicles with CACC [6]. The link is also the weakness. Latency and packet loss reach the controller, and under persistent loss CACC degrades to ACC and needs a larger distance to stay string stable, which motivated a graceful degradation that estimates the predecessor's acceleration with onboard sensors [17].

The ring road showed that a line needs no leader to amplify, as the second history segment tells.

> ⧖ HISTORY **A Jam With No Bottleneck, 2008**
>
> An experiment reported in 2008 put 22 cars on a circuit of 230 m. The drivers were asked to cruise at about 30 km/h and to follow the car ahead safely. Small fluctuations in the gaps grew until several cars stopped for a moment, and a jam formed with no bottleneck at all. Cars inside the jam stood still while those outside moved at about 40 km/h, and the jam traveled backward at roughly 20 km/h, the common speed of a jam on real highways [8]. A later field experiment on a 260 m track with 21 or 22 vehicles found that such waves emerged consistently and that controlling the velocity of a single vehicle in the flow dampened them [10].

What the application adds to these experiments is a chain on which the frequency-domain criterion and the measurement that the field studies made can be read together, one car at a time.

14.3 The Line, Its Sensors and Its Radio

The simulated system is one straight lane of identical cars, drawn in Figure 14.1. The cars are numbered from the front, car 1 being the lead, and a line holds $N = 10$ by default and up to 24. Two further slots hold a car that cuts in from the next lane and a slower car that may pull in ahead of the lead, so the state vector has 26 slots of three states each, 78 in all, and a slot that is not on the road has zero derivatives. On the ring there is no lead, and car 1 follows car N a lap ahead. Table 14.1 classifies the signals as controlled (CV), manipulated (MV) and disturbance (DV) variables, states, measurements and outputs, with SP for a setpoint.

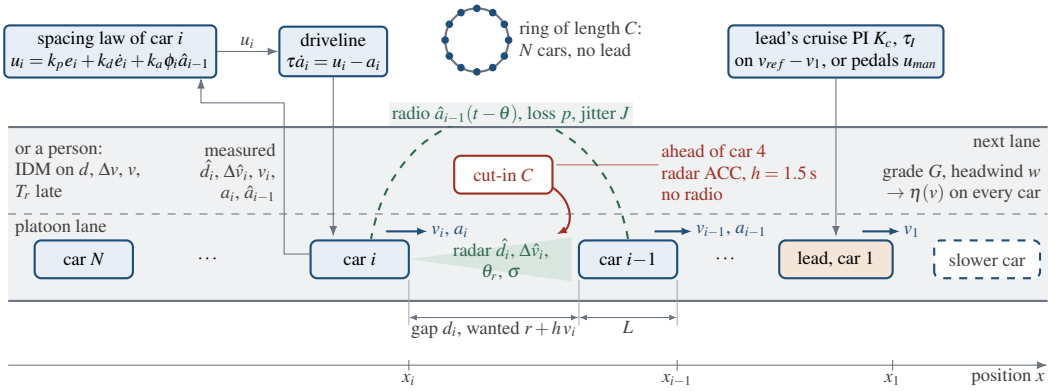


Figure 14.1 The line of cars, its sensors and its radio, not to scale. Car i has its front bumper at x_i, length L, speed v_i and delivered acceleration a_i, which follows the command u_i through the driveline lag τ. Its radar measures the gap d_i to the car ahead and the gap's rate of change as $\hat{d}_i$ and $\Delta\hat{v}_i$, its own speed and acceleration are measured exactly, and the radio delivers the acceleration of car $i-1$ after a delay θ with loss p and jitter J. The spacing law wants a gap of $r + h v_i$. The lead runs a cruise PI on its speed reference v_{ref}, which follows the set speed v_{set} through the chosen profile, or takes the operator's pedals, a car may cut in ahead of car 4, a slower car may pull in ahead of the lead, grade G and headwind w act on every car, and on a ring of length C the N cars follow one another with no lead.

Table 14.1 Signals of the Platoon by Role, With Their Defaults

Symbol	Role	Default	Range
v_1	CV of the lead's cruise loop	SP 72 km/h	10–130
d_i	CV of follower i	SP $r + h v_i$	≥ 0
u_1	MV, the lead's command	held 50 ms	-6 to 2 m/s^2
u_i	MV, a follower's command	held 50 ms	-6 to 2 m/s^2
u_{man}	MV in MANUAL, the lead's pedals	0	-6 to 2 m/s^2
x_i, v_i, a_i	states of 26 slots	laid out	—
G, w	DVs, grade and headwind	0, 0	± 6 %, ± 15 m/s
—	DVs, tap, hard stop, cut-in, radio burst, slower car	per case	events
$\hat{d}_i$, $\Delta\hat{v}_i$	measured, radar	50 ms late	σ 0.05
v_i, a_i	measured, wheel speed and accelerometer	exact	—
$\hat{a}_{i-1}$	measured, radio	50 ms late	p, J
γ, $\lvert\Gamma\rvert_{max}$	outputs, measured and predicted amplification	—	—
q, d_{min}	outputs, lane flow, smallest gap	—	—

14.3.1 The Cars

Every car is a point mass whose delivered acceleration follows its command with a first-order lag, as shown in Equation (14.1).

$$\dot{x}_i = v_i, \qquad \dot{v}_i = a_i + \eta(v_i), \qquad \tau \dot{a}_i = u_i - a_i \tag{14.1}$$

where u_i is the command that the car's controller or driver holds over each 50 ms control period, already clipped to $[a_{min}, a_{max}] = [-6, 2]$ m/s^2, $\tau = 0.4$ s is one lag for throttle and brake alike, and η is the disturbance on $\dot{v}$ that grade and headwind add, Equation (14.2). A car whose speed has reached zero stays stopped while $a_i + \eta < 0$, because brakes hold and nothing reverses.

$$\eta(v) = -\frac{g\,G}{100} - k_w\left[(v+w)\lvert v+w\rvert - v\lvert v\rvert\right] \tag{14.2}$$

where $g = 9.81\,\mathrm{m/s^2}$, G is the grade in percent with uphill positive, w is the headwind and $k_w = 2.6 \times 10^{-4}\,\mathrm{m^{-1}}$ is the drag coefficient $\frac{1}{2}\rho C_d A/m$ that turns the headwind into a deceleration. The bracket subtracts the drag of still air, so on a level road in still air no road load acts and a coasting car keeps its speed. The gap to the car ahead is Equation (14.3).

$$d_i = x_{p(i)} - x_i - L + C_i \tag{14.3}$$

where $p(i)$ is the car ahead of car i, $L = 4.5\,\mathrm{m}$, and C_i equals the ring length C for car 1 of a ring, whose car ahead is car N a lap further on, and zero otherwise. Nothing overtakes, so the order never changes. A gap that reaches zero is a collision. The two cars share their momentum as equal masses in a perfectly inelastic impact, are placed touching, brake at a_{min} to a standstill and stay out of the run, which pauses at the first one. The speed floor, the collision test, the cut-in car's lane change, the slower car leaving and every measured quantity are handled between integration steps, so the derivative function itself contains no switch.

14.3.2 Spacing Policies and the Following Law

A radar follower wants a gap that is either fixed or grows with its own speed, as shown in Equation (14.4).

$$d_i^* = \begin{cases} r + h\,v_{set} & \text{constant spacing} \\ r + h\,v_i & \text{constant time gap} \end{cases}, \qquad e_i = d_i - d_i^* - \delta_i \tag{14.4}$$

where $r = 3\,\mathrm{m}$ is the standstill gap, h the time gap, v_{set} the lead's set speed and $\delta_i \le 0$ the supervisor's allowance after a cut-in (Section 14.5.3). Constant spacing holds the gap that the time gap would give at the set speed, 33 m at 72 km/h for $h = 1.5\,\mathrm{s}$, whatever the car's own speed. Radar ACC and cooperative ACC run Equation (14.5) on the error.

$$u_i = k_p e_i + k_d \dot{e}_i + k_a \phi_i \hat{a}_{i-1}, \qquad \dot{e}_i = \Delta\hat{v}_i - h\,a_i \tag{14.5}$$

where $\Delta\hat{v}_i$ is the radar's measure of the rate at which the gap opens, $v_{p(i)} - v_i$, a_i is the car's own accelerometer reading of $\dot{v}_i$, $k_p = 2\,\mathrm{s^{-2}}$ and $k_d = 1\,\mathrm{s^{-1}}$ are the time-gap gains, $k_a = 1$ is the radio feedforward gain of cooperative ACC and zero for radar ACC, and $\phi_i \in [0, 1]$ is the supervisor's fade of the radio term. For constant spacing the $h\,a_i$ term is absent, because the wanted gap does not move with the car's speed. Under a time gap it is the derivative of the speed-dependent part of the wanted gap, so the law damps itself with the car's own acceleration. Constant spacing runs its own gains, 0.3 and 1.0, because the time-gap gains leave that loop with a damping ratio of 0.054 (Section 14.4). An ACC also respects a set speed of its own, $1.2\,v_{set}$, by taking the smaller of u_i and $\frac{1}{2}(1.2\,v_{set} - v_i)$.

14.3.3 People

A follower may instead be a person. Each driver follows the intelligent driver model [18] on the gap, the approach rate and their own speed, all seen a reaction delay T_r late, as shown in Equation (14.6).

$$\begin{aligned} u_i(t) &= a_m \left[1 - \left(\frac{v}{v_0} \right)^4 - \left(\frac{s^*}{s} \right)^2 \right]_{t - T_r} \\ s^* &= s_0 + v\,T_g + \frac{v\,\Delta v}{2\sqrt{a_m b}} \end{aligned} \tag{14.6}$$

where $s = d_i$, $v = v_i$, $\Delta v = v_i - v_{p(i)}$ is the approach rate, $v_0 = 120\,\mathrm{km/h}$ is the desired speed, $T_g = 1.6\,\mathrm{s}$ the desired time gap, $a_m = 0.73\,\mathrm{m/s^2}$ the maximum acceleration, $b = 1.67\,\mathrm{m/s^2}$ the comfortable deceleration and $s_0 = 2\,\mathrm{m}$ the jam distance, values the application takes from that paper. Each driver records the true gap, approach rate and speed at every 50 ms tick and acts on the record $T_r = 0.70\,\mathrm{s}$ old, fourteen ticks. The equilibrium gap $(s_0 + v T_g)/\sqrt{1 - (v/v_0)^4}$ is 36.44 m at 72 km/h and 15.36 m at 30 km/h, so 22 such drivers hold 30 km/h on a ring of $22(L + 15.36) = 437.0\,\mathrm{m}$, and on the 230 m circuit of the ring experiments [8] they would settle at 8.90 km/h. The reaction delay is a calibrated teaching value, because at 0.8 s the model collides (Section 14.6).

14.3.4 Radar and Radio

The radar and the radio are the same sampled link seen two ways, the application's channel model, which now lives in the framework as `shared/src/framework/core/channel.js`. Every 50 ms the sender stamps a sample with the time it was taken, and the link delivers or drops it as shown in Equation (14.7).

$$\begin{aligned} t_{arr,k} &= t_k + \theta + J\nu_k, \quad \text{dropped if } \mu_k < p \text{ or } t_k < t_{burst} \\ \hat{y}(t) &= y(t_{k^*}), \quad k^* = \arg\max_k \{t_k : t_{arr,k} \le t\} \end{aligned} \tag{14.7}$$

where θ is the fixed delay, J the largest jitter, p the loss probability, t_{burst} the end of a burst that drops every packet, and ν_k and μ_k two uniform deviates that every packet draws from its link's own seeded generator whatever the settings, so a link at 10 % loss and the same link at 20 % lose the same packets first. The receiver keeps the newest packet that has arrived and discards one overtaken by a newer one. A sample is used at the first 50 ms tick after it arrives, so the radar's 50 ms latency is served exactly and a latency of 120 ms becomes 150 ms, which the application's own check confirms. The radar's range and range rate carry Gaussian noise of 0.05 m and 0.05 m/s from a third generator of each car. The radio carries the predecessor's measured acceleration with a default delay of 50 ms. A link counts as stopped after 0.5 s, ten periods, without an arrival, and a burst drops every packet of every link for 6 s. The car that cuts in and the slower car ahead of the lead have no radio. Volume 5, Chapter 8 meets the same seeded delay, jitter and loss on the link of a remote manipulator.

14.3.5 Lead, Events and the Ring

The lead's set speed is 72 km/h unless a case names another. A profile moves its reference, holding steady cruise, running a stop and go that holds the set speed for 20 s, slows at 1.5 m/s^2 to a third of it, holds 20 s and returns at 0.8 m/s^2, repeating every 100 s, or adding a sine of ± 1 m/s at a chosen period. A brake tap overrides the cruise control with 3 m/s^2 of braking for 1.5 s. Scripted and learner taps hold 30 control periods, 1.5 s, with endpoint comparisons tolerant of floating-point roundoff. The commanded speed decrement is 4.5 m/s. A hard stop brakes at 5 m/s^2 to a standstill and holds 4 s. A cut-in places a car with no radio in the middle of the gap ahead of car 4, at car 4's speed and only if that gap exceeds a car length plus 1 m and holds 0.5 m in a stop at the line's present deceleration answered 0.45 s later at half the braking limit, running radar ACC at $h = 1.5$ s with $k_p = 1\,\mathrm{s}^{-2}$ and $k_d = 1\,\mathrm{s}^{-1}$. A slower car pulls in 3 m plus 1 s of travel ahead of the lead, holds 70 % of the lead's speed v_S with the command $\frac{1}{2}(v_S - v)$, and leaves the lane after 30 s. On the ring, $N = 22$ drivers start evenly spaced on $C = 437$ m at the speed their policy settles at for that spacing, each moved by at most 10 cm by a seeded uniform draw.

14.3.6 What Is Measured

The flow is computed from the platoon's own spacing, as shown in Equation (14.8).

$$q = 3600\,\frac{\bar{v}}{\bar{\ell}}, \qquad q_h = \frac{3600}{h + (L + r)/v} \tag{14.8}$$

where q is in vehicles per hour, $\bar{v}$ is the mean speed of the cars on the road, $\bar{\ell}$ their mean front-to-front spacing, and q_h the flow of a time-gap policy at speed v. At 72 km/h a 1.5 s gap gives 1920 veh/h and a 0.6 s gap 3692 veh/h. The lane-flow tile of the scorecard averages q over the last 30 s, sampled every 0.5 s. The amplification is measured as a ratio of root-mean-square accelerations, Equation (14.9).

$$\gamma_i = \sqrt{\frac{\int a_i^2 \,\mathrm{d}t}{\int a_{p(i)}^2 \,\mathrm{d}t}}, \qquad \gamma = \max_i \gamma_i \tag{14.9}$$

where both integrals run over the same window and a here is the kinematic acceleration $\dot{v}$. A pair counts once the car ahead has passed back at least 0.5 m^2/s^3, an RMS of 0.1 m/s^2 for 50 s, and only

Table 14.2 Model Constants and Where Each Is Set: V in the Show Model Dialog, A Among the Advanced Controls, H Hidden From the Student, C in the Source

Symbol	Meaning	Value	Class
Cars			
L	car length	4.5 m	V
a_{max}, a_{min}	acceleration and braking limits	2, –6 m/s^2	V
k_w	drag per unit mass, headwind only	2.6×10^{-4} m^{-1}	V
τ	driveline lag	0.4 s	A
—	cut-in car's speed relative to car 4	0	H
Sensors and radio			
θ_r, σ	radar latency, noise on range and rate	50 ms, 0.05	A
θ, p, J	radio delay, loss, jitter	50 ms, 0, 0	A
—	link timeout, burst length	0.5, 6 s	C
People, intelligent driver model			
v_0, T_g	desired speed and time gap	120 km/h, 1.6 s	A
a_m, b	maximum acceleration, comfortable braking	0.73, 1.67 m/s^2	A
s_0	jam distance	2 m	A
T_r	reaction delay	0.70 s	A
Line, ring and measurement			
N	cars in the line or on the ring	10, or 22	A
C	ring length	437 m	A
—	ring start jitter	±10 cm	C
—	energy threshold, quiet time	0.5 m^2/s^3, 15 s	C
Numerics			
Δt, T	integration step, control period	10, 50 ms	C
—	log rate, seed	10 Hz, 20261001	C

while the whole line has run its policy for 15 s, untouched by a cut-in's gap regrowth or a fallback's change of gap. A pair starts again when a car's predecessor changes, the pair behind the slower car is never counted, and on the ring neither is car 1's. The RMS ratio is the gain that the criterion of Section 14.4 bounds, an $\mathscr{L}_2$ notion of string stability [6], and Section 14.4.6 shows why a ratio of peaks is not used.

14.3.7 Constants, Start and Integration

Table 14.2 gathers the model's constants, and the settings a student can move appear in Table 14.3. Every case on the open road starts with the line at its policy's gap for the lead's set speed and every car at that speed with zero acceleration, an equilibrium of Equations (14.1) to (14.5) to within 4×10^{-16} by the application's check, and each sensor starts as though it had been tracking the car ahead, with samples already in flight. Classical fourth-order Runge-Kutta integrates the 78 states on a fixed 10 ms step, every controller and sensor runs at 20 Hz, and the historian records at 10 Hz, with controllers, scripted events and the historian on the engine's schedule of Section 1.6.1. The default seed is 20261001, and each car's sensors draw from generators of their own, so a case replays bit for bit and the losses of one link do not change when another car's settings do. Halving the step changes the measured amplification of the 0.6 s case and of the human case by a relative 1.3×10^{-6}. The application passes its 53 verification checks.

14.3.8 Assumptions

One lane, longitudinal motion only. A cut-in joins the line at once in the middle of the gap, and its lane change is only drawn. Grade and headwind act on every car at once, as a long uniform slope and a steady wind would.

Identical cars with one lag. Every car has the same 0.4 s lag for throttle and brake and the same hard limits. A real driveline brakes faster than it accelerates, and its lag changes with speed, gear and load.

No road load in still air. The drag of still air and rolling resistance are left out, so the lead's cruise control has no steady load to hold and a coasting car keeps its speed. Aerodynamic drag and rolling resistance would add a steady load that grows with speed on every car.

People who look one car ahead. The drivers are identical, noiseless and a pure delay late, and they watch only the car ahead. Real drivers anticipate and watch several cars ahead, and without that this model collides from 0.8 s of delay.

An ideal radar and an independent radio. The radar sees exactly the car ahead at any range. The radio's losses are independent, the clocks of the cars agree, and loss does not grow with the number of cars talking, as it would on a congested channel.

A criterion from one linearization. The criterion is computed at the line's mean follower speed and describes small waves. For the radar and radio laws, which are linear, it holds until a limit is reached, while for people a large wave grows faster than the linear peak suggests.

Teaching values. A collision is perfectly inelastic, and nothing about injury or damage is modeled. Every threshold, the challenge's among them, is a teaching value, and the application names ISO 15622 for adaptive cruise control and ISO 20035 for cooperative adaptive cruise control as the sources of real requirements.

14.4 The Criterion From the Cars' Own Law

The application never writes down a model of a follower. It linearizes the function that the cars execute, so the criterion and the chain cannot disagree about what the law is.

14.4.1 Linearization

A follower's command is one function of five measured quantities. Central differences of that function about the operating point give five gains, g_s on the gap, g_{dv} on its opening rate, g_v on the car's own speed, g_{ai} on its own acceleration and g_a on the predecessor's broadcast acceleration. With $A_i(s)$ the transform of a_i, the gap and the opening rate written through the two accelerations as $(A_{i-1}-A_i)/s^2$ and $(A_{i-1}-A_i)/s$, each input arriving late by its own delay and the command held over the control period, the ratio $\Gamma(s)=A_i(s)/A_{i-1}(s)$ is Equation (14.10).

$$\begin{aligned}\Gamma(s) &= \frac{N(s)}{D(s)}, \qquad H(s)=\frac{1-e^{-sT}}{sT}\\ N(s) &= H\left[e^{-\theta_d s}(g_s+g_{dv}s)+e^{-\theta_c s}g_a s^2\right]\\ D(s) &= \tau s^3+s^2+H\left[e^{-\theta_d s}(g_s+g_{dv}s)\right.\\ &\qquad \left.-e^{-\theta_v s}g_v s-g_{ai}s^2\right]\end{aligned} \tag{14.10}$$

where $T=50\,\text{ms}$, θ_d is the radar's latency rounded up to whole ticks or the driver's delay, θ_v the driver's delay on their own speed and zero for an ACC, and θ_c the radio's effective delay. A lost packet leaves the last good one in use for a geometric number of extra periods, with mean $p/(1-p)$, so θ_c is the mean of $T\lceil(\theta+Jv)/T\rceil$ over the jitter plus $Tp/(1-p)$, and 20 % loss adds 12.5 ms. For radar ACC the gains are the law's own, $g_s=k_p$, $g_{dv}=k_d$, $g_v=-k_p h$, $g_{ai}=-k_d h$ and $g_a=k_a$, and the units make Γ dimensionless. The line is string stable when $|\Gamma(j\omega)|\le 1$ at every frequency [2, 9], as in Section 1.4.10. Because a lasting change of speed is passed on whole, $\Gamma(0)=1$ for every policy, so a stable line has its largest $|\Gamma|$ at zero frequency and an unstable one has a peak above one somewhere. The application finds the peak on 161 logarithmically spaced frequencies from 0.003 to 30 rad/s, refined by 30 golden-section steps, and counts the line stable when the peak is at most $1+10^{-6}$.

The application's own check drives the lead with a sine of 0.2 m/s and requires the ratio passed from car 2 to car 3 to match Equation (14.10) within 1 % for three configurations. Repeated at the peak frequency of five unstable configurations and at 1 rad/s for two stable ones, the two agree within

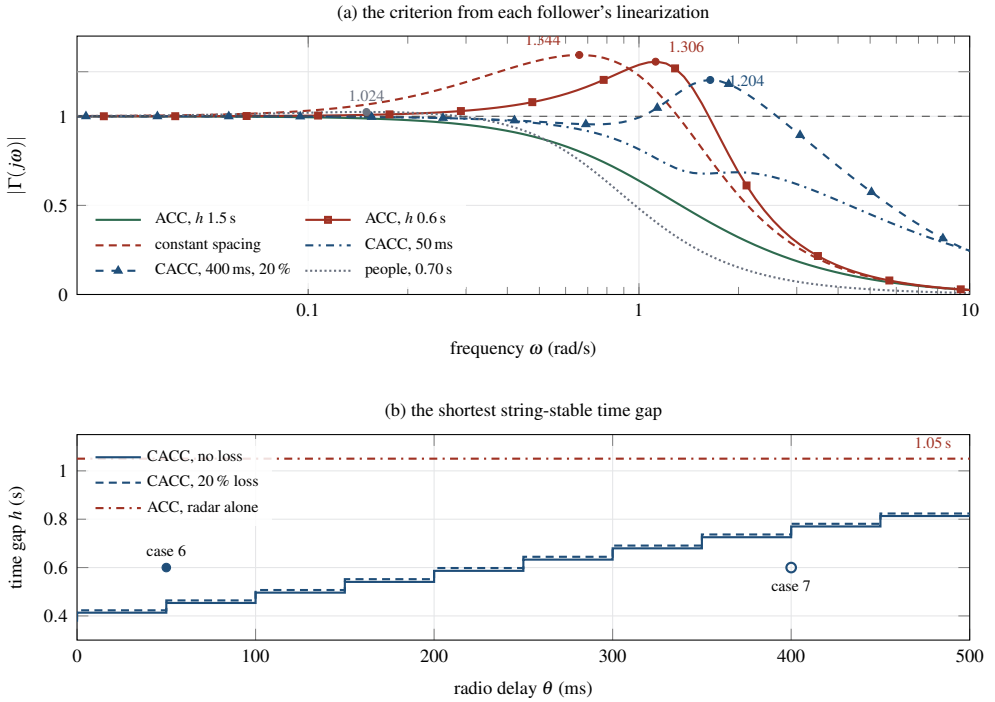


Figure 14.2 The string-stability criterion at 72 km/h, from the application's own linearization. (a) $|\Gamma(j\omega)|$ for radar ACC at time gaps of 1.5 and 0.6 s, constant spacing with its gains of 0.3 and 1.0, cooperative ACC at 0.6 s with a 50 ms radio and with 400 ms and 20 % loss, and drivers 0.70 s late, with the unstable peaks marked. Radar ACC at 1.5 s is solid, at 0.6 s solid with squares, constant spacing densely dashed, cooperative ACC with the fast radio dash-dotted and with the slow one dashed with triangles, and the drivers dotted. (b) The shortest time gap at which cooperative ACC with $k_a = 1$ is string stable, rising in steps of one 50 ms tick of radio delay, solid without loss and dashed at 20 % loss, against the 1.05 s that radar ACC needs, dash-dotted. Case 6, the filled circle, lies in the stable region above the steps and case 7, the open circle, below them.

0.19 %, and within 0.27 % for a stable line at 30 % packet loss. Figure 14.2(a) draws the criterion for six followers at 72 km/h.

14.4.2 Radar ACC With a Time Gap

With no delays and no hold, Equation (14.10) reduces for radar ACC to the transfer function of Equation (14.11).

$$\Gamma(s) = \frac{k_d s + k_p}{\tau s^3 + (1 + k_d h)s^2 + (k_d + k_p h)s + k_p} \tag{14.11}$$

The application's check confirms that its linearization reproduces it within a relative 4.1×10^{-11}. With no lag, $\tau = 0$, it is Equation (1.11) with h in place of t_g. With N and D now the numerator and denominator of Equation (14.11), string stability asks $|D|^2 - |N|^2 \geq 0$ at every frequency, and the difference is the polynomial of Equation (14.12).

$$\begin{gathered} |D|^2 - |N|^2 = c_2\omega^2 + c_4\omega^4 + \tau^2\omega^6 \\ c_2 = k_p(k_p h^2 - 2), \qquad c_4 = (1 + k_d h)^2 - 2\tau(k_d + k_p h) \end{gathered} \tag{14.12}$$

Near zero frequency $|\Gamma|^2 \approx 1 + \omega^2(2/k_p - h^2)$, so the line is stable only for $h \geq \sqrt{2/k_p}$, 1.00 s at $k_p = 2\,\mathrm{s}^{-2}$. The polynomial is non-negative at every frequency exactly when $c_2 \geq 0$ and either $c_4 \geq 0$

or $c_4^2 \le 4\tau^2 c_2$, and at the default gains this full condition gives the same 1.0000 s, because $c_4 = 1.6$ there. The radar's 50 ms and the hold move the boundary to 1.050 s. At k_p of 1 and 3 the bound is 1.414 s and 0.816 s, and 1.463 s and 0.868 s with the delays. The time gap helps twice. The wanted gap grows with the car's own speed, which supplies the $k_p h$ term of the s coefficient, and the rate of the error carries the car's own acceleration, which supplies the $k_d h$ term of the s^2 coefficient.

The denominator of Equation (14.11) is also the characteristic polynomial of one follower behind a car at constant speed, and it says nothing alarming. At $h = 1.5$ s its complex pair has a damping ratio of 0.968, and at 0.6 s the roots are $-0.684 \pm 1.197\,j$ and -2.632 s^{-1}, a damping ratio of 0.496 at a natural frequency of 1.378 rad/s. The same loop has its $|\Gamma|$ peak of 1.306 at 1.12 rad/s, near that natural frequency, and stays above one from zero up to 1.63 rad/s. The peak is 1.0035 at 1.0 s, 1.084 at 0.8 s and 1.306 at 0.6 s, as shown in Figure 14.2(a). Raising k_p to 3 moves the peak to 1.45 rad/s and leaves it at 1.305 while the damping ratio falls to 0.456, and raising k_d to 2 lowers the peak to 1.089 while the damping ratio rises to 0.773.

Pause. Try Exercise 14.1 before reading on.

14.4.3 Constant Spacing

Without the $h\,v_i$ term the h in the coefficients disappears, and the transfer function becomes $(k_d s + k_p)/(\tau s^3 + s^2 + k_d s + k_p)$, with $c_2 = -2k_p$. Near zero frequency $|\Gamma|^2 \approx 1 + 2\omega^2/k_p$, above one for every positive gain, the limitation that the literature traces to a complementary sensitivity integral [4]. Two integrators lie between a car's acceleration and its gap, and a fixed gap asks the car to match every speed change of the car ahead at once. The shipped gains of 0.3 and 1.0 give each car a damping ratio of 0.816 and the line a peak of 1.344 at 0.66 rad/s. The time-gap gains of 2 and 1 leave the loop damped at only 0.054, with a peak of 131.7. Over 90 gain pairs, k_p from 0.05 to 3 and k_d from 0 to 3, the smallest peak is 1.050 at $k_p = 0.05$ and $k_d = 1$, the lower end of the spacing-gain slider, a loop with a real pole at –0.053 s^{-1}, so slow that its gap drifts.

Pause. Try Exercise 14.2 before reading on.

14.4.4 Cooperation and Delay

The radio adds $-k_a\omega^2$ to the real part of the numerator and changes the low-frequency coefficient to $c_2 = k_p\big(k_p h^2 - 2(1-k_a)\big)$, so near zero frequency $|\Gamma|^2 \approx 1 + \omega^2\big(2(1-k_a)/k_p - h^2\big)$. At $k_a = 1$ the bracket is negative at any gap, and the bound moves to higher frequencies, where the lag and the delays decide it. With no delays at all the shortest stable gap is 0.342 s, the radar's latency and the hold raise it to 0.376 s with an instant radio, and the default 50 ms radio to 0.413 s. Each further 100 ms of radio delay costs about 0.09 s of gap, to 0.541 s at 200 ms, 0.725 s at 400 ms and 0.813 s at 500 ms, and 20 % loss adds 0.009 to 0.012 s more, 0.737 s at 400 ms. Every delay is served at a 50 ms tick, so the boundary rises in steps of one tick, as shown in Figure 14.2(b). The feedforward gain has an optimum near one. At 50 ms the boundary is 0.758 s at $k_a = 0.5$, 0.505 s at 0.8, 0.413 s at 1.0 and 0.464 s at 1.2.

14.4.5 Drivers Seen Through the Linearization

For a driver the gains are the slopes of Equation (14.6) at its equilibrium, all delayed by T_r. At 72 km/h they are $g_s = 0.0349\,\text{s}^{-2}$, $g_{dv} = 0.339\,\text{s}^{-1}$ and $g_v = -0.0787\,\text{s}^{-1}$, and the peak is 1.024 at 0.15 rad/s, a period of 42 s. A driver 0.70 s late is thus barely string unstable, and only for slow waves, while reaction delays of 0.5, 0.6, 0.75, 0.8 and 1.0 s give peaks of 1.017, 1.020, 1.027, 1.031 and 1.084. At 30 km/h, the speed of the ring, the slopes grow to $g_s = 0.0947\,\text{s}^{-2}$ and $g_v = -0.153\,\text{s}^{-1}$ and the peak to 1.198 at 0.43 rad/s. These are small-wave numbers. The braking term of Equation (14.6) grows with the square of s^*/s, and s^* grows with the approach rate, so a driver who is late and close brakes far harder than the linearization predicts.

Table 14.3 Controller Settings, Defaults and Ranges

Symbol	Setting	Default	Range
Lead			
v_{set}	set speed	72 km/h	10–130
K_c, τ_I	cruise PI gain and integral time	0.8 s^{-1}, 10 s	0.1–3, 1–60
—	speed profile	cruise	3 profiles
P	sine period	6 s	3–60
Followers			
—	policy	ACC, time gap	6 policies
h	time gap	1.5 s	0.3–2.5
k_p, k_d	spacing and rate gains	2, 1	0.05–3, 0–3
r	standstill gap	3 m	1–10
Radio and supervisor			
k_a	radio feedforward gain	1	0–1.5
h_{fb}	fallback time gap	1.5 s	0.6–3
ρ	gap regrowth after a cut-in	1 m/s	0–3
—	radio fade, widening, narrowing	0.5, 3, 6 s	fixed
—	MPC preview	plan	3 settings

Units of k_p and k_d are s^{-2} and s^{-1}. The six policies are people, ACC with constant spacing, ACC with a time gap, CACC, CACC with fallback and MPC (Section 14.5.4).

14.4.6 Which Ratio the Criterion Bounds

Equation (14.10) bounds the ratio of RMS values. By Parseval's theorem the energy of a_i is the integral of $|\Gamma|^2$ times the energy spectrum of $a_{p(i)}$, at most $|\Gamma|^2_{max}$ times the energy of $a_{p(i)}$. The ratio of peak accelerations is bounded instead by the integral of the absolute impulse response, which exceeds $|\Gamma(0)| = 1$ whenever the response changes sign, so a string-stable line can pass back larger peaks. With the lead of case 4 switched to the stop and go profile from 72 km/h, its scripted tap included, ACC at 1.2 s passes back peaks 1.0130 times larger at an RMS ratio of 0.9904 over 100 s, and cooperative ACC at 0.6 s 1.0164 at 0.9902. The tap's own peak hides the effect at 1.5 s (0.9869 at 0.9845), but a model run without the tap gives 1.0096 at 0.9867 there. The measured amplification of Equation (14.9) is therefore the $\mathscr{L}_2$ quantity, the one the criterion can be checked against [6].

Pause. Try Exercise 14.3 before reading on.

14.5 Cruise Control, Following Laws and the Supervisor

Every controller runs every 50 ms in a fixed order. Each car with a radio broadcasts its acceleration to the car behind it, each radar takes its sample, everything that has arrived is delivered, the lead's cruise control runs, every follower runs its law under the platoon supervisor, and the criterion of Equation (14.10) is recomputed when a setting it depends on changes, every 5 s, or when the speed of a line of people drifts by 0.25 m/s. Each command is held until the next tick, and no controller reads another car's state except through its sensors. The settings are listed in Table 14.3.

14.5.1 The Lead's Cruise Control

The lead's cruise control is the framework's velocity-form PID of Section 1.6.2, here a proportional-integral loop on the speed reference whose output is the acceleration command, as shown in Equation (14.13).

$$\begin{aligned} e_k &= v_{ref,k} - v_{1,k}, \qquad S_k = S_{k-1} + K_c e_k - P_{k-1} \\ \tilde{u}_k &= S_k + I_{k-1}, \qquad u_{1,k} = \min\left(a_{max}, \max(a_{min}, \tilde{u}_k)\right) \\ I_k &= I_{k-1} + \left[\frac{K_c}{\tau_I} e_k + \frac{u_{1,k} - \tilde{u}_k}{\tau_I}\right] T \end{aligned} \tag{14.13}$$

where $P_{k-1} = K_c e_{k-1}$ is the proportional term formed at the previous scan, $K_c = 0.8\,s^{-1}$, in m/s^2 per m/s, $\tau_I = 10$ s, the tracking time of the back-calculation equals τ_I, T is the 50 ms control period,

and there is no derivative action and no rate limit. These are Equations (1.13) to (1.15) without the derivative term, and Volume 2, Chapter 2 tunes the same loop on a single car. The unclamped running sum S_k telescopes to $K_c e_k$, the clip acts on the emitted output only, and the anti-windup acts on the integral only, so the law is the positional PI with back-calculation on its integral, at the limits as well. Fed the same errors through a run whose set-speed steps hold the command at a limit for 153 of 1800 scans, the two give identical commands. A tap or a hard stop replaces the command while it lasts. When it ends, the PI restarts from zero output and the reference restarts at the speed the car is doing, climbing back to the profile at 1 m/s^2, as a production cruise control resumes from coasting. In MANUAL the operator's pedal value u_{man} is the command and the integral is set each tick so that the PI's output equals it, so a switch of mode in either direction moves nothing. Behind a slower car the lead takes the smaller of the PI's command and that of a radar ACC at 1.5 s on the slower car.

14.5.2 The Followers

Each follower evaluates Equation (14.5), or Equation (14.6) for a person, on what its sensors delivered at this tick, clips the result to the acceleration limits and holds it. A radar car whose first sample of a new target has not yet arrived keeps its last command, and a car that has collided brakes at a_{min}. A person's law reads the record kept 14 ticks earlier in a history of 64 ticks, 3.2 s. The rate of the error comes from the radar's range rate and the accelerometer, so no signal is differentiated and no derivative filter is needed. The followers carry no integral term, so nothing winds up against the acceleration limits and no anti-windup is needed. The steady gap error is nevertheless zero behind a car at constant speed, because two integrations lie between a car's command and its gap. Every policy is direct acting on the gap. A gap longer than wanted or a faster car ahead makes it accelerate, and a shorter gap or a slower car ahead makes it brake, which the application's direction check confirms for all four.

14.5.3 What the Supervisor Asks For

The supervisor shapes what each follower is asked for. When a new car appears ahead, the follower accepts the gap it was left and grows it back at a limited rate, as shown in Equation (14.14).

$$\delta_i(t_0) = \min\left(0,\ \hat{d}_i - d_i^*\right), \qquad \dot{\delta}_i = \rho \ \text{ while } \delta_i < 0 \tag{14.14}$$

where t_0 is the tick at which the radar's first sample of the new car arrives and $\rho = 1\,\mathrm{m/s}$, so the car opens its gap gradually instead of braking for the whole shortfall at once, and the car that cut in does the same toward the car ahead of it. A cooperative follower watches the silence since its newest packet. Past 0.5 s the link counts as stopped, ϕ_i fades from one to zero over 0.5 s, and in the degraded mode the time gap in force moves toward the fallback gap $h_{fb} = 1.5\,\mathrm{s}$, as shown in Equation (14.15).

$$\dot{h}_{eff} = \begin{cases} \min\left(0.1,\ (h_{fb} - h_{eff})/3\,\mathrm{s}\right) & \text{link stopped} \\ \max\left(-0.05,\ (h - h_{eff})/6\,\mathrm{s}\right) & \text{packets back} \end{cases} \tag{14.15}$$

where the rates are seconds of time gap per second and h_{eff} replaces h in Equations (14.4) and (14.5). When packets return, ϕ_i fades back to one over 0.5 s and the gap narrows again. Both moves are first-order approaches with a rate limit, so neither ends in a corner that would leave the cars at the back braking out of a closing speed. A predecessor with no radio at all, such as the car that cut in, is treated like a stopped link, so the car behind it falls back to h_{fb}.

14.5.4 Distributed Model Predictive Control

A sixth policy, model predictive control (MPC), never the default, replaces Equation (14.5) in every follower by a receding-horizon optimization [19] (Section 1.6.3), the distributed arrangement in which each car plans its own command and passes the plan to the car behind [20, 21]. At every control period $T = 50\,\mathrm{ms}$ each follower plans its command over the next $N_p = 60$ steps, 3 s or more than seven driveline lags, held in eight blocks of $b_j = 1, 1, 2, 3, 5, 8, 15$ and 25 steps, applies the first

move, holds it for one period and plans again. Its model is the car of Equation (14.1) written in the frame of the gap, Equation (14.16).

$$\begin{aligned} &\dot{d} = \Delta v, \quad \dot{\Delta v} = a_p - (a+\beta), \quad \dot{v} = a+\beta, \\ &\bar{\tau}\dot{a} = u - a, \quad \dot{\beta} = 0 \end{aligned} \tag{14.16}$$

where Δv is the opening rate, a_p the acceleration of the car ahead, which enters as a preview over the horizon, β an input disturbance on $\dot{v}$ that stands for grade, headwind and any lag other than the declared one, and $\bar{\tau} = 0.4\,\mathrm{s}$ the declared lag, fixed when the application loads, so that a lag set on the plant is model mismatch. The model is written from the plant's equations, not identified, and sampled exactly by a zero-order hold over T, which a fourth-order Runge-Kutta integration reproduces to 4×10^{-16}. Each plan solves Equation (14.17).

$$\begin{aligned} \min_{u^b, s} \; & \sum_{k=1}^{N_p} \left(Q_e e_k^2 + Q_v \Delta v_k^2 \right) + P_e e_{N_p}^2 + P_v \Delta v_{N_p}^2 \\ & + S \sum_{j=0}^{7} \left(u_j^b - u_{j-1}^b \right)^2 + R \sum_{k=0}^{N_p - 1} \left(u_k + \hat{\beta} \right)^2 + \rho s^2 + \sigma s \\ \text{s.t.} \; & e_k = d_k - (r + h_{eff} v_k + \delta_i), \\ & d_k - \tfrac{1}{2} h_{eff} v_k \geq \tfrac{1}{2} r + \delta_i - s, \quad s \geq 0, \\ & a_{min} \leq u_j^b \leq a_{max}, \quad \left| u_j^b - u_{j-1}^b \right| \leq j_{max} T \, c_j \end{aligned} \tag{14.17}$$

where u_j^b is the command of block j, u_k the command in force at step k, u_{-1}^b the move in force at the decision, e_k the gap error of the time-gap policy with the supervisor's h_{eff} and δ_i, and s one slack for the whole horizon. The weights are $Q_e = Q_v = 1$, $P_e = P_v = 10$, $S = 0.5$ and $R = 0.2$ in SI units, with $\rho = 10^5$ and $\sigma = 10^4$ on the slack. The R term prices the acceleration $u + \hat{\beta}$ that the command would sustain, so holding speed against a grade costs nothing. The acceleration limits are hard and read as the clip of Section 14.5 reads them, the floor at half the time-gap gap is soft, and $j_{max} = 10\,\mathrm{m/s^3}$ is a comfort limit on the jerk, hard, with $c_0 = 1$ on the first move and $c_j = (b_{j-1} + b_j)/2$ between blocks, the change that a ramp at the limit makes between two block centers. The set-speed cap and the braking of a collided car stay in force above the plan.

The ramp form corrects the first build, which applied one limit of $0.5\,\mathrm{m/s^2}$ at every block boundary, as the shared solver does whatever a block's length, so a plan believed its command could change by at most $4\,\mathrm{m/s^2}$ in 3 s and braking hard committed it to braking for seconds. In a hard stop at 8 s on case 10's line car 2 then commanded only -3.63 to $-4.83\,\mathrm{m/s^2}$ from 8.8 to 10.4 s behind a lead at $-4.2\,\mathrm{m/s^2}$ to -5.0 (Section 14.6.6). With the ramp form it does as well as with no jerk limit, and the largest move on case 10's tap is $0.37\,\mathrm{m/s^2}$ per period.

A steady-state Kalman filter [22], its gain computed once, estimates the five states, augmented with the gap and its rate one period earlier because the radar's samples arrive one period late, from the radar's gap and rate, the wheel speed and the accelerometer, which reads $a + \beta$, with variances of 0.05^2, 0.05^2, 10^{-6} and 10^{-4} in SI units and a process noise covariance of $\mathrm{diag}(10^{-5}, 4 \times 10^{-3}, 10^{-6}, 10^{-3}, 10^{-4}, 10^{-9}, 10^{-9})$. The disturbance state makes the controller offset free. With a lag of 0.6 s, a 3 % grade and an 8 m/s headwind on the plant, car 2's $\hat{\beta}$ settles at $-0.3960\,\mathrm{m/s^2}$ against the plant's -0.3942, and the worst gap error is 3.3 cm against 3.2 cm at nominal conditions, the radar's noise.

With the plan preview each car broadcasts ten numbers, its eight block commands, its driveline state and $\hat{\beta}$, from which the car behind reconstructs the predicted acceleration, an extension of the radio protocol under this policy only. With the acceleration preview the one acceleration the radio carries today is held for 0.4 s and then taken as zero, the like-for-like comparison. With the preview off nothing from the radio is used, and the supervisor widens every gap to h_{fb} as for a stopped link.

The lead has no plan, so car 2 always previews one acceleration. The plan is a forecast, and its acceleration 1 s ahead is in error by 0.20 m/s^2 RMS on the tap of case 10, 0.60 m/s^2 under the 400 ms, 20 % radio of case 7 and 0.62 m/s^2 in the hard stop.

The controller uses a plan only when the solver returns the status optimal, the plan, predictions, cost and slack are finite, every planned command lies within the limits and the first planned command equals the move applied. Otherwise Equation (14.5) with $k_a = 1$ drives that car, with a bias from the last move that decays with a 1 s time constant, so the hand-over moves nothing, until a later plan is accepted. Under the plan k_p, k_d and k_a have no effect, while h, r, h_{fb} and ρ stay in force.

With no constraint active, the first move and the broadcast plan are affine in the estimate, the last command and the preview, with gains taken from the solver itself. Because a plan rather than one number passes from car to car, the criterion replaces Equation (14.10) by propagating each car's acceleration and plan along the line frequency by frequency, through the plant at the lag as set, the estimator and the delays, and reports the worst car. It matches simulated 6 s sines within 2 % by the application's check, and it holds while nothing binds, as on the tap of case 10, where the floor's slack stays at zero, but not in a hard stop, where the braking limit binds.

The weights were chosen once, on the 0.3 s line of case 6, as those with the least $|\Gamma|$ peak whose hard stop from cruise leaves at least 2 m. A rule on the tap alone would have picked $R = 3$ and $S = 2$, 0.943 on the tap of case 10, whose linearization peaks at 1.119 at 0.4 s and which keeps 0.11 m in the hard stop at 8 s. In Chrome 157 each period's nine solves take 2.2 % of a processor core at real time, 0.12 ms a solve, and 41 % at twenty times. At sixty times the page painted 14.5 frames a second, so a frame guard stops a frame's stepping after 25 ms, skipping no decision and leaving every logged row unchanged, and the page runs at 37.6 times real time and 33 frames a second.

14.5.5 Simulation, Control and Analysis Models

Equations (14.1) to (14.7) make up the simulation model, nonlinear through the clips, the set-speed cap, the stopped-car rule, the collisions and the driver model. Under the five shipped policies no controller holds a model of the line, since the spacing laws are model-free and the lead's PI knows nothing of the followers, while the predictive controller of Section 14.5.4 carries its own model of one car at the declared lag, Equation (14.16), and sees the plant only through the instruments of Equation (14.7). Equation (14.10), evaluated at the line's mean follower speed with the delays as the cars experience them, is an analysis model used only to display the criterion and its map, and no model is identified from data.

14.6 Eleven Runs Along the Line

Table 14.4 lists the cases. Every case begins from the defaults of Tables 14.2 and 14.3, ten cars at 72 km/h on radar ACC at 1.5 s, and departs from them as its row gives. The figures that follow were produced by the engine alone, with the default seed and a span of 90 s unless another is named. Each variant is the case with the setting changed at $t = 0$ before Run, and a time gap or policy chosen there lays the line out at the gap it wants. Reset keeps the student's settings and replays the case's scripted events, so the same setting carried through Reset gives a run identical to the last digit on the open road. In cases 2 to 7 the brake tap is the first disturbance the line meets. Until it comes, the drivers of case 2 hold their start exactly, and the radar cars of cases 3 to 7 stay within 0.16 km/h of it, moved only by the radar's noise. Table 14.5 collects the results, and Figure 14.3 sets the predicted peak beside the measured amplification and the lane flow for the main policies.

14.6.1 People and the Reaction Delay

In case 1 the operator holds the lead's pedals while the scenario taps the brake at 6 s, 3 m/s^2 for 1.5 s, which commands a 4.5 m/s speed decrement. In MANUAL nothing restores it, and with no road load the lead coasts at 55.8 km/h. Each driver sees the change 0.70 s late and answers harder than the one ahead, so the dip deepens down the line. Car 2 bottoms out at 53.3 km/h, car 5 at 42.8, car 8 at 14.3 and car 9 at 0.9. Car 10 reaches zero at 27.14 s, 21.14 s after the tap. The cars pass back up to 1.29

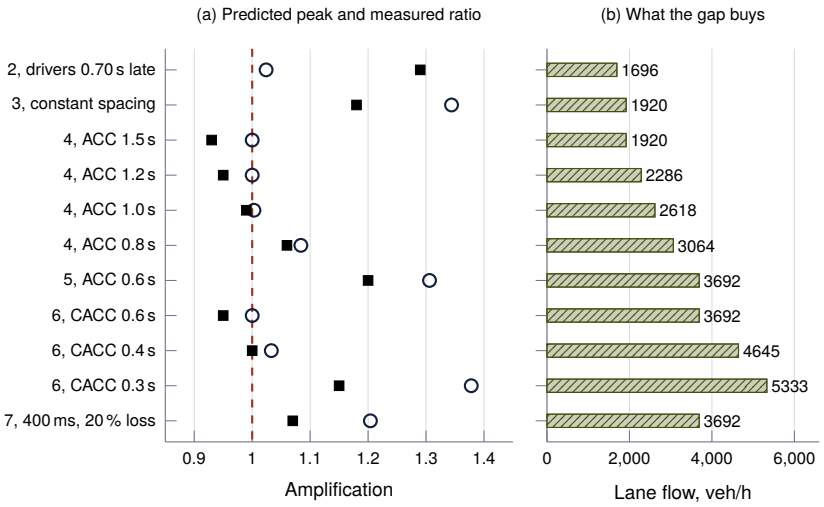


Figure 14.3 Rows of Table 14.5. (a) The criterion's peak $|\Gamma|_{max}$ (open circles) and the measured amplification γ (filled squares), with the string-stability bound of one dashed. (b) The lane flow over the last 30 s. Notice that each shorter time gap of radar ACC buys flow, reaching the bound at 1.0 s and passing it at 0.8 s, while the radio holds 0.6 s within it at the flow of case 5, and that the drivers measure 1.29 where the criterion predicts only 1.024.

Table 14.4 The Ten Cases and the Challenge

#	Change from the reference line	Teaching point
1	people 0.70 s late, lead in MANUAL, tap at 6 s	the accordion
2	people 0.70 s late, tap at 5 s	delay as a cliff
3	constant spacing at 33 m, gains 0.3 and 1.0, tap	no gain suffices
4	time gap 1.5 s, tap	stability and its price
5	time gap 0.6 s, tap	well damped, unstable
6	CACC at 0.6 s, 50 ms radio, tap	stability regained
7	CACC with fallback at 0.6 s, 400 ms, 20 % loss, tap, burst at 40 s	the network's price
8	CACC with fallback at 0.6 s, cut-in at 10 s	shaping the request
9	22 people on a 437 m ring at 30 km/h	the phantom jam
10	MPC at 0.4 s, broadcast plan, tap	planning against CACC
11	challenge, 30 km/h, 0.6 s, 10 % loss, cut-in, stop and go	graceful degradation

times the RMS acceleration they receive, at car 6, and the closest call is 2.09 m, car 10 behind car 9. Reset replays the tap, so a student's own pedal input adds to it. Braking at 0.5 m/s^2 from 10 to 19 s takes the lead down to 39.6 km/h, and car 10 still reaches zero at 27.14 s. Accelerating at 1.5 m/s^2 from 7.6 to 10.7 s restores about 72 km/h and keeps car 10 above 42.8 km/h.

Case 2 gives the lead its cruise control, which brings it back from a low of 57.3 km/h at 7.1 s. Car 10 bottoms out at 28.3 km/h at 26.7 s, and the measured ratio rises from 0.80 at car 2 to 1.29 at car 10. The reaction delay is a cliff. At 0.5 s the line measures 0.99 and car 10 stays above 60.5 km/h. At 0.6 s it measures 1.03, and at 0.75 s car 10 falls to 0.5 km/h. At 0.8 s car 10 runs into car 9 at 9.7 km/h after 26.62 s. At 1.0 s four collisions follow, the first car 7 into car 6 at 26.0 km/h after 20.95 s. Over the same range of delay the linear peak rises only from 1.017 to 1.084, while the measured ratio is already 1.29 at 0.70 s. The tap is not a small wave, and the driver model's braking is nonlinear in the approach rate.

Table 14.5 Measured Amplification, Car 10's Lowest Speed, Smallest Gap and Lane Flow. Runs Last 90 s Except Where Named

Case and configuration	$\|\Gamma\|_{max}$	γ (car)	v_{10} (km/h)	d_{min} (m)	q (veh/h)
1, as shipped	1.024	1.29 (6)	0.0	2.09	1462
2, T_r 0.70 s	1.024	1.29 (10)	28.3	20.55	1696
2, T_r 0.5 s	1.017	0.99 (10)	60.5	29.63	1757
2, T_r 0.8 s	1.031	1.35 (7)	0.0	collision	412
2, T_r 1.0 s	1.084	1.44 (5)	0.0	4 collisions	255
3, as shipped	1.344	1.18 (3)	6.3	21.88	1920
3, k_p 0.05, k_d 1, 120 s	1.050	1.02 (10)	56.5	29.04	1920
3, k_p 2, k_d 1	131.7	4.23 (2)	0.0	8.54	973
4, h 1.5 s	1.000	0.93 (10)	66.8	28.18	1920
4, h 1.2 s	1.000	0.95 (10)	64.4	22.64	2286
4, h 1.0 s	1.003	0.99 (10)	61.2	18.97	2618
4, h 0.8 s	1.084	1.06 (6)	54.5	14.82	3064
5, h 0.6 s	1.306	1.20 (3)	36.3	8.22	3692
5, k_p 3	1.305	1.21 (3)	37.9	8.95	3692
5, k_d 2	1.089	1.06 (6)	53.2	11.35	3692
6, CACC 0.6 s	1.000	0.95 (10)	63.7	12.81	3692
6, h 0.4 s	1.033	1.00 (5)	58.1	9.40	4645
6, h 0.3 s	1.378	1.15 (3)	46.2	6.87	5333
6, k_a 0.5	1.048	1.05 (8)	57.2	12.25	3692
7, as shipped, 120 s	1.204	1.07 (5)	37.6	12.12	3692
7, h 0.8 s, 120 s	1.000	0.93 (10)	43.9	16.37	3063
7, plain CACC, 120 s	1.204	1.07 (5)	52.2	12.12	3692
8, regrowth 1 m/s, 80 s	1.000	—	59.1	5.24	3117
8, regrowth 0, 80 s	1.000	—	36.5	5.24	3117
9, ring, 300 s	1.198	1.03 (6)	—	—	810
9, ring on ACC 1.5 s, 300 s	1.000	—	—	—	1494

$\|\Gamma\|_{max}$ at the start of the run, v_{10} the lowest speed of car 10, q over the last 30 s. A run with a collision pauses at the first one, and its row is read after Run is pressed again.

14.6.2 Spacing Policy and Gains

Case 3 holds a fixed 33 m gap with gains that damp each car at 0.816. After the tap car 10 bottoms out at 6.3 km/h at 17.6 s, a dip of 65.7 km/h behind the lead's 14.7, and the line measures 1.18 at car 3, every car above one. The best pair of the grid, which the spacing-gain slider reaches at its lower end, still measures 1.02 over 120 s, and the time-gap gains measure 4.23 at car 2, with six cars standing still at once. This is the honestly infeasible case, because no gain on the grid or off it brings the low-frequency term of Section 14.4 below one.

Case 4 asks for 3 m plus 1.5 s of travel, the same 33 m at 72 km/h. The tap now dies out on its way back. Car 10's low is 66.8 km/h at 20.4 s against the lead's 57.3, and the measured ratio rises from 0.68 at car 2 to 0.93 at car 10, every car passing back less than it receives. The price is flow, 1920 veh/h. Shortening the gap buys flow with amplification. At 1.2 s the line measures 0.95 at 2286 veh/h, at 1.0 s, with its peak of 1.003 just past the boundary, 0.99 at 2618 veh/h, and at 0.8 s it amplifies, 1.06 at 3064 veh/h.

Case 5, at 0.6 s, is the chapter's central demonstration, shown in Figure 14.4. Each car on its own is damped at 0.496, while the criterion peaks at 1.306 at 1.12 rad/s, a period of 5.6 s. Car 10 bottoms out at 36.3 km/h at 16.1 s and the line measures 1.20 at car 3. The tap takes the lead's acceleration to −2.93 m/s^2, the middle car's, car 6, to −3.20 m/s^2 at 10.80 s and the last car's to −4.02 m/s^2 at 14.60 s. The last car then commands its 2 m/s^2 limit from the control period at 16.20 to the one at 21.50 s, its gap peaks at 27.23 m at 28.84 s, it reaches its own set speed of 86.4 km/h at 31.85 s, and it brakes again, to −2.87 m/s^2 at 33.05 s, as it closes the gap, a second disturbance that the catching up creates by itself. The lane would carry 3692 veh/h, and the line cannot hold together at that density. A spacing gain of 3 moves the peak without lowering it (1.305, measured 1.21), and a rate gain of 2 lowers it without removing it (1.089, measured 1.06), because at the top of the slider, $k_p = 3$, the

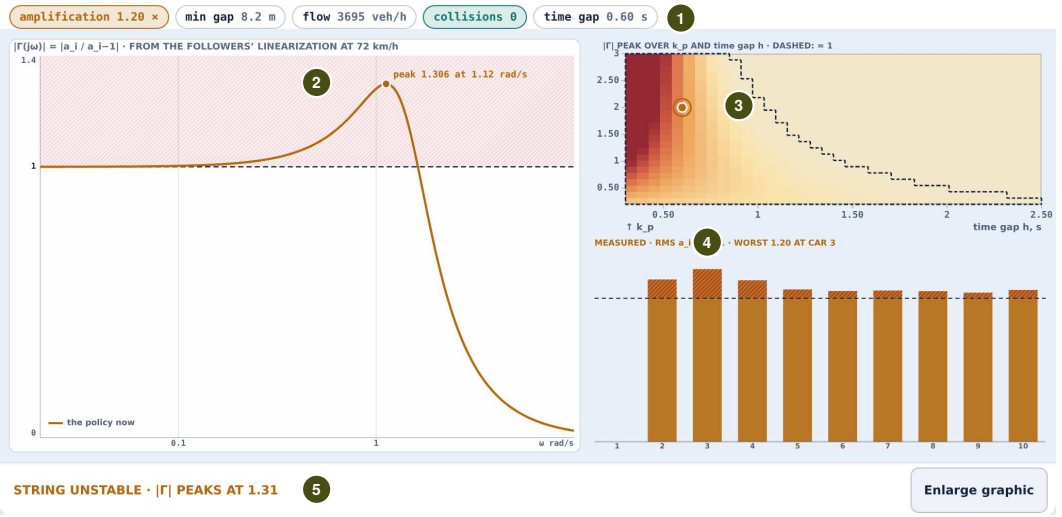


Figure 14.4 Case 5, the gap shrunk to 0.6 s, at $t = 60.00$ s in the $|\Gamma(j\omega)|$ view, run in AUTO with no control changed, cropped to the plant view. The badges (1) read amplification 1.20×, min gap 8.2 m, flow 3695 veh/h, collisions 0 and time gap 0.60 s. The left panel draws $|\Gamma(j\omega)|$ from the followers' linearization at 72 km/h, with its peak of 1.306 at 1.12 rad/s (2) marked in the hatched region above one, and the map beside it places the operating point (3), $k_p = 2$ and $h = 0.6$ s, left of the dashed contour where the peak equals one. Below the map every measured bar (4) from car 2 to car 10 is above one, the worst 1.20 at car 3, and the status line (5) reads String unstable, $|\Gamma|$ peaks at 1.31. The run's trends and scorecard are cut, and Table 14.5 gives its measured values.

boundary $\sqrt{2/k_p}$ is still 0.82 s.

14.6.3 The Radio and Its Failures

Case 6 keeps the 0.6 s gap and the gains and adds the radio at $k_a = 1$. The tap dies out, car 10 bottoming out at 63.7 km/h at 12.9 s and every car passing back at most 0.95, and the lane still carries 3692 veh/h, 92 % more than the 1.5 s ACC of case 4. At 0.4 s, just past the 0.413 s boundary, the peak is 1.033 and the line measures 1.00 at 4645 veh/h. At 0.3 s it amplifies again, 1.378 and measured 1.15, and half the feedforward gain leaves a peak of 1.048 at 0.6 s, measured 1.05.

> **TRY IT IN THE APP** **6 · Cooperative ACC at 0.6 s**
>
> Load the scenario, press Run and watch the speed bars after the brake tap at 5 s. Car 10 bottoms out at 63.7 km/h, where the radar-only line of case 5 took it down to 36.3 km/h, and every car passes back at most 0.95 of the acceleration it receives. The gap and the gains are those of case 5. Only the radio is new.

Case 7 slows the radio to 400 ms with one packet in five lost, an effective 412.5 ms. The radio term now arrives too late to help at 0.6 s. The peak is 1.204 at 1.64 rad/s, the line measures 1.07 at car 5, the stable gap moves out to 0.737 s, and at 0.8 s the line measures 0.93 at 3063 veh/h. Random loss of one packet in five almost never stops a link, because a stop needs ten silent periods in a row, a chance of 0.2^{10}, about 10^{-7}, at each packet, and none of the nine links stops before the burst at 40 s, while all nine stop during it, at the control periods at 40.85 and 40.90 s. Car 6's link counts as stopped at 40.90 s, its time gap widens to 1.160 s, the burst ends at 46 s and the link counts as live again at 46.50 s, and the gap is back within 0.01 s of 0.6 s by 72.0 s. The fallback is not free. When every gap widens at once the last car must fall back by all of them, and car 10 slows from 72 to 37.6 km/h at

47.5 s, deeper than its low of 52.2 km/h after the tap. Plain CACC, which rides out the silence on radar alone at 0.6 s with no wave to amplify, keeps car 10 above 71.1 km/h after 38 s.

> ✦ DID YOU KNOW? **A Fallback That Deepens the Dip**
>
> From the simulator: in case 7 the fallback widens every gap when the radio burst silences all nine links, and car 10's lowest speed in the run is 37.6 km/h. Plain cooperative ACC rides out the same silence on radar alone, and its lowest, 52.2 km/h, comes from the brake tap (Table 14.5).

14.6.4 A Cut-In

In case 8 a car without a radio cuts into the middle of the 15 m gap ahead of car 4 at 10 s, leaving 5.24 m on either side. With the gap regrowing at 1 m/s, car 4 brakes at no more than 2.02 m/s^2 and the cars behind it slow to 59.0 to 59.1 km/h and recover. Asked for the whole gap at once, car 4 faces a shortfall of 9.88 m and brakes at 5.96 m/s^2, close to its limit, the car that cut in brakes at 5.70 m/s^2 for its own 1.5 s gap, car 4 drops to 26.9 km/h and the cars behind it to between 30.7 and 36.5 km/h. Either way the cut-in costs flow, from 3691 veh/h before it to 3117 veh/h over the last 30 s, because car 4 now keeps 1.5 s behind a car with no radio. Nothing in the loop differs between the two runs except what the supervisor asks of it.

14.6.5 The Ring

Case 9 puts 22 drivers on the 437 m ring at 30 km/h, each displaced by at most 10 cm. Linearized there, the drivers peak at 1.198, far above their 1.024 on the open road at 72 km/h. The spread between the fastest and slowest car grows from 0.4 km/h at 30 s to 2.3 at 60 s and 11.3 at 90 s, the first car stops at 98.7 s, and from then on a stop-and-go wave runs backward around the ring with up to six cars standing at once. The ring carries 810 veh/h over the last 30 s of five minutes, against 1510 veh/h at the uniform start. With a reaction delay of 0.6 s the first stop comes at 132.0 s, and with 0.8 s at 75.8 s. Radar ACC at 1.5 s on every car keeps the same start uniform, with a spread of 0.1 km/h at 1494 veh/h. The jam also lowers the line's capacity. In a model run of 16 cars whose scripted tap is replaced by a hard stop at 5 s, these drivers discharge from the standstill at 997 veh/h, against the 1766 veh/h that their equilibrium carries at its best speed, 65.9 km/h, a drop of 44 %, while radar ACC at 1.5 s discharges at 1880 veh/h against 1920, a drop of 2 %. Fired by hand at 7, 30 or 60 s of case 2 with 16 cars, after its tap, the hard stop gives 939 to 1183 veh/h for the drivers, and 1880 again for the ACC of case 4. The ramp-metering application of Volume 2, Chapter 8 declares a capacity drop of 10 %, between these two measurements of this model.

> TRY IT IN THE APP **9 · The Ring Road: A Jam From Nothing**
>
> Press Run on the Time–space view at 20× and watch for the first red stripe. The first car stops at 98.7 s, and a stop-and-go wave then runs backward around the ring. Press Reset, set the Follower policy to ACC · time gap and run again: the spread of speeds stays at 0.1 km/h.

14.6.6 Planning Against Cooperative ACC

Case 10 runs case 6's cooperative line at 0.4 s, just past the 0.413 s boundary of cooperative ACC, under the controller of Section 14.5.4 with the plan preview and a tap at 5 s. Table 14.6 compares the structures on several cases with the same seed, noise and radio in every row and the controller's weights never retuned. Cooperative ACC retuned once by the rule that chose those weights, on the same 0.3 s line, takes $k_p = 0.5$ and $k_d = 3$ (peak 1.031, 4.57 m left in a hard stop at 3 s), and the strongest classical form found feeds those gains the commanded acceleration of the car ahead, which leads the delivered one by the lag.

On the tap of case 10 the controller measures 0.980, worst at car 2, against 1.003 for cooperative ACC, its linearization stays at one where that of cooperative ACC peaks at 1.033 at 1.38 rad/s, and its RMS command jerk is 1.4 m/s^3 against 3.3. With only the one acceleration the radio carries today

Table 14.6 Cooperative ACC Against the Predictive Controller: Measured Amplification γ, Criterion Peak $|\Gamma|_{max}$, RMS Command Jerk J, Lane Flow q and Smallest Gap d_{min}

	10, h 0.4 s, tap			6, 0.6 s	6 at 0.3 s		7, 400 ms, 20 %[a]		11[a]	hard stop, d_{min} (m)	
Structure	γ	$\|\Gamma\|_{max}$	J (m/s^3)	γ	γ	$\|\Gamma\|_{max}$	γ	q (veh/h)	γ	at 8 s	at 30 s
CACC, shipped gains	1.003	1.033	3.3	0.950	1.152	1.378	1.072	3339	0.989	2.92	2.93
CACC, commanded acceleration	0.968	1.033	8.1	0.959	0.976	1.378	1.001	3347	0.992	3.01	3.04
CACC retuned, k_p 0.5, k_d 3	0.975	1.000	4.7	0.941	1.018	1.031	0.967	3306	0.989	3.01	3.02
retuned, commanded acceleration	0.954	1.000	12.7	0.946	0.961	1.031	0.977	3306	0.988	3.22	3.20
MPC, plan preview	0.980	1.000	1.4	0.939	1.028	1.051	0.987	3371	0.985	2.95	2.97
MPC, acceleration preview	0.992	1.000	1.3	0.935	1.045	1.051	1.022	3342	0.988	2.98	2.96
MPC, preview off[b]	—	1.000	0.9	—	—	1.000	—	1920	0.974	3.33	3.48

Runs of 60, 60, 60, 80 and 110 s, default seed. Hard stops on case 10's line, runs of 40 and 60 s. [a]Cooperative rows with the fallback. Every challenge row passes except preview off (1497 veh/h). [b]With the preview off the supervisor widens every gap to $h_{fb} = 1.5$ s, the tap's pairs fall inside the quiet time of Equation (14.9) and γ has no value, and the criterion describes the law at 1.5 s. Every MPC run solved 7,200 to 19,799 plans, none rejected, at 0.88 to 0.90 ms for each period's nine solves in Node.

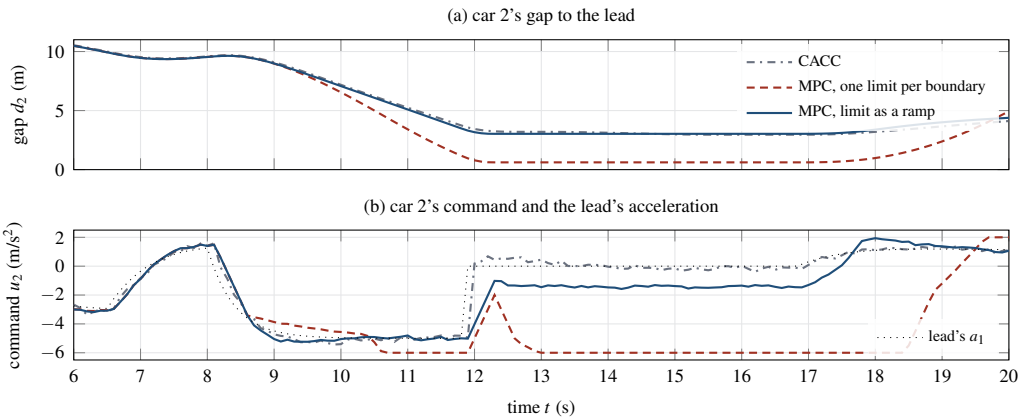


Figure 14.5 Car 2 in a hard stop fired at 8 s on the 0.4 s line of case 10. (a) Its gap and (b) its command beside the lead's acceleration (dotted), under cooperative ACC (dash-dotted, smallest gap 2.96 m), the predictive controller's first build, one jerk limit per block boundary (dashed, 0.61 m), and the controller with the limit as a ramp (solid, 3.03 m).

it measures 0.992, and 0.935 against 0.939 on case 6, so most of the gain is the controller, not the message. The plan counts where the radio is slow, 0.987 against 1.022 in case 7. Car 2, which always previews the lead's one acceleration, is the worst car under the plan in cases 7 and 10. With the preview off the widened gap leaves no wave to grow, but the lane carries 1920 veh/h instead of 4648, and with the widening removed, a bound kept out of the application, the controller measures 1.116.

Cooperative ACC retuned once by the controller's own rule beats it on amplification on case 10, 0.975 against 0.980, at 0.3 s, 1.018 against 1.028 with a peak of 1.031 against 1.051, and on case 7, 0.967 against 0.987. With the commanded acceleration it reaches 0.954 on case 10, the best of any row, at an RMS jerk of 12.7 m/s^3, nine times the controller's. The controller's advantage is therefore not its amplification but the smoothest commands of the table, 0.9 to 1.4 m/s^3 on case 10 against 4.7 for the retuned law, limits held as constraints and a plan shown to the car behind. Every row but the radar-only controller passes the challenge, at 0.985 to 0.992.

The hard stop is the case the first build lost, as shown in Figure 14.5. Fired at 8 s, it left 0.61 m under the plan preview and 0.46 m under the acceleration preview, a stop from cruise at 30 s left 0.30 m and a collision, and the radar-only controller collided nine times. The ramp form keeps 2.95, 2.98 and 3.33 m. The same stop was then fired at 17 instants from 1 to 30 s on four lines, each run continued 25 s after it. On cases 10, 6 and 7 and at 0.3 s no car under the controller collides at any

Table 14.7 The Challenge of 110 s in Five Configurations

Configuration	γ (car)	d_{min} (m)	q (veh/h)	Target
as shipped, ACC 0.6 s	1.13 (10)	1.75	2283	not met
ACC 1.5 s	0.98 (10)	5.50	1497	not met
CACC 0.6 s	1.04 (4)	1.75	2281	not met
CACC with fallback	0.99 (8)	1.75	2158	met
MPC, plan preview	0.99 (10)	1.75	2156	met

instant with any preview, the smallest gaps being 2.31, 0.46 and 2.01 m under the plan, acceleration and radar-only previews, while cooperative ACC at 0.3 s collides at four instants. The first build had collided at 9 and 13 instants under the plan preview, at 0.3 s and in case 7, at 8 to 14 on three lines under the acceleration preview and at 15 or 16 on every line with the preview off. On the challenge's line a stop before 5 s now leaves no collision, the cut-in being refused. A stop at 5 to 12 s still does under every structure (radar-only 5 to 6.5 s), as the radio-less car that cut in, still regrowing its gap, hits car 3 at up to 4.69 m/s, and none collides without the cut-in.

14.6.7 Stop and Go With a Cut-In

The challenge runs ten cars at 30 km/h and 0.6 s for 110 s with one packet in ten lost. At 5 s a car without a radio cuts in ahead of car 4 with 1.75 m on either side, and at 40 s the lead begins a stop and go to 10 km/h for 20 s. The target asks for no collision, a measured amplification of at most 1.00, no gap below 1.5 m and at least 2000 veh/h over the last 30 s. Table 14.7 compares five configurations. As shipped, radar ACC fails on amplification rather than on safety. The line measures 1.13 at car 10, cars 8, 9 and 10 come to a standstill at 52.1, 52.6 and 53.2 s, the smallest gap is the cut-in's 1.75 m, and the flow is 2283 veh/h. A 1.5 s gap is stable, at 0.98, but carries 1497 veh/h. Plain CACC is stable everywhere except behind the car that cut in, where car 4 runs on radar alone at 0.6 s and measures 1.04. CACC with the fallback, at the shipped gains, widens only car 4's gap to 1.5 s and passes, measuring 0.99 with no gap below 1.75 m at 2158 veh/h. This is the challenge's suggested setting. Its measured amplification of 0.989 leaves margins of 0.011, 0.25 m and 158 veh/h on the three targets, whether Run follows the suggestion directly or after Reset. Four other seeds of the radar noise and the radio losses give 0.988 to 0.991 for it, and 1.122 to 1.133 for the shipped radar ACC, so the verdict does not depend on the seed. The controller of Section 14.5.4 passes at 0.985, because the supervisor widens car 4's gap behind the car with no radio as it does for cooperative ACC.

14.7 Technology in Practice

What Measures It

Each follower of this chapter reads three things: the gap and its rate of change from a radar, its own speed and acceleration, and the acceleration of the car ahead over a radio. In service the gap comes from a millimeter-wave radar at the front of the car, which measures distance and radial velocity directly [23]. The radio is a vehicle-to-vehicle link whose messages carry position, speed and braking status [24]. The chapter's radar is a sampled link 50 ms late with 0.05 m and 0.05 m/s of noise and perfect target selection, and its radio is the same link model with a delay, a jitter and a loss probability. Figure 14.6 draws one follower's loop on that hardware.

What Moves It

The final element is the car's own throttle and brakes, which the platoon controller reaches only through the vehicle's low-level controller [25]. The chapter lumps that controller, the engine and the brakes into one first-order lag of 0.4 s, the same for throttle and brake, and clips the command to -6 to 2 m/s^2. A real driveline brakes faster than it accelerates, as the Assumptions note.

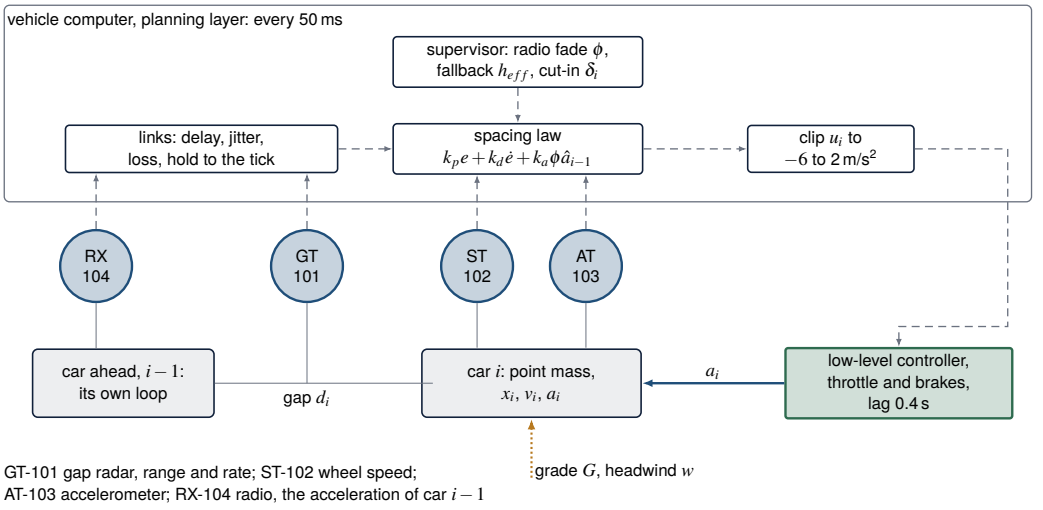


Figure 14.6 One follower's loop on vehicle hardware. The gap radar GT-101 and the radio receiver RX-104 reach the controller through the sampled links, the wheel speed ST-102 and the accelerometer AT-103 directly. The spacing law runs under the supervisor's fade, fallback and cut-in allowance, and its clipped command goes to the low-level controller of the throttle and brakes, the chapter's 0.4 s lag. Notice that the car ahead enters twice, through the gap the radar sees and through the acceleration the radio reports.

> **TECHNOLOGY** **Sensor: A Millimeter-Wave Radar at the Bumper**
>
> An automotive radar measures the distance and the radial velocity of a target directly, and with a suitable antenna system its angle. Its functions grew from distance warning to adaptive cruise control and emergency braking. Most new sensors work at 76 to 81 GHz and send fast sequences of frequency chirps, whose echoes a two-dimensional Fourier transform turns into a map of range and velocity [23]. A federal rulemaking notice gave a perception range of up to 200 m for a 77 GHz radar, with a field of view of 18 degrees [24]. The chapter's radar is 50 ms late, carries a little noise and always sees the right car.

Where the Controller Runs

In cooperative adaptive cruise control research the vehicle is organized in sensing, planning and actuation phases. Data from the radio unit and the onboard sensors are gathered on the vehicle's data bus, and the planning phase holds the high-level controller where the longitudinal control law runs. In one test fleet the production adaptive cruise control stayed installed, the cooperative law ran on a separate prototyping computer connected to the data bus, and the driver could switch between the two in real time [25]. The chapter's spacing law, the supervisor's fade, fallback and regrowth, and the predictive controller all belong to that planning layer and run every 50 ms. A proposed safety standard asked each car to broadcast ten times a second, every 100 ms [24], twice the period of the chapter's radio.

> **TECHNOLOGY** **Final Element: Throttle and Brakes Behind a Low-Level Controller**
>
> A platoon controller plans an acceleration and leaves its execution to a low-level controller, which converts the target into throttle and brake actions [25]. The delay of this actuation significantly limits the shortest inter-vehicle distance at which a cooperative platoon stays string stable [25]. In this chapter one lag of 0.4 s stands for all of it, and the bound on the time gap moves with it, because the lag enters the τs^3 term of the car-to-car transfer function. Its limits are the other half of the final element: in case 5 the last car commands its 2 m/s^2 limit from 16.20 to 21.50 s while it catches up.

What Is Modernizing It

Three developments are changing the hardware around this loop. The radio band is changing. The federal regulator kept the upper 30 MHz of the 5.9 GHz band for intelligent transportation systems and opened the lower 45 MHz to unlicensed use [26], and it required the transition from dedicated short-range communications to cellular vehicle-to-everything technology, with rules for the new devices and a timeline for retiring the old [27]. The radar is moving to multiple-input multiple-output antennas, digital beamforming and digital modulation, toward images of the whole scene rather than single reflections [23]. And the automated car is becoming an actuator of the traffic flow itself. On a ring track a single automated vehicle dampened the stop-and-go waves of about twenty human drivers, which suggests that flow control may be possible with a few such vehicles, less than 5 %, long before most cars are automated [10]. The chapter's ring of case 9 shows the same wave with every car under one policy.

> **IN PRACTICE** **What a Proposed Rule Asked of the Radio**
>
> A rulemaking notice of 2017 proposed Federal Motor Vehicle Safety Standard No. 150 for vehicle-to-vehicle radios. Each car would broadcast a basic safety message with its position, speed and braking status ten times a second under non-congested conditions, each transmission shifted at random by up to 5 ms to ease channel congestion. The message had to be received anywhere within at least 300 m with a packet error rate below 10 % [24]. The chapter's radio sends every 50 ms, and case 7 loses one packet in five, twice that error rate.

14.8 Limitations

The cars are identical point masses in one lane with one lag for throttle and brake, no road load and no tires, and a cut-in is instantaneous. The drivers are identical and noiseless and look one car ahead, so the reaction delay of 0.70 s is a value calibrated to stay short of the model's collisions at 0.8 s, not a measured human figure. The radar has perfect target selection, and the radio's losses are independent of how many cars share the channel. The criterion comes from one linearization at the mean follower speed, which describes small waves, and for people it understates how fast a large wave grows.

> **OPEN CHALLENGE** **Many Radars on One Road**
>
> Vehicles now carry up to twelve radars each, and no rules yet govern how sensors share the 76 to 77 GHz band. A simulation of a dense six-lane highway found many transmitters along the road reaching each receiver, which can lower sensitivity and accuracy and in rare cases blind the sensor. Detecting and repairing disturbed signals, active avoidance and cooperation between sensors are being studied [23]. Progress would be a radar whose gap and closing rate stay trustworthy in the densest traffic, where a platoon runs its shortest gaps. The chapter's radar has no interference.

The predictive controller's model is one car at the declared 0.4 s lag, which meets any other lag, grade or wind only through β and omits the stopped-car rule and collisions. Its preview cannot foresee the lead's brake, in error by 0.62 m/s^2 RMS one second ahead in a hard stop, and its criterion holds only while no limit binds.

One limitation concerns use rather than physics. Reset keeps every setting the student changed and replays the case's scripted events, its tap included, while reloading the scenario restores the case's own settings. A new time gap or policy lays the line out at the gap it wants only while the run stands at $t = 0$. A time gap moved while the line is running acts on the running line instead. For radar ACC the closing-up then enters the measurement, which reads 1.39 instead of 0.99 for 1.0 s chosen at 2 s, and for cooperative ACC the supervisor's gradual move pauses it. On the ring a policy chosen at $t = 0$ keeps the start speed of the drivers, 30.0 km/h, while one carried through Reset starts the ring at its own, 29.7 km/h, and the quoted figures of case 9 reproduce either way.

Every threshold is a teaching value. The challenge's 1.00 and 1.5 m, the 0.5 s link timeout and the fallback gap are not calibrations of any vehicle, and no number here describes how a production system performs.

OPEN CHALLENGE **A Radio Channel Shared by Every Car**

Every nearby car talks on the same channel, so a dense platoon is also a crowded channel. The 2017 notice called for research to identify the congestion that could impair safety applications and to develop mitigation. The strategies it weighed lower the number of messages a device sends each second, and the power or range of each one, when the device detects congestion [24]. A cooperative follower could then hear the car ahead less often just when its gap is shortest. Progress would be cooperative following that stays string stable under message rates that adapt to traffic density. In this chapter the losses do not depend on how many cars share the channel.

14.9 Conclusion

A line of individually well-damped cars can turn a brake tap into a stop, and this simulator measures by how much and why. At a 0.6 s time gap each radar follower is damped at 0.50 while the criterion peaks at 1.306, and car 10 dips to 36.3 km/h behind a lead that dipped to 57.3. Constant spacing amplifies for every gain, the best of 90 pairs peaking at 1.050. A time gap is stable above $\sqrt{2/k_p}$, 1.05 s with the radar's latency and the hold, at a cost in flow, 1920 veh/h at 1.5 s, and the predecessor's acceleration over a 50 ms radio regains 3692 veh/h at 0.6 s until 400 ms of delay and 20 % loss move the boundary to 0.74 s. Drivers 0.70 s late, barely unstable by the linearization at 1.024, collide once the delay reaches 0.8 s, and on a ring of 22 the first car stops 98.7 s after a uniform start. The challenge improves from 1.13 with three cars stopped to 0.99 at 2158 veh/h. A predictive controller in every follower measures 0.980 at a 0.4 s gap against 1.003 for cooperative ACC at 1.4 against 3.3 m/s^3 of command jerk, but the classical law retuned by its rule reaches 0.975, so it adds smoothness and constraints rather than amplification. The model is a single lane of identical, lagged cars whose drivers watch one car ahead. Further work would add mixed vehicles and drivers who watch several cars ahead, channel loss that grows with the number of cars, and a fallback that estimates the predecessor's acceleration from the radar. The platoon makes string stability a quantity that can be read off a chain of loops rather than inferred from one of them.

KEY POINTS

- A line of cars that are each well damped can still turn a brake tap into a much deeper slowdown at the back of the line.
- Constant spacing amplifies for every gain, and a time gap is string stable above $\sqrt{2/k_p}$, 1.05 s with the radar's latency and the hold, at a cost in lane flow.
- The predecessor's acceleration over a 50 ms radio regains 3692 veh/h at 0.6 s until 400 ms of delay and 20 % loss move the boundary to 0.74 s.
- Drivers 0.70 s late are barely unstable by the linearization, yet they collide once the delay reaches 0.8 s.

Where else this applies. Volume 9, Chapter 5 tells each gate's predictive controller the latest plan of the gate below, as each car here passes its plan to the car behind.

References

[1] L. Peppard, "String stability of relative-motion PID vehicle control systems," *IEEE Transactions on Automatic Control*, vol. 19, no. 5, pp. 579–581, 1974.

[2] D. Swaroop and J. K. Hedrick, "String stability of interconnected systems," *IEEE Transactions on Automatic Control*, vol. 41, no. 3, pp. 349–357, 1996.

[3] G. Gunter, D. Gloudemans, R. E. Stern, S. McQuade, R. Bhadani, M. Bunting, M. L. Delle Monache, R. Lysecky, B. Seibold, J. Sprinkle, B. Piccoli, and D. B. Work, "Are commercially implemented adaptive cruise control systems string stable?" *IEEE Transactions on Intelligent Transportation Systems*, vol. 22, no. 11, pp. 6992–7003, 2021.

[4] P. Seiler, A. Pant, and K. Hedrick, "Disturbance propagation in vehicle strings," *IEEE Transactions on Automatic Control*, vol. 49, no. 10, pp. 1835–1841, 2004.

[5] G. Gunter, D. Gloudemans, R. E. Stern, S. McQuade, R. Bhadani, M. Bunting, M. L. Delle Monache, R. Lysecky, B. Seibold, J. Sprinkle, B. Piccoli, and D. B. Work, "Are commercially implemented adaptive cruise control systems string stable?" arXiv preprint 1905.02108v1, 2019. [Online]. Available: https://arxiv.org/abs/1905.02108v1

[6] J. Ploeg, N. van de Wouw, and H. Nijmeijer, "$\mathscr{L}_p$ string stability of cascaded systems: Application to vehicle platooning," *IEEE Transactions on Control Systems Technology*, vol. 22, no. 2, pp. 786–793, 2014.

[7] H. L. Lee, V. Padmanabhan, and S. Whang, "Information distortion in a supply chain: The bullwhip effect," *Management Science*, vol. 43, no. 4, pp. 546–558, 1997.

[8] Y. Sugiyama, M. Fukui, M. Kikuchi, K. Hasebe, A. Nakayama, K. Nishinari, S.-i. Tadaki, and S. Yukawa, "Traffic jams without bottlenecks—experimental evidence for the physical mechanism of the formation of a jam," *New Journal of Physics*, vol. 10, no. 3, p. 033001, 2008.

[9] G. J. L. Naus, R. P. A. Vugts, J. Ploeg, M. J. G. van de Molengraft, and M. Steinbuch, "String-stable CACC design and experimental validation: A frequency-domain approach," *IEEE Transactions on Vehicular Technology*, vol. 59, no. 9, pp. 4268–4279, 2010.

[10] R. E. Stern, S. Cui, M. L. Delle Monache, R. Bhadani, M. Bunting, M. Churchill, N. Hamilton, R. Haulcy, H. Pohlmann, F. Wu, B. Piccoli, B. Seibold, J. Sprinkle, and D. B. Work, "Dissipation of stop-and-go waves via control of autonomous vehicles: Field experiments," *Transportation Research Part C: Emerging Technologies*, vol. 89, pp. 205–221, 2018.

[11] W. Levine and M. Athans, "On the optimal error regulation of a string of moving vehicles," *IEEE Transactions on Automatic Control*, vol. 11, no. 3, pp. 355–361, 1966.

[12] R. E. Fenton and R. J. Mayhan, "Automated highway studies at the Ohio State University-an overview," *IEEE Transactions on Vehicular Technology*, vol. 40, no. 1, pp. 100–113, 1991.

[13] S. E. Shladover, "PATH at 20—history and major milestones," *IEEE Transactions on Intelligent Transportation Systems*, vol. 8, no. 4, pp. 584–592, 2007.

[14] D. Swaroop, J. K. Hedrick, C. C. Chien, and P. Ioannou, "A comparision of spacing and headway control laws for automatically controlled vehicles," *Vehicle System Dynamics*, vol. 23, no. 1, pp. 597–625, 1994.

[15] R. Rajamani, H.-S. Tan, B. K. Law, and W.-B. Zhang, "Demonstration of integrated longitudinal and lateral control for the operation of automated vehicles in platoons," *IEEE Transactions on Control Systems Technology*, vol. 8, no. 4, pp. 695–708, 2000.

[16] S. E. Shladover, D. Su, and X.-Y. Lu, "Impacts of cooperative adaptive cruise control on freeway traffic flow," *Transportation Research Record: Journal of the Transportation Research Board*, vol. 2324, no. 1, pp. 63–70, 2012.

[17] J. Ploeg, E. Semsar-Kazerooni, G. Lijster, N. van de Wouw, and H. Nijmeijer, "Graceful degradation of cooperative adaptive cruise control," *IEEE Transactions on Intelligent Transportation Systems*, vol. 16, no. 1, pp. 488–497, 2015.

[18] M. Treiber, A. Hennecke, and D. Helbing, "Congested traffic states in empirical observations and microscopic simulations," *Physical Review E*, vol. 62, no. 2, pp. 1805–1824, 2000.

[19] D. Q. Mayne, J. B. Rawlings, C. V. Rao, and P. O. M. Scokaert, "Constrained model predictive control: Stability and optimality," *Automatica*, vol. 36, no. 6, pp. 789–814, 2000.

[20] W. B. Dunbar and D. S. Caveney, "Distributed receding horizon control of vehicle platoons: Stability and string stability," *IEEE Transactions on Automatic Control*, vol. 57, no. 3, pp. 620–633, 2012.

[21] Y. Zheng, S. E. Li, K. Li, F. Borrelli, and J. K. Hedrick, "Distributed model predictive control for heterogeneous vehicle platoons under unidirectional topologies," *IEEE Transactions on Control Systems Technology*, vol. 25, no. 3, pp. 899–910, 2017.

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

[23] C. Waldschmidt, J. Hasch, and W. Menzel, "Automotive radar—from first efforts to future systems," *IEEE Journal of Microwaves*, vol. 1, no. 1, pp. 135–148, 2021.

[24] National Highway Traffic Safety Administration, "Federal motor vehicle safety standards; V2V communications, notice of proposed rulemaking, docket no. NHTSA-2016-0126," Federal Register, vol. 82, no. 8, pp. 3854–4019, 12 January 2017, 2017. [Online]. Available: https://www.govinfo.gov/content/pkg/FR-2017-01-12/pdf/2016-31059.pdf

[25] Z. Wang, G. Wu, and M. J. Barth, "A review on cooperative adaptive cruise control (CACC) systems: Architectures, controls, and applications," in *2018 21st International Conference on Intelligent Transportation Systems (ITSC)*, 2018, pp. 2884–2891.

[26] Federal Communications Commission, "Use of the 5.850–5.925 GHz band, final rule," Federal Register, vol. 86, pp. 23281–23299, 3 May 2021, 2021. [Online]. Available: https://www.govinfo.gov/content/pkg/FR-2021-05-03/pdf/2021-08802.pdf

[27] ——, "Use of the 5.850–5.925 GHz band, final rule," Federal Register, vol. 89, no. 240, pp. 100838–100856, 13 December 2024, 2024. [Online]. Available: https://www.govinfo.gov/content/pkg/FR-2024-12-13/pdf/2024-28980.pdf

14.10 Exercises

14.1 Concept check · 2 min

In case 5 each radar follower behind a car at constant speed has the closed-loop roots $-0.684 \pm 1.197\,j$ and $-2.632\,\mathrm{s}^{-1}$, a damping ratio of 0.496, yet the brake tap takes car 10 from 72 km/h down to 36.3 km/h. Which statement explains this?

(a) The damping ratio leaves out the radar's latency, and with the latency included each follower's own loop is unstable

(b) The tap drives every follower onto its acceleration limits, and the line amplifies only because the limits clip the commands

(c) A damping ratio of 0.496 overshoots by about 17 %, and nine such overshoots compound into the dip at car 10

(d) The car-to-car gain $|\Gamma|$ exceeds one from zero up to 1.63 rad/s, peaking near the loop's own natural frequency

14.2 Concept check · 2 min

For radar ACC with constant spacing, Section 14.4 finds $|\Gamma(j\omega)|^2 \approx 1 + 2\omega^2/k_p$ near zero frequency. What follows for the follower's gains?

(a) No pair of positive gains makes the line string stable, and the best of 90 pairs on the grid still peaks at 1.050

(b) A spacing gain above about $2\,\mathrm{s}^{-2}$ makes the low-frequency term negative and the line string stable

(c) The rate gain decides it, and any k_d larger than k_p brings the peak below one

(d) The line becomes string stable once the radar's latency is set to zero, since the expansion holds only with delay

14.3 Concept check · 2 min

The app measures the amplification of a car as the ratio of its RMS acceleration to that of the car ahead, Equation (14.9), and not as a ratio of peak accelerations. Why?

(a) Peak accelerations carry the radar's noise of 0.05 m/s, while RMS values average that noise away

(b) The criterion bounds the RMS ratio by Parseval's theorem, while a string-stable line can pass back larger peaks

(c) A peak ratio cannot be formed during a cut-in, so an RMS ratio taken over the whole run replaces it

(d) Peak ratios exceed one only on the ring road, so they cannot tell the open-road cases apart

14.4 Theory and calculation · 7 min

Radar ACC with a time gap and no delays has the transfer function of Equation (14.11), with $\tau = 0.4\,\mathrm{s}$ and $k_d = 1\,\mathrm{s}^{-1}$.

(a) Use the coefficient c_2 of Equation (14.12) to find the shortest string-stable time gap for k_p of 1, 2 and $3\,\mathrm{s}^{-2}$. *Checkpoint:* 1.414 s, 1.000 s and 0.816 s.

(b) At $h = 1.0\,\mathrm{s}$ and $k_p = 2\,\mathrm{s}^{-2}$, evaluate c_2 and c_4, and say whether the full condition of Section 14.4 holds. *Checkpoint:* $c_2 = 0$, $c_4 = 1.6$; it holds, on the boundary.

(c) At $h = 0.6\,\mathrm{s}$ with $k_p = 2$, evaluate $|\Gamma(j\omega)|$ at 1 rad/s, and find the k_p that the low-frequency bound would need at this gap. *Checkpoint:* 1.213; $k_p \geq 5.56\,\mathrm{s}^{-2}$.

14.5 Theory and calculation · 6 min

A radar follower at $h = 0.6\,\mathrm{s}$ with $k_p = 2\,\mathrm{s}^{-2}$ and $k_d = 1\,\mathrm{s}^{-1}$ follows a car at constant speed. Its characteristic polynomial is the denominator of Equation (14.11).

(a) Write the polynomial with numbers, and check that $s = -2.632$ is nearly a root. *Checkpoint:* $0.4s^3 + 1.6s^2 + 2.2s + 2$; the residual is about 3×10^{-4}.

(b) Divide out the real root and find the natural frequency and the damping ratio of the remaining pair. *Checkpoint:* $s^2 + 1.368s + 1.899$; 1.378 rad/s and 0.496.

(c) The $|\Gamma|$ peak lies at 1.12 rad/s. Find the period of a wave at that frequency, and say what case 5's sine profile at a period of 5.6 s probes. *Checkpoint:* 5.61 s.

14.6 Theory and calculation · 6 min

The lead holds 72 km/h. Use $L = 4.5\,\mathrm{m}$ and $r = 3\,\mathrm{m}$.

(a) Find the wanted gap and the flow q_h of Equation (14.8) at time gaps of 1.5, 0.6 and 0.4 s. *Checkpoint:* 33, 15 and 11 m; 1920, 3692 and 4645 veh/h.

(b) Case 7's radio has $\theta = 400\,\mathrm{ms}$ and $p = 20\,\%$ with no jitter. Find the effective delay θ_c the criterion uses. What does a radar latency of 120 ms become in the app? *Checkpoint:* 412.5 ms; 150 ms.

(c) A link counts as stopped after ten periods without an arrival. Find the chance that ten packets in a row are lost at $p = 0.2$, and the expected number of stops from random loss alone on nine links in the 40 s before case 7's burst, counting one chance at each packet. *Checkpoint:* 1.0×10^{-7}; about 7×10^{-4}.

14.7 App challenge · 8 min

Load scenario 4, *Constant time gap at 1.5 s*, at the default seed. Run to $t = 90\,\mathrm{s}$ at a playback speed of 5, pause, read the Measured amplification and Lane flow tiles, and export the CSV. Then press Reset, set Time gap to 1.0 s before Run, and run to 90 s again. Repeat with a time gap of 0.8 s. Car 10's lowest speed is the minimum of the column `v_last_mps` (the Slowest any car went tile shows the lead here). Tolerances: 0.01 on amplification, 5 veh/h on flow, 0.2 km/h on speed.

(a) Read the three values at 1.5 s. *Checkpoint:* 0.93 at car 10, 1920 veh/h, 66.8 km/h.

(b) Read them at 1.0 and 0.8 s, with the $|\Gamma|$ peak tile. *Checkpoint:* 0.99 at car 10, 2618 veh/h, 61.2 km/h, peak 1.003; 1.06 at car 6, 3064 veh/h, 54.5 km/h, peak 1.084.

(c) The peak at 1.0 s is above one, yet the line measures 0.99. Explain, using the boundary of the app's criterion. *Checkpoint:* The boundary is 1.050 s, and the excess is tiny and at 0.37 rad/s.

14.8 App challenge · 8 min

Load scenario 5, *The gap shrunk to 0.6 s*, at the default seed and run to 90 s at a playback speed of 5. Read the Measured amplification, $|\Gamma|$ peak and Slowest any car went tiles. Then load scenario 6, *Cooperative ACC at 0.6 s*, run to 90 s and read the same tiles, the Lane flow tile, and car 10's lowest speed from the column `v_last_mps` of the CSV. Finally press Reset, set Time gap to 0.4 s before Run, and run to 90 s. Tolerances as in the previous exercise.

(a) Read case 5. *Checkpoint:* 1.20 at car 3; 1.306; 36.3 km/h.

(b) Read case 6 at 0.6 s, and say what changed between the two cases. *Checkpoint:* 0.95 at car 10; 1.000; 3692 veh/h; car 10 at 63.7 km/h.

(c) Read case 6 at 0.4 s and relate the result to the cooperative boundary. *Checkpoint:* 1.00 at car 5; 1.033; 4645 veh/h.

14.9 App challenge · 9 min

Load the challenge, *Challenge, stop and go with a cut-in*, at the default seed. The run stops itself at 110 s, and a playback speed of 5 takes 22 s of wall time. Before each later run press Reset, choose the Follower policy, then press Run. Leave every other setting as loaded and do not reveal the suggestion. Read the Measured amplification, Smallest gap, Lane flow and Challenge target tiles (tolerance 0.01, 0.01 m, 5 veh/h).

(a) Run as loaded, on ACC with a time gap. Which target fails? *Checkpoint:* 1.13 at car 10, 1.75 m, 2283 veh/h; amplification fails.

(b) Run with the policy CACC. *Checkpoint:* 1.04 at car 4, 1.75 m, 2281 veh/h; not met.

(c) Run with CACC + fallback. Then press Reset, set the policy back to ACC · time gap and run once more. What does the last run show? *Checkpoint:* 0.99 at car 8, 1.75 m, 2158 veh/h, met; the last run repeats (a).

14.10 Application area · 4 min

In case 8 a car with no radio cuts into the 15 m gap ahead of car 4 of a cooperative line at 0.6 s and 72 km/h. With the gap regrowing at 1 m/s, car 4 brakes at no more than 2.02 m/s^2. Asked for the whole gap at once, it brakes at 5.96 m/s^2. The lane carries 3691 veh/h before the cut-in and 3117 veh/h over the last 30 s.

(a) Find the fraction of the flow lost, and say why the flow falls although no car collides. *Checkpoint:* 15.6 %.

(b) What should a fleet engineer specify for the supervisor of a short-gap platoon, and why must the platoon plan for cut-ins at all? *Checkpoint:* A rate-limited regrowth of the gap, and a fallback gap behind any car without a radio.

14.11 Application area · 3 min

After a hard stop in a model run of 16 cars, drivers 0.70 s late discharge from the standstill at 997 veh/h, against the 1766 veh/h that their equilibrium carries at its best speed. Radar ACC at 1.5 s discharges at 1880 veh/h against 1920 veh/h. The ramp-metering app declares a capacity drop of 10 %.

(a) Find both drops. *Checkpoint:* 44 % and 2 %.

(b) What does the comparison suggest to a traffic engineer, and what in Section 14.8 keeps it from being a forecast? *Checkpoint:* A queue of these drivers discharges far below capacity, and the model's drivers are identical and look one car ahead.

14.12 Review the agent · 5 min

Asked how far a brake tap travels down case 2's line of drivers, an AI assistant answers: "The linearized criterion peaks at 1.024 for drivers 0.70 s late, so each car passes on at most 2.4 % more than it receives. Over nine cars the dip grows at most $1.024^9 = 1.24$ times. The lead dips 14.7 km/h, from 72 to 57.3, so car 10 dips at most 18.2 km/h and stays above 53.8 km/h. Reaction delays near 0.7 s are therefore safe in a ten-car line." Find the flaw, and say how you would check the claim in the app and what the check shows.