INTERACTIVE CONTROL ENGINEERING
Principles and Applications

Volume 20 Motion Control, Robotics, and Autonomous Systems

CHAPTER 9

Warehouse Robot Simulator for Stopping Distance, Blind Corners and Localization Error

Pages 211–234

Edited by John D. Hedengren

How to cite this chapter. "Warehouse Robot Simulator for Stopping Distance, Blind Corners and Localization Error," in *Motion Control, Robotics, and Autonomous Systems*, John D. Hedengren, Ed., vol. 20 of *Interactive Control Engineering: Principles and Applications*. Advanced Process Solutions, LLC, 2026, ch. 9, pp. 211–234.

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

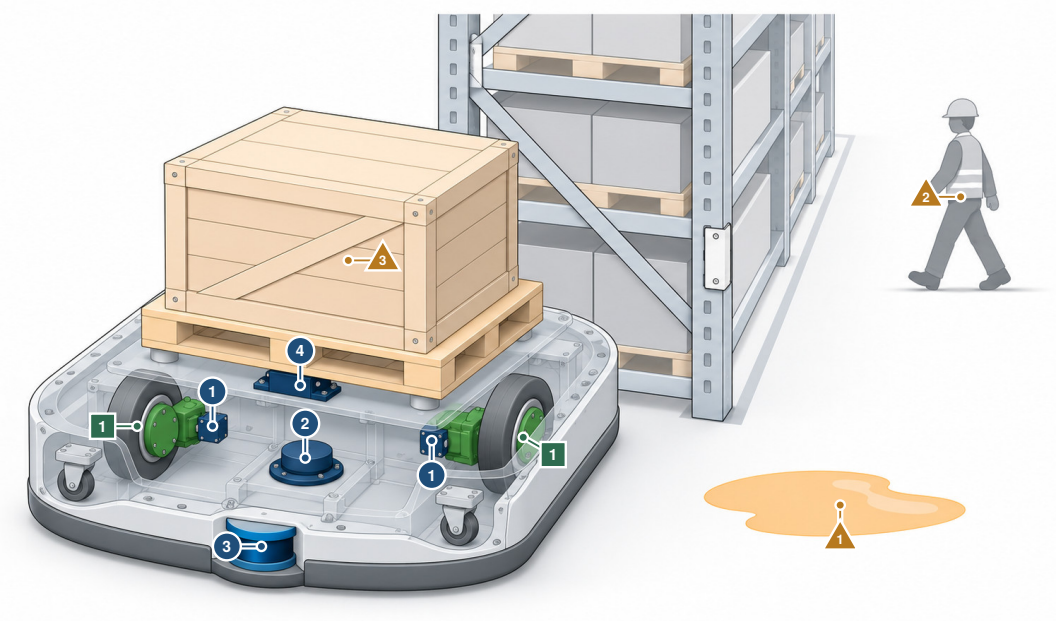
The app for this chapter
https://apmonitor.com/apps/warehouse-robot/

The volume online: contents, other chapters and citation records
https://apmonitor.com/apps/books/v20-motion-robotics/#warehouse-robot

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

CHAPTER 9

Warehouse Robot Simulator for Stopping Distance, Blind Corners and Localization Error



The controller moves 1 left and right wheel motors, one torque each

The controller reads 1 wheel encoders, which count the turns 2 yaw gyro 3 laser scanner: obstacles and reflectors 4 load cell under the pallet

What disturbs it 1 wet patch that lengthens the stop 2 worker hidden by the rack end 3 pallet mass

Notice that the rack end hides the cross-aisle from the low scanner, so the robot has to slow for space it cannot see, while the load and the wet patch both lengthen its stop. Illustration, not to scale.

IN PLAIN TERMS A Robot That Has to Stop in Time

The main idea. A robot carries pallets down aisles where people also walk. It drives only as fast as lets it stop inside the space it can see, using its own guesses of where it is and how hard it can brake.

The challenge. Each guess can be wrong. A heavy load or a wet floor makes the stop longer, a tire a little too big makes it lose its place, and the end of a shelf can hide a worker who steps out.

Why it matters. A loaded robot that trusts one wrong guess can hit a person, so each guess must be checked and corrected.

THREE WAYS TO READ THIS CHAPTER

Explore	11 min	the first page, the history segment of Section 9.2, the fact on the test pieces (Section 9.3), the fact on blind corners, the try-it segments of scenarios 2 and 5 (Section 9.5), and the key points
Study	50 min	the whole chapter with its three pause points and the twelve exercises at its end
Practice	15 min	the first page, the measured cases (Section 9.5), Technology in Practice (Section 9.6), Limitations (Section 9.7) and the open challenge

The control question A warehouse robot chooses its speed from estimates of its position, its braking and the space it can see. What does each estimate cost when it is wrong, and which corrections let a loaded robot finish a shift among people without a collision or a clearance violation?

Learning objectives

- Compute the braking capability of the robot from its motor and floor limits, and the stopping band and envelope speed that follow.
- Explain why the covariance of dead reckoning does not cover a wheel-radius error, and how reflector updates remove it.
- Tell an avoidable collision from an unavoidable one by the entry distance, the true stopping distance and what the scanner could see.
- Relate the gap a robot leaves behind an obstacle to the lag of its speed loop, and check the relation against the speed gain.
- Configure the robot to pass the graded shift, and show that each of the four corrections is needed.

Prerequisites Sections 1.4.2 (differential drive), 1.4.5 (path following), 1.4.6 (dead reckoning and the Kalman filter), 1.4.9 (stopping distance and the speed envelope) and 1.6.2 (the wheel loops).

The app https://apmonitor.com/apps/warehouse-robot/

See also Volume 5, Chapter 13 for an extended Kalman filter that applies late fixes behind a chi-square gate while its drift stays invisible aboard. Chapter 15 for a speed law beside a person who walks in, and Volume 7, Chapter 4 for braking at the friction peak.

9.1 Introduction

Most feedback loops are judged by how closely a measurement follows its setpoint, and on a process unit that judgment is usually sound because the measurement is the controlled variable. On an autonomous pallet robot the same judgment misleads, because the robot never measures the three quantities its safety rests on. Its position is an estimate assembled from wheel encoders, a gyro and occasional sightings of reflectors. Its braking capability is a belief that holds only for the pallet and the floor it was computed for. Its free space is the part of the aisle its scanner happens to see, and a rack end hides the rest. In the shipped scenarios a collision or a missed delivery is never a tuning failure. It is an assumption that was wrong at the one moment it mattered, and the tracking error recorded until then gives no warning of it.

Automated guided vehicles have moved loads through factories and distribution centers for decades, and fleets of mobile robots now carry shelves, totes and pallets through warehouses in which people also work [1–3]. The safety requirements of such driverless industrial trucks, and the tests that verify them, are standardized [4]. Behind the requirement to stop in time sits one control problem, a speed limit computed from the robot's own estimates of its state and of its capabilities.

The simulator of this chapter addresses the belief that a controller acts on the plant. It acts on an estimate of the plant and on a model of the boundary it must not cross, so a loop with good tracking can drive a loaded robot into a cart it saw in good time. A block diagram hides this, because it draws a single arrow from the plant to the controller. The simulator keeps three things apart instead, the true robot, what its sensors report and what its controller believes, draws all three on one map and scores every run against the truth.

The class of problem is general. Limiting a velocity so that the plant can always stop inside the space it has verified is the basis of collision avoidance with admissible velocities [5], of the analysis of inevitable collision states [6] and of its extension to sensors with a limited field of view [7], and it is written as a control barrier function in safety filters whose worked examples include an adaptive cruise controller held at a safe following distance [8]. A process unit driven toward a constraint that only an inferential measurement reports, a dynamically positioned vessel holding station on an observer, and a lane-keeping loop tracking a road it can only estimate have the same shape. In each of

them the controller is only as safe as its model of its own limits and its estimate of its own state [9], and on a smooth floor the largest errors in that estimate are systematic rather than random [10].

However, estimation, planning and speed supervision are usually taught as separate subjects, each with a clean example of its own, while the failures that matter sit at the interfaces between them. This chapter tests the assumptions of a warehouse robot one at a time, on a simulator whose controller reads only measurements, and measures what each costs when it is wrong. The differential-drive model, the extended Kalman filter, pure pursuit, A* and the stopping-distance rule are all established methods. The simulator joins them in a 13-state robot with slip-based traction, a sensor layer that is the only path from the simulated truth to the controller, and a referee that scores collisions, clearances and deliveries from the truth and classifies collisions using a modeled stopping-distance test. It carries 35 automatic checks, one of which confirms that moving the true pose without changing the measurements leaves the torque commands unchanged to the last bit. Eight cases each break one assumption, and a graded shift makes each of four fixes necessary.

Section 9.2 traces where the problem came from. Section 9.3 sets out the robot, its floor, the warehouse and the sensors, and Section 9.4 the estimators, the stopping envelope, the route follower, the wheel loops and the referee, in enough detail to rebuild both. Section 9.5 reports what each case measures, from the price of caution on a clear route to the graded shift. Section 9.6 names the devices of real robots, Section 9.7 states what the model leaves out, and Section 9.8 draws the lessons together.

9.2 From Guide Wires to Reflectors and Robot Fleets

Driverless transport entered factories in the 1950s as towing vehicles that followed a guide wire, and for decades afterward the route of an automated guided vehicle was built into the floor as a buried wire, a painted line or a magnetic strip [11]. Guidance was a sensing problem solved by construction. What remained were the problems of a fleet on a fixed network, and a survey of that literature organizes it around guide-path layout, the number of vehicles, dispatching, routing, the positioning of idle vehicles and battery management [1]. These are questions of traffic rather than of the motion of one vehicle.

Free navigation replaced the wire with a map. Matching range measurements to geometric beacons of known position, and correcting a dead-reckoned pose with an extended Kalman filter at every match, localized a mobile robot in 1991 [9]. The same structure scaled to industrial vehicles. A 17.5-ton container carrier for port terminals drove at up to 6 m/s on millimeter-wave radar sightings of surveyed beacons and held its position to better than 3 cm [12]. The dead reckoning between sightings was characterized in the same years. On a smooth indoor floor the dominant errors of a differential-drive robot are systematic, an unequal pair of wheel diameters and an uncertain effective wheelbase, and a calibration run can measure and remove both [10].

> ⧖ HISTORY **The Bumper Came First (1953 to 2015)**
>
> Automatic guided vehicle systems date from 1953, and their original device for finding an obstacle was a collapsible contact bumper, still fitted to many vehicles. The bumper is instrumented to trigger a safety-rated stop when it strikes something. The United States standard ANSI/ITSDF B56.5 of 2012 asks that such a bumper be fail-safe, push with no more than 134 N, and stop the vehicle within its collapsible range, before the vehicle itself strikes the obstacle. Non-contact sensing has since spread, yet in 2015 the standards still set performance measures with no test methods to check them, and NIST began to develop those methods [13].

Warehouses then brought robots and people into the same aisles. Systems in which hundreds of mobile robots carry storage pods to picking stations were operating in distribution facilities by 2008 [2], and robotized handling now includes shuttle-based storage and robotic mobile fulfillment for e-commerce operations [3]. A robot that shares an aisle with people has to detect them and stop

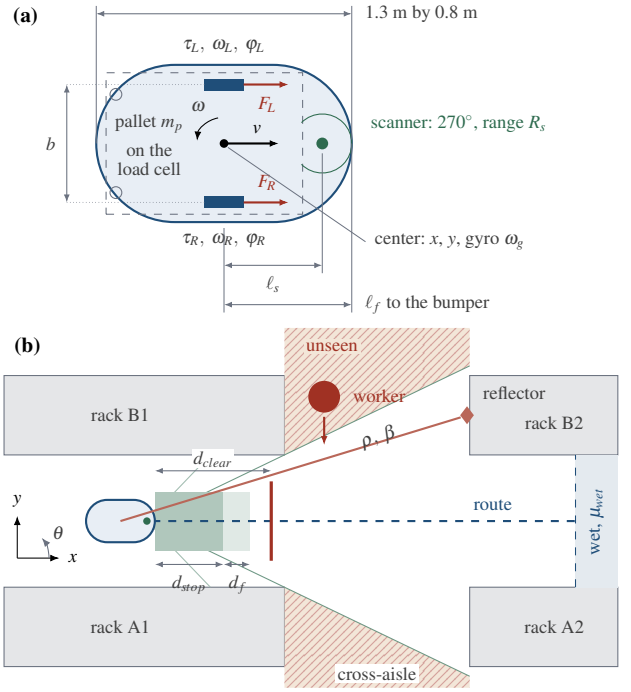


Figure 9.1 The robot and a blind cross-aisle. (a) The robot heading right, with its pallet m_p on a load cell, drive wheels of radius r on track b with torques τ_L, τ_R, speeds ω_L, ω_R, encoder angles φ_L, φ_R and traction forces F_L, F_R, the speed v and yaw rate ω at the center (x, y) where the gyro reads ω_g, and the scanner ℓ_s and the bumper ℓ_f ahead of the center. (b) Aisle AB at the first cross-aisle, map geometry to scale, heading θ. The stopping band $d_{stop} + d_f$ must fit inside the clear distance d_{clear}, which here ends where the map says the scanner cannot see past the rack ends (hatched). A worker stands in that space, a reflector on rack B2 is seen at range ρ and bearing β, and the wet patch of friction μ_{wet} begins at $x = 19\,\text{m}$.

in time, and the safety requirements and verification methods for driverless industrial trucks were set out in a standard published in 2020 [4].

The theory of safe motion grew alongside. The dynamic window approach restricts a robot to velocities from which it can stop before the nearest obstacle on its path [5]. Inevitable collision states extend the idea to any state from which no future control avoids a collision [6], and passive safety for a robot with a limited field of view asks only that the robot be at rest whenever a collision occurs, which obliges it to allow for obstacles hidden outside what it has seen [7]. Route search on a map with A* [14] and geometric trackers that steer toward a point ahead on the route [15] complete the layered structure the simulator implements, a planner above a path follower above two wheel loops.

9.3 A Pallet Robot on a Mapped Floor

The simulated system is one differential-drive pallet robot on a level warehouse floor 32 m by 16 m. Two driven wheels on a 0.60 m track carry 60 % of the weight of the robot and its pallet, and casters carry the rest. The footprint, pallet included, is a stadium 1.3 m long and 0.8 m wide whose bumper lies 0.65 m ahead of the center, and the range scanner sits 0.5 m ahead of the center. The robot, its forces and sensors, and the geometry of a blind cross-aisle are shown in Figure 9.1.

Table 9.1 Classification of the Robot's Signals

Symbol	Role	Unit	Default or range
v	CV, read as v_{enc}	m/s	SP $v_{ref} \le v_c$
θ	CV, read as $\hat{\theta}$	rad	SP θ_{ref}
τ_m	MV, mean wheel torque	N m	$\lvert\tau_m\rvert \le \tau_{max} - \lvert\tau_d\rvert$
τ_d	MV, differential wheel torque	N m	$\lvert\tau_d\rvert \le \tau_{max}$
x, y, θ	states, pose	m, rad	floor 32 by 16 m
v, ω	states, speed and yaw rate	m/s, rad/s	—
ω_L, ω_R	states, wheel speeds	rad/s	—
τ_L, τ_R	states, motor torques	N m	± 40
φ_L, φ_R	states, wheel angles	rad	counted
E	state, battery energy used	J	of 4.8 kWh
—	state, contact flag	—	0 or 1
m_p	DV, pallet mass	kg	0 to 1200, 300
ε_r	DV, wheel radius error	%	-5 to $+5$, 0
μ	DV, wet patch in aisle AB	—	0.7 dry, 0.1 wet
—	DV, reflectors hidden	—	x 7 to 24 m
b_g	DV, gyro bias	deg/s	-1 to $+1$, 0.05
—	DV, people, carts, pallets	—	per scenario
—	DV, worker or pallet ahead	—	events

9.3.1 Signals and Their Roles

The variables are classified in Table 9.1. Two loops act on one robot. The speed loop regulates the encoder speed with the mean wheel torque, and the heading loop regulates the estimated heading with the differential wheel torque. The positions of people and objects, the floor friction and the true pose are disturbances or states the controller never reads directly.

9.3.2 Body, Wheels and Traction

The body moves only along its heading, because the drive wheels do not skid sideways, and its planar motion is the differential-drive form of Equation (1.1) [16] with rolling resistance and caster scrub added, Equation (9.1).

$$\begin{aligned} \dot{x} &= v\cos\theta, \quad \dot{y} = v\sin\theta, \quad \dot{\theta} = \omega \\ m\dot{v} &= F_L + F_R - C_{rr}\, mg \tanh(v/v_0) \\ I\dot{\omega} &= (F_R - F_L)\, b/2 - c_t\, mg \tanh(\omega/\omega_0) \end{aligned} \tag{9.1}$$

where $m = m_r + m_p$ is the robot plus the pallet, $I = m_r(0.35\,\mathrm{m})^2 + m_p(1.2^2 + 0.8^2)/12$ in kg m^2 treats the pallet as a 1.2 m by 0.8 m slab, C_{rr} is the rolling resistance, c_t is a scrub arm for the casters, and $v_0 = 0.02\,\mathrm{m/s}$ and $\omega_0 = 0.02\,\mathrm{rad/s}$ smooth both resistances through zero. Each wheel carries its motor and gearbox reflected to the shaft, and each motor delivers its commanded torque through a first-order lag, Equation (9.2).

$$J_w\dot{\omega}_i = \tau_i - r_i F_i - b_w\omega_i, \quad T_m\dot{\tau}_i = \tau_{c,i} - \tau_i, \quad \dot{\varphi}_i = \omega_i \tag{9.2}$$

for $i \in \{L, R\}$, where $\tau_{c,i}$ is the torque command, T_m the motor lag and b_w the bearing drag. The wheel angles φ_i are carried as states because they are what the encoders count.

Traction is carried by the contact slip, the difference between the rim speed and the ground speed under each wheel, through a smooth friction curve, Equation (9.3).

$$\begin{aligned} s_{L,R} &= r_{L,R}\,\omega_{L,R} - (v \mp \omega b/2) \\ F_i &= \mu N \tanh(s_i/s_0), \qquad N = \chi mg/2 \end{aligned} \tag{9.3}$$

where $r_L = r_R = r(1 + \varepsilon_r/100)$ is the true radius, χ is the share of the weight on the drive wheels and μ is the friction under the robot's center. With the default 300 kg pallet the 59 N per wheel that

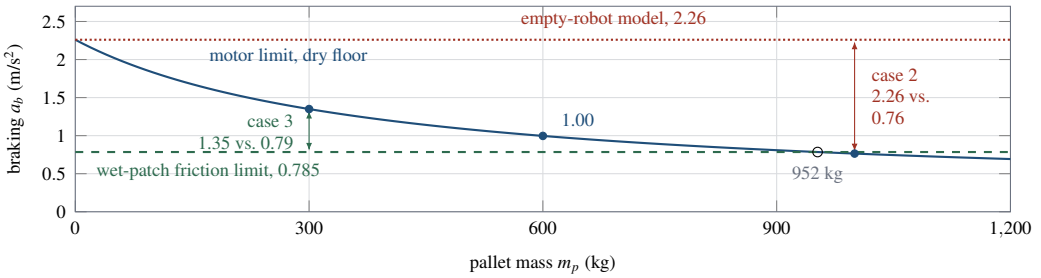


Figure 9.2 Braking capability against pallet mass from Equation (9.4). The motor limit (solid) falls with the load, the wet patch caps braking at 0.785 m/s^2 (dashed), and the empty-robot model believes 2.26 m/s^2 whatever the load (dotted). Filled dots mark three pallet masses and the open circle the crossing of the two limits. The arrows mark the gaps between belief and truth that cases 2 and 3 expose.

rolling resistance needs costs 5 mm/s of slip on the dry floor, a third of a percent at the 1.5 m/s cruise speed, so the encoders track the ground closely. On the wet patch the traction limit μN falls from 1236 N to 177 N per wheel, a hard stop locks the wheels while the robot slides, and the encoders stop measuring the robot. The curve rises linearly and saturates at μN, without the peak and fall that measured force–slip characteristics show [17]. Volume 7, Chapter 4 regulates a braking wheel at that peak.

9.3.3 Braking Capability

A flat-out stop is limited by the motors or by the floor, whichever runs out first, as shown in Equation (9.4). It is the braking capability A_b of Section 1.4.9, written a_b here.

$$a_b(m,\mu) = \min\left(\frac{2\tau_{max}/r + C_{rr} m g}{m + 2J_w/r^2},\ (\mu\chi + C_{rr})\, g \right) \tag{9.4}$$

The first term adds the reflected inertia of the rolling wheels to the mass the torque has to stop, and the second is the friction under the drive wheels plus rolling resistance. On the dry floor the motors always bind, because friction would allow 4.32 m/s^2. The capability is 2.26 m/s^2 empty, 1.35 m/s^2 with the default pallet, 1.00 m/s^2 with 600 kg and 0.76 m/s^2 with 1000 kg, while the wet patch caps it at 0.785 m/s^2. The two limits cross at a pallet of 952 kg, above which a wet floor no longer lengthens the stop, as shown in Figure 9.2. A verification check brakes the model flat out from 1.5 m/s with 0 and 1000 kg pallets, and the stopping distances agree with $v^2/(2a_b)$ to 0.8 %.

The battery supplies the motors through a drive efficiency, recovers part of the braking energy, and carries a copper loss and a hotel load for the computer, the scanner and the lights, Equation (9.5).

$$\dot{E} = P_h + \sum_i \left[k_{Cu}\tau_i^2 + \begin{cases} \tau_i\omega_i/\eta_d, & \tau_i\omega_i \geq 0 \\ \eta_r\,\tau_i\omega_i, & \tau_i\omega_i < 0 \end{cases} \right] \tag{9.5}$$

9.3.4 The Warehouse, Its Traffic and Contact

The map the controller knows holds the walls, six racks, five stations and 49 surveyed reflectors, listed in Table 9.2. Every rack is a solid box to the scanner, so each place where a cross-aisle meets an aisle is a blind intersection. Each reflector sits 5 cm proud of a rack end, a rack face or a wall and can be seen only from in front of its face.

The world adds what the map does not hold. Workers are circles of 0.3 m radius and carts are boxes 1.0 m by 0.7 m, and each moves along a polyline at a fixed speed, starting at a fixed time or when the robot's true front crosses a trigger line. A pallet can be set down in an aisle, the wet patch

Table 9.2 The Warehouse Map, in Meters

Feature	Extent
Floor and walls	x 0 to 32, y 0 to 16, walls 0.5 thick
Racks A1, A2	x 4.5 to 13.5 and 17.0 to 26.0, y 4.0 to 5.5
Racks B1, B2	the same x, y 8.0 to 9.5
Racks C1, C2	the same x, y 12.0 to 13.5
Aisles	y 0.5 to 4.0, 5.5 to 8.0 (AB), 9.5 to 12.0 (BC), 13.5 to 15.5
Cross-aisles	x 0.5 to 4.5, 13.5 to 17.0, 26.0 to 31.5
Dock D, stations	D (2.25, 6.75), S1 (28.75, 6.75), S2 (28.75, 14.5), S3 (28.75, 2.25), S4 (2.25, 14.5)
Reflectors	six per rack (both ends, and at 28 and 72 % of each long face), two on the west wall, three on the east wall, four each on the north and south walls
Wet patch	x 19 to 23, y 5.5 to 8.0
Dropout zone	reflectors with x 7 to 24

can be switched on, and stock can hide the reflectors in the dropout zone. Agent motion follows time once its fixed or robot-triggered start is set, so a seeded run with the same controls replays exactly, and people and carts never yield. Contact between the stadium and a wall, a rack or an object ends the run with the robot stopped where it touched. Contact with a person or a cart counts only while the robot moves faster than 0.05 m/s, because a scripted worker who walks into a standing robot is not a collision the robot caused.

> ✦ DID YOU KNOW? **Two Cylinders Stand In for a Person**
>
> The driverless-truck safety standards test with two cylinders: one 200 mm across and 600 mm long lying in the path, and one 70 mm across and 400 mm tall standing in it. They can be taken for a torso or upper leg on the floor and for a lower leg, although the standards do not say so [13].

9.3.5 Sensors

A single function reads the true state and the world and returns a measurement record, and all random draws happen there, from a seeded generator. The encoders report whole counts, the gyro adds a bias and white noise, and the load cell reports the pallet mass exactly, as shown in Equation (9.6).

$$\begin{aligned} n_i &= \lfloor \varphi_i N_{enc} / 2\pi \rfloor, \qquad \omega_g = \omega + b_g + \sigma_g w \\ \rho_j &= \| \mathbf{p}_j - \mathbf{p}(t-\delta) \| + \sigma_\rho w \\ \beta_j &= \angle\big(\mathbf{p}_j - \mathbf{p}(t-\delta)\big) - \theta(t-\delta) + \sigma_\beta w \end{aligned} \tag{9.6}$$

where each w stands for a fresh draw from the unit normal distribution, $\mathbf{p}$ is the position of the center and $\mathbf{p}_j$ that of reflector j. The range scanner casts 109 rays over 270°, 2.5° apart, from its mount out to R_s, returns the first wall, rack, person, cart or pallet on each ray with noise σ_s, and runs at 10 Hz. A reflector is identified when it lies within R_s and farther than 0.3 m from the scanner, inside the field of view and in line of sight, is seen at less than 75° from the normal of its face, and is outside the dropout zone when that disturbance is on. Its range and bearing are taken from the center, reported δ late, and arrive with the reflector's identity, so there is no association problem.

9.3.6 Parameters, Numerics and the Initial State

The parameters are listed in Table 9.3. The reader can change every one marked adjustable, and three are hidden.

The 13 states are integrated by the classical fourth-order Runge–Kutta method on a 2 ms step. The step is set by the traction, not by the motion. Linearizing Equations (9.1) to (9.3) about a steady roll at 1.5 m/s gives two slip modes at $-812\ \mathrm{s}^{-1}$ and $-794\ \mathrm{s}^{-1}$ with a 1200 kg pallet on the dry floor, $-350\ \mathrm{s}^{-1}$

Table 9.3 Robot, Floor and Sensor Parameters

Symbol	Meaning	Value	Class
m_r	robot mass, empty	300 kg	adj., 150 to 600
J_w	wheel and motor inertia	0.4 kg m^2	adj.
r	wheel radius believed	0.10 m	adj.
b	track	0.60 m	adj.
τ_{max}	torque limit, each wheel	40 N m	adj., 10 to 80
T_m	motor torque lag	0.05 s	adj.
b_w	bearing drag	0.05 N m s	hidden
μ_{dry}, μ_{wet}	floor friction	0.7, 0.1	adj.
χ	weight on the drive wheels	0.6	adj.
s_0	traction slip scale	0.1 m/s	adj.
C_{rr}	rolling resistance	0.02	adj.
c_t	caster scrub arm	0.003 m	hidden
v_0, ω_0	smoothing through zero	0.02 m/s, rad/s	fixed
P_h	hotel load	120 W	adj.
η_d, η_r	drive efficiency, recovery	0.85, 0.25	adj.
k_{Cu}	copper loss	0.05 W/(N m)^2	hidden
—	battery capacity	4.8 kWh	adj.
N_{enc}	encoder counts per turn	2048	adj.
σ_g	gyro noise per sample	0.003 rad/s	adj.
R_s	scanner range	8 m	adj., 3 to 15
σ_s	scanner range noise	0.02 m	adj.
σ_ρ, σ_β	reflector noise	0.03 m, 0.01 rad	adj.
δ	reflector delay	0.1 s	adj.
ℓ_s, ℓ_f	scanner and bumper offsets	0.5 m, 0.65 m	fixed

with the default pallet and $-44\ \mathrm{s}^{-1}$ on the wet patch. The method is stable on the negative real axis only for $|\lambda|\Delta t < 2.79$, so a 10 ms step would diverge under a heavy pallet and 2 ms leaves a margin of 1.7. Halving the step to 1 ms moves the end-of-step detection of the case 1 delivery by 0.001 s, while the controller-stamped mission time agrees to numerical precision and the final position differs by less than a micrometer. The wheel loops and the estimators run every 20 ms, navigation, the scanner and the supervisor every 100 ms, and the historian logs at 10 Hz. Each physics step follows the order of Section 1.6.1, with a contact check after the integration. Every scenario starts at rest at the dock, at 2.25 m, 6.75 m heading east with no torque, where every mechanical rate is zero to 10^{-12} and only the hotel load drains the battery. The robot then stands 2 s while its pallet is loaded.

9.3.7 Assumptions

The robot is planar and rigid on a level floor, and its drive wheels never skid sideways. That holds at warehouse speeds on a dry floor. It fails in a hard turn on the wet patch, where a real robot would also slide laterally.

Traction follows a smooth curve with no peak. The braking force therefore never falls as slip grows, and with locked wheels near standstill it fades, which slightly lengthens a stop at the friction limit.

Rolling resistance and caster scrub are smoothed through zero, and below 0.02 m/s and 0.02 rad/s they act as stiff viscous dampers. The last degrees of a turn in place are therefore overdamped (Section 9.4.4).

The controller's braking model is the plant's own Equation (9.4) with the nominal parameters, the exact pallet mass from the load cell and, when the wet zone is mapped, the true wet friction. A real robot must estimate all three and cover what it gets wrong with margin.

The scanner is one horizontal plane in which racks are solid boxes and people are circles, and reflectors arrive with their identities. A real scanner sees under shelving and misses objects below its

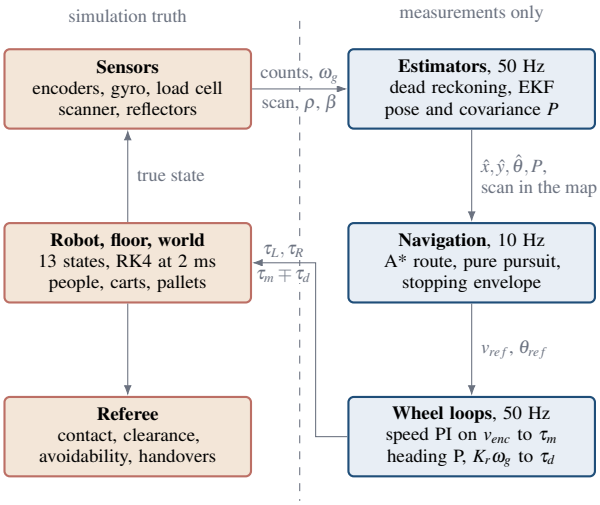


Figure 9.3 The layers of the controller. Only the sensor layer reads the simulation truth, and the referee scores the run from the truth without feeding anything back. The encoders, gyro, load cell and wheel loops run at 50 Hz, and the scanner, reflectors, navigation and supervisor at 10 Hz.

plane, and a wrong association between a return and a reflector would bias the estimate in a way the noise model does not describe.

People and carts follow fixed scripts and never yield, which makes every run reproducible and puts the whole responsibility for avoiding contact on the robot. There is one robot, a known map and no fleet traffic management.

9.4 A Controller That Reads Only Measurements

The controller is layered as shown in Figure 9.3. Its decision code reads the measurement record, the settings, the robot's nominal parameters and its own state, and nothing else. A verification check runs two identical controllers on the same measurements while the true pose of one is moved by 0.6 m, –0.3 m and 0.2 rad, and the torque commands and speed setpoints agree to the last bit over ten control steps.

9.4.1 Dead Reckoning and the Extended Kalman Filter

Two estimators run side by side on the same measurements, and a switch chooses which one the robot steers by. Section 1.4.6 treats dead reckoning and the Kalman filter in general. Both estimators here turn the count increments Δn_i of one 20 ms period into a distance, with the believed radius r, and a heading change, as shown in Equation (9.7).

$$\begin{aligned} \Delta s &= r\,(\Delta\varphi_L + \Delta\varphi_R)/2 \\ \Delta\theta &= \begin{cases} r(\Delta\varphi_R - \Delta\varphi_L)/b & \text{dead reckoning} \\ \omega_g\,\Delta t_c & \text{EKF} \end{cases} \end{aligned} \tag{9.7}$$

with $\Delta\varphi_i = 2\pi\Delta n_i / N_{enc}$ and $\Delta t_c = 20\,\text{ms}$. Each estimate $\hat{\mathbf{x}} = [\hat{x}, \hat{y}, \hat{\theta}]$ starts at the dock with standard deviations of 0.02 m in position and 0.5° in heading, and is propagated at the midpoint heading $\theta_m = \hat{\theta} + \Delta\theta/2$ with its covariance, Equation (9.8).

$$\begin{aligned} &\hat{x} \leftarrow \hat{x} + \Delta s\cos\theta_m, \quad \hat{y} \leftarrow \hat{y} + \Delta s\sin\theta_m, \quad \hat{\theta} \leftarrow \hat{\theta} + \Delta\theta \\ &P \leftarrow FPF^T + G\,\text{diag}(q_s, q_\theta)\,G^T \end{aligned} \tag{9.8}$$

where F is the identity with $-\Delta s \sin\theta_m$ and $\Delta s \cos\theta_m$ in its third column, G maps the distance and heading increments into the pose, and the odometry error is a random walk in distance traveled rather than in time, $q_s = (\sigma_{10}^2/10\,\mathrm{m})|\Delta s| + 10^{-10}\,\mathrm{m}^2$. The heading noise is $q_\theta = 3\times10^{-5}|\Delta s|$ in rad^2 for dead reckoning and $10^{-4}\Delta t_c + (0.01\,\omega_g \Delta t_c)^2$ for the filter. Neither model knows about a wrong radius, a skid or the gyro bias.

Every reflector identified in a scan updates the filter in turn. The innovation is formed against the estimate stored for the moment of measurement, δ earlier, and the correction is applied to the present estimate, as shown in Equation (9.9).

$$\begin{aligned} \nu &= \left[\rho - \hat{\rho},\ \mathrm{wrap}(\beta - \hat{\beta})\right]^T, \quad S = HPH^T + R \\ K &= PH^T S^{-1}, \quad \hat{\mathbf{x}} \leftarrow \hat{\mathbf{x}} + K\nu \\ P &\leftarrow (I - KH)\,P\,(I - KH)^T + KRK^T \end{aligned} \tag{9.9}$$

where $\hat{\rho}$, $\hat{\beta}$ and the Jacobian H are formed from the stored pose, but P is the current covariance. This delayed-pose innovation is an approximation rather than a full out-of-sequence covariance update, and $R = k_\rho^2\,\mathrm{diag}(\sigma_\rho^2, \sigma_\beta^2)$ with an assumed noise multiplier $k_\rho = 1$. A sighting is rejected when $\nu^T S^{-1} \nu$, the squared Mahalanobis distance of the innovation, exceeds 13.8, the 99.9 % point of the chi-square distribution with two degrees of freedom. When two scans that each hold two or more reflectors have every sighting rejected, with none accepted in between, the estimate rather than the reflectors is taken to be wrong, and the position variances grow by 0.25 m^2 and the heading variance by $(5^\circ)^2$ so the next scan can pull the estimate back. The ellipse drawn in the scene has the 95 % radius $r_{95} = \sqrt{5.991\,\lambda_{max}}$, with λ_{max} the larger eigenvalue of the position block of P. The dead-reckoning estimate never receives a reflector update. The underwater vehicle of Volume 5, Chapter 13 applies its late acoustic fixes behind the same kind of chi-square gate.

9.4.2 The Stopping Envelope

The supervisor sets the speed so that the reaction distance plus the braking distance, with the braking the robot believes it has, fits inside the room available. Inverting that condition gives the envelope speed of Equation (9.10), which is Equation (1.17) of Chapter 1 written with the belief a_b.

$$v t_r + \frac{v^2}{2a_b} \le d \iff v \le V(d) = a_b\left(\sqrt{t_r^2 + 2d/a_b} - t_r\right) \tag{9.10}$$

with $V(d) = 0$ for $d \le 0$. The belief a_b is Equation (9.4) evaluated with $m = m_r$ for the empty-robot model or $m = m_r + m_p$ from the load cell for the payload-aware model, and with μ_{wet} when the mapped wet zone, grown by 0.5 m, lies within 4 m ahead on the route. The envelope treats an obstacle as standing still. Chapter 15 adds the distance a person walks while the machine reacts and brakes.

Each scan is placed in the map through the estimate the robot steers by. A return within 0.55 m of the route and at most 8 m ahead along it is an obstacle in the path, and the nearest sets d_{obs}. In the cautious mode two further limits apply. A return that lies farther than $\max(0.3, r_{95} + 0.1)$ meters from mapped racking, and so is something the map does not know, counts as an obstacle d_{side} when it is within 0.5 m of that strip and ends the clear distance when it is within 1.5 m. Space the scanner has not seen ends it as well. Along the route, from the bumper to 7.5 m ahead in 0.25 m steps, every 0.5 m map cell within 2.05 m of the route point is tested, and a free cell at least 0.3 m ahead of the center counts as unseen unless all four of its corners, inset 5 cm, lie within $R_s - 0.1$ m of the scanner with no mapped rack in the sightline. The first route point with such a cell sets d_{occ}. The definition comes from the map rather than from the rays, because a gap between rays is not a blind corner and the shadow of something already in the path is not either. The speed target is the smallest of the limits in Equation (9.11).

$$\begin{aligned} v^* = \min\big(&v_s,\ V(d_{obs} - d_f),\ V(d_{side} - d_f), \\ &\max[v_{cr}, V(d_{occ})],\ v_{dock}\big) \end{aligned} \tag{9.11}$$

where $v_{cr} = 0.3\,\mathrm{m/s}$ is the creep speed, which stops within 0.12 m with the default pallet and t_r. Its positive floor relaxes Equation (9.10) when the unseen gap approaches zero, so the implemented policy is not a guarantee of stopping inside all visible space. The terms in d_{side} and d_{occ} are dropped when unseen space is treated as clear. The static limit v_s is the least of the cruise speed v_c, a turning limit $v_c\,\mathrm{clamp}(1.15 - |\alpha|/60^\circ, 0.25, 1)$, zero while the robot turns in place, a corner limit $\sqrt{0.6^2 + 2(0.5)\max(0, d_{corner} - 0.3)}$ for the next vertex turning more than 0.4 rad, and 0.4 m/s while r_{95} exceeds the confidence limit r_c. The docking limit is $v_{dock} = \max(0.06, \sqrt{2(0.35)(d_{goal} - 0.03)})$ for $d_{goal} > 0.06\,\mathrm{m}$ and zero below it. Between scans the stored distances shrink by the distance the encoders report. The setpoint v_{ref} follows a falling v^* at once and a rising one at $0.6\,\mathrm{m/s^2}$. An emergency stop brakes both wheels flat out when $v_{enc} > 0.1\,\mathrm{m/s}$ and $d_{obs} < 0.1\,v_{enc} + v_{enc}^2/(2a_b)$, and releases after the encoder speed has stayed below 0.02 m/s for 1 s. The clear distance $d_{clear} = \min(d_{obs}, d_{occ})$ and the stopping distance d_{stop} at the present speed are drawn in the scene and logged.

9.4.3 Route, Pure Pursuit and the Mission

Routes come from A* [14] on a grid of 64 by 32 cells of 0.5 m. A cell is passable when its center lies at least $0.4\,\mathrm{m} + c_p$ from every wall, rack and blocked cell, each step costs its length times $1 + 2(1.25 - \min(c, 1.25))/1.25$ for a clearance c in meters so that the route keeps to the middle of an aisle, diagonal steps may not cut a corner, and the route is then straightened wherever a chord keeps 1.0 m of clearance. The path follower is pure pursuit in its carrot form, one of the path-following laws of Section 1.4.5. The heading setpoint is the bearing from the estimated position to the route point one lookahead distance L_d beyond the robot's projection, Equation (9.12),

$$\theta_{ref} = \mathrm{atan2}(y_\ell - \hat{y},\, x_\ell - \hat{x}), \qquad \alpha = \mathrm{wrap}(\theta_{ref} - \hat{\theta}) \tag{9.12}$$

and a heading loop closes on it, rather than a curvature being fitted through the point [15]. Within 0.5 m of the goal the setpoint becomes the heading of the last route segment, and when $|\alpha|$ exceeds 50° the robot turns in place until it falls below 12°. Chapter 2 follows a road with the Stanley law, which steers on the heading error and the cross-track error of the front axle.

The mission loads the pallet for 2 s, drives each leg, and hands over for 2 s at each station. Arrival is declared when fewer than 0.06 m remain on the route and the encoder speed is below 0.03 m/s, and the handover proceeds when the estimate is within 0.1 m of the station, or after two re-approaches. When $d_{obs} - d_f < 0.4\,\mathrm{m}$ with the robot stopped, the robot waits. After the waiting time t_w it marks the map cells of the returns in its path as blocked, on its own copy of the map, and plans again, and a blocked cell that a later scan sees straight through is cleared.

9.4.4 The Wheel Loops as Implemented

Both wheel loops use the velocity-form PID routine of Section 1.6.2, Equations (1.19) to (1.22), in which only the emitted output is clamped and anti-windup by back-calculation, with the tracking time $T_t = \tau_I$, changes the integral term alone. Both loops run with $\tau_D = 0$, no output slew and the action $a = +1$, reverse acting in the convention of Section 1.3, so that a positive error calls for more torque.

The speed loop has $\mathrm{SP} = v_{ref}$, $\mathrm{PV} = v_{enc} = \Delta s/\Delta t_c$, $K_c = K_v = 150\,\mathrm{N\,m}$ per m/s, $\tau_I = 3\,\mathrm{s}$ and $u_b = 0$, and its output is the mean torque τ_m. The heading loop has $\mathrm{SP} = \theta_{ref}$ and $\mathrm{PV} = \hat{\theta}$, both unwrapped, $K_c = K_\theta = 100\,\mathrm{N\,m/rad}$, no integral by default (the setting of 60 s means off), and its yaw-rate damping enters as a live output bias $u_b = -K_r\,\omega_g$ with $K_r = 60\,\mathrm{N\,m\,s/rad}$. Its output is the differential torque τ_d. The rate term is a derivative on the measurement taken from a second sensor rather than by differencing. The two outputs share each wheel's limit with turning served first, the moment-first allocation of Section 1.4.7, Equation (9.13).

$$\begin{aligned} |\tau_d| &\le \tau_{d,max}, & |\tau_m| &\le \tau_{max} - |\tau_d| \\ \tau_{c,L} &= \tau_m - \tau_d, & \tau_{c,R} &= \tau_m + \tau_d \end{aligned} \tag{9.13}$$

so a robot braking hard can still steer. A traction limiter sets $\tau_{d,max} = 0.3\,\tau_{max}$ for 0.5 s whenever the yaw rate implied by the encoders differs from the gyro by more than 0.25 rad/s, which is the signature

of a spinning wheel. An emergency stop replaces both loops with $\tau_d = 0$ and $\tau_m = -\tau_{max}\,\mathrm{sign}(v_{enc})$ until the encoder speed falls below 0.02 m/s. A robot standing still with a zero setpoint holds zero torque, so the integral cannot creep it.

In the physical units of the plant, the heading loop applies $K_\theta b/r = 600\,\mathrm{N\,m}$ of yaw torque per radian and $K_r b/r = 360\,\mathrm{N\,m\,s}$ per radian of damping. With the default pallet the rigid-body loop, whose yaw inertia includes $7.2\,\mathrm{kg\,m^2}$ reflected from the wheels, would have a natural frequency of 2.5 rad/s and a damping ratio of 0.75, and with 1000 kg of 1.66 rad/s and 0.50. A 20° heading step at rest, with the mission held in its completed phase and the normal sensors and wheel loops active, reaches 90 % in 1.34 s with the default pallet and 1.63 s with 1000 kg, with no overshoot in either case. The smoothed caster scrub supplies the missing damping, and with it removed the same step overshoots by 18.3 % with 1000 kg. The speed gain places the proportional pole of the speed loop at $2K_v/(r\,m_{eff}) = 4.4\,\mathrm{s^{-1}}$ with the default pallet and $2.2\,\mathrm{s^{-1}}$ with 1000 kg, where $m_{eff} = m + 2J_w/r^2$.

In manual mode the operator sets τ_m and τ_d directly, by slider, by holding the arrow keys (which add $0.6\tau_{max}$ and $0.5\tau_{max}$) or by timed touch buttons. The estimators, the scene and the scorecard keep running, the speed setpoint tracks the encoder speed, and each loop tracks the operator's torque through its integral term, as Section 1.6.2 describes. A return to automatic remains smooth when its reference and bias agree with the tracked values. A new speed target, heading target or torque allocation can change the first automatic command. No speed limit applies in manual mode, but the emergency stop does. It brakes both wheels flat out over the operator's torque until the robot has stood still for 1 s. A steady 20 N m push toward the pallet of case 7 is stopped from 3.18 m/s, and because the push resumes each time the stop releases, the robot still touches the pallet at 0.07 m/s after three more stops, below the 0.1 m/s at which the stop is armed.

9.4.5 Three Models and a Referee

Equations (9.1) to (9.6), together with the world of Table 9.2, form the simulation model. The controller's internal models are separate and simpler. Odometry uses the believed radius and track, the estimators assume a random walk in distance and white reflector noise, the supervisor uses Equation (9.4) with a mass and a friction it is told, and the planner uses the map without the objects the world adds. No model is identified from data anywhere in the application.

The referee scores the run from the truth and feeds nothing back, the check of a monitor against the truth that Section 1.4.9 describes. A handover counts as a delivery when the true center is within 0.3 m of the station, and a clearance violation is counted each time a person or an object comes within 0.2 m of the outline while the robot moves faster than 0.05 m/s. A collision is scored as unavoidable only when the object gave no warning, having been seen at rest and started moving within the last 2 s or appeared less than 0.25 s before, and entered the strip the robot sweeps closer than the true stopping distance of Equation (9.14).

$$d_{enter} < v\,T_{min} + \frac{v^2}{2\,a_b(m_r + m_p,\ \mu_{true})}, \qquad T_{min} = 0.02\,\mathrm{s} + T_m. \tag{9.14}$$

The referee adds the configured motor lag to the 0.02 s control period, giving 0.07 s at default settings, with no scan period counted. This optimistic stopping-distance classification is a simulation rule, not a proof of inevitability for every controller or traffic trajectory. An object that walks out of space the scanner could not see is classified as avoidable by definition. That convention does not establish a safety guarantee for the implemented positive-creep policy.

The settings are listed in Table 9.4, and each scenario starts from these defaults and changes the settings and disturbances for its lesson.

Table 9.4 Controller Settings, With Ranges

Symbol	Meaning	Default	Range
v_c	cruise speed	1.5 m/s	0.2 to 2.0
K_v	speed gain, N m per m/s	150	5 to 400
$\tau_{I,v}$	speed integral time	3 s	0.2 to 20
L_d	lookahead distance	1.2 m	0.4 to 4
K_θ	heading gain	100 N m/rad	5 to 400
K_r	yaw-rate damping	60 N m s/rad	0 to 200
$\tau_{I,\theta}$	heading integral time	off (60 s)	1 to 60
d_f	following distance	0.5 m	0.1 to 2.0
t_r	reaction time assumed	0.3 s	0.05 to 1.5
—	braking model	payload-aware	or empty robot
—	unseen space	slow for it	or treat as clear
—	wet-floor zone	not mapped	or mapped
—	estimator	EKF	or encoders only
σ_{10}	odometry drift, m per 10 m	0.1	0.01 to 0.5
r_c	confidence limit	0.6 m	0.1 to 2.0
k_ρ	reflector noise assumed	1	0.5 to 10
—	replan around blockages	on	or off
t_w	wait before replanning	3 s	1 to 30
c_p	planning clearance	0.35 m	0.1 to 0.6

9.5 Eight Broken Assumptions and a Graded Shift

Table 9.5 gives the scenario sequence, and Table 9.6 collects what each run measured. Every number was read from the running engine at the default seed, 20261117. Runs that completed or collided ended 1 s after the last handover or collision. Unfinished variants were observed to the stated horizon. The engine was advanced at every 2 ms physics step. Learner variants were applied before Run. Mission times are counted from leaving the dock, 2 s after the start, and the running application is shown during the graded shift in Figure 9.5. Figure 9.4 draws two of the outcomes as bars.

The native completion workflow is summarized in Table 9.5. Cases 1–7 retain their completed observations across Reset. Their finished-run checks accept either mission completion or a collision, and do not require the safer variant to deliver successfully. Cases 1–6 also verify settings against the logged trajectory, so changing an option after a finished run does not earn that comparison. Case 7's second check requires replanning off, elapsed time at least 60 s, active blockage waiting for at least 20 s, and neither a collision nor mission completion. Case 8 requires its observations within one attempt. The challenge judges actual outcomes and preserves the case's model and disturbances.

9.5.1 The Price of Caution

On the clear route the fused estimate stays within 0.034 m of the true position, 0.011 m root-mean-square, and the robot hands over 3 cm from Station 1 in 31.6 s. It does not hold its cruise speed. Twice the map tells the supervisor that the racks hide space beside the route, at the cross-aisle and where aisle AB opens into the station area, and unseen space limits the speed from $x = 10.7$ m to 13.0 m and again from 23.2 m to 25.4 m, where the robot creeps at 0.3 m/s until it can see around the corner. Treating unseen space as clear delivers in 22.9 s and 7.9 kJ against 10.4 kJ, so caution at two blind corners costs 8.7 s on a 26.5 m route. The dead-reckoning estimate running beside the filter ends 0.14 m off after the same 26.5 m.

Table 9.5 Scenarios and Completion Requirements

Scenario	Kind	Completion requirement
1. Clear route	Experiment	Finish runs with unseen space cautious and treated as clear. Reset between runs and set the option before Run. The log must show a finite unseen gap for caution and the 8 m cap throughout the clear-space run.
2. Heavier pallet	Demonstration	Finish both empty-model and payload-aware runs with the 1000 kg pallet. Logged assumed deceleration must stay above 2 and below 1 m/s^2, respectively.
3. Slippery patch	Demonstration	Finish unmapped and mapped wet-floor runs. Assumed deceleration stays above 1.3 m/s^2 when unmapped and falls below 1 m/s^2 at least once when mapped.
4. Wheel-radius error	Demonstration	Finish encoder-only and EKF runs with 2.5 % radius error. The selected-position log must equal the corresponding estimator log throughout each run.
5. Occluded intersection	Demonstration	Finish clear-space and cautious runs with the worker. The unseen-gap log distinguishes the two policies as in case 1.
6. Unexpected cart	Demonstration	Finish at cruise speeds 1.5 and 0.8 m/s. Each setting must remain unchanged throughout its logged run. Read the referee's collision classification.
7. Blocked aisle	Experiment	Finish with replanning on and at least one replan. Reset and switch replanning off. Reach 60 s while still waiting behind the blockage for at least 20 s, without collision or completed mission.
8. Landmark dropout	Experiment	Accumulate 2 s with low-confidence speed limiting. Past $x = 23$ m, satisfy the covariance bound $5.991[\max(P_{xx}, P_{yy}) + \lvert P_{xy}\rvert] < 0.01$ m^2, then finish with one successful delivery.
9. Mixed shift	Challenge	At mission end, deliver to S1, S2 and S4 within 0.3 m, with no misses, collision or clearance violation and minimum moving clearance at least 0.2 m. Keep model parameters and disturbances as given.

Table 9.6 Outcomes of Every Case and Variant

Case and condition	Result
1, slow for unseen space	31.6 s, 3 cm from S1, 10.4 kJ
1, unseen treated as clear	22.9 s, 7.9 kJ
2, empty-robot braking	E-stop at 18.22 s, cart hit at 0.74 m/s
2, payload-aware	stops 0.25 m short, 40.3 s
2, payload-aware, K_v 100 / 200 / 250	stops 0.15 / 0.39 / 0.48 m short
3, wet zone not mapped	stops 0.07 m short (violation), 36.6 s
3, wet zone mapped	stops 0.41 m short, 36.7 s
4, encoders only	0.44 m past S1 (missed), 45.1 s
4, EKF	2 cm from S1, 30.9 s
5, unseen treated as clear	E-stop at 10.18 s, worker hit at 0.72 m/s
5, slow for unseen space	never within 0.92 m, 31.6 s
6, cruise 1.5 m/s	cart hit at 0.71 m/s, unavoidable
6, unseen treated as clear	cart hit at 1.24 m/s, unavoidable
6, cruise 0.8 m/s	stops 0.52 m short, 74.1 s
7, replanning on	replans at 22.6 s, route 11.4 m longer, 57.5 s
7, replanning off	still waiting at 100 s
8, reflectors hidden	r_{95} 1.17 m, error 0.30 m, 42.6 s
9, as shipped	worker hit at 1.47 m/s at 11.1 s
9, suggested	3 of 3 in 120.6 s, 0.27 m clearance, 48.4 kJ
9, suggested but unseen clear	worker hit at 1.47 m/s
9, suggested but empty-robot braking	pallet hit at 0.25 m/s
9, suggested but wet zone not mapped	0.11 m to the cart (violation)
9, suggested but encoders only	S1 missed, unfinished at 170 s

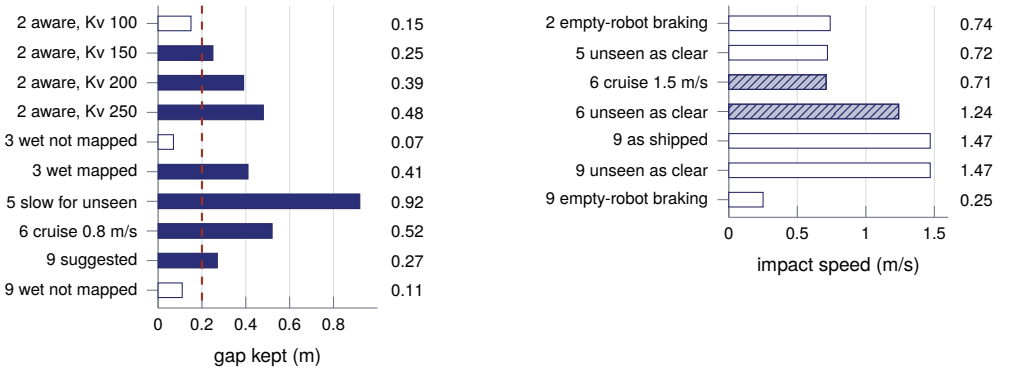


Figure 9.4 Two outcomes of Table 9.6 drawn as bars, labeled by case. Left, the gap each run kept to the cart, worker or pallet ahead, with the 0.2 m clearance as a dashed line: solid bars keep it and white bars violate it. Right, the impact speed of each run that collided, with the two collisions the referee scored unavoidable hatched. Notice that a slower speed loop or an unmapped wet patch eats the gap that the following distance was meant to keep, and that every collision but those two comes from a wrong belief the robot could have corrected.

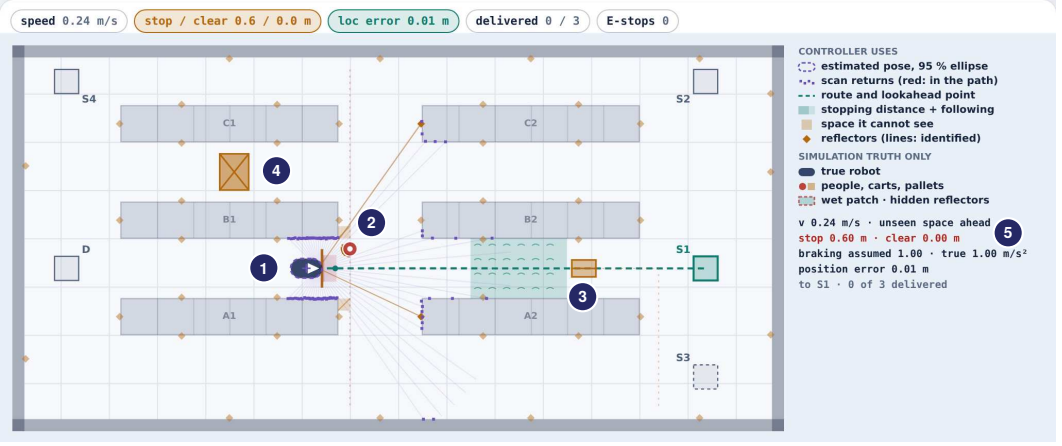


Figure 9.5 The graded shift with the suggested configuration at $t = 11.00$ s, shortly before the shipped configuration's collision at 11.09 s, cropped to the plant view. The 600 kg robot (1, solid, with its estimate dashed) approaches the first cross-aisle at 0.24 m/s. Shaded amber cells mark unseen space beside the rack ends, and the supervisor requests a 0.30 m/s creep. The worker (2) is approximately 1.05 m from the robot outline. The cart (3) waits beyond the wet patch and a pallet (4) blocks aisle BC. The scene readout (5) gives a displayed clear distance of zero and a stopping band of approximately 0.60 m, including the 0.50 m following distance, and the assumed and true braking both read 1.00 m/s^2.

DID YOU KNOW? What Two Blind Corners Cost

From the simulator: slowing for unseen space costs 8.7 s on the clear route, 31.6 s against 22.9 s (Table 9.6). With a worker in the cross-aisle the cautious robot still delivers in 31.6 s, while the robot that treats unseen space as clear hits the worker at 0.72 m/s.

9.5.2 Braking for the Load and for the Floor

Case 2 carries 1000 kg and plans with the empty robot's 2.26 m/s^2, where the load leaves 0.76 m/s^2. From 1.5 m/s the robot believes it needs 0.50 m of braking and needs 1.47 m. It brakes far too late, the emergency stop fires at 18.22 s, and flat-out braking is still only 0.76 m/s^2, so it hits the parked cart at 0.74 m/s at 18.61 s. With the payload-aware model the stopping band at 1.5 m/s is 1.92 m plus the following distance, braking starts earlier, and the robot stops 0.25 m from the cart, then follows it when it moves off and delivers in 40.3 s.

> **TRY IT IN THE APP** **2 · Heavier Pallet**
>
> Run the case as shipped: the robot plans its stops with the empty robot's braking and hits the parked cart at 0.74 m/s. Press Reset, set Braking model to Payload-aware before pressing Run, and run again. The robot starts braking earlier, stops 0.25 m from the cart and delivers in 40.3 s.

The 0.25 m is not the 0.5 m following distance, because the supervisor plans as though the speed loop tracked its setpoint without lag. The loop's lag is measurable directly. With the payload-aware model the robot stops 0.15 m from the cart at a speed gain of 100 N m per m/s, after an emergency stop and with a clearance violation, 0.25 m at the default 150, 0.39 m at 200 and 0.48 m at 250. An outer loop's plan is only as good as the inner loop that has to follow it.

Pause. Try Exercise 9.1 before reading on.

Case 3 stops on a wet patch that allows 0.785 m/s^2 while the robot plans with the dry-floor 1.35 m/s^2. The difference eats 0.43 m of the 0.5 m following distance, and the robot comes to rest 0.07 m from the cart, inside the 0.2 m every person and object must keep, a clearance violation. The locked wheels slip at up to 0.77 m/s, and the fused estimate falls up to 0.28 m from the truth, both in the logged record. Between 21.2 s and 26.5 s, while the wheels lock during the stop and spin during the restart on the patch, 33 reflector sightings fail the gate, the covariance is re-inflated twice, and the traction limiter acts for 1.78 s in all. With the wet zone mapped the robot plans with wet friction near the patch and stops 0.41 m short. The unmapped and mapped runs deliver in 36.6 s and 36.7 s, respectively, so a successful delivery alone does not distinguish their clearance safety. The maximum localization error is 0.32 m at the control rate, greater than the 0.28 m peak in the exported record.

9.5.3 Where the Robot Believes It Is

In case 4 a wheel 2.5 % larger than the controller believes turns $1/1.025$ times as much over the same ground. Dead reckoning therefore understates rim travel by 2.44 %, with additional contact slip. The true-to-believed rim-travel ratio is 1.025 in a verification check. The encoder-only robot believes it has arrived when it is 0.44 m past Station 1, a missed delivery. Its covariance models random drift rather than a bias, and at the handover the 95 % ellipse has semi-axes of 1.20 m across the aisle and only 0.40 m along it, so the truth lies outside it at a squared Mahalanobis distance of 7.78 against 5.99. The growing ellipse also passes the 0.6 m confidence limit, and the robot drives 22.1 s at the low-confidence speed and takes 45.1 s. The fused estimate, corrected by reflectors that do not care what radius the wheels have, hands over 2 cm from the mark in 30.9 s, while the dead reckoning running beside it ends 0.51 m off.

Pause. Try Exercise 9.2 before reading on.

In case 8 stock hides every reflector between $x = 7$ m and 24 m, and the last sighting before the zone is accepted with the robot at $x = 4.4$ m. The 95 % radius grows from 0.03 m as the robot leaves the dock to 1.17 m, while the true error reaches 0.30 m. Inside the zone the claimed radius stays at least 0.11 m above the true error, so the claim is conservative. The robot drives 14.8 s at the low-confidence speed and delivers in 42.6 s, against 31.6 s on the clear route. The first reflector past the zone, seen with the robot at $x = 23.4$ m, takes the radius from 1.17 m to 0.35 m and the true error from 0.30 m to 0.05 m in one update, and within 2 s the radius is 0.04 m and the error 0.01 m.

9.5.4 What the Scanner Cannot See

In case 5 a worker walks down the cross-aisle close to the rack end. Treating unseen space as clear, the robot meets the worker at speed. The worker enters its path 0.57 m ahead of the bumper while the robot is doing 1.43 m/s and needs 0.86 m to stop, the emergency stop fires at 10.18 s, and the robot hits the worker at 0.72 m/s at 10.81 s. The collision is scored as avoidable, because the worker came out of space the robot could not see. Slowing for unseen space, the robot creeps around the corner, never comes within 0.92 m of the worker, and delivers in 31.6 s, the same time as the empty route. The result depends on timing in one direction only. Moving the worker's start between 2.0 and 8.0 s in 0.1 s steps, the reckless policy collides or violates the clearance at 18 of the 61 starts, all between 3.8 and 5.5 s, and passes cleanly at the rest, while the cautious policy keeps at least 0.52 m at all 61 starts and never takes longer than 34.3 s.

> **TRY IT IN THE APP** **5 · Occluded Intersection**
>
> Load the case and press Run. Unseen space is set to Treat as clear, so the robot holds its speed toward the cross-aisle. Watch the green stopping band ahead of it as the worker steps out from behind the rack end. The emergency stop fires at 10.18 s, and the robot still hits the worker at 0.72 m/s.

In case 6 a cart stands in an empty floor bay of rack A2, a bay the map shows as racking, and is shoved into the aisle as the robot passes. The supervisor starts braking 0.25 s after the shove, as soon as the cart stands clear of the rack face and becomes something the map does not know beside the route, and the emergency stop follows 0.2 s later. The cart still enters the path 0.40 m ahead while the robot is doing 1.17 m/s, and the true stopping distance of Equation (9.14) is then 0.59 m. The robot hits it at 0.71 m/s and the referee records an unavoidable collision. Treating unseen space as clear, the robot brakes only when the cart is in its path and hits it at 1.24 m/s, again unavoidable because the cart entered 0.40 m ahead of a robot that needed 0.93 m. At a cruise speed of 0.8 m/s the cart appears at the same place, the robot stops 0.52 m short, replans around it and delivers in 74.1 s.

Pause. Try Exercise 9.3 before reading on.

9.5.5 A Map That Is Out of Date

In case 7 a pallet left in aisle AB stops the robot at the following distance. After waiting 3 s it marks the pallet's face on its own map and replans at 22.6 s onto a route 11.4 m longer, back to the cross-aisle, along aisle BC and down to Station 1, and it delivers in 57.5 s against 31.6 s on the clear route. With replanning off it waits behind the pallet, and 100 s after the start it is still there. The waiting time trades a detour around a person who would have moved on against idling behind a pallet that will not.

9.5.6 The Graded Shift

The challenge combines a 600 kg pallet, tires 2 % larger than the controller believes, the blind-corner worker, a cart parked just past the wet patch and a pallet blocking aisle BC, with deliveries to Stations 1, 2 and 4. Its native completion target is evaluated when the mission finishes or a collision occurs, with no fixed deadline. It requires all three handovers within 0.3 m, no missed delivery, no collision, no clearance violation and a worst moving clearance of at least 0.2 m. The model parameters and disturbance settings must retain the scenario values during the attempt. As shipped, with empty-robot braking, unseen space treated as clear, encoder-only steering and the wet zone unmapped, the robot hits the worker at 1.47 m/s 11.1 s into the shift. With the four assumptions corrected and every gain at its default, it delivers at 39.0 s, 51.1 s and 122.6 s after the start, 2 to 3 cm from each mark, replans around the pallet at 81.0 s onto a route 14.8 m longer, and uses 48.4 kJ. Its closest approach, while moving, is 0.27 m to the cart as it stops on the wet patch at 22.8 s. The moment of the shipped collision is shown for the corrected configuration in Figure 9.5.

Each correction is necessary on its own, as a verification check confirms by restoring one wrong assumption at a time. Treating unseen space as clear hits the worker at 1.47 m/s. Empty-robot braking delivers twice and then hits the aisle-BC pallet at 0.25 m/s. An unmapped wet patch delivers all three pallets but stops 0.11 m from the cart, a clearance violation. Steering by the encoders misses Station 1 by 0.34 m, delivers to Station 2 0.14 m off, and has not finished after 170 s. The suggested configuration passed at all 61 worker start times between 2.0 and 8.0 s, with a smallest clearance of 0.27 m. The longest mission in that sweep took 151.5 s after leaving the dock, demonstrating why the completion rule is distinct from the 150 s verification-run horizon. Applying the suggestion before Run or applying it and then pressing Reset gives identical default-seed results. The scorecard scores each constraint as its worst case over the run, so a clean finish cannot hide an earlier violation.

9.6 Technology in Practice

The robot of Section 9.3 reads wheel encoders, a gyro, a load cell and one range scanner, and it moves on two driven wheels. Figure 9.6 draws its loops with real devices in the place of the schematic's symbols.

What Measures It

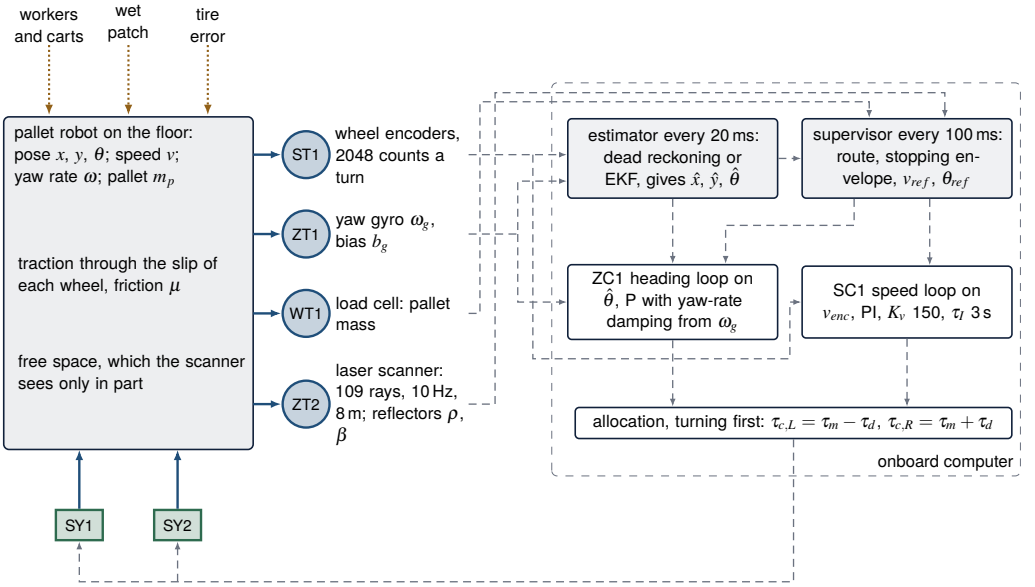
Autonomous mobile robots carry laser scanners, cameras, accelerometers, gyroscopes and wheel encoders. The laser scanner gives a precise cloud of distances around the robot, and a camera adds a wide view in which obstacles can be recognized [18]. The scanner of Equation (9.6), 109 rays over 270° at 10 Hz out to 8 m, stands for such a laser scanner. In this chapter the same scanner finds obstacles and sights the surveyed reflectors, as the navigation sensor of a NIST test vehicle measures the distance and angle to fixed reflective markers and computes the vehicle's position and heading from them [19]. The missed returns and phantom points of a real scanner lie outside the model.

> **TECHNOLOGY** **Sensor: Laser Scanner**
>
> A laser scanner times a pulse of light out to a surface and back. The range is the flight time times the speed of light, halved, so one nanosecond of timing is 150 mm of range, and a mirror sweeps the single beam across the scene. Its accuracy depends on the reflectivity, color and gloss of the target and on the angle at which the beam meets it. A beam that falls on two surfaces at once can report a phantom point between them, and a matte black surface can return nothing [20]. A transparent curtain is usually seen as an obstacle that stops the vehicle [21]. On a trend a missed return is a clear distance that stays long while something stands in the path, and a curtain is a stop with nothing there.

What Moves It

Each wheel of Equation (9.2), a torque command through a 0.05 s lag up to 40 N m, stands for an electric motor and gearbox on a driven wheel, in the two-wheel layout with free supporting wheels that most guided vehicles use [18]. The traction curve of Equation (9.3) carries the slip that a real wheel shows on a slippery floor or a metal plate, and the payload of cases 2 and 9 the load that slows a real vehicle's acceleration [21]. The wheel-radius error of case 4 stands for a diameter that wear changes and that is never known exactly [19].

> **TECHNOLOGY** **Final Element: Driven Wheels**
>
> Most guided vehicles have either one steerable traction wheel in front or two independently driven wheels with free-turning supporting wheels, a simple layout of low cost [18]. An encoder on each driven wheel converts revolutions into distance through a scale, and the wheel's diameter is never known exactly and changes slightly with wear [19]. In NIST tests on ramps a heavy payload slowed acceleration and cost grip: the vehicle slipped on a metal floor transition, would not run faster than 500 mm/s with 54 kg aboard, and its controller reported the travel complete before the vehicle reached the goal [21]. On a trend, slip is distance that the encoders count and the vehicle never travels.



SY1, SY2: left and right wheel drives, an electric motor and gearbox on each driven wheel, 40 N m each through a 0.05 s torque lag; casters carry the rest of the weight. The emergency stop brakes both flat out.

Figure 9.6 The loops of the chapter with real devices in place of the schematic's symbols. ST1 counts the turns of the wheels, ZT1 is the yaw gyro, WT1 the load cell under the pallet and ZT2 the laser scanner, which both finds obstacles and sights the reflectors. In the onboard computer (dashed frame) the estimator feeds the supervisor and the heading loop ZC1, the supervisor sets the setpoints of ZC1 and of the speed loop SC1, and their torques are shared between the wheel drives SY1 and SY2. Notice that the speed limit rests on three beliefs that no sensor checks: the estimated pose, the braking computed from the load cell, and the free space beyond what the scanner sees.

Where the Controller Runs

The wheel loops and the estimators of this chapter run every 20 ms, and the scanner, the navigation and the supervisor every 100 ms. In industrial control, supervisory and lower-level loops run with cycle times from milliseconds to minutes [22], and the robot's two layers sit at its fast end. A guided vehicle typically has its path planned and its motion controlled off board, while a mobile robot typically does both on board [23]. An autonomous robot can also negotiate with machines and other systems itself, where a guided-vehicle system leaves scheduling, routing and dispatching to a central unit [18]. The structure of this chapter lives in the robot's own computer: the speed and heading loops are two velocity-form PID blocks of the fast layer, and the estimator and the stopping envelope sit in the layer above them.

What Is Modernizing It

Guidance has moved through mechanical, optical, inductive, inertial and laser stages to vision-based systems that map the facility as they go and need no reference points installed in advance. Lithium-ion batteries and wireless charging relieve the limited capacity and long charging times that held back earlier vehicles [18]. Performance standards for autonomous industrial vehicles, written since 2014, now cover navigation, obstacle detection and environmental conditions, and docking and position measurement are being added [21]. A robot that maps as it goes must also match each return to a landmark, a problem that the identified reflectors of this chapter leave out (Section 9.3.7).

IN PRACTICE **Calibrating the Wheel Scale**

A NIST guideline calibrates a guided vehicle by driving it along a straight line, measuring the distance actually traveled, and comparing it with the distance the control system reports for each wheel. The new encoder scale is the old one times the ratio of the two distances. Steering offsets and the mounting angle of the navigation sensor are calibrated as well, and correcting that angle cut the lateral error by more than an order of magnitude. Before calibration the test vehicle stopped up to about 25 mm apart when sent to one point from opposite directions, although its navigation could repeat a position to 4 mm [19]. The tire error of case 4 is what the scale corrects.

9.7 Limitations

The robot never skids sideways and its traction has no peak, so lateral sliding on the wet patch and any benefit of modulating the slip at the friction limit lie outside the model [17]. The smoothed caster scrub also overdamps a turn in place near rest.

The controller knows its load and, when the zone is mapped, the wet friction exactly, and its supervisor ignores the lag of the speed loop, which uses 0.25 m of the 0.5 m following distance at the default gain. A deployed robot has to estimate both quantities and budget for the lag, so the margins measured here are the best case.

Perception is idealized, with a single scan plane, reflectors that carry their identities, and a map that is exact apart from the objects a scenario adds.

The positive creep floor can carry the robot beyond the strict visible-space stopping envelope. The timing sweeps report success for their tested scripts, not a safety guarantee for arbitrary traffic. The traffic is scripted and never yields, and each shipped case uses one timing. The reckless policy fails case 5 only for worker starts between 3.8 and 5.5 s, so a single run can show that a policy is unsafe but never that it is safe. One robot on a known map, with no fleet traffic and no mapping of its own, completes the scope.

OPEN CHALLENGE **Capabilities That Must Hold Together**

Performance standards for autonomous industrial vehicles have been written since 2014 and cover navigation, obstacle detection and environmental conditions. NIST experiments found gaps where capabilities combine. A vehicle that meets a transparent curtain usually takes it for an obstacle and stops, so its safety sensors may be switched off or ignored to let it pass, a risk if a person stands behind. A payload that the vehicle carries on a level floor can defeat it on a ramp [21]. Progress would be tests that rate a vehicle with its load, its floor and its obstacles together, as the graded shift of Section 9.5 combines them.

9.8 Conclusion

A warehouse robot's safety rests on what it believes about its position, its braking and the space around it, and the simulator measures what each belief costs when it is wrong. Planning a 1000 kg pallet's stops with the empty robot's 2.26 m/s^2 of braking, where the load leaves 0.76 m/s^2, hits a parked cart at 0.74 m/s despite an emergency stop, while payload-aware braking stops 0.25 m short. An unmapped wet patch leaves 0.07 m to a cart, a mapped one 0.41 m. Encoders with a 2.5 % tire error hand the pallet over 0.44 m past the station with the truth outside the robot's own ellipse, and the fused estimate hands it over 2 cm from the mark. Treating a blind corner as clear hits a worker at 0.72 m/s, and slowing for unseen space keeps 0.92 m at a cost of 8.7 s. A cart shoved out 0.40 m ahead of a robot that needs 0.59 m is scored as unavoidable. The graded shift fails at 11.1 s as shipped and passes with four corrected assumptions, three deliveries in 120.6 s and 0.27 m of minimum clearance. The model omits lateral skid, a friction peak, perception errors and people who yield, and a fuller model would add an estimated load and friction, reflector association errors and a second robot. The simulator makes the assumptions behind a mobile robot's speed visible, measurable and open to correction.

KEY POINTS

- A warehouse robot picks its speed from beliefs about its position, its braking and the space it can see, and the simulator scores each run against the truth.
- Planning a 1000 kg pallet's stops with the empty robot's braking hits a parked cart at 0.74 m/s despite an emergency stop, while payload-aware braking stops 0.25 m short.
- With a 2.5 % tire error the encoders hand the pallet over 0.44 m past the station, and the fused estimate 2 cm from the mark.
- Treating a blind corner as clear hits a worker at 0.72 m/s, and slowing for unseen space keeps 0.92 m at a cost of 8.7 s.

Where else this applies. Chapter 15 limits a machine's speed beside a person who walks in, and Volume 5, Chapter 13 steers on a filter whose drift nothing aboard can see.

References

[1] I. F. A. Vis, "Survey of research in the design and control of automated guided vehicle systems," *European Journal of Operational Research*, vol. 170, no. 3, pp. 677–709, 2006.

[2] R. D'Andrea and P. Wurman, "Future challenges of coordinating hundreds of autonomous vehicles in distribution facilities," in *2008 IEEE International Conference on Technologies for Practical Robot Applications*, 2008, pp. 80–83.

[3] K. Azadeh, R. De Koster, and D. Roy, "Robotized and automated warehouse systems: Review and recent developments," *Transportation Science*, vol. 53, no. 4, pp. 917–945, 2019.

[4] *Industrial trucks - Safety requirements and verification - Part 4: Driverless industrial trucks and their systems*, BSI British Standards Standard BS EN ISO 3691-4:2020, Jun. 2020.

[5] D. Fox, W. Burgard, and S. Thrun, "The dynamic window approach to collision avoidance," *IEEE Robotics & Automation Magazine*, vol. 4, no. 1, pp. 23–33, 1997.

[6] T. Fraichard and H. Asama, "Inevitable collision states—a step towards safer robots?" *Advanced Robotics*, vol. 18, no. 10, pp. 1001–1024, 2004.

[7] S. Bouraine, T. Fraichard, and H. Salhi, "Provably safe navigation for mobile robots with limited field-of-views in dynamic environments," *Autonomous Robots*, vol. 32, no. 3, pp. 267–283, 2012.

[8] A. D. Ames, X. Xu, J. W. Grizzle, and P. Tabuada, "Control barrier function based quadratic programs for safety critical systems," *IEEE Transactions on Automatic Control*, vol. 62, no. 8, pp. 3861–3876, 2017.

[9] J. J. Leonard and H. F. Durrant-Whyte, "Mobile robot localization by tracking geometric beacons," *IEEE Transactions on Robotics and Automation*, vol. 7, no. 3, pp. 376–382, 1991.

[10] J. Borenstein and L. Feng, "Measurement and correction of systematic odometry errors in mobile robots," *IEEE Transactions on Robotics and Automation*, vol. 12, no. 6, pp. 869–880, 1996.

[11] G. Ullrich, *Automated Guided Vehicle Systems: A Primer with Practical Applications*. Berlin, Heidelberg: Springer, 2015.

[12] H. F. Durrant-Whyte, "An autonomous guided vehicle for cargo handling applications," *The International Journal of Robotics Research*, vol. 15, no. 5, pp. 407–440, 1996.

[13] R. J. Norcross, R. V. Bostelman, and J. A. Falco, "Automated guided vehicle bumper test method development," National Institute of Standards and Technology, Gaithersburg, MD, NIST Interagency/Internal Report 8029, 2015.

[14] P. Hart, N. Nilsson, and B. Raphael, "A formal basis for the heuristic determination of minimum cost paths," *IEEE Transactions on Systems Science and Cybernetics*, vol. 4, no. 2, pp. 100–107, 1968.

[15] J. Wit, C. D. Crane, III, and D. Armstrong, "Autonomous ground vehicle path tracking," *Journal of Robotic Systems*, vol. 21, no. 8, pp. 439–449, 2004.

[16] G. Campion, G. Bastin, and B. d'Andréa-Novel, "Structural properties and classification of kinematic and dynamic models of wheeled mobile robots," *IEEE Transactions on Robotics and Automation*, vol. 12, no. 1, pp. 47–62, 1996.

[17] H. B. Pacejka, "Tire characteristics and vehicle handling and stability," in *Tire and Vehicle Dynamics*. Elsevier, 2012, pp. 1–58.

[18] G. Fragapane, R. de Koster, F. Sgarbossa, and J. O. Strandhagen, "Planning and control of autonomous mobile robots for intralogistics: Literature review and research agenda," *European Journal of Operational Research*, vol. 294, no. 2, pp. 405–426, 2021.

[19] S. Legowik, R. V. Bostelman, T. Hong, and E. R. Messina, "Guideline for automatic guided vehicle calibration," National Institute of Standards and Technology, Gaithersburg, MD, NIST Interagency/Internal Report 8168, 2017.

[20] W. C. Stone, M. Juberts, N. Dagalakis, J. Stone, and J. Gorman, "Performance analysis of next-generation LADAR for manufacturing, construction, and mobility," National Institute of Standards and Technology, Gaithersburg, MD, NIST Interagency/Internal Report 7117, 2004.

[21] S. Yoon, R. Bostelman, and A. Virts, "Experiments to test the A-UGV capabilities standard," National Institute of Standards and Technology, Gaithersburg, MD, NIST Interagency/Internal Report 8407, 2021.

[22] K. Stouffer, M. Pease, C. Tang, T. Zimmerman, V. Pillitteri, S. Lightman, A. Hahn, S. Saravia, A. Sherule, and M. Thompson, "Guide to operational technology (OT) security," National Institute of Standards and Technology, Gaithersburg, MD, Special Publication 800-82r3, 2023.

[23] R. V. Bostelman and T. H. Hong, "Review of research for docking automatic guided vehicles and mobile robots," National Institute of Standards and Technology, Gaithersburg, MD, NIST Interagency/Internal Report 8140, 2016.

9.9 Exercises

9.1 ○ Concept check · 2 min

In scenario 2, Heavier pallet, as shipped, the robot carries a 1000 kg pallet but plans its stops with the braking of the empty robot. Its emergency stop fires at 18.22 s, and it still hits the parked cart at 0.74 m/s. Why does the emergency stop not prevent the collision?

(a) The emergency stop removes the heading torque only, so the mean wheel torque keeps pushing the robot on

(b) The emergency stop releases once the speed falls below its arming speed, well before the robot reaches the cart

(c) The wheels lock on the floor, the encoders read zero, and the stop releases while the robot still slides

(d) The stop fires on the believed braking, so it fires late, and flat-out braking with the load is only 0.76 m/s^2

9.2 ○ Concept check · 1 min

In scenario 4, Wheel-radius error, the robot steers by its encoders alone while its tires are 2.5 % larger than the controller believes. At the handover its 95 % ellipse is 1.20 m across the aisle and 0.40 m along it, and the true position lies outside it. Which statement explains this?

(a) The covariance models random drift in distance, while a radius error adds the same bias to every meter

(b) The ellipse covers 95 % of outcomes, so the truth falls outside it on one run in twenty by chance

(c) The gyro bias turns the estimate across the aisle, and the encoder model leaves the gyro out

(d) The reflector updates pull the estimate the wrong way because they use the believed wheel radius

9.3 ○ Concept check · 2 min

The referee scores the cart collision of scenario 6, Unexpected cart, as unavoidable. It scores the worker collision of scenario 5, Occluded intersection, with unseen space treated as clear, as avoidable. What sets the difference?

(a) The referee forgives contact with an object but never contact with a person

(b) The cart gave no warning and entered closer than the true stopping distance, while the worker came out of space the robot could not see

(c) The robot hit the cart below the speed at which the referee counts contact as a collision

(d) The emergency stop fired before the cart collision and not before the worker collision

9.4 Theory and calculation · 6 min

The robot has a torque limit $\tau_{max} = 40\,\mathrm{N\,m}$ per wheel, wheel radius $r = 0.10\,\mathrm{m}$, wheel and motor inertia $J_w = 0.4\,\mathrm{kg\,m^2}$, empty mass $m_r = 300\,\mathrm{kg}$, rolling resistance $C_{rr} = 0.02$ and $\chi = 0.6$ of the weight on the drive wheels. Use $g = 9.81\,\mathrm{m/s^2}$ and Equation (9.4).

(a) Compute the braking capability on the dry floor ($\mu = 0.7$) with no pallet and with a 1000 kg pallet. Which limit binds? *Checkpoint:* $2.26\,\mathrm{m/s^2}$ and $0.76\,\mathrm{m/s^2}$, both motor-limited

(b) From 1.5 m/s, compute the braking distance and the stopping band $v\,t_r + v^2/(2a_b)$ with $t_r = 0.3\,\mathrm{s}$, for the loaded robot with the empty-robot belief and with its true braking. *Checkpoint:* 0.50 m and 1.47 m; bands 0.95 m and 1.92 m

(c) On the wet patch ($\mu = 0.1$), compute the friction limit and the pallet mass above which the wet floor no longer shortens the stop. *Checkpoint:* $0.785\,\mathrm{m/s^2}$, 952 kg

9.5 Theory and calculation · 6 min

Use the envelope speed of Equation (9.10) with $t_r = 0.3\,\mathrm{s}$ and the braking values of the previous exercise ($2.26\,\mathrm{m/s^2}$ empty, $0.76\,\mathrm{m/s^2}$ with a 1000 kg pallet, $1.35\,\mathrm{m/s^2}$ with the default 300 kg pallet).

(a) A robot carrying 1000 kg has a clear distance $d = 1.5\,\mathrm{m}$. Compute $V(d)$ with the empty-robot belief and with the true braking. Which speed does the robot drive at with a cruise speed of 1.5 m/s and the empty-robot belief? *Checkpoint:* 2.01 m/s and 1.30 m/s; it drives at 1.5 m/s

(b) At that speed, how far beyond the 1.5 m does the robot need to stop with its true braking? *Checkpoint:* 0.42 m

(c) With the default pallet, how far does the creep speed $v_{cr} = 0.3\,\mathrm{m/s}$ need to stop? Explain why the term $\max[v_{cr}, V(d_{occ})]$ of Equation (9.11) is not a guarantee of stopping inside visible space. *Checkpoint:* 0.12 m

9.6 Theory and calculation · 5 min

The heading loop applies $K_\theta b/r = 600\,\mathrm{N\,m}$ of yaw torque per radian and $K_r b/r = 360\,\mathrm{N\,m\,s}$ per radian per second of damping. The yaw inertia of the robot is $I = m_r(0.35\,\mathrm{m})^2 + m_p(1.2^2 + 0.8^2)/12$, with $m_r = 300\,\mathrm{kg}$, plus the wheel inertia reflected to yaw, $2J_w(b/2)^2/r^2$ with $J_w = 0.4\,\mathrm{kg\,m^2}$, $b = 0.60\,\mathrm{m}$ and $r = 0.10\,\mathrm{m}$.

(a) Compute the reflected inertia and the natural frequency and damping ratio of the rigid-body heading loop with the default 300 kg pallet. *Checkpoint:* $7.2\,\mathrm{kg\,m^2}$; 2.50 rad/s, 0.75

(b) Repeat for a 1000 kg pallet, and give the overshoot a step would show for that damping ratio. *Checkpoint:* 1.66 rad/s, 0.50; about 16 %

(c) The measured 20° step with 1000 kg shows no overshoot, and 18.3 % once the caster scrub is removed. What does this say about the rigid-body estimate?

9.7 App challenge · 8 min

Load scenario 2, Heavier pallet, as shipped (braking model Empty robot, 1000 kg pallet), default seed, at speed 2.

(a) Press Run and let the run stop itself. Read the collision note: the time and the speed of the impact. Read the E-stops tile. *Checkpoint:* 18.6 s, 0.74 m/s (within 0.01 m/s), 1 E-stop

(b) Press Reset, set Braking model to Payload-aware before Run, and run until the run stops itself. Read the Min clearance and Mission time tiles. *Checkpoint:* 0.25 m (within 0.01 m), 40.3 s (within 0.1 s)

(c) Use the bands of the theory exercise on braking to explain why the first run hits the cart even though the second leaves a gap.

9.8 App challenge · 8 min

Load scenario 5, Occluded intersection, as shipped (Unseen space Treat as clear), default seed, at speed 2.

(a) Press Run and let the run stop itself. Read the collision note and say how the referee records the collision. *Checkpoint:* 0.72 m/s at 10.8 s, avoidable

(b) Press Reset, set Unseen space to Slow for it before Run, and run until the run stops itself. Read the Min clearance and Mission time tiles. *Checkpoint:* 0.92 m (within 0.01 m), 31.6 s (within 0.1 s)

(c) The worker entered the path 0.57 m ahead of the bumper while the robot did 1.43 m/s. Use $t_r = 0.3$ s and the default braking of 1.35 m/s^2 to estimate the stopping band at that speed, and compare it with the 0.86 m the chapter reports for the true stop. *Checkpoint:* about 1.19 m

9.9 App challenge · 8 min

Load the challenge, Challenge · a mixed shift, default seed. Change no model parameter or disturbance. Use speed 5 for the long runs.

(a) Press Run as shipped and let the run stop itself. Read the collision note. *Checkpoint:* worker hit at 1.47 m/s at 11.1 s

(b) Press Reset. Before Run, set Braking model to Payload-aware, Unseen space to Slow for it, Estimator to Fused (EKF) and Wet-floor zone to Mapped. Run until the run stops itself and read the Deliveries, Min clearance, Mission time and Energy tiles. *Checkpoint:* 3 of 3, 0.27 m, 120.6 s, 48.4 kJ

(c) Press Reset, set Wet-floor zone back to Not mapped, and run again. Read the same tiles and the Violations tile. Does the attempt pass? *Checkpoint:* 3 of 3, 0.11 m, 1 violation, fails

9.10 Application area · 3 min

Scenario 1, Clear route, delivers to Station 1 in 31.6 s using 10.4 kJ when the robot slows for unseen space, and in 22.9 s using 7.9 kJ when it treats unseen space as clear.

(a) By what percentage does caution raise the cycle time and the energy of this delivery? *Checkpoint:* 38 % and 32 %

(b) A site manager asks to set the fleet to treat unseen space as clear to recover the time. Name two changes to the site or the sensing that recover time at the two blind corners without that risk.

9.11 Application area · 4 min

The controller reads the pallet mass exactly from its load cell. On a real robot the load cell can read a 1000 kg pallet as 900 kg. Use Equation (9.4) with the parameters of the theory exercise on braking.

(a) Compute the braking the robot then believes and the braking distance it plans from 1.5 m/s, against the true 1.47 m. *Checkpoint:* 0.81 m/s^2, 1.39 m, 0.08 m short

(b) Which direction of load-cell error is unsafe, and how should the supervisor budget for it?

9.12 Review the agent · 5 min

An AI assistant reviews scenario 2, Heavier pallet, with the payload-aware braking model. It writes: "The robot stops 0.25 m short of the cart although the following distance is 0.5 m. The braking model therefore overestimates the braking by about a factor of two. Halve the assumed braking capability to restore the full following distance." Find the flaw, and say how you would check your answer in the app.