INTERACTIVE CONTROL ENGINEERING
Principles and Applications

Volume 2 Practical PID Control

CHAPTER 8

Ramp Metering Simulator for Holding a Freeway Merge Short of Breakdown

Pages 196–223

Edited by John D. Hedengren

How to cite this chapter. "Ramp Metering Simulator for Holding a Freeway Merge Short of Breakdown," in *Practical PID Control*, John D. Hedengren, Ed., vol. 2 of *Interactive Control Engineering: Principles and Applications*. Advanced Process Solutions, LLC, 2026, ch. 8, pp. 196–223.

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

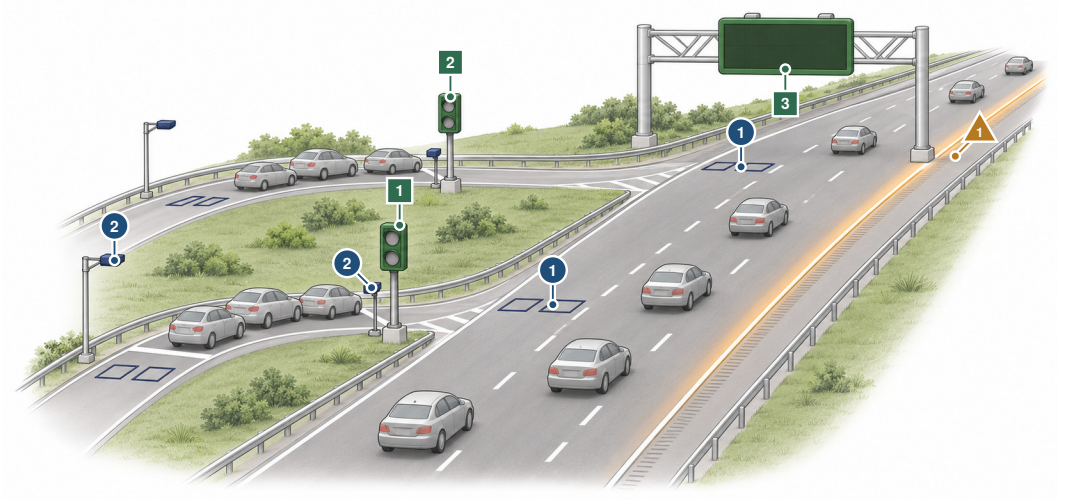
The app for this chapter
https://apmonitor.com/apps/ramp-metering/

The volume online: contents, other chapters and citation records
https://apmonitor.com/apps/books/v02-practical-pid/#ramp-metering

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

CHAPTER 8

Ramp Metering Simulator for Holding a Freeway Merge Short of Breakdown



The controller moves 1 the signal on the upstream ramp 2 the signal on the merge ramp 3 the speed limit sign over the freeway

The controller reads 1 loops in the road at each merge, which sense how much of the time a car is over them

2 counters where cars join each ramp and where they leave its signal

What disturbs it 1 a crash past the merge that closes part of the road

Notice the cars held at both signals: the ramps store the wait that keeps the freeway flowing past the merges. Illustration, not to scale.

IN PLAIN TERMS **Holding a Freeway Merge Short of a Jam**

The main idea. Where an entrance ramp joins a freeway, the road can carry only so many cars. A signal on the ramp lets cars in at a pace that a controller sets, so that the merge stays just below its limit.

The challenge. If the merge gets too full it jams, and a jammed merge passes about a tenth fewer cars until the jam is gone. Every car held back waits on the ramp, and a ramp has room for only so many.

Why it matters. A few minutes of too much traffic can cost more than an hour of jam.

THREE WAYS TO READ THIS CHAPTER

Explore	12 min	the first page, the fact on the count of ramp meters (Section 8.1), the two history segments (Section 8.2), the try-it segments of cases 1 and 4 and the fact on the fixed plan (Section 8.5), and the key points
Study	50 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 8.5), Technology in Practice (Section 8.6), Limitations (Section 8.7) and the open challenges

The control question A freeway merge passes the most traffic just before it breaks down, and the queue it then forms discharges a tenth less until it is gone. How far below that fold should an integral controller hold the merge, and what happens when the ramp storage that pays for the margin runs out?

Learning objectives

- Compute the critical occupancy, the dropped capacity and the recovery density of a merge from its fundamental diagram.
- Explain how a capacity drop with hysteresis turns minutes of overload into more than an hour of breakdown.
- Show that ALINEA is integral control on occupancy, and predict its decay per report from the static gain of the merge.
- Choose an occupancy setpoint with a margin below the fold, and weigh the time saved on the freeway against the ramp queue.
- Judge what a queue override, coordination and a detector watchdog each protect, and what each costs.

Prerequisites Sections 1.4.2 and 1.6.2 (velocity form), 1.4.9 (windup and back-calculation) and 1.4.11 (a loop that acts once per report). Section 1.6.3 is needed only for the predictive comparison.

The app https://apmonitor.com/apps/ramp-metering/

See also Volume 7, Chapter 4 for a setpoint at the peak of the plant's map and Volume 7, Chapter 2 for folds with ignition and extinction. Volume 4, Chapter 2 for a protective override beside a folded map and Volume 4, Chapter 14 for a finite store spent to defend a target.

8.1 Introduction

A lane of freeway carries more vehicles as its density rises, up to a critical density, and fewer as the density rises further, so the steady relation between flow and density folds over at a peak. The kinematic wave theory built on that relation treats traffic as a fluid whose changes travel as waves and whose abrupt changes are shock waves [1, 2], and the cell transmission model represents highway traffic in cells consistently with that hydrodynamic theory [3]. Ramp metering puts a traffic signal on the on-ramp that releases its vehicles at a rate the controller chooses, and an overview of the field explains why it can lead to a dramatic amelioration of traffic conditions on freeways [4]. ALINEA, a local feedback strategy, reported easy application, flexibility and high efficiency in its field applications in Paris and Amsterdam [5].

The mechanism is not obvious, because every vehicle held on the ramp still has to pass the same merge. On the average Tuesday of the simulator the demand at the merge exceeds its 6000 veh/h for 17 min in all, from 7:45 to 8:05, yet with the signal dark the merge breaks down at 7:48 and stays in the capacity drop for 104.9 min, and the morning costs 1840.7 veh·h. ALINEA holding an occupancy of 15.5 % keeps the merge flowing to 10:00 for 1529.4 veh·h, 16.9 % less, while the ramp queue reaches 163.9 vehicles, 103.9 of them past the ramp's storage. The simulator answers the objection that a meter only moves delay from the freeway to the ramp. It moves delay, but the merge it protects keeps a tenth of its capacity that a breakdown would throw away, and the storage that holds the moved delay is finite.

> ✦ DID YOU KNOW? **About 2,160 Ramp Meters**
>
> In 2002 the United States had about 2,160 ramp meters, and about 2,370 were expected by the end of 2005. Metering systems ran in 26 metropolitan areas. Each meter is a signal on an entrance ramp with its controller in a cabinet beside the ramp [6].

The same shape recurs outside traffic. When a computer metered the traffic entering the Lincoln Tunnel in 1967, congestion was already known to reduce the tunnel's throughput appreciably [7]. In the first congestion collapse of the early Internet, in October 1986, the data throughput between two

sites fell from 32 kbps to 40 bps [8], and Chapter 7 holds the queue of such a link with a PI controller. An axial compressor at its stall line falls into surge, or into rotating stall at a substantially reduced flow and pressure ratio [9], and Volume 4, Chapter 2 keeps a centrifugal compressor short of surge with a protective override. In each the operating point should sit near the peak of a map that folds, a step past the peak costs a lower branch, and whatever is held back must wait somewhere. None of the three survives a linearization of the merge around free flow, where the release simply adds to the occupancy.

However, the evaluations of metering cited here were made in the field, at ramps in Paris and Amsterdam and during a metering holiday in the Twin Cities [5, 10], where no morning can be run twice, and a student cannot replay one morning with one setting changed and read the minutes in the drop, the ramp queue and the total time spent side by side. This chapter takes the problem apart into the margin below the fold, the storage behind the meter, the measurement it trusts and the actuator it uses, each with its own setting on a simulated merge whose morning replays identically. The cell transmission model, the capacity drop, ALINEA and its extensions are established methods, and the chapter applies them as published.

Section 8.2 traces freeway control from the kinematic wave to coordinated metering. Section 8.3 gives the model in full, enough to rebuild it without the source: a capacity drop held by a per-cell latch with hysteresis on the cell transmission model, detectors and counters that report once a minute, and a score integrated from the model's own state. Section 8.4 gives the meters as implemented. They read only the minute reports of occupancy and vehicle counts, and ALINEA runs on the velocity-form PID of Section 1.6.2. Section 8.5 reports the measured mornings, ten cases and a graded morning, each isolating one part of the problem and each credited by its own completion rule. Section 8.6 describes the devices of a real meter, Section 8.7 states the limitations, Section 8.8 gathers the results, and the exercises follow in Section 8.9.

8.2 From Kinematic Waves to Ramp Meters

The theory came first. In 1955 a functional relation between flow and concentration on crowded roads, which by then had experimental backing, was used to deduce how changes in the distribution of traffic propagate, to show that vehicles entering a region of higher concentration must reduce speed rather suddenly at a shock wave, and to estimate the extent and duration of the hold-up at a bottleneck too narrow to admit the increased flow [1]. A year later the same picture was reached by replacing individual vehicles with a continuous density and an empirical relation between speed and density, and a traffic signal was found to cause a threshold effect, minor disturbances in light traffic that suddenly build to large values once a critical density is exceeded [2].

Control followed within a decade. A 1965 volume based on papers presented at an interdisciplinary clinic on instrumentation requirements for traffic control systems, held in New York City in December 1963, included experimentation with manual and automatic ramp control [11]. The history box below follows the first ramp meters and a tunnel metered by a computer [7]. By 1973 automatic freeway surveillance and control systems had been introduced in several cities of the United States over the preceding decade, and mathematical models of freeway traffic that allow on-ramp metering were being developed to evaluate and optimize them [12].

> ⧖ HISTORY **From a Police Officer to Feedback, 1963 to 1972**
>
> The first ramp meters went up in 1963 on an expressway in Chicago, and a traffic enforcement officer in the field controlled them by hand. Experiments followed in Detroit and Los Angeles [6]. During most of 1967 an experimental system counted the vehicles in three sections of one lane of the Lincoln Tunnel from detectors at four points, on a computer forty miles away, and reduced the traffic let in whenever congestion threatened the tunnel's throughput. It gave longer periods of sustained high flow [7]. In 1970 the first two meters near St. Paul ran on fixed time, and two years later they were upgraded to respond to traffic [6], the step from case 2 to case 3 of this chapter.

The strategies then grew from fixed-time plans, through traffic-responsive regulators, to nonlinear optimal control [4]. ALINEA, Asservissement Linéaire d'Entrée Autoroutière, a local feedback strategy, had by 1997 been applied at single and multiple ramps of the Boulevard Périphérique in Paris and on the A10 West freeway in Amsterdam [5], and it uses occupancy measured in the merging area, at most a few hundred meters downstream of the ramp [13]. Metering an on-ramp was studied as a way to increase the capacity of an isolated merge [14], and the capacity drop has been observed at freeway merges under ramp control and integrated into a first-order model [15]. The cell transmission model [3] was given an asymmetric form for the optimal metering of freeway ramps [16].

Successors extended the local law. A linked strategy, responsive to traffic, coordinates the metering of two or more consecutive on-ramps with ALINEA at the local level, uses only real-time measurements, and in simulation reached the efficiency of proactive optimal control [17]. When the bottleneck lies further downstream than the merge, ALINEA may give a poorly damped loop, and its proportional-integral extension, PI-ALINEA, restores satisfactory performance [13]. Variable speed limits act on aggregate flow in a way similar to ramp metering at potentially active bottlenecks and can be integrated with coordinated metering [18]. Ramp meters have also been put on trial, with evidence from a metering holiday in the Twin Cities [10].

> ⧖ HISTORY **A Metering Holiday, 2000**
>
> In 2000 a small but vocal group of citizens opposed to ramp metering prompted the Minnesota legislature to have the meters of the Twin Cities switched off for six weeks and to fund an independent evaluation. Traffic was measured for six weeks with the meters and six weeks without. With the meters off, the volumes on the freeway mainline fell by nine percent, freeway speeds fell by 14 percent, travel times became nearly twice as unpredictable, and crashes on the freeways and ramps rose by 26 percent. The system's benefits were put at about 40 million dollars, 15 times its costs [6]. The agency then moved from metering every rebuilt freeway by default to a more performance-based approach [6].

8.3 The Freeway and Its Model

The road is six kilometers of three-lane freeway, traffic from left to right, divided into $N = 12$ cells of $L = 0.5\,\text{km}$, as shown in Figure 8.1. The merge ramp joins cell $m = 8$ at 4.0 km and the upstream ramp, closed unless a case opens it, joins cell $m_1 = 4$ at 2.0 km. Speed limits can hold cells 6 and 7, over the kilometer before the merge, and an incident halves cell 10, a kilometer past it. The signals are classified in Table 8.1. The integrated state holds 39 numbers, 24 density slots of which the 12 cells use 12, the three queues q_0, q_1 and q, and 12 registers that the detectors, counters and score accumulate. Everything discrete, the capacity-drop latches, the incident's end, the queue accounting and the end of the morning at 10:00, is updated after each integration step.

8.3.1 Cells and the Fundamental Diagram

Each cell balances the vehicles that enter and leave it, with time in hours, as shown in Equation (8.1).

$$\lambda L \frac{\mathrm{d}k_i}{\mathrm{d}t} = f_{i-1} - f_i + r\,\delta_{i,m} + r_1\,\delta_{i,m_1} \tag{8.1}$$

where k_i is the density per lane in veh/km, $\lambda = 3$ the lanes, f_i the flow in veh/h from cell i to cell $i+1$, and $\delta_{i,j}$ the Kronecker delta. The flow between two cells is the smaller of what the upstream cell can send and what the downstream one can receive, as shown in Equation (8.2).

$$\begin{aligned} S_i &= \lambda \min\left(v_i k_i,\ (1-\delta\,\ell_i)\,Q_i\right) \\ R_i &= \lambda \max\left(0,\ \min(Q_i,\ w(k_{jam}-k_i))\right) \\ f_i &= \min(S_i,\ R_{i+1}), \qquad f_{N-1} = S_{N-1} \\ f_{-1} &= \min\left(\lambda Q,\ q_0/\Delta t + D_m,\ R_0\right) \end{aligned} \tag{8.2}$$

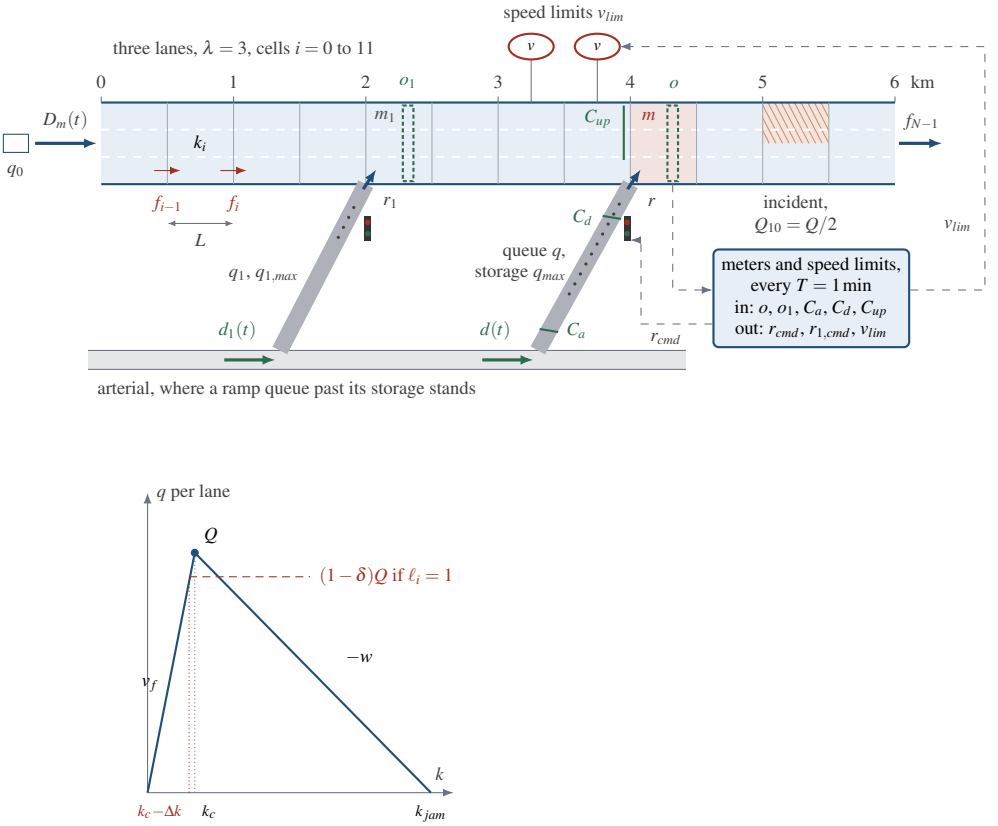


Figure 8.1 The six kilometers in plan, to scale along the road and not across it. Cell i of length L and λ lanes holds the density k_i and passes the flow f_i to cell $i+1$, the mainline demand $D_m(t)$ enters cell 0 through the queue before the map q_0, and f_{N-1} leaves at 6 km. The merge ramp joins the merge cell m with the release r from a queue q of storage q_{max} that the arrivals $d(t)$ feed, and a queue past the storage stands on the arterial. The upstream ramp is the same with r_1, q_1, $q_{1,max}$ and $d_1(t)$ at cell m_1. Loop detectors in cells m and m_1 report the occupancies o and o_1, counters at each ramp's entrance and stop line give C_a and C_d, and a counter at the upstream face of the merge cell gives the mainline count C_{up}. Every T the controller sets r_{cmd}, $r_{1,cmd}$ and the speed limit v_{lim} over cells 6 and 7. The incident halves the capacity of cell 10. The inset is the fundamental diagram per lane, with free speed v_f, wave speed w, capacity Q at the critical density k_c, jam density k_{jam}, the dropped capacity $(1-\delta)Q$ that a cell sends while its latch ℓ_i is set, and the density $k_c - \Delta k$ below which the latch releases.

where S_i and R_i are the sending and receiving flows, v_i and Q_i the free speed and capacity of cell i now, $\ell_i \in \{0,1\}$ its capacity-drop latch, δ the drop, w the congestion wave speed, k_{jam} the jam density, f_{-1} the inflow to cell 0 from the mainline demand D_m of Equation (8.8) and the queue before the map q_0, and Δt the integration step, so that $q_0/\Delta t$ releases that queue in one step. The nominal diagram is a triangle whose three declared numbers close it, as shown in Equation (8.3).

$$Q = v_f k_c = w\,(k_{jam} - k_c) \tag{8.3}$$

so that $k_c = Q/v_f$ and $k_{jam} = k_c + Q/w$. With $Q = 2000$ veh/h per lane, $v_f = 80$ km/h and $w = 16$ km/h the critical density is 25 veh/km and the jam density 150 veh/km per lane. Holding the capacity, a critical density of 25 and a jam density of 150, the triangle closes only at these two speeds, and a free speed of 100 km/h would need a capacity of 2500 veh/h per lane. Rain at a fraction ρ

Table 8.1 Signals of the Freeway by Role, With Their Defaults

Symbol	Role	Default	Range
o	CV, merge occupancy (reported)	SP 15.5 %	SP 12–20
o_1	CV, upstream occupancy	SP 15.5 %	SP 12–20
r	MV, merge ramp release	ALINEA	240–1800 veh/h
r_1	MV, upstream ramp release	ALINEA	240–1800
v_{lim}	MV, speed limit, cells 6, 7	off	20–80 km/h
k_i	states, 12 cell densities	free flow	0–150 veh/km
q, q_1	states, ramp queues	0	storage 60, 150
q_0	state, queue before the map	0	—
—	12 registers, counts and score	0	—
D_m, d, d_1	DV, demand by weekday	Tuesday	5 days
—	DV, scaling, surge	0	±10, 0–20 %
—	DV, incident at 5 km	off	20 min event
—	DV, an event lets out	off	+300 veh/h, 30 min
—	DV, rain, diverted traffic	0	20 %, 600 veh/h
—	DV, detector noise, freeze	2 %, off	0–10 %
C_a, C_d	measured, ramp counts	exact	—
C_{up}	measured, mainline count	exact	—
TTS	score, total time spent	—	veh·h

lowers v_i and Q_i by the same share, so the critical density stays where it was. The incident halves Q_{10} for 20 min. A speed limit v_{lim} below the free speed keeps the congested branch and moves the peak, $v_i = v_{lim}$ and $Q_i = \min\left(Q_i,\ v_{lim} w k_{jam}/(v_{lim}+w)\right)$, which is 1565 veh/h per lane at 30 km/h. Each cell's critical density is $k_{c,i} = Q_i/v_i$.

8.3.2 The Capacity Drop

After every step each cell's latch follows the rule of Equation (8.4).

$$\ell_i \leftarrow \begin{cases} 1, & \ell_i = 0 \text{ and } k_i > k_{c,i} \\ 0, & \ell_i = 1 \text{ and } k_i < k_{c,i} - \Delta k \\ 0, & \text{a speed limit holds cell } i \\ \ell_i, & \text{otherwise} \end{cases} \tag{8.4}$$

with $\delta = 10\,\%$ and $\Delta k = 3$ veh/km per lane, declared teaching values. The gap between the two densities is the hysteresis of the drop. A latched merge sends at most $(1-\delta)\lambda Q = 5400$ veh/h instead of 6000, and it recovers only once its density falls below 22 veh/km per lane, where free flow carries 5280 veh/h. In free flow no cell reaches its critical density, so the latch acts at the head of a queue, and the merge's latch is the one the morning is scored on. The incident cell latches too, so while it is blocked it discharges 2700 veh/h, not the 3000 its halved capacity would allow. The exemption under a speed limit is the premise of mainstream control, that vehicles held to a limit accelerate to it rather than out of a jam.

8.3.3 Ramps, Queues and the Merge

Each ramp's queue gains its arrivals and loses its release, and the release is limited four ways, as shown in Equation (8.5).

$$\begin{aligned} \frac{\mathrm{d}q}{\mathrm{d}t} &= d - r, \qquad \frac{\mathrm{d}q_0}{\mathrm{d}t} = D_m - f_{-1} \\ r &= \max\left(0,\ \min\left(r_{cmd},\ r_{max},\ q/\Delta t + d,\ \tfrac{1}{3}R_m\right)\right) \end{aligned} \tag{8.5}$$

where r_{cmd} is the meter's command and $r_{max} = 1800\,\text{veh/h}$ a two-lane signal's highest rate, and q_1 and r_1 obey the same equations with d_1 and R_{m_1}. The ramp enters on top of the mainline inflow, up to a third of the merge cell's receiving flow, which the application describes as the merge of the asymmetric cell transmission model [16]. In a congested merge whose ramp has vehicles waiting that share is a quarter of what the merge passes, and when mainline and ramp together bring more than the merge can pass, the merge cell fills past k_c, which is the breakdown. A queue longer than its storage is not blocked. The extra vehicles stand on the arterial, are counted as spillback, and still count in the total time spent.

8.3.4 Detectors, Counts and the Score

A loop detector reports the share of time a vehicle is over it, the occupancy, which in steady traffic is the density times the effective vehicle length. Every reporting period T the merge station reports the period's mean, as shown in Equation (8.6).

$$\begin{aligned} o(k) &= \max\left(0,\ \frac{L_{eff}}{1000}\,\bar{k}_m(k)\,\left(1+\sigma_n \nu_k\right)\right) \\ L_{eff} &= L_{veh} + g_{det} \end{aligned} \tag{8.6}$$

where $\bar{k}_m(k)$ is the mean density of the merge cell over report k, $L_{veh} = 4.5\,\text{m}$ the mean vehicle length, $g_{det} = 2.0\,\text{m}$ the detection zone, $\sigma_n = 2\,\%$ the noise and ν_k a standard normal deviate from the seeded generator. The critical occupancy is $L_{eff} k_c/1000 = 16.25\,\%$, one percentage point of merge occupancy is 369 veh/h of free-flow discharge, and the default setpoint of 15.5 % stands for 23.85 veh/km per lane and 5723 veh/h. The upstream station reports o_1 the same way. A frozen detector repeats its last reading. The counters at each ramp's entrance and stop line give the arrivals $\hat{d} = 60\,\Delta C_a/T$, the release $\hat{r} = 60\,\Delta C_d/T$ and the queue $\hat{q}_k = \max(0,\ \hat{q}_{k-1} + \Delta C_a - \Delta C_d)$, and the mainline counter gives the flow into the merge cell $\hat{q}_{up} = 60\,\Delta C_{up}/T$, with T in minutes. The meters read nothing else from the road.

The score is the total time spent of Equation (8.7).

$$\text{TTS}(t) = \int_{6:00}^{t} \left(\sum_{i=0}^{N-1} \lambda L k_i + q_0 + q_1 + q \right) dt \tag{8.7}$$

It is integrated as a state, from the model's own vehicles rather than from any detector, and every vehicle-hour on the freeway, the ramps and the arterial counts the same.

8.3.5 The Morning's Demand

The demand is declared, not routed. Each stream s, the mainline, the merge ramp and the share of the mainline that uses the upstream ramp when it is open, follows Equation (8.8).

$$D_s(t) = P_s(t - \tau_d)\,\kappa_d \left(1 + \frac{a}{100}\right)\left(1 + \frac{u}{100}\right)\left(1 + \varepsilon_{s,d}(t)\right) \tag{8.8}$$

where P_s is the average weekday of Table 8.2, interpolated linearly in time, κ_d and τ_d the scale and shift of weekday d, a the demand scaling and u the surge, both in percent, and $\varepsilon_{s,d}$ a first-order autoregressive fluctuation of 2 % standard deviation and 6 min correlation time, drawn once a minute from a generator seeded by the weekday and interpolated between minutes. Tuesday is the average day. Monday is 1.5 % lighter and 6 min later, Wednesday 1.5 % heavier and 4 min earlier, Thursday 3 % heavier and 8 min earlier, and Friday 4.5 % lighter and 12 min later. When the upstream ramp is open its stream is taken out of the mainline rather than added, so the merge sees the same morning. Diverted traffic and an event letting out add to the merge ramp. On Tuesday the mainline and the merge ramp bring 4396 veh/h at 6:00, peak at 6147 veh/h at 7:51 and are above 6000 veh/h for 17 min.

Table 8.2 The Average Weekday, veh/h, at the Knots of $P_s(t)$

Clock	6:00	6:20	6:50	7:20	7:45	8:05	8:30	8:55	9:25	10:00
Mainline	3975	4020	4415	4770	4900	4900	4725	4370	4020	3800
Merge ramp	520	540	720	920	1050	1050	915	720	595	520
Upstream	380	390	470	560	600	600	560	480	420	380

The upstream row is part of the mainline, used only while that ramp is open.

Table 8.3 Model Constants. Class V Is a Slider in the Show Model Dialog, H a Hidden Parameter and C Fixed in the Code

Symbol	Meaning	Value	Class
Fundamental diagram and breakdown			
v_f	free-flow speed	80 km/h	V
Q	capacity per lane	2000 veh/h	V
w	congestion wave speed	16 km/h	V
λ	lanes	3	H
k_c, k_{jam}	critical and jam density, per lane	25, 150 veh/km	derived
δ, Δk	capacity drop, recovery margin	10 %, 3 veh/km	V
Geometry and numerics			
L, N	cell length, cells	0.5 km, 12	V, at reset
Δt	Euler step	2 s	C
—	ramp share of R_m	1/3	C
—	incident: cell 10, capacity, duration	$Q/2$, 20 min	C
Ramps			
q_{max}, $q_{1,max}$	storage, merge and upstream ramp	60, 150 veh	V
r_{min}, r_{max}	lowest and highest release	240, 1800 veh/h	V
Detectors and demand			
L_{veh}, g_{det}	vehicle length, detection zone	4.5, 2.0 m	V
T	reporting period	1 min	V
—	fluctuation, deviation and correlation	2 %, 6 min	C
—	noise seed, weekday seeds	20261001, 101–505	C

8.3.6 Initial State and Numerics

The constants are collected in Table 8.3, and the settings a student changes in Table 8.4. Every case starts at 6:00 in steady free flow at the opening demand, each cell at the density $\min(\phi_i/\lambda,\ 0.999Q)/v_f$ for the opening flow ϕ_i through it, 302.6 vehicles in all, with every queue empty, and the run pauses at 10:00. The states advance by explicit Euler at $\Delta t = 2$ s, a step that the application supplies in place of the engine's default (Section 1.6.1), because the cell transmission model is defined as that forward step, and the Courant number $v_f \Delta t / L$ is 0.089. Each step follows the order of Figure 1.10, a scripted event due at its start, then a due report and the meters' action, the Euler update, the latches with the queue accounting, the outputs, and every 10 s a log row. A row of the exported file holds the merge occupancy reported and true, the releases commanded and made, the arrivals, the queues, the speed limit, the merge's discharge and capacity, the drop flags, twelve segment densities, the running total time spent, and the disturbances, parameters and settings in force. Halving the step moves the total time spent of cases 1 and 3 by 0.005 % and 0.003 %, and halving the cell length, to 24 cells, moves it by 0.080 % and −0.027 % and moves the breakdown of case 1 by 14 s. The application passes its 58 verification checks, among them vehicle conservation to 10^{-9} of the arrivals and a jam front that travels upstream at w within 5 %.

8.3.7 Assumptions

A first-order model. Densities and flows are cell averages on a triangular diagram, so every wave inside a queue travels at w, fronts are smeared over a cell, and the stop-and-go oscillations that real queues carry are absent.

Table 8.4 Meter Settings, Defaults and Ranges

Symbol	Setting	Default	Range
—	merge metering strategy	ALINEA	4 choices
$\hat{o}$	merge occupancy setpoint	15.5 %	12–20
K_R	ALINEA gain, veh/h per % per min	70	0–200
K_P	PI-ALINEA gain, veh/h per %	0	0–200
r_{man}	manual release	1800 veh/h	0–1800
$\hat{C}$	capacity estimate, demand-capacity	5800 veh/h	4500–6500
—	stuck-detector watchdog	off	on, off
—	queue override	off	on, off
q_{lim}	override limit, share of storage	85 %	30–100
—	upstream ramp	closed	open
$\hat{o}_1$, K_{R1}	upstream setpoint and gain	15.5 %, 70	as above
—	coordination	off	on, off
—	variable speed limits	off	on, off
$\hat{o}_V$	speed-limit occupancy setpoint	16.0 %	12–20
K_V	speed-limit gain, km/h per % per min	2	0–20
v_{man}	manual speed limit	80 km/h	20–80

A deterministic drop. The capacity drop is a latch with two declared numbers. The application's own assumptions add that real breakdowns near capacity are probabilistic rather than certain past a density, and that the size of the drop differs between sites and mornings.

A lumped merge. Weaving, lane changing and the length of the acceleration lane are lumped into the ramp's one-third share of the merge cell's receiving flow.

Declared demand. Nobody diverts when a ramp queue grows or the freeway jams, and nobody leaves earlier the next day, so metering cannot change the demand it serves.

Point queues and exact counts. A ramp queue is a number against a storage, the arterial's own traffic is not modeled, the counters miss nothing and the detector's only fault is its noise or a freeze.

Obedient speed limits. Drivers hold the limit exactly, and a held cell is exempt from the drop, the most optimistic assumption in the model.

Teaching values. Every threshold is a teaching value. The application states that real capacities, critical occupancies, storages and queue limits come from the operating agency's own detector data and design standards, and it names the Highway Capacity Manual and the Ramp Management and Control Handbook of the Federal Highway Administration as such sources.

8.4 The Meters as Implemented

The controller reads the minute reports of Section 8.3, the period means of occupancy and the ramp counts, and never a density or a queue. The stations report and the meters act once per period $T = 1$ min, at the start of the step that begins on each whole minute. A priming call at 6:00 issues the first commands from the initial state, and from 6:01 on every report covers exactly the 30 integration steps of the minute just ended. Table 8.4 lists the settings.

8.4.1 ALINEA on the Shared Controller

ALINEA raises the release when the merge occupancy is below its setpoint and lowers it when it is above, as shown in Equation (8.9).

$$r(k) = \hat{r}(k-1) + K_R T \left(\hat{o} - o(k)\right) - K_P \left(o(k) - o(k-1)\right) \tag{8.9}$$

clipped to $[r_{min}, r_{max}] = [240, 1800]$ veh/h, with the occupancies in percent and $K_P = 0$ by default, while $K_P > 0$ gives PI-ALINEA [13]. The previous rate $\hat{r}(k-1)$ is the release the stop-line counter measured, clipped to the same range, not the last command, because with no queue waiting the ramp releases only what arrives, and a law that started from its own command would wind up far above it. The law runs on the velocity-form PID of Section 1.6.2, executed once per report, as shown in

Equation (8.10). It is Equations (1.23) to (1.26) with no derivative term and no bias, with s_k for S_k, c_k for v_k, the report period T for the control period and K_i for K_c/τ_I.

$$\begin{aligned} e_k &= \hat{o} - o(k) \\ s_k &= s_{k-1} + K_c\,(e_k - e_{k-1}) \\ c_k &= s_k + I_{k-1} \\ u_k &= \min\left(r_{max},\ \max(r_{min},\ c_k)\right) \\ I_k &= I_{k-1} + K_i\, e_k\, T + (u_k - c_k)\, T/T_t \end{aligned} \tag{8.10}$$

where e_k is the error in percent, s_k accumulates the proportional increments without limit, c_k is the demanded rate, u_k the emitted rate, I_k the integral, which also carries the initial rate because the bias is zero, and $T_t = T$ the tracking time of the back-calculation. The integral used at report k was accumulated through report $k-1$, so the present error rides on the proportional gain. ALINEA sets $K_c = K_P + K_R T$ and $K_i = K_R$, and before each report it shifts I_{k-1} by $\hat{r}(k-1) - u_{k-1}$, the difference between the counted release and the loop's last output. Away from the limits Equation (8.10) is then Equation (8.9) exactly, which the application's checks confirm at reporting periods of 1 and 2 min. The application passes $T_t = T$ rather than the default of Section 1.6.2, the integral time $K_c/K_i = T + K_P/K_R$, so at a limit the back-calculation absorbs the whole clip in one period and the two laws agree for any K_P. PI-ALINEA with $K_P = 50$, driven to its floor by two fabricated reports at 30 %, answers a report back at the setpoint with 965 veh/h, as its law gives, and the checks confirm it. There is no derivative term and no rate limit on the release, whose only limits are the clip to [240, 1800] veh/h, and the loop raises the release when the occupancy is below its setpoint, the direction the application's sign checks confirm.

In free flow a point of merge occupancy is 369 veh/h of discharge, so each report moves the occupancy by about $K_R/369 = 0.19$ of its error at the default gain. In a model run on the average day's plateau, with the detector noise at zero, the demand fluctuation removed (which no setting of the application does) and a queue on the ramp, a setpoint step from 15.5 to 15.0 % at 7:52 decays by a factor of 0.791 per report against the 0.810 that the static gain predicts, a time constant of 4.3 min, and by 0.901 against 0.905 at $K_R = 35$. The loop is integral on a plant whose gain near the setpoint is nearly static, so in steady flow the occupancy settles at $\hat{o}$ whatever the demand, and the meter finds the merge's room without knowing its capacity.

8.4.2 Queue Override and Coordination

The override releases at least what keeps the counted queue inside its limit at the next report, whatever the occupancy says, as shown in Equation (8.11).

$$r(k) \geq \hat{d}(k) - \left(q_{lim}\, q_{max} - \hat{q}(k)\right) \frac{60}{T} \tag{8.11}$$

with $q_{lim} = 85\,\%$, 51 of the 60 vehicles of storage, and T in minutes. When it binds, the integral is shifted to the rate applied. The upstream meter runs its own ALINEA on cell m_1 whenever the merge meter is metering, with the same override on its storage of 150. With coordination, which the application calls HERO-style after the linked strategy of [17], the merge ramp's counted queue asks the upstream ramp to store, as shown in Equation (8.12).

$$r_1(k) \leq \hat{d}_1(k) - \left(q_1^* - \hat{q}_1(k)\right) \frac{60}{\tau_c} \tag{8.12}$$

where $\tau_c = \max(T, 5\,\text{min})$ and the target q_1^* is the upstream ramp's override limit, 127.5 vehicles, from the report at which $\hat{q}$ reaches 60 % of its storage until it falls below 30 %. After that the target falls by 2 % of the upstream storage, 3 vehicles, a minute, so the held vehicles reach the merge as a trickle. The upstream ramp's own override still wins.

8.4.3 Speed Limits, Fixed Plan and Demand-Capacity

The speed limit over cells 6 and 7 is a third integral loop on the merge occupancy, $v(k) = v(k-1) + K_V T(\hat{o}_V - o(k))$, between 20 and 80 km/h and changing by at most 10 km/h a minute, on the same shared PID with $K_c = K_V T$ and $K_i = K_V$. Its tracking time is T, so the integral absorbs the clip and the rate limit in one period and $v(k-1)$ is the limit last applied. Its setpoint is its own. Below the meter's it takes over from the meter, and above it, it acts only once the meter can no longer hold the merge. The fixed-time plan replays a forecast of the average day, in quarter hours, as shown in Equation (8.13).

$$r_{plan}(t) = 60\,\mathrm{round}\Big(\frac{0.97\lambda Q - P_m(t_s)}{60}\Big) \tag{8.13}$$

clipped to $[r_{min}, r_{max}]$, where t_s is the middle of the quarter hour, which gives 1800 veh/h at 6:00, 900 veh/h at 7:45 and 1800 again from 9:15. The demand-capacity rule releases $\hat{C} - \hat{q}_{up}$, the capacity estimate less the counted mainline flow into the merge cell, while the occupancy is below the critical 16.25 %, and r_{min} above it. A watchdog, when it is on, treats a reading that three reports in a row have repeated as a failed detector and falls back to the fixed plan.

In manual mode each actuator takes its station, r_{man}, the upstream release and v_{man}, while every PID tracks the applied value, so that switching back to automatic moves no actuator, the bumpless transfer of Section 1.4.10. The simulation model is Equations (8.1) to (8.8). Nothing in the meters is identified from data. The classical controllers' internal models are the plan's forecast of the average day, the capacity estimate $\hat{C}$ of 5800 veh/h, the critical occupancy of the nominal diagram in the demand-capacity rule, and the counted queues. ALINEA uses no model at all beyond the sign of its gain.

8.4.4 Predictive Metering With Counted Inventories

The predictive option follows the arrangement of Section 1.6.3, and this subsection gives what is specific to the merge. Predictive traffic coordination [19] motivates two nominal cells, $L_a = 2$ km and $L_b = 0.5$ km, with counted queues. Existing detectors and counters supply occupancy and flows. Nominal $v_f = 80$ km/h, $Q = 2000$ veh/h per lane, $w = 16$ km/h and $\lambda = 3$ stay fixed. Occupancy gives k_a, k_b through the nominal effective vehicle length. Their balances are

$$\begin{aligned} 60\lambda L_a \dot{k}_a &= d_0 + r_1 - F, \quad 60\dot{q}_1 = d_1 - r_1 \\ 60\lambda L_b \dot{k}_b &= F + r - S(k_b), \quad 60\dot{q} = d - r \\ F &= \min(S(k_a), R(k_b)) \\ S(k) &= \lambda \min(v_f k, Q) \\ R(k) &= \lambda \max(0, \min(Q, w(k_{jam} - k))). \end{aligned} \tag{8.14}$$

Time is in minutes, flows in veh/h. Background d_0 subtracts counted upstream release from mainline flow into the merge, clipped nonnegative despite their travel-time mismatch. Priming uses nominal sending at measured upstream occupancy less public average ramp demand. Normalized states and inputs are $x = (k_a/k_c, k_b/k_c, q_1/100, q/100)$ and $u = (r_1, r)/1000$. Central differences of width 10^{-5} at the measured state and exact zero-order-hold discretization give

$$x_{j+1} = A_d x_j + B_d u_{b(j)} + E_d w_j + c_d + (b_a, b_b, 0, 0)^{\mathsf{T}}, \tag{8.15}$$

where $w_j = (d_0, d_1, d)$, c_d retains the affine remainder and $b(j)$ indexes a move block. Nominal lags are 1.5/0.375 min, without explicit delay or data fitting. Analytic discretization agrees within 5.1×10^{-12}. Noise-free one-report reconstruction using realized counts for evaluation gives merge RMSE 0.03257/2.73173/0.03839/5.55719/2.01376 percentage points for ALINEA/override/coordination/incident/frozen detector. This does not measure forecast accuracy.

Density reconstruction innovations e_i are clipped to ± 0.5 normalized, then $b_i^+ = b_i + (1 - e^{-T/5})(e_i - b_i)$ at every report, including classical control and Manual. Queues remain counted balances. Bias correction addresses offset [20], without Kalman filtering or guaranteed offset removal.

Optimization begins at priming. The detector period T_{det} (minutes) sets $T = \max(1, T_{det})$, nominally 1 min. Thirty intervals in blocks $(1,1,1,1,2,2,3,4,6,9)$ give 20 release decisions and three horizon-wide slacks. The initial release pair is selected [21] from Equation (8.16).

$$\begin{aligned} \min \quad & \sum_{j=1}^{30} \ell^{\mathsf{T}} x_j + \sum_{h=1}^{10} \|\Delta u_h\|_2^2 + s^{\mathsf{T}} W s + v^{\mathsf{T}} s + \tfrac{10^{-8}}{2} \|z\|_2^2 \\ \text{s.t.} \quad & \text{Equation (8.15)}, \quad x_j \geq 0, \quad s \geq 0 \\ & r_{min}/1000 \leq u_{h,i} \leq r_{max}/1000 \\ & |\Delta u_{h,i}| \leq 1.8, \quad x_{j,2} \leq 0.94 + s_o \\ & x_{j,3} \leq 0.85 q_{1,max}/100 + s_1 \\ & x_{j,4} \leq 0.85 q_{max}/100 + s_2. \end{aligned} \tag{8.16}$$

Here $\ell = T(2.5, 0.625, 100/60, 100/60)$ converts predicted inventory to vehicle-hours, $s = (s_o, s_1, s_2)$, $W = \text{diag}(20000, 5000, 5000)$, $v = (100, 20, 20)$ and z contains all releases and slacks. Penalty weights price the normalized variables in vehicle-hours. Move differences start from counted releases clipped to the command box. The QP adds the ridge shown (excluded from reported J), with 600 iterations allowed. Closed upstream input is fixed at zero. Current command limits set the hard release boxes. Cell/queue nonnegativity is hard. Queue rows cumulatively constrain releases by forecast arrivals, while the 15.275 % occupancy and 85 % current-storage ceilings are soft. The 1.8 change cap spans the ordinary release range and imposes no physical slew restriction. The reduced inventory cost omits most road cells and the origin queue, so it differs from scored TTS.

Preview On adds the public profile's change from now to each midpoint to measured demand. Error scales the future profile before subtracting its present average, up to ± 20 %. Off holds measured counts. Events and seeded fluctuations remain unknown. The fixed plan uses that public profile in quarter-hour slots. ALINEA/demand-capacity lack preview, making adjacent Off rows the present-information comparison.

A requested allowance differs from available flow in Equation (8.5). With a 600 veh/h floor and about 380 veh/h upstream arrivals, hard inventories make 11/11 attempts infeasible over 10 min. Invalid status, numbers, shapes, inventories or future bounds reject a plan. First failure holds commands for one report, then tracked ALINEA acts. The watchdog immediately selects the fixed plan. Protections retain priority and valid plans recover control. Inactive/watchdog-skipped attempt fields are zero, with observer history/cost current.

8.5 Ten Mornings and a Challenge

Table 8.5 lists the cases. All of them start from the defaults of Table 8.4 on the average Tuesday, with the changes that define a case written in its row. Every number below was measured headlessly, on the engine with no interface stepped as the browser steps it, with any change made at 6:00 before the run starts, at the default seed and over the whole morning from 6:00 to 10:00 unless another window is named, and the main results are collected in Table 8.7, three of whose columns Figure 8.2 draws. A setting chosen before the run and the same setting carried through a Reset give one run. The detector noise comes from the default seed of Section 1.6.1, so a case replays exactly, and where ten other noise seeds change a comparison the text says so. At its fastest speed the application plays a whole morning in 0.8 s.

Table 8.6 gives each case under the name the application shows, with the requirement for completing it under the rules of Section 1.6.4. A change that a second morning asks for must be in place

Table 8.5 The Ten Cases and the Challenge

#	Change from the defaults	Teaching point
1	signal dark	the breakdown and its drop
2	fixed-time plan, five weekdays	a forecast, right on average
3	the defaults, ALINEA at 15.5 %	the margin and its queue
4	$\hat{o} = 16.5\,\%$	a setpoint past the fold
5	queue override at 85 %	storage is the budget
6	incident at 7:40, 20 min	the honest infeasible case
7	override, detector freezes at 7:20	integral action on a frozen reading
8	override, upstream ramp open, coordination	borrowed storage
9	signal dark, speed limits at 15.2 %	the mainline as storage
10	predictive metering against ALINEA, both ramps	model and forecast error
11	demand +1 %, both ramps, fixed plan	graded morning

Table 8.6 Scenarios and Completion Requirements

#	Scenario	Kind	Required to complete
1	No metering: the breakdown and the shockwave	Demonstration	Run the dark-signal morning through the breakdown of the merge to the 10:00 end.
2	Fixed-time metering, day by day	Experiment	Run the fixed-time plan to 10:00 on Tuesday, then a second morning to 10:00 on Thursday.
3	ALINEA just under the critical occupancy	Experiment	Run ALINEA at $\hat{o} = 15.5\,\%$ to 10:00, then a second morning to 10:00 at $\hat{o} = 15.0\,\%$.
4	The setpoint a little too high	Demonstration	Run to 10:00 with $\hat{o}$ of at least 16.4 % and two or more breakdowns, then a second morning to 10:00 at $\hat{o} = 16.0\,\%$.
5	The queue override	Demonstration	Run to 10:00 with the override on, a merge ramp storage of at most 60 vehicles and the merge broken down, then a second morning to 10:00 with a storage of 115 to 125 vehicles.
6	An incident past the merge	Demonstration	Run to 10:00 with ALINEA and no override, through the breakdown the incident causes, then a second morning to 10:00 with the override on.
7	A stale detector	Demonstration	Run past 8:00 without the watchdog until the merge breaks down, then a second morning past 8:00 with the watchdog on before 7:19, so that the meter falls back to the fixed plan after the 7:20 freeze.
8	Two ramps and coordination	Experiment	Run to 10:00 with coordination on, then a second morning to 10:00 with coordination off.
9	Speed limits against metering	Experiment	Run to 10:00 with the speed limits at $\hat{o}_V = 15.2\,\%$, then a second morning to 10:00 with $\hat{o}_V$ between 14.5 and 15.0 %.
10	MPC versus feedback on the weekday forecast	Experiment	Run the predictive controller to 10:00 with both ramps open, then repeat with ALINEA after Reset.
11	Challenge, a busier Tuesday with both ramps open	Challenge	By 10:00, with 1 % more demand and both ramps open as given, a total time spent below 1650 veh·h, neither ramp queue ever past its storage, and at most 3 min in the capacity drop.[a]

[a] As shipped (fixed plan at the merge ramp, ALINEA upstream) the morning fails all three parts, with 1833.4 veh·h, 103.2 min in the drop and queues of 66.9 of 60 and 187.8 of 150 vehicles. The suggested ALINEA with the override and coordination passes with 1551.9 veh·h (98.1 veh·h below the limit), no breakdown and queues of 51.2 of 60 and 127.4 of 150.

within its first 10 min. The challenge is judged on the case as given, so a changed model parameter or disturbance voids the attempt, while the strategies, the override, coordination and the speed limits are the reader's to choose. The footnote of the table gives its outcome as shipped and under the suggested configuration.

Table 8.7 Measured Results Over the Morning, 6:00 to 10:00. Queue Is the Worst Merge Ramp Queue Against Its Storage of 60

Case and configuration	TTS, veh·h	Breaks	Drop, min	Queue
1, no metering	1840.7	7:48	104.9	0
2, fixed plan, Tuesday	1438.5	—	0	59.3
2, fixed plan, Thursday	2667.7	7:22	158.3	82.2
3, ALINEA at 15.5 %	1529.4	—	0	163.9
3, $\hat{o}$ 15.0 %	1835.7	—	0	365.2
3, $\hat{o}$ 16.0 %	1425.9	—	0	50.2
3, demand-capacity	1487.9	—	0	125.5
4, $\hat{o}$ 16.5 %	1731.5	4 times	40.4	341.3
4, $\hat{o}$ 16.2 %	1597.0	twice	19.3	257.7
5, override at 51	1766.0	7:57	89.1	51.2
5, override, storage 120	1481.0	—	0	102.2
6, incident, ALINEA	3836.5	7:41	102.9	1056.7
6, incident, no metering	3846.5	7:40	139.5	79.4
7, frozen detector	1840.7	7:48	104.9	0.2
7, with the watchdog	1434.7	—	0	51.1
8, coordination	1459.9	—	0	51.1
8, without coordination	1716.2	7:57	88.9	51.2
9, speed limits at 15.2 %	1511.0	—	0	0
9, speed limits at 14.5 %	2095.6	—	0	0

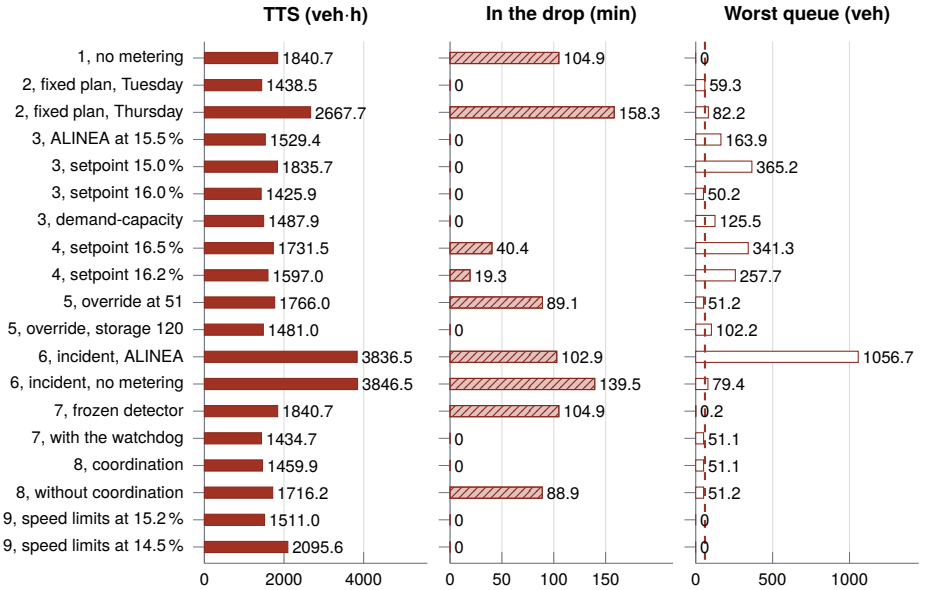


Figure 8.2 Three columns of Table 8.7 over the whole morning: total time spent (solid bars), minutes in the capacity drop (hatched) and the worst merge ramp queue (open), with the storage of 60 vehicles dashed. Notice that the meters that keep the merge out of the drop hold a queue on the ramp instead, that the speed limits of case 9 store the delay on the mainline with no queue at all, and that the incident of case 6 costs more than any other morning whatever the meter does.

8.5.1 The Breakdown and the Drop That Holds It

With the signal dark, the merge passes 5815 veh/h on average from 7:37 to 7:46, at most 5890, and 5950 at 7:47. At 7:48 its density passes 25 veh/km per lane and it latches, and from then on it discharges 5400 veh/h, as shown in Figure 8.3(b). The queue grows for as long as the demand exceeds 5400, which it does until 8:40, and its tail reaches 3.5 km behind the merge at 8:44. The

merge regains its capacity only at 9:33, when its density falls below 22 veh/km per lane, so 17 min of overload cost 104.9 min in the drop. On the diagram the merge cell leaves the free branch at its peak, holds the density of the queue behind it, up to 59.1 veh/km per lane, while it sends the dropped 1800 veh/h per lane, and returns to the free branch only below the recovery density, as shown in Figure 8.3(a). The merge cell sits off the triangle there because a cell's outflow in the model is its sending flow, and the head of a queue sends its dropped capacity.

> **TRY IT IN THE APP** **1 · No Metering: The Breakdown and the Shockwave**
>
> Load the case and press Run. Watch the road at the merge, at the 4 km mark: a little before 8:00 it turns from pale to red as the merge breaks down. Keep the run going to 10:00 and read the Capacity drop tile. The merge stayed in the drop for 104.9 min, although the demand was above capacity for only 17.

The queue's tail is a shock. From 7:58 to 8:10 the free flow arriving at it carries 1634 veh/h per lane at 21.0 veh/km per lane and the queue just inside it passes 1459 at 57.1 veh/km, and the jump condition, the difference of the flows over the difference of the densities, gives 4.8 km/h upstream against the 5.0 km/h that the time–space record shows between 1 and 2 km behind the merge. The congestion wave speed of 16 km/h is the speed of changes inside the queue, not of its tail. The morning costs 1840.7 veh·h.

Pause. Try Exercise 8.1 before reading on.

8.5.2 A Plan Built for the Average Day

The fixed-time plan releases, in each quarter hour, what should fill the merge to 97 % of its capacity on the average day. On the Tuesday it was built for it holds the merge, with a queue of 59.3 at the worst, and the morning costs 1438.5 veh·h against 1840.7. On every other day its forecast misses. Wednesday and Thursday come earlier and heavier, the merge breaks down at 7:33 and 7:22, two minutes after and in the same minute as with no meter, and on both days the queue puts 22 vehicles on the arterial, so Wednesday saves 37.0 veh·h and Thursday nothing. On Monday and Friday the merge would have flowed anyway and the plan only holds vehicles, 44.9 and 30.9 at the worst, for 25.7 veh·h and 15.5 veh·h more than no meter. Over the five days the mean falls from 1899.5 to 1819.9 veh·h, 4.2 %, and the week's net saving of 398.0 veh·h is less than the 402.2 that Tuesday saves alone. ALINEA on the same five days averages 1579.0 veh·h, because a meter that measures the merge does not need to know which morning it is.

> **DID YOU KNOW?** **The Plan That Wins on Its Own Day**
>
> From the simulator: on the Tuesday it was built for, the fixed-time plan costs 1438.5 vehicle-hours, less than the 1529.4 of ALINEA at its default setpoint. On Thursday, which comes earlier and heavier, the same plan lets the merge break down at 7:22, and the morning costs 2667.7 (Table 8.7).

8.5.3 Integral Control Under the Fold

ALINEA at 15.5 % never lets the merge break down. From 7:30 to 8:30 the merge passes 5719 veh/h on average at a true occupancy of 15.49 %, against the 5400 a broken-down merge discharges, and the release bottoms at 655 veh/h at 7:55. By 9:01 the metered merge has passed 321 more vehicles than the unmetered one, and by 10:00 both have passed the same 20 706, so what metering saves is time, 311.3 veh·h of it. The cost is the ramp queue, which passes its storage of 60 at 7:57, peaks at 163.9 vehicles at 8:27 with 103.9 of them on the arterial, and is empty only at 9:02, as shown in Figure 8.3(c). The margin cuts both ways. A setpoint of 15.0 % holds back vehicles the merge had room for, the queue reaches 365.2 and the morning costs 1835.7 veh·h, nearly what no meter costs, while 16.0 % keeps the queue at 50.2 for 1425.9 veh·h without a breakdown this Tuesday. That last margin is thinner than it looks. On each of ten other detector-noise seeds 16.0 % breaks the merge down for

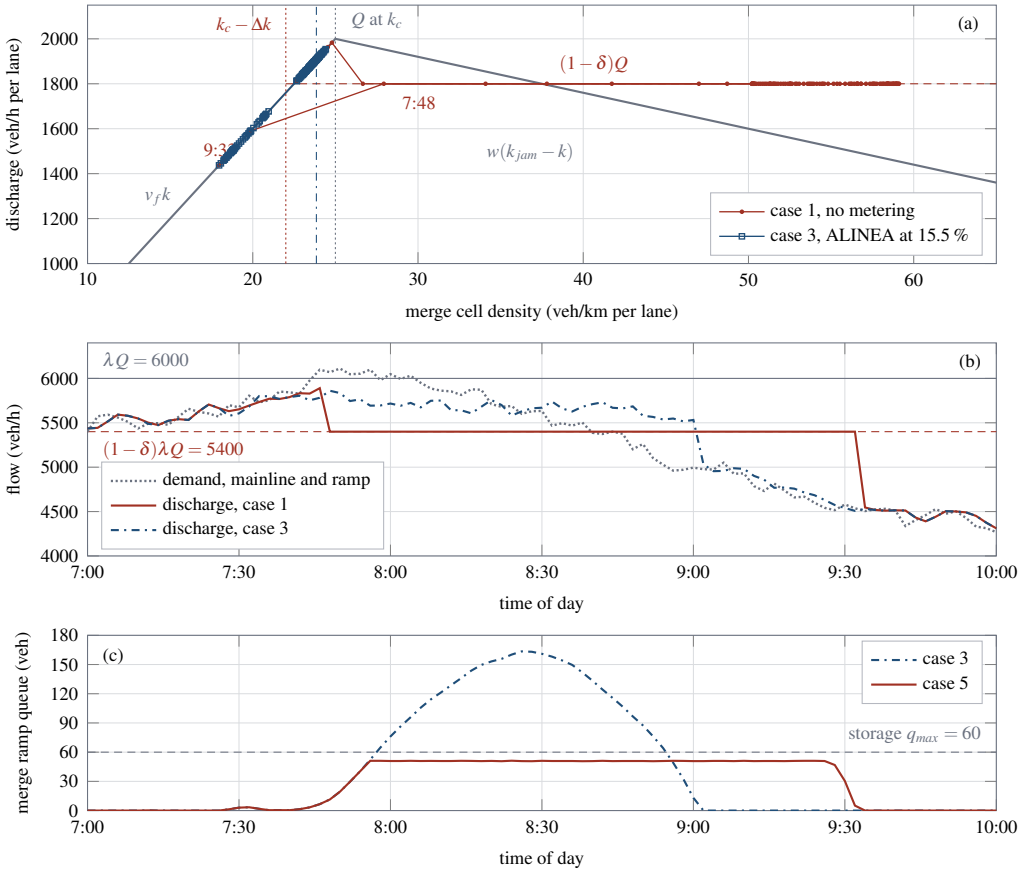


Figure 8.3 Breakdown and its prevention on the average Tuesday, from the engine. (a) The merge cell's density against its discharge per lane once a minute from 7:00 to 10:00, case 1 with filled circles and case 3 with open squares, and the setpoint's 23.85 veh/km as the dash-dotted vertical line. (b) The merge discharge of case 1 (solid) and case 3 (dash-dotted) against the demand of mainline and merge ramp (dotted), every 2 min. (c) The merge ramp queue that ALINEA costs, without the override (case 3, dash-dotted) and with it (case 5, solid).

12.7 to 24.1 min, while 15.5 % never does. The gain matters less than the setpoint, 1526.6 veh·h at $K_R = 35$ and 1537.8 at 140, and PI-ALINEA with $K_P = 50$ gives 1530.1. The demand-capacity rule with its 5800 veh/h estimate gives 1487.9 veh·h with a queue of 125.5, right for as long as its estimate is.

8.5.4 A Setpoint Past the Fold

At 16.5 %, a quarter point past the critical occupancy, the setpoint asks for 25.4 veh/km per lane, a density at which the merge cannot stay. It breaks down at 7:48, the true occupancy reaches 25.4 %, and the integral action drives the release to its 240 veh/h floor from 7:50 to 8:01. The merge regains its capacity at 8:11, once its density is below 22, and the integral action, still aiming past the fold, pushes it over again, four times in all, from 7:48 to 8:11, 8:20 to 8:26, 8:39 to 8:45 and 9:03 to 9:08, 40.4 min in the drop. Meanwhile the meter holds the ramp hard, the queue reaches 341.3, 281.3 of them on the arterial, and the morning costs 1731.5 veh·h. Ten other detector-noise seeds give three or four breakdowns. At 16.2 %, just below the critical occupancy, the merge still breaks down twice, for 19.3 min. An integral loop aimed past the peak of its map has no steady state to settle in, because

each recovery puts the merge back on the free branch, from which the integral steers it over the edge again. Volume 7, Chapter 4 treats a loop whose setpoint sits at the peak itself, where the process gain passes through zero and reverses.

> **TRY IT IN THE APP** **4 · The Setpoint a Little Too High**
>
> Load the case, whose setpoint of 16.5 % lies a quarter point past the critical occupancy, and run it to 10:00 on the Diagram view. Watch the merge point climb toward the setpoint and fall off the fold, four times in all. Then press Reset, set the setpoint to 16.0 % and run again: on this Tuesday the merge never breaks down.

8.5.5 Storage Is the Budget

The queue override keeps every vehicle off the arterial. The queue reaches its limit of 51 at 7:56, the override releases what arrives, and the merge breaks down at 7:57, nine minutes later than with no meter, as Table 8.7 lists and Figure 8.4 shows. It stays in the drop until 9:26, 89.1 min, and the morning costs 1766.0 veh·h, which gives back 236.6 of the 311.3 veh·h that metering saved, 76.0 %. A meter moves delay in time by storing it, and when the storage runs out the breakdown it postponed arrives. Volume 4, Chapter 14 meets the same budget as a vault of reserves with a floor. With the limit at the full storage the breakdown comes at 7:58. On this seed the storage needed lies between 105 and 110 vehicles, since with 105 the merge breaks down at 8:04 and with 110 it never does, and with 120 the queue reaches its new limit of 102 only at 8:07, the true occupancy peaks at 16.23 % against the critical 16.25, and the morning costs 1481.0 veh·h. The threshold moves with the detector noise. Over ten other seeds a storage of 110 still breaks down on five and 120 on one.

Pause. Try Exercise 8.3 before reading on.

8.5.6 An Incident No Meter Can Absorb

At 7:40 a crash at 5.0 km halves cell 10 for 20 min. The merge had been passing 5741 veh/h in the ten minutes before, the blocked cell latches and discharges 2700, and every strategy breaks down between 7:40:30 and 7:40:36, no metering, the fixed plan, ALINEA with and without the override, demand-capacity and speed limits at 15.2 %. From 7:45 to 7:58 the merge discharges 3038 veh/h on average, and at 8:00 the section holds 1470 vehicles whatever the meter has done. No setting of the meter can keep the merge flowing upstream of a cell that passes less than the mainline alone brings, because the queue is made of mainline traffic. What differs is the recovery. ALINEA without the override holds its ramp at the 240 veh/h floor from 7:43 to 9:23 and gives the merge to the mainline queue, so the merge regains its capacity at 9:23 after 102.9 min, while with no meter it is still in the drop at 10:00, 139.5 min on. It does so by holding up to 1056.7 vehicles at 9:26, 996.7 of them on the arterial, and the total time spent hardly moves, 3836.5 veh·h against 3846.5, every strategy within 0.56 % of the others. The speed limits recover first, at 8:29, but fall back into the drop seven more times, 61.9 min in all, and the demand-capacity rule recovers at 9:23 and falls back eight more times. Over ten other noise seeds ALINEA's morning costs 3834.7 to 3852.9 veh·h and the merge recovers at 9:23 on every one. The freeway recovers sooner and the morning does not.

8.5.7 A Frozen Detector

In case 7 the merge detector freezes at 7:20 at a reading of 14.86 %, 0.64 % below the setpoint. The frozen error is added every minute, 44.9 veh/h on top of whatever the stop line counted, so the release rides just above the arrivals, 1100 against 1063 veh/h at 7:49, the queue never holds more than 0.2 vehicles, and the meter has switched itself off with nothing on its card to show it. The merge breaks down at 7:48 and the morning costs 1840.7 veh·h, within 0.01 veh·h of no meter. The watchdog trips at the 7:22 report, the third to repeat the frozen reading, falls back to the fixed plan, which holds this Tuesday, and the morning costs 1434.7 veh·h. The frozen value is one draw of the detector noise. Over ten other seeds it lies between 14.75 and 15.69 %, and the merge breaks down on every one. Only the reading above the setpoint, 15.69 %, lets a queue form, which the override holds at its limit



Figure 8.4 Case 5, The queue override, at 10:00, where the application paused itself at the end of the morning, in the Time–space view with the 4.0 h trend window. The figure is a composite: the density map and the ramp-queue strip above the trends, with three tiles of the Learn panel moved into the empty band at the right. (1) The density map reads the breakdown as a wedge growing back from the merge at 4 km, and (2) its tail is labeled “queue tail 4 km/h upstream”. (3) The strip beneath holds the merge ramp queue against its dashed storage of 60. (4) In the trends a dashed vertical line marks the breakdown, which the run records at 7:57 (1.94 h). Before it the reported merge occupancy scatters, by the run’s record from 14.3 to 16.6 %, about its 15.5 % setpoint from 7:00. By the same record it stays above 30 % in the reports from 8:00 to 9:23 and falls back below the setpoint after the capacity returns at 9:26. (5) The tiles give 89.1 min in the drop, the queue held at its limit of 51, and the total time spent of 1766 veh·h, which the last trend shows with 85 on the ramps.

of 51 while the morning costs 1966.2 veh·h. The watchdog holds the merge on all ten. A freeze fired at 7:40 on the settings of case 5, which are those of case 7 without its scripted freeze, holds a reading of 15.29 %, breaks the merge down at 7:48 without the watchdog, 1841.1 veh·h, and costs 1431.7 with it.

🕮 **Pause.** Try Exercise 8.2 before reading on.

8.5.8 Borrowed Storage

Case 8 opens the upstream ramp with its storage of 150 and turns coordination on. The merge ramp’s counted queue passes 36, 60 % of its storage, at the 7:54 report and the upstream meter starts to store, its queue peaks at 123.7 at 8:21 and is empty again at 9:04, the merge ramp never holds more than its limit of 51, the merge never breaks down, and the morning costs 1459.9 veh·h. Without coordination the upstream meter holds nothing until the jam reaches it, the merge ramp’s override gives way, and the merge breaks down at 7:57 for 88.9 min and 1716.2 veh·h. Without the override the merge ramp

Table 8.8 The Graded Morning in Seven Configurations

Configuration	TTS	Drop	Spill, veh	Met
as shipped, fixed plan	1833.4	103.2	6.9, 37.8	no
ALINEA	1569.9	0	161.9, 0	no
ALINEA, override	1894.7	102.9	0, 0	no
ALINEA, coordination	1619.6	0	36.7, 0	no
ALINEA, override, coordination	1551.9	0	0, 0	yes
demand-capacity, both	1534.9	0	0, 0	yes
speed limits at 15.2 %	1550.9	0	0, 0	yes

TTS in veh·h, drop in minutes, spill as the most vehicles past the storage of the merge and upstream ramps.

spills 13.6 vehicles onto the arterial for 1516.3 veh·h. Coordination adds storage and not capacity. On Thursday both ramps fill and the merge breaks down at 7:53 anyway, for 2262.6 veh·h.

8.5.9 The Mainline as Storage

In case 9 the signal is dark and the speed limits aim at 15.2 %. They first come down at 7:11, stay below 80 km/h from 7:25 until 9:15 and reach 28.2 km/h at 8:23. The merge never breaks down, no vehicle waits on the ramp, and the morning costs 1511.0 veh·h, slightly less than ALINEA's 1529.4, with the delay spread over every mainline driver instead of concentrated on the ramp. That ranking lies within the noise. On seven of ten other detector-noise seeds the limits let the merge break down for 3.7 to 4.5 min and the morning costs up to 1641.5 veh·h, while ALINEA never breaks down and costs 1517.8 to 1559.2. The window is narrow. At 15.5 % the merge breaks down for 5.1 min, at 15.0 % the morning costs 1608.2 veh·h, at 14.7 % 1860.9 and at 14.5 % 2095.6, the last two more than no control. A setpoint below 14.625 % holds a steady merge under the 5400 a broken-down merge passes. The limits also act early, from 6:59 at 14.5 %, 49 min before an unmetered merge breaks down, and by 9:30 that merge has passed 183 fewer vehicles than with no control. The first setting that costs less than no control is 14.8 %, at 1762.9 veh·h, and it does so on all ten other seeds. With ALINEA and the override first and the limits at 16.0 % behind them, the limits act once the override has given up, 11.5 min in the drop instead of 89.1, and 1595.7 veh·h.

8.5.10 The Graded Morning

The challenge runs Tuesday with 1 % more traffic on every stream and both ramps open, and asks by 10:00 for a total time spent below 1650 veh·h, neither ramp queue ever past its storage, and at most 3 min in the drop. Table 8.8 compares the configurations. As shipped, with the fixed plan at the merge ramp and ALINEA at the upstream ramp, the merge breaks down at 7:53 for 103.2 min, both ramps spill, and the morning costs 1833.4 veh·h. Each piece fails alone. ALINEA alone holds the merge but puts up to 161.9 vehicles on the arterial, the override without coordination breaks the merge down at 7:52, and coordination without the override still spills 36.7. ALINEA with the override and coordination meets the target with 1551.9 veh·h, the merge ramp at 51.2 and the upstream ramp at 127.4, and it keeps meeting it with the demand scaled up by as much as 2 %, twice the challenge's, and fails at 2.5 %, while the shipped plan already fails at 0.5 %. Speed limits alone at 15.2 % also meet it, with 1550.9 veh·h and nobody on a ramp, but fail at 2 %. The detector noise separates the two answers. Over ten other noise seeds the three-part answer meets the target on all ten, with the demand scaled up by 1.5 % on nine and by 2 % on only two, while the speed limits alone meet it on one, because the merge breaks down for 3.1 to 4.1 min or the total passes 1650 veh·h on the other nine.

Table 8.9 Frozen-Setting Comparisons, 240 Minutes at Each Case's Seed and 2 s Physics Step

Case	Structure	TTS (veh·h)	Delay (veh·h)	Drop (min)	$q^{obs}_{max}/q^{obs}_{1,max}$ (veh)	IAE (%-min)	TV_r (veh/h)	Sat. (%)	Attempts/rejects	Mean QP (μs)
Case 5	ALINEA	1766.034	365.191	89.07	51.21/0.00	1999.8	6966	38.3	0/0	0
Case 5	Demand-capacity	1764.888	364.046	89.23	51.21/0.00	2004.1	5381	37.9	0/0	0
Case 5	Fixed plan	1434.689	33.847	0.00	51.08/0.00	403.1	1858	4.6	0/0	0
Case 5	MPC On	1845.881	445.039	99.77	51.21/0.00	2165.2	7006	40.8	240/23	584
Case 5	MPC Off	1851.596	450.753	100.40	51.21/0.00	2174.5	8312	37.9	240/17	312
Case 5	MPC, +20%	1843.509	442.667	101.40	51.16/0.00	2167.2	9700	37.5	240/4	362
Case 5	Fixed model rule	1864.211	463.369	98.23	51.21/0.00	2174.9	6925	42.5	0/0	0
Challenge	As shipped	1833.372	468.038	103.20	66.95/187.80	2263.5	1800	0.0	0/0	0
Challenge	ALINEA	1551.877	186.542	0.00	51.21/127.41	349.8	6590	6.7	0/0	0
Challenge	Demand-capacity	1534.883	169.549	0.00	51.09/127.33	355.0	10841	1.7	0/0	0
Challenge	Fixed plan	1833.372	468.038	103.20	51.08/127.61	2258.8	1903	6.7	0/0	0
Challenge	MPC On	1560.782	195.447	0.00	50.99/127.70	349.7	6392	1.3	240/2	486
Challenge	MPC Off	1769.882	404.547	67.97	51.16/127.70	1408.3	11046	23.3	240/4	361
Challenge	MPC, +20%	1950.138	584.803	99.47	51.21/127.58	2071.3	10352	36.3	240/0	453
Challenge	Fixed model rule	1588.991	223.656	0.00	51.21/127.67	327.1	6308	15.0	0/0	0
Two ramps	ALINEA	1459.917	108.100	0.00	51.07/123.65	379.9	6484	1.3	0/0	0
Two ramps	Demand-capacity	1475.952	124.135	0.00	47.68/116.47	380.4	8932	0.0	0/0	0
Two ramps	Fixed plan	1494.885	143.068	0.00	51.08/126.78	380.6	1858	4.6	0/0	0
Two ramps	MPC On	1469.389	117.572	0.00	29.47/120.62	376.4	6996	0.0	240/0	467
Two ramps	MPC Off	1446.587	94.771	0.00	10.92/116.26	381.7	6830	0.0	240/0	378
Two ramps	MPC, +20%	1834.909	483.092	90.23	51.21/127.61	1912.9	10455	33.8	240/0	469
Two ramps	Fixed model rule	1545.069	193.253	0.00	51.21/127.46	340.0	6780	7.5	0/0	0

The two-ramp row is an extra run that repeats case 8's physical conditions. TV and saturation are the shared merge-command metrics. IAE retains the ALINEA reference, not the MPC ceiling. Delay subtracts an arrival-matched free-travel reference from full-road TTS. Physical storage exceedance is zero except the shipped challenge (17.90/13.60 min, merge/upstream). Soft 85% bands can still be exceeded. Timing is host dependent.

8.5.11 Measured Predictive Comparison

Case 10 repeats case 8 under the predictive controller, then compares ALINEA after Reset. Runs share 240 min, seeds, 2 % noise, arrivals and protections, including the challenge's suggested controls. MPC fixes margin 0.94, move price 1 and zero forecast error throughout. One predetermined classical rule rounds $0.94(16.25)$ to 15.3 % at both meters. Nominal gain 0.00270833 percentage point/(veh/h) and 5 min correction give $K_R = 73.846$, rounded to 75 at both meters, with $K_P = 0$. No case-specific tuning follows.

The predictive controller costs 1469.389 veh·h, against 1459.917 veh·h for ALINEA and 1475.952 veh·h for the demand-capacity rule. With the forecast Off it costs 1446.587 veh·h, so the forecast adds 22.801 veh·h. A +20 % forecast error gives 1834.909 veh·h and 90.23 min drop. The fixed classical rule costs 1545.069 veh·h, also losing to shipped ALINEA. Figure 8.5 compares trajectories.

Case 5 MPC loses, 1845.881 veh·h against ALINEA's 1766.034 veh·h, with 99.77/89.07 min drop. The local cost omits the latch and protective command changes. Challenge MPC passes at 1560.782 veh·h, zero drop and 50.99/127.70-vehicle queues, losing to ALINEA/demand-capacity at 1551.877/1534.883 veh·h. Its soft storage band is exceeded by 0.198 vehicle for 2.43 min, despite passing the physical-storage grade.

Fault replicas add 300 veh/h ramp demand at minutes 90–120, persistent 15 % rain at minute 90, or detector freeze at minute 80, with watchdog enabled for every law. Event MPC/ALINEA/demand-capacity cost 2086.678/2132.295/2081.616 veh·h, and rain costs 4451.582/4467.444/4461.352 veh·h, all failing the challenge criteria. Frozen MPC costs 1495.001 veh·h, with 82 attempts, zero rejects and 158 min watchdog fallback. All 47 rows have paired half-step runs. Most critically, noisy challenge MPC-Off changes from 1769.882 veh·h, 67.97 min drop and fail at 2 s to 1493.373 veh·h, zero drop and pass at 1 s. Small report/control differences precede a discontinuous latch crossing. Neither this trace nor separate changed-input diagnoses isolate a sole cause or prove convergence. Removing

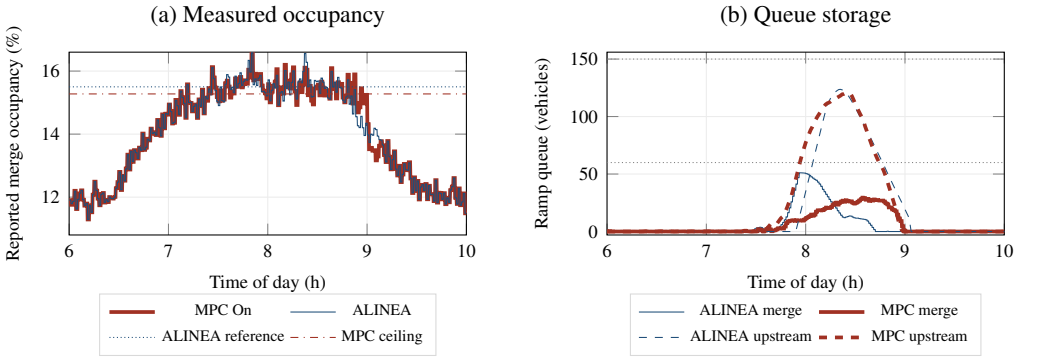


Figure 8.5 Two-ramp comparison from the original noisy runs. (a) Reported merge occupancy under ALINEA (thin line) and predictive metering with forecast On (thick line). The ALINEA reference is 15.5 % (dotted), while the predictive ceiling is 15.275 % (dash-dotted). (b) Counted merge queues (solid) and upstream queues (dashed) under the same laws, thin for ALINEA and thick for the predictive controller. Neither run enters the capacity drop or exceeds physical storage, yet their full-road TTS differs by 9.472 veh·h.

detector noise alone gives 1479.561/1479.596 veh·h at the two steps. Additionally suppressing demand fluctuation gives 1505.590/1505.584 veh·h. Both changed-input pairs pass without replacing the noisy result.

At one simulated minute per second, 20 warm Chrome decisions average 1730 μs for the complete controller (0.173% of one core). Solver/model/observer/decision means are 740/620/295/1395 μs. At 300 min/s, guarded and unguarded playback reach 298.24 and 295.37 min/s at 59.98 and 60.07 frames/s, with zero budget stops. These samples establish no pacing benefit. The 25 ms frame-work guard retains reached decisions and physics steps but cannot bound a single slow solve. Forced guard stops retain 240 solves and exact non-timing records. Same-Chrome speed/guard comparisons are byte-identical, while Node/Chrome differ in exported TTS by up to 10^{-5} veh·h.

8.6 Technology in Practice

The meters of Section 8.4 read a minute report of the merge occupancy and the counts at each ramp, and act through the release of a ramp signal. Figure 8.6 draws the merge loop with the devices a real ramp carries in place of the schematic's symbols.

What Measures It

The occupancy comes from inductive loops cut into each lane of the merge, the detectors with which ramp metering is mainly carried out [22]. Equation (8.6) stands for that loop: the period mean of the cell's density times a vehicle length plus a detection zone, with 2 % noise, and a freeze that repeats the last reading. Loops cannot measure a ramp queue directly, only estimate it [22], and the exact counters of the model are one of the assumptions of Section 8.3.7.

> **TECHNOLOGY** **Sensor: The Inductive Loop Detector**
>
> A loop is one or more turns of wire in a slot sawn into the pavement, wired to an electronics unit in the roadside cabinet that excites it at 10 to 200 kHz. A vehicle over the loop induces eddy currents that lower its inductance, and the unit senses a fall in frequency and signals a vehicle present. Loops are the common standard for accurate occupancy, while density and speed must be inferred [23]. Pavement wear, poor installation and weather degrade them. Some units flag a step of 25 % in inductance as a loop failure, and an open loop that calls constantly can be taken off line [23].

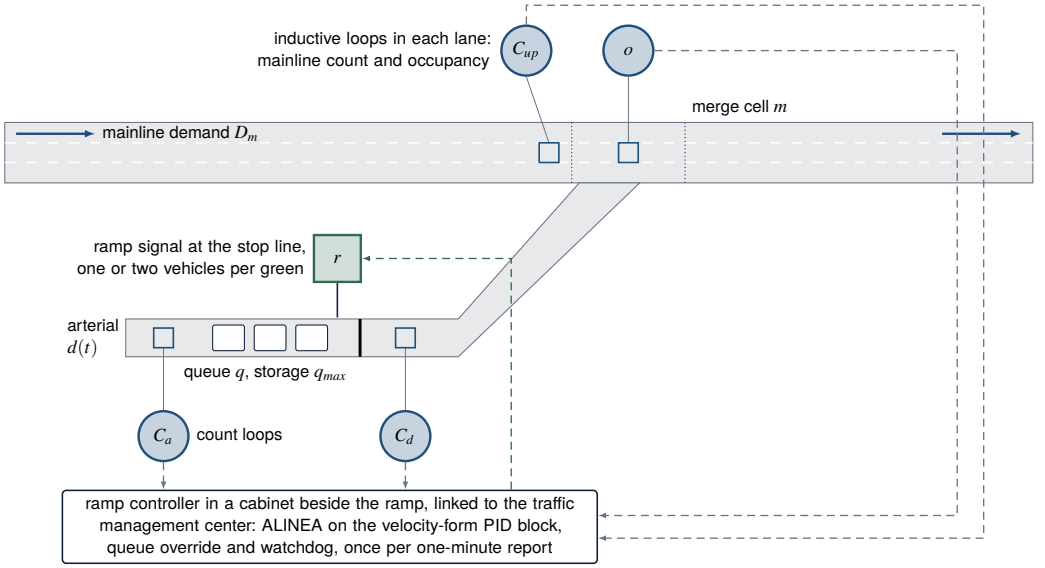


Figure 8.6 The merge meter with the devices of a real ramp in place of the schematic's symbols. Loops cut into the pavement give the merge occupancy o and the mainline count C_{up}, count loops on the ramp give C_a and C_d, and a controller in the roadside cabinet sets the release r of the ramp signal. Notice that the meter never measures the queue it is holding: it counts what enters and what leaves the ramp.

What Moves It

The final element is the ramp signal. The model's release is a continuous flow clipped to [240, 1800] veh/h, the floor being the lowest practical rate of a signal and the ceiling that of a two-lane meter. A real meter releases whole vehicles, one or two per green.

> **TECHNOLOGY** **Final Element: The Ramp Signal**
>
> The meter is a traffic signal at a stop line. With one vehicle per green, a demand detector before the stop line calls the green and a passage detector just beyond it ends it, and a single lane passes about 900 vehicles an hour. A minimum cycle of 4 s, 2.5 s of red and 1.5 s of green, gives 1000 an hour, and the lowest practical rate is 240 an hour, a 15 s cycle. For a ramp volume of 1600 to 1800 an hour the guidance is two metered lanes, each releasing two vehicles per green [6]. The release is therefore a train of single vehicles, and the cycle sets its floor and its ceiling.

Where the Controller Runs

The ramp controller sits in a cabinet beside the ramp with its load switches and loop amplifiers. It is a microprocessor-based traffic signal controller that also acts as a data station, and it reaches the traffic management center by telephone line, fiber, microwave or radio [6]. Field algorithms work on smoothed reports: running averages over 5 min, a one-minute occupancy per lane in one agency's local algorithm, and filtered 30 s data in a zone algorithm [6, 23]. The meters of this chapter act once per one-minute report. The detector handbook describes ALINEA as a recursive law that holds the mainline occupancy at a preselected value, with a rate that violates a limit adjusted to a feasible value [23]. In this chapter that recursion is the velocity-form block of Equation (8.10), and the adjustment is its clip with back-calculation.

> **IN PRACTICE** **A Fixed Plan Behind the Feedback**
>
> Traffic-responsive metering often does five to ten percent better than a pre-timed plan, yet the low cost of pre-timed systems makes them attractive backups when the primary approach fails. Queue storage is sized by rules of thumb, one agency's being ten percent of the peak-hour volume before metering. When the queue outgrows it, the rate is raised, which costs the freeway but protects the arterial, and the queue must not be dumped onto the freeway all at once. A queue detector near the surface street raises the rate, and a second one further up raises it sharply [6]. The watchdog of case 7 and the override of case 5 are these two practices.

What Is Modernizing It

Three developments touch this loop. Over-roadway sensors, video image processors and frequency-modulated microwave radar, report lane occupancy without a cut in the pavement, and radar of that kind also sees stopped vehicles [23]. Messages broadcast by connected vehicles add data from which a robust filter estimates the ramp queue, with a bound on its long-term error even when the share of connected vehicles is unknown [22]. Learning-based meters trained on loop data have been tested against false data injected into their inputs, because a richer input is also a larger surface for attack [24]. The frozen detector of case 7 is the plain form of that risk.

8.7 Limitations

The cases inherit the assumptions of Section 8.3.7. With a first-order model and a deterministic latch the merge breaks down at one density and never by chance, every wave inside a queue moves at one speed and no stop-and-go wave forms, so the minutes in the drop belong to the declared δ and Δk. Volume 7, Chapter 14 shows how a line of cars amplifies a brake tap into such a wave. With declared demand no driver avoids a long ramp queue, and with point queues and exact counts the override always sees the true queue. The total time spent weighs every vehicle-hour the same, so fuel, emissions, reliability and fairness between ramp and mainline drivers are outside the score. Every capacity, storage, threshold and target is a teaching value chosen for the lesson, and no number here is a measurement of a real road.

> **OPEN CHALLENGE** **Where Is the Fold Today?**
>
> A local meter holds its setpoint just below the critical occupancy, and that number is hard to know. A study of ramp metering calls the critical density most difficult to estimate correctly online, and the calibration of the fundamental diagram far from obvious [25]. In this chapter a setpoint a quarter point too high breaks the merge down four times. Progress would be a meter that finds its margin from its own data and keeps it as weather, traffic mix and the morning move the fold.

The predictive approximation omits the capacity-drop latch, downstream receiving restrictions, distributed transport, origin queue, speed-limit dynamics and substep empty-queue release cap. Reconstruction using realized counts does not establish multi-step demand accuracy. The 600 veh/h floor experiment exposes request/flow infeasibility rather than permitting negative inventory. The losses in case 5 and the challenge, the wrong-forecast failure and the noisy refinement reversal prevent a general advantage or convergence claim.

> **OPEN CHALLENGE** **A Fair Share of the Wait**
>
> Conventional metering often spreads delay very unevenly across on-ramps, which undermines its acceptance, and strategies seen as unfair meet signal violation and public opposition [26]. The Twin Cities shutdown of 2000 began with such opposition [6]. The score of this chapter weighs every vehicle-hour the same. Progress would be a metering law, simple enough for existing controllers, that keeps the merge flowing while it bounds the wait at the worst-served ramp.

8.8 Conclusion

A freeway merge is a plant whose steady-state map folds over, and the simulator of this chapter makes the fold, the drop below it and the meter that avoids it measurable on a morning that replays exactly. With the signal dark, 17 min of overload hold the merge in the capacity drop for 104.9 min and the morning costs 1840.7 veh·h. ALINEA at 15.5 % keeps the merge out of the drop for 1529.4 veh·h, while its ramp queue reaches 163.9 vehicles against a storage of 60. A setpoint past the fold breaks the merge down four times, the override that protects the arterial hands back 76.0 % of the saving, coordination stores the difference upstream for 1459.9 veh·h, a frozen detector quietly switches the meter off, and an incident that halves the road defeats every strategy, whose totals agree within 0.56 %. The graded morning improves from 1833.4 veh·h with a breakdown and two spilling ramps to 1551.9 with neither. The model is first order with a deterministic drop and declared demand. A fuller model would add a probabilistic breakdown, demand that diverts from long queues, a second-order model with stop-and-go waves, a chain of merges metered together, and setpoints chosen against the spread of the detector noise. The simulator treats a freeway as a feedback problem whose setpoint is a margin and whose actuator spends a budget.

The predictive comparison costs 1469.389 veh·h against ALINEA's 1459.917 veh·h, with 1446.587 veh·h when the forecast is withheld, demonstrating that a nominal constrained plan need not improve the scored morning.

KEY POINTS

- A merge passes the most traffic just before it breaks down, and once it breaks down its queue discharges a tenth less until it is gone.
- ALINEA is integral control on the merge occupancy: in steady flow it settles at its setpoint whatever the demand, without knowing the merge's capacity.
- A setpoint past the fold has no steady state, so the integral action breaks the merge down again after each recovery.
- A meter moves delay in time by storing it, and when the storage runs out the breakdown it postponed arrives.

Where else this applies. Volume 4, Chapter 2 holds a compressor short of its own fold with a protective override.

References

[1] M. J. Lighthill and G. B. Whitham, "On kinematic waves II. A theory of traffic flow on long crowded roads," *Proceedings of the Royal Society of London. Series A. Mathematical and Physical Sciences*, vol. 229, no. 1178, pp. 317–345, 1955.

[2] P. I. Richards, "Shock waves on the highway," *Operations Research*, vol. 4, no. 1, pp. 42–51, 1956.

[3] C. F. Daganzo, "The cell transmission model: A dynamic representation of highway traffic consistent with the hydrodynamic theory," *Transportation Research Part B: Methodological*, vol. 28, no. 4, pp. 269–287, 1994.

[4] M. Papageorgiou and A. Kotsialos, "Freeway ramp metering: an overview," *IEEE Transactions on Intelligent Transportation Systems*, vol. 3, no. 4, pp. 271–281, 2002.

[5] M. Papageorgiou, H. Hadj-Salem, and F. Middelham, "ALINEA local ramp metering: Summary of field results," *Transportation Research Record: Journal of the Transportation Research Board*, vol. 1603, no. 1, pp. 90–98, 1997.

[6] L. Jacobson, J. Stribiak, L. Nelson, and D. Sallman, "Ramp management and control handbook," Federal Highway Administration, U.S. Department of Transportation, Washington, DC, Tech. Rep. FHWA-HOP-06-001, Jan. 2006.

[7] D. C. Gazis and R. S. Foote, "Surveillance and control of tunnel traffic by an on-line digital computer," *Transportation Science*, vol. 3, no. 3, pp. 255–275, 1969.

[8] V. Jacobson, "Congestion avoidance and control," *ACM SIGCOMM Computer Communication Review*, vol. 18, no. 4, pp. 314–329, 1988.

[9] E. M. Greitzer, "Surge and rotating stall in axial flow compressors—part I: Theoretical compression system model," *Journal of Engineering for Power*, vol. 98, no. 2, pp. 190–198, 1976.

[10] D. Levinson and L. Zhang, “Ramp meters on trial: Evidence from the Twin Cities metering holiday,” *Transportation Research Part A: Policy and Practice*, vol. 40, no. 10, pp. 810–828, 2006.

[11] A. D. May, “Experimentation with manual and automatic ramp control,” in *Traffic Control: Theory and Instrumentation*, T. R. Horton, Ed. Springer US, 1965, pp. 157–208.

[12] L. Isaksen and H. J. Payne, “Freeway traffic surveillance and control,” *Proceedings of the IEEE*, vol. 61, no. 5, pp. 526–536, 1973.

[13] Y. Wang, E. B. Kosmatopoulos, M. Papageorgiou, and I. Papamichail, “Local ramp metering in the presence of a distant downstream bottleneck: Theoretical analysis and simulation study,” *IEEE Transactions on Intelligent Transportation Systems*, vol. 15, no. 5, pp. 2024–2039, 2014.

[14] M. J. Cassidy and J. Rudjanakanoknad, “Increasing the capacity of an isolated merge by metering its on-ramp,” *Transportation Research Part B: Methodological*, vol. 39, no. 10, pp. 896–913, 2005.

[15] A. Srivastava and N. Geroliminis, “Empirical observations of capacity drop in freeway merges with ramp control and integration in a first-order model,” *Transportation Research Part C: Emerging Technologies*, vol. 30, pp. 161–177, 2013.

[16] G. Gomes and R. Horowitz, “Optimal freeway ramp metering using the asymmetric cell transmission model,” *Transportation Research Part C: Emerging Technologies*, vol. 14, no. 4, pp. 244–262, 2006.

[17] I. Papamichail and M. Papageorgiou, “Traffic-responsive linked ramp-metering control,” *IEEE Transactions on Intelligent Transportation Systems*, vol. 9, no. 1, pp. 111–121, 2008.

[18] R. C. Carlson, I. Papamichail, M. Papageorgiou, and A. Messmer, “Optimal motorway traffic flow control involving variable speed limits and ramp metering,” *Transportation Science*, vol. 44, no. 2, pp. 238–253, 2010.

[19] A. Hegyi, B. De Schutter, and H. Hellendoorn, “Model predictive control for optimal coordination of ramp metering and variable speed limits,” *Transportation Research Part C: Emerging Technologies*, vol. 13, no. 3, pp. 185–209, 2005.

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

[21] D. Mayne, J. Rawlings, C. Rao, and P. Scokaert, “Constrained model predictive control: Stability and optimality,” *Automatica*, vol. 36, no. 6, pp. 789–814, 2000.

[22] Y. Tang, K. Ozbay, and L. Jin, “Robust queue length estimation for ramp metering in a connected vehicle environment,” arXiv preprint arXiv:2305.17921, 2023.

[23] L. A. Klein, M. K. Mills, and D. R. P. Gibson, “Traffic detector handbook: Third edition, volume I,” Federal Highway Administration, U.S. Department of Transportation, McLean, VA, Tech. Rep. FHWA-HRT-06-108, Oct. 2006.

[24] D. Liu, L. Liu, and L. D. Han, “Analyzing robustness of the deep reinforcement learning algorithm in ramp metering applications considering false data injection attack and defense,” arXiv preprint arXiv:2301.12036, 2023.

[25] C. Join, H. Abouaïssa, and M. Fliess, “Ramp metering: modeling, simulations and control issues,” arXiv preprint arXiv:2111.06610, 2021.

[26] K. Riehl, O. Alami Badissi, A. Kouvelas, and M. A. Makridis, “C-EQ-ALINEA: Distributed, coordinated, and equitable ramp metering strategy for sustainable freeway operations,” arXiv preprint arXiv:2601.06311, 2026.

8.9 Exercises

8.1 ○ Concept check · 2 min

With the signal dark, the demand at the merge exceeds its capacity of 6000 veh/h for 17 min, yet the merge stays in the capacity drop for 104.9 min. Which explanation is right?

(a) The detectors report once a minute, so the breakdown is seen too late for the queue to clear

(b) The congestion wave moves at 16 km/h, so the tail of the queue needs that long to reach the merge

(c) The broken-down merge discharges 5400 veh/h, so its queue grows until the demand is below that and drains only afterward

(d) The demand stays above the capacity of 6000 veh/h until 9:33, the minute at which the merge recovers

8.2 Concept check · 2 min

In scenario 7 the merge detector freezes at 7:20 on a reading of 14.86 %, below the 15.5 % setpoint, and the watchdog is off. What does ALINEA do through the rest of the peak?

(a) It holds the release at its 7:20 value, because a constant reading gives a constant output

(b) It ramps the release up to the 1800 veh/h ceiling within minutes and winds up there

(c) It drives the release down to the 240 veh/h floor, and the ramp queue spills onto the arterial

(d) It adds one fixed increment to the counted release each minute, so the release rides just above the arrivals

8.3 Concept check · 1 min

Scenario 5 adds the queue override, with its limit at 51 vehicles, to the ALINEA meter of scenario 3. What does the override do on the default morning?

(a) It keeps every vehicle off the arterial, and the merge breaks down at 7:57 and loses most of what metering saved

(b) It keeps every vehicle off the arterial at no cost, because ALINEA still holds the occupancy at its setpoint

(c) It lets the queue pass the storage for a few minutes, and in exchange the merge never breaks down

(d) It delays the breakdown until the peak has passed, so the total time spent equals that of scenario 3

8.4 Theory and calculation · 5 min

The merge has three lanes with a capacity of 2000 veh/h per lane, a free speed of 80 km/h and a congestion wave speed of 16 km/h. The effective vehicle length at the detector is 6.5 m.

(a) Find the critical density and the jam density per lane. *Checkpoint:* 25 and 150 veh/km.

(b) Find the critical occupancy, and the free-flow discharge of the merge that one percentage point of occupancy stands for. *Checkpoint:* 16.25 % and 369 veh/h.

(c) Find the density and the discharge at the default setpoint of 15.5 %, and the margin between that discharge and the capacity. *Checkpoint:* 23.85 veh/km per lane, 5723 veh/h and 277 veh/h.

8.5 Theory and calculation · 8 min

A queue waits on the ramp and the merge is in free flow. Treat the merge as a static plant: the occupancy of report k, in percent, is $o(k) = \big(q_{up} + r(k-1)\big)/g$ with $g = 369\,\text{veh/h}$ per point and a constant mainline flow q_{up}. ALINEA follows Equation (8.9) with $K_P = 0$ and $T = 1\,\text{min}$.

(a) Show that the error $e(k) = \hat{o} - o(k)$ shrinks by a fixed factor per report. Find the factor and the time constant at the default $K_R = 70$. *Checkpoint:* 0.810 per report and 4.8 min.

(b) Find the gain that removes the error in one report and the gain above which the error grows. Find the factor at the top of the slider range, $K_R = 200$. *Checkpoint:* 369 and 738 veh/h per percent per minute, and 0.458.

(c) Give the proportional gain and the integral time of the PID of Equation (8.10) that reproduce ALINEA at $K_R = 70$, and PI-ALINEA at $K_R = 70$ with $K_P = 50$. *Checkpoint:* $K_c = 70$ with 1 min, and $K_c = 120$ with 1.71 min.

8.6 Theory and calculation · 6 min

The merge ramp stores 60 vehicles and its override limit is 85 % of the storage. Reports arrive every minute, and the entrance counter reads 1050 veh/h of arrivals.

(a) Find the least release that Equation (8.11) demands at a counted queue of 45, 51 and 55 vehicles. *Checkpoint:* 690, 1050 and 1290 veh/h.

(b) ALINEA holds the release at 655 veh/h, its lowest value in scenario 3. Find the time for an empty ramp to reach the override limit. *Checkpoint:* 7.7 min.

(c) The upstream ramp stores 150 vehicles. Find the counted merge-ramp queues at which coordination starts and ends, the target the upstream ramp fills to, and the time the target takes to fall to zero afterward. *Checkpoint:* 36 and 18 vehicles, 127.5 vehicles and 42.5 min.

8.7 App challenge · 8 min

Load scenario 1, No metering: the breakdown and the shockwave. Use the default seed, and run each morning to 10:00, where the app pauses.

(a) Change no setting. Read the Capacity drop and Total time spent tiles. *Checkpoint:* 104.9 min and 1841 veh·h.

(b) Press Reset, open Show Model, set Recovery hysteresis below the critical density to 0 and run again. Read the same two tiles. *Checkpoint:* 104.5 min and 1841 veh·h.

(c) Reload the scenario, set Capacity drop once a queue forms to 0 % in Show Model and run again. Read the Total time spent tile, and say which of the two parameters holds the merge in the drop. *Checkpoint:* 1405 veh·h, within 1.

8.8 App challenge · 8 min

Load scenario 3, ALINEA just under the critical occupancy. Use the default seed and run each morning to 10:00. Make each change after Reset and before Run. Read the Total time spent, Capacity drop and Ramp queue, worst tiles.

(a) Run the scenario as loaded, then with the Merge occupancy setpoint at 15.0 %, then at 16.0 %. *Checkpoint:* 1529, 1836 and 1426 veh·h, with queues of 164, 365 and 50 vehicles and no breakdown.

(b) Keep 16.0 %, switch to Advanced and set Demand scaling to 2 % under the disturbances. Run again. *Checkpoint:* 1773 veh·h and 24.5 min in the drop.

(c) Keep the demand and return the setpoint to 15.5 %. Run again, and choose a setpoint for a site whose demand is known only to within 2 %. *Checkpoint:* 1725 veh·h, no breakdown, 284 vehicles.

8.9 App challenge · 8 min

Load scenario 5, The queue override. Use the default seed and run each morning to 10:00. Make each change after Reset and before Run.

(a) Change no setting. Read the Total time spent, Capacity drop and Ramp queue, worst tiles. *Checkpoint:* 1766 veh·h, 89.1 min and 51 of 60.

(b) Open Show Model and set Merge ramp storage to 105 vehicles, run, then set it to 110 vehicles and run again. Read the same tiles. *Checkpoint:* 1708 veh·h with 74.6 min, then 1474 veh·h with no breakdown.

(c) Scenario 3 needed a queue of 164 vehicles to hold the merge without the override. Explain why a storage of 110 is enough here, and say whether a designer should rely on it. *Checkpoint:* The true occupancy peaks at 16.23 %.

8.10 Application area · 3 min

At 7:40 a crash one kilometer past the merge halves the capacity of one cell for 20 min. The mainline alone brings about 4900 veh/h.

(a) Find the capacity of the blocked cell, and its discharge once a queue stands behind it. *Checkpoint:* 3000 and 2700 veh/h.

(b) Estimate the vehicles that must queue during the 20 min if the ramp is closed completely, and say what metering can and cannot do here. *Checkpoint:* About 730 vehicles.

8.11 Application area · 4 min

Table 8.8 shows two configurations that meet the graded morning at nearly the same total time spent: ALINEA with the override and coordination at 1551.9 veh·h, and speed limits alone at 1550.9 veh·h.

(a) Say who bears the delay under each configuration. *Checkpoint:* Ramp drivers in one, mainline drivers in the other.

(b) Choose one to commission, using the chapter's results at higher demand and at other detector-noise seeds. *Checkpoint:* Ten of ten seeds against one of ten.

(c) Name two things that matter to a road agency and that the total time spent leaves out. *Checkpoint:* The Limitations section names four.

8.12 Review the agent · 5 min

An AI assistant is asked for the best setpoint of the ALINEA meter and replies as follows. "The merge carries the most at its critical occupancy of 16.25 %. ALINEA is an integral controller, and integral action removes the error, so the occupancy will sit at the setpoint it is given. Set the setpoint to 16.2 %. The merge will then pass 5982 veh/h instead of 5723 veh/h, the queue will shrink, and the margin of 0.05 points keeps the merge on the free branch." Find the flaw in the reply and say how it could be checked.