INTERACTIVE CONTROL ENGINEERING

Motion Control, Robotics, and Autonomous Systems

Interactive Control Engineering: Principles and Applications

Edited by John D. Hedengren

Control technology

1 Process Dynamics, Modeling, and Identification
2 Practical PID Control
3 Feedforward, Cascade, and Ratio Control
4 Override, Split-Range, and Constraint Control
5 Delays, Slow Measurements, and State Estimation
6 Multivariable and Plantwide Control
7 Nonlinear Dynamics, Oscillations, and Stabilization
8 Batch, Run-to-Run, and Hybrid Control
9 Model Predictive Control and Economic Optimization

Industry applications

10 Control in Upstream and Midstream Oil and Gas
11 Control in Petroleum Refining
12 Control in Petrochemicals and Polymers
13 Control in Chemical Plants
14 Control in Semiconductor Manufacturing
15 Control in Mining and Mineral Processing
16 Control in Iron, Steel, and Metals Production
17 Control of Energy Conversion and Storage
18 Control of Nuclear and Fusion Systems
19 Control of Biological Systems and Human Performance
20 Motion Control, Robotics, and Autonomous Systems

INTERACTIVE CONTROL ENGINEERING
Principles and Applications

Volume 20

Motion Control, Robotics, and Autonomous Systems

Tracking, navigation, contact and allocation

John D. Hedengren
Editor

Advanced Process Solutions, LLC

Interactive Control Engineering: Principles and Applications
Volume 20: Motion Control, Robotics, and Autonomous Systems
Edited by John D. Hedengren

First edition, 2026

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

Produced by Advanced Process Solutions, LLC.

No part of this publication may be reproduced, stored in a retrieval system, or transmitted in any form or by any means without the prior written permission of the publisher, except for brief quotations in a review or scholarly work.

The simulators described in this book are teaching models. Their equations, parameters and results are simplified for instruction and are not validated for the design, operation or safety assessment of any real plant, vehicle or medical device. The editor and the publisher accept no liability for use of this material outside that purpose.

Names of products and companies are used for identification only and may be trademarks of their owners.

Every chapter has a companion simulator that runs in a web browser at https://apmonitor.com/apps/. This volume online, with chapter files and citation records: https://apmonitor.com/apps/books/v20-motion-robotics/

How to cite this volume. John D. Hedengren, Ed., *Motion Control, Robotics, and Autonomous Systems*, vol. 20 of *Interactive Control Engineering: Principles and Applications*. Advanced Process Solutions, LLC, 2026.

Typeset in Times and Helvetica-family type with LaTeX and the memoir class.

Preface

This book is one of 20 volumes of *Interactive Control Engineering: Principles and Applications*. Each chapter is about one plant, machine or system and one control lesson, and each has a simulator that runs in a web browser with nothing to install. The simulators are collected at https://apmonitor.com/apps/. Nine volumes are organized by control technology and eleven by industry or application area; the last pages of this volume list them all.

How to read this book with the apps

Chapter 1 teaches the theory this volume relies on and explains, once, the simulator framework that the apps share. Each chapter after it opens with the control question it answers, its learning objectives, what to read first, and the address of its app, also given as a code to scan. The chapter then develops the model, the controllers and the measured results. Every number reported was measured on the app, so a reader can open the app, load the named scenario and reproduce it.

Every chapter ends with exercises. Each is labeled with its kind and the time it should take, and none takes more than ten minutes. Concept checks and calculations can be done on paper. App challenges name a scenario, the settings and the run length. One exercise in each chapter, labeled *Review the agent*, presents a tuning, a model or a claim proposed by an AI assistant and asks what is wrong with it. Short answers to the odd-numbered exercises are at the back of the book, and worked solutions are available to instructors.

How AI agents were used

AI agents did much of the development work behind this series. Under the editor's direction they wrote and tested the simulators, drafted the papers from which the chapters are drawn, measured the reported results on the running apps, converted the papers to book form, and drafted the introductory and closing chapters and the exercises. Other agents then checked that work: they re-measured reported numbers, audited what each app displays against the state of its model, reviewed the papers as referees, and solved the exercises without seeing the answers. The editor set the scope, chose the plants and the lessons, reviewed the results and is responsible for the content.

That way of working is also a subject of the book. The last chapter of each volume describes what the agents got wrong while this volume was made and how the errors were caught, and it sets out what the editor believes this means for engineers: more of the routine work can be delegated, and more of the engineer's value lies in checking, supervising, choosing the problem and knowing the process.

About this volume

This volume is about machines that move. Chapters 2 to 8 cover vehicles on the road, in the air, on the water and in orbit. Chapters 9 to 15 cover robots that pick, play, weld, assemble, operate and work beside people. The chapters share four tasks: following a path, knowing where the machine is, allocating effort among actuators, and making contact safely.

Volume 7 covers the stabilization that these systems assume. Each chapter names the technology volume that teaches its main technique.

John D. Hedengren
2026

Contents

Preface v

Notation x

1 Introduction: Controlling Machines That Move 1
- 1.1 The Question This Volume Answers . . . 1
- 1.2 A Short History . . . 3
- 1.3 Notation and Conventions . . . 4
- 1.4 Motion Under Uncertainty . . . 4
- 1.5 Instruments and Final Elements in This Volume . . . 13
- 1.6 The Simulator Framework . . . 15
- 1.7 Working With AI Agents While Studying . . . 19
- 1.8 Chapter Roadmap . . . 19
- 1.9 Exercises . . . 22

2 Lane Keeping on a Kinematic Bicycle: Speed Dependence of a Stanley Path-Following Loop 24
- 2.1 Introduction . . . 25
- 2.2 Background . . . 26
- 2.3 Vehicle and Road Model . . . 28
- 2.4 Dynamic Analysis . . . 32
- 2.5 Controller Design . . . 35
- 2.6 Results . . . 38
- 2.7 Technology in Practice . . . 46
- 2.8 Limitations . . . 48
- 2.9 Conclusion . . . 48
- 2.10 Exercises . . . 50

3 Drone Taxi Simulator for Load Feedforward and Wind Estimation Between Buildings 54
- 3.1 Introduction . . . 55
- 3.2 Hover, City Wind and the Passenger Multirotor . . . 56
- 3.3 A Planar Air Taxi Among Rooftops . . . 57
- 3.4 Three Loops and Three Levels of Knowledge . . . 62
- 3.5 What a Loop Must Be Told . . . 66
- 3.6 Scenarios and Measured Results . . . 67
- 3.7 Technology in Practice . . . 72
- 3.8 Limitations . . . 74
- 3.9 Conclusion . . . 74
- 3.10 Exercises . . . 76

4 Lift-Plus-Cruise VTOL Simulator for Control Allocation and Measured Transition Gates 79
- 4.1 Introduction . . . 80
- 4.2 Transition Flight From Lift Engines to Electric Rotors . . . 81
- 4.3 A Planar Lift-Plus-Cruise Aircraft . . . 82
- 4.4 Where the Wing Can Take the Weight . . . 87
- 4.5 Loops, Allocation and Supervisor . . . 88
- 4.6 Measured Cases . . . 92
- 4.7 Technology in Practice . . . 97

4.8 Limitations . . . 99
4.9 Conclusion . . . 99
4.10 Exercises . . . 101

5 Flying a Foiling Boat Between a Slamming Hull and a Ventilating Foil 104
5.1 Introduction . . . 105
5.2 Hydrofoils, Their Surface and Their Controllers . . . 106
5.3 The Boat and the Sea as Modeled . . . 107
5.4 Dynamic Character of the Flying Boat . . . 112
5.5 Implemented Control System . . . 113
5.6 Eleven Cases, Measured . . . 116
5.7 Technology in Practice . . . 122
5.8 Limitations . . . 124
5.9 Conclusion . . . 124
5.10 Exercises . . . 126

6 Choosing What Not to Reject in a Dynamic Positioning Simulator 130
6.1 Introduction . . . 131
6.2 From Anchors to a Kalman Filter . . . 133
6.3 Vessel, Environment and Thruster Model . . . 133
6.4 Estimator, Controllers and Allocation . . . 139
6.5 Scenarios . . . 145
6.6 Results . . . 145
6.7 Technology in Practice . . . 151
6.8 Limitations . . . 153
6.9 Discussion . . . 153
6.10 Conclusion . . . 154
6.11 Exercises . . . 155

7 CubeSat Attitude Simulator Whose Wheels Keep Every Torque They Absorb 158
7.1 Introduction . . . 159
7.2 From Spinning Satellites to Three-Axis CubeSats . . . 160
7.3 The Spacecraft, Its Orbit and Its Instruments . . . 162
7.4 What the Inertia and the Field Permit . . . 166
7.5 The Flight Software as Written . . . 167
7.6 Ten Cases and a Five-Hour Challenge, Measured . . . 172
7.7 Technology in Practice . . . 179
7.8 Limitations . . . 181
7.9 Conclusion . . . 182
7.10 Exercises . . . 183

8 A Low Earth Orbit Simulator Flown on Mean Elements and Collision Probability 187
8.1 Introduction . . . 188
8.2 Drag, Mean Orbits and Conjunction Assessment . . . 189
8.3 Satellite, Atmosphere and Debris . . . 190
8.4 The Orbit the Planners See . . . 194
8.5 Navigation, Screening and Collision Probability . . . 195
8.6 Implemented Planners and Lower Layer . . . 196
8.7 Cases and Measured Outcomes . . . 199
8.8 Technology in Practice . . . 203
8.9 Limitations . . . 206
8.10 Conclusion . . . 206

8.11 Exercises . . . 208

9 Warehouse Robot Simulator for Stopping Distance, Blind Corners and Localization Error 211
9.1 Introduction . . . 212
9.2 From Guide Wires to Reflectors and Robot Fleets . . . 213
9.3 A Pallet Robot on a Mapped Floor . . . 214
9.4 A Controller That Reads Only Measurements . . . 219
9.5 Eight Broken Assumptions and a Graded Shift . . . 223
9.6 Technology in Practice . . . 228
9.7 Limitations . . . 230
9.8 Conclusion . . . 230
9.9 Exercises . . . 232

10 Robot Conveyor Picking Simulator for Late Camera Reports and Moving-Target Interception 235
10.1 Introduction . . . 236
10.2 Picking Products From a Moving Belt . . . 237
10.3 The Picking Cell . . . 238
10.4 Four Budgets for One Grasp . . . 242
10.5 Implemented Controller . . . 244
10.6 Scenarios and Measured Results . . . 247
10.7 Technology in Practice . . . 251
10.8 Limitations . . . 253
10.9 Conclusion . . . 253
10.10 Exercises . . . 255

11 A Table Tennis Robot That Must Swing Before It Knows Where 258
11.1 Introduction . . . 259
11.2 Robots at the Table . . . 260
11.3 The Table and Two Arms as Modeled . . . 261
11.4 Half a Second, Spent Twice . . . 265
11.5 The Pipeline as Coded . . . 266
11.6 Nine Cases and a Game, Measured . . . 269
11.7 Technology in Practice . . . 273
11.8 Limitations . . . 275
11.9 Conclusion . . . 276
11.10 Exercises . . . 277

12 Self-Regulating Arc and Distance-Keyed Seam Preview in a Robotic Welding Simulator 280
12.1 Introduction . . . 281
12.2 Arc Sensing, Seam Vision and the Adaptive Cell . . . 282
12.3 The Cell as Modeled . . . 283
12.4 The Arc's Own Loop . . . 289
12.5 Implemented Controller . . . 290
12.6 Measured Cases . . . 293
12.7 Technology in Practice . . . 298
12.8 Limitations . . . 300
12.9 Conclusion . . . 300
12.10 Exercises . . . 302

13 Robotic Peg-in-Hole Simulator for Compliance, Jamming and Force Control Through Contact 306

13.1 Introduction . . . 307
13.2 Force Control and Compliant Assembly . . . 308
13.3 Peg, Part and Contact Model . . . 309
13.4 Three Controllers on One Wrench Law . . . 312
13.5 What Sets the Contact Force . . . 315
13.6 Scenarios and Measured Results . . . 319
13.7 Technology in Practice . . . 322
13.8 Limitations . . . 324
13.9 Conclusion . . . 324
13.10 Exercises . . . 326

14 Surgical Robot Simulator for Control Through a Fixed Port onto Moving Tissue 329
14.1 Introduction . . . 330
14.2 From Stereotaxy to Supervised Suturing . . . 331
14.3 Instrument, Tissue and Patient as Simulated . . . 332
14.4 The Controllers as Coded . . . 338
14.5 Twelve Cases, Measured . . . 340
14.6 Technology in Practice . . . 346
14.7 Limitations . . . 348
14.8 Conclusion . . . 349
14.9 Exercises . . . 350

15 Collaborative Robot Speed and Separation Simulator With a Truth Audit 354
15.1 Introduction . . . 355
15.2 From Fenced Cells to Shared Workspaces . . . 356
15.3 The Shared Cell . . . 357
15.4 The Budget and the Law . . . 362
15.5 The Controller as Implemented . . . 365
15.6 Eight Cases and a Shift . . . 367
15.7 Technology in Practice . . . 372
15.8 Limitations . . . 374
15.9 Conclusion . . . 375
15.10 Exercises . . . 376

16 Summary and Outlook 380
16.1 What the Chapters Showed . . . 381
16.2 Choosing a Method . . . 384
16.3 Where Practice Is Heading . . . 385
16.4 Engineering With AI Agents . . . 387
16.5 Gaps in Motion Control and Autonomy . . . 389
16.6 Exercises . . . 394

Answers to Selected Exercises 395

Glossary 400

Index of Control Topics 403

Index of Applications 407

About the Editor 410

Volumes in This Series 411

Notation

The chapters share the symbols below. A symbol used by one chapter only is defined where it first appears, and each chapter states its own units.

Symbol	Meaning
SP	setpoint, the value a controller is asked to hold
PV	process variable, the measurement a controller acts on
CV	controlled variable
MV	manipulated variable; also OP, the controller output
DV	disturbance variable
e	control error, SP $-$ PV
K_c	controller gain
τ_I, τ_D	integral (reset) time and derivative time
K_p	process gain, the steady-state change in PV per unit change in MV
τ_p	process time constant
θ_p	process dead time
τ_c	closed-loop time constant chosen in a tuning rule
Δt	controller sample time
s	Laplace variable
x, u, y, d	state, input, output and disturbance of a state-space model
A, B, C	matrices of the linear model $\dot{x} = Ax + Bu$, $y = Cx$
Q, R	weights on outputs (or states) and on input moves in an optimal controller
$\hat{x}$	an estimate of x

Abbreviations used throughout: PID, proportional-integral-derivative; PI, proportional-integral; FOPDT, first-order plus dead time; MPC, model predictive control; MHE, moving horizon estimation; IAE, integral of the absolute error; RGA, relative gain array. Times are in seconds unless a chapter says otherwise. Controller output and valve position are in percent of span. A scenario or case number refers to the numbered list in the chapter's app.

CHAPTER 1

Introduction: Controlling Machines That Move

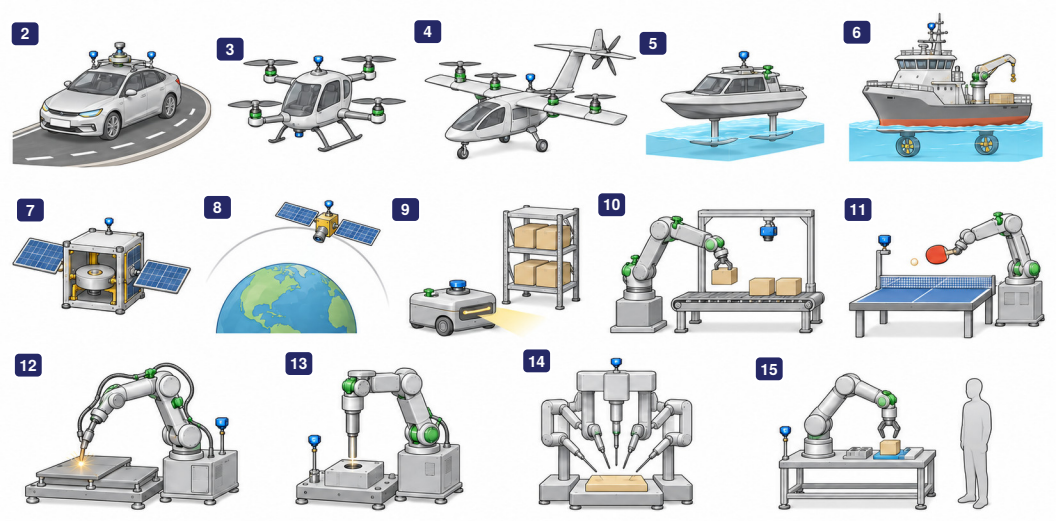

The chapters 2 lane keeping 3 drone taxi 4 lift-plus-cruise transition 5 foiling boat 6 dynamic positioning 7 CubeSat attitude 8 orbit maintenance 9 warehouse robot 10 conveyor picking 11 table tennis robot 12 arc welding 13 peg-in-hole assembly 14 surgical robot 15 collaborative robot

Notice that the top row moves through road, air and sea, the middle row through orbit and a warehouse to the first work cells, and the bottom row brings the arm into contact with parts, tissue and a person. Illustration, not to scale.

IN PLAIN TERMS **Guiding Machines That Move Through the World**

The main idea. A car keeps to its lane, a ship holds still over one spot of the seabed, and a robot arm picks a box off a moving belt. Each machine is pushed by motors, propellers or wheels, and it must know where it is to know how hard to push.

The challenge. Its position comes from sensors that drift, report late or go blind, and its load changes from one trip to the next. It must also be able to stop in time, even when what it believes about itself is wrong.

Why it matters. Good motion control lets machines carry goods, fly, sail and work safely beside people.

THREE WAYS TO READ THIS CHAPTER

Explore	7 min	This page, Section 1.2 with its timeline and two History boxes, Figure 1.2, the Did you know? boxes in Sections 1.4.6, 1.4.7 and 1.4.9, and Key points in Section 1.8. No equations.
Study	45 min	The whole chapter in order, working each number of the example C-1, then Exercises 1 to 10.
Practice	5 min	This page, Section 1.4.10 with Table 1.2, Section 1.5 with Figures 1.7 and 1.10, and Section 1.8.

1.1 The Question This Volume Answers

A machine that moves is controlled through geometry and inertia. A car keeps its lane by steering, a multirotor moves sideways only by tilting its thrust, a vessel holds station over a wellhead on thrusters alone, and a satellite turns its camera with wheels that keep every torque they absorb. A robot arm meets a product on a moving belt, presses a peg into a hole or works beside a person. In each case the quantity that matters is a distance to a path, a position over a pad, a force at a contact or a separation from a person. It is computed from estimates rather than read from one instrument, and the actuators

reach it through one or two integrations and a mass. This volume answers one question: how does a moving machine complete its task when its position, its load and its surroundings are known only approximately, and how is it kept from doing harm when they are wrong?

Eight control problems define the field, and every chapter meets several of them.

- *The error is geometric, and speed is a gain.* A turning command reaches a lateral position through two integrations whose coefficient is the speed, so a loop tuned at one speed meets a different plant at another.
- *Inertia changes with the load.* A payload changes the mass, the rotational inertia and the braking, and it changes at every stop.
- *Loops are nested.* A position loop commands a velocity or an attitude, which commands a force or a torque, and each inner loop must be faster than the loop it serves.
- *Motion is planned inside limits.* A reference must respect speed, acceleration, jerk and torque, and an interception must arrive at a place, at a time and at a speed.
- *The state is estimated, and often late.* Position comes from fusing encoders, inertial sensors, beacons and cameras. Reports arrive old, and part of a measured motion can be real and still not an error.
- *Actuators are shared, limited or directional.* There are more effectors than demands, effectors whose authority changes with speed, and actuators that act only across a field.
- *Contact turns position into force.* Once the tool touches the part, its stiffness converts every position error into a force.
- *A machine must be able to stop in time.* Its speed is bounded by what it believes about its braking and the space it can see, and a monitor must hold that bound against people who move.

Fourteen chapters answer the question on fourteen machines, each with an app that the reader runs in a browser. Chapters 2 to 8 move vehicles through a medium: a car in its lane, a passenger multirotor between rooftops, a lift-plus-cruise aircraft handing its weight from rotors to wing, a foiling boat, a vessel on dynamic positioning, a small satellite pointing a camera and a satellite holding its orbit among debris. Chapter 9 drives a pallet robot among people. Chapters 10 to 13 put an arm in a work cell, picking from a belt, returning a table tennis ball, welding a seam and inserting a peg. Chapters 14 and 15 close the volume on robots that work on or beside a person.

A reader who has worked through the volume can do the following on a moving machine that is new to the reader.

- Write the kinematics and the rigid-body balances of a vehicle or an arm and say what a payload changes.
- Design a cascade of position, velocity and force loops, and feed it the acceleration and the forces it can be told about.
- Generate a trajectory inside speed, acceleration and jerk limits, and say what smoothness costs in time.
- Make a path-following law independent of speed, and find the offset it leaves on a curve.
- Fuse a drifting sensor with a noisy one, remove a sensor bias and correct a report for its age.
- Allocate demands among limited actuators by priority or by weights.
- Set an impedance, predict a landing force and find the delay margin of a force loop.
- Turn a stopping distance into a speed limit near an obstacle or a person.

This chapter develops each of these on one small example and names the chapter that uses it. PID forms and tuning are taken as known from Volume 2, feedforward and cascade from Volume 3, allocation and overrides from Volume 4, estimation and Kalman filtering from Volume 5, unstable plants and state-space design from Volume 7, trajectories and switching from Volume 8 and predictive control from Volume 9. Standard texts treat the robotics [1], the marine craft [2] and the spacecraft [3] at length.

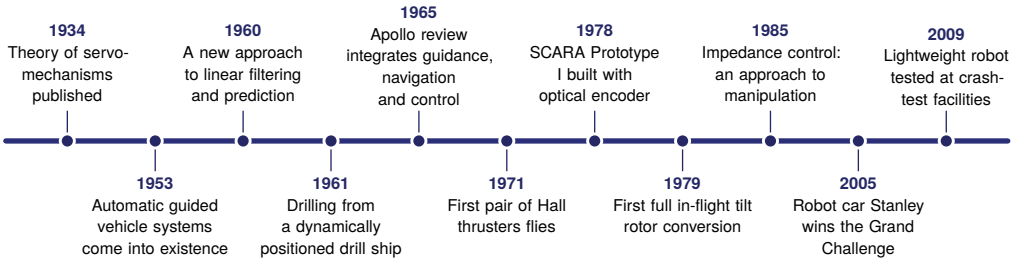

Figure 1.1 Milestones of motion control, robotics and autonomous systems. The dates are from [4] (1934), [5] (1953), [6] (1960), [7] (1961), [8] (1965), [9] (1971), [10] (1978), [11] (1979), [12] (1985), [13] (2005) and [14] (2009). Notice that the vehicles come first and that the arms which touch parts and people come late. Events are evenly spaced, not to a time scale.

1.2 A Short History

The milestones of Figure 1.1 follow the machines of this volume, from the theory of the servo to arms tested for the harm they could do to a person.

A theory of servo-mechanisms was published in 1934 [4], and a servo that makes a shaft or a body follow its command is the motion loop of Section 1.4.3. Automatic guided vehicle systems have existed since 1953 [5], the forerunners of the pallet robot of Chapter 9. In 1960 a new approach to linear filtering and prediction [6] and a guidance system for the rendezvous of satellites [15] gave navigation two tools that Section 1.4.6 and Chapter 8 still use. In March and April of 1961 a drill ship held itself over deep water by dynamic positioning, with radar reflectors and acoustic transducers on buoys as its position reference [7], the problem of Chapter 6.

> ⧖ HISTORY **1961 to 1965: Gyros Aligned by Stars**
>
> The plan of late 1961 was to fly the lunar spacecraft within two to three years, which left no time for a new inertial system, so its design was based on an existing guidance system. Three integrating gyros held a stable member still in space while the spacecraft turned around it, and three accelerometers on that member measured the changes of velocity that the computer added to the state. The gyros drift, so in flight the member was realigned by sighting stars, and the most critical parameter was the gyro bias drift. A body-mounted design was examined and set aside, partly because no such system had yet left the laboratory. A review of 1965 integrated the guidance, navigation and control functions so that an alternate mode remained after a failure [8]. The star tracker and the gyro bias of Chapter 7 are the same pair.

A pair of Hall thrusters first flew in 1971, and more than 300 more had flown by 2018 [9], the kind of thruster that Chapter 8 fires. The first prototype of an assembly arm that yields sideways and is stiff vertically was built in October 1978, with a direct-current motor and an optical encoder as its servo elements [10]. A tilt rotor research aircraft made its first full conversion from helicopter to airplane in flight on 24 July 1979 [11], the transition of Chapter 4. Robot table tennis goes back to a challenge of 1983 whose simplified rules allowed for the technology of the time [16]. Impedance control was put forward in 1985 as an approach to manipulation [12], and a robot with improved positioning accuracy for brain surgery guided by computed tomography was reported in 1988 [17]. In 2009 a report described a lightweight arm impact-tested in automobile crash-test facilities to quantify the risk of injury it posed [14], the question behind Chapter 15.

> ⧖ HISTORY **2005: A Desert Robot Built in Layers**
>
> The robot car that won the desert race of October 2005 ran about 30 software modules in parallel, in six layers: sensor interface, perception, control, vehicle interface, user interface and global services. The sensor layer received each laser at 75 Hz, the camera at about 12 Hz, the satellite receiver and its compass at 10 Hz and the inertial unit and the vehicle bus at 100 Hz. No master process kept the modules in step. Every datum was time-stamped, and the stamps were used whenever sources were combined. A filter estimated the position, orientation and velocities, a planner set the trajectory, and two closed-loop controllers followed it, one for the steering and one for the brake and throttle [13]. Figure 1.2 has the same layers.

1.3 Notation and Conventions

Positions are in meters, angles in radians or degrees as each chapter gives them, forces in newtons and torques in newton meters. Speed is v, heading ψ, yaw rate ω and path curvature κ. A reference is marked with a subscript r and an estimate with a hat. The framework's h is its integration step and its a the controller action, and neither is used for anything else here. This chapter therefore writes an acceleration limit as A, a braking capability as A_b and a jerk limit as $j_{\max}$. Most chapters count time in seconds. Chapter 11, Chapter 12 and Chapter 13 quote their fast events in milliseconds, Chapter 7 counts orbits and minutes, and Chapter 8 counts hours and days.

Loop action follows Section 1.6.2: $a = +1$ is reverse acting, for a process whose measurement rises with the manipulated variable, and $a = -1$ is direct acting. A joint servo whose torque rises when the angle falls short of its reference is therefore reverse acting, as are a wheel speed loop, a force loop that moves its reference deeper when the force reads low and the altitude loop of a multirotor. The position loop of the multirotor in Chapter 3 is direct acting, because a positive tilt accelerates the vehicle toward negative x, and so is the pitch loop of the foiling boat in Chapter 5.

The chapters keep their own symbols, and several collide. κ is the path curvature in Chapter 2 and the tilt stiffness $\ell K_{c,\theta}$ in Chapter 3. λ is the rate of arc self-regulation in Chapter 12, the chamfer force factor in Chapter 13, a damping factor in Chapter 11 and a forgetting factor in Chapter 14. μ is a friction coefficient in Chapter 9 and Chapter 13, the reduced mass in Chapter 15, the thrust margin in Chapter 3 and the lift margin in Chapter 4. J is the inertia matrix in Chapter 7 and the Jacobian in Chapter 10 and Chapter 11. K_p is a pointing gain in Chapter 7 and a joint stiffness in Chapter 10, never the process gain K_p. a is an acceleration in Chapter 2, the semimajor axis in Chapter 8, the braking deceleration in Chapter 15 (written a_b for the braking capability in Chapter 9) and the lift-curve slope in Chapter 4, where a_z and a_x are demanded accelerations, none of them the action. h is the ride height in Chapter 5, the restart hysteresis in Chapter 15, the servo period in Chapter 10 and the height above the surface in Chapter 4. θ is a tilt or a pitch in several chapters, the heading in Chapter 9 and the crossing angle of a debris pass in Chapter 8. φ is the wand angle in Chapter 5, a wheel angle in Chapter 9 and the hole tilt in Chapter 13, while Chapter 7 writes the rotation vector ϕ. n is the orbital mean motion in the two space chapters. K is an observer bandwidth in Chapter 3, an observer gain in Chapter 6 and a Kalman gain in Chapter 9.

1.4 Motion Under Uncertainty

1.4.1 The Flow Scheme

A moving machine has no pipes, but its signals follow one path in every chapter, shown in Figure 1.2. A mission layer or a supervisor chooses what to do next and when to switch modes. A planner turns that into a path or a trajectory in time. Motion loops, usually nested, compute the forces and torques that follow it. An allocation shares those demands among the actuators, which act on a rigid body that may be in contact with something. Sensors report on the body, and a navigation layer fuses their reports into the estimates that every other layer uses. A safety monitor watches the estimates and can stop the machine. The figure marks the layer where each chapter's main lesson lies and, in the bottom row, the medium in which each chapter's machine moves.

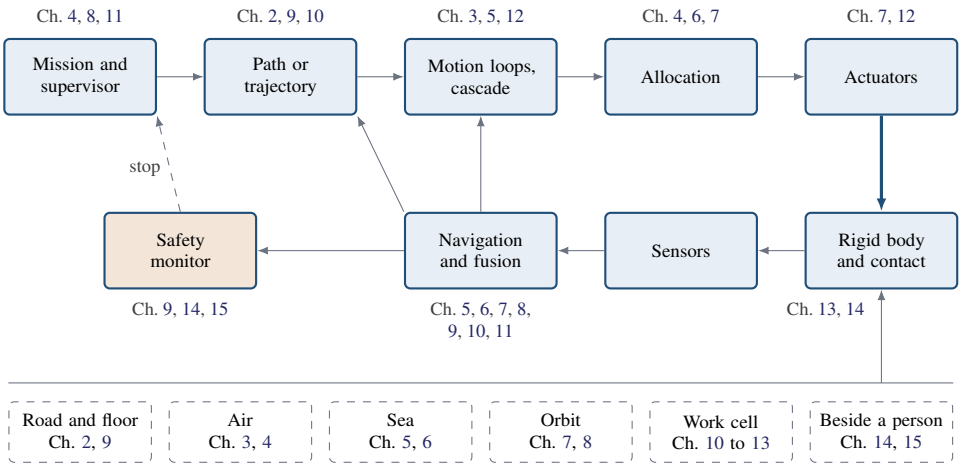

Figure 1.2 The Flow Scheme of a Moving Machine. Solid Blocks Are the Layers Every Chapter Has, With the Chapters Whose Main Lesson Lies in Each Layer Next to It. The Safety Monitor Reads the Estimates and Can Stop the Machine. Dashed Blocks in the Bottom Row Place Each Chapter's Machine in the Medium It Moves In.

The techniques that the chapters share are developed below on one running example, a small transport cart called C-1. It has two drive wheels on a track of $b = 0.5\,\mathrm{m}$, each limited to 50 N of force at the floor, and casters elsewhere. It weighs 50 kg empty and 100 kg with its payload, so its largest acceleration is $2\,\mathrm{m/s^2}$ empty and $1\,\mathrm{m/s^2}$ loaded. It cruises at 1 m/s, and its controller executes every 10 ms. Every number of C-1 is declared here or where it is used and is computed by the script kept with this chapter's record. The numbers are chosen for clear arithmetic and are not taken from any one chapter.

1.4.2 Kinematics and Rigid-Body Dynamics

A wheel that rolls without slipping sideways constrains the body to move along its heading [18]. With the midpoint of the drive axle at (x, y), the planar motion of C-1 is Equation (1.1).

$$\frac{\mathrm{d}x}{\mathrm{d}t} = v\cos\psi, \qquad \frac{\mathrm{d}y}{\mathrm{d}t} = v\sin\psi, \qquad \frac{\mathrm{d}\psi}{\mathrm{d}t} = \omega, \qquad m\frac{\mathrm{d}v}{\mathrm{d}t} = F_R + F_L, \qquad I\frac{\mathrm{d}\omega}{\mathrm{d}t} = \frac{b}{2}(F_R - F_L) \quad (1.1)$$

where F_R and F_L are the wheel forces, m the mass and I the yaw inertia. The first three relations are the kinematics. They have no mass in them, and they make the lateral position the second integral of the turning command with the speed as a coefficient. The last two are the rigid-body balances. The payload doubles m and raises I from 3.00 to 4.33 $\mathrm{kg\,m^2}$, so the same wheel forces give half the acceleration and 69 % of the yaw acceleration. Chapter 9 drives a pallet robot on these equations with wheel slip added. Chapter 2 replaces the yaw-rate command by a steering angle, $\psi = (v/L)\tan\delta$, which puts the speed into the heading equation as well.

An arm obeys the same balances in joint coordinates, as shown in Equation (1.2).

$$M(q)\,\ddot{q} + C(q,\dot{q})\,\dot{q} + F(\dot{q}) = \tau, \qquad M \leftarrow M + m_p J^{\top} J \quad (1.2)$$

where q holds the joint angles, M is the inertia matrix, $C\dot{q}$ the velocity terms, F the friction, τ the joint torques and J the tip Jacobian [1]. A held payload m_p adds $m_p J^{\top} J$ to the inertia, so it changes every entry of M and the coupling between the joints. Chapter 10 checks that term to 10^{-12} and shows the torque budget a payload left out of the controller's model breaks. Chapter 15 drives a planar two-link arm of the same form, and Chapter 11 inverts the kinematics of a six-axis arm by damped least squares.

A body that turns in three dimensions needs one more idea. Any change of attitude is a rotation by one angle about one axis, the eigenaxis, and its rotation vector ϕ points along that axis with the angle as its length. A pointing law that commands a rate $-k\phi$ turns the body about the eigenaxis, the shortest path [19]. A rate limit keeps it there only if it limits the length of the command and not its parts. For an error of 1.0 and 0.25 rad about two axes, whose eigenaxis lies 14.0° from the first, clipping each component at 0.05 rad/s turns the body about an axis 45.0° away. Limiting the length keeps 14.0°. Chapter 7 flies a 90.45° slew through 90.62° with the length limited and through 105.28° with three separate axes. Its reaction wheels add a momentum budget. A wheel exchanges angular momentum with the body and cannot shed it, so a drag torque that never stops fills one wheel by 37.85 % of its capacity every orbit.

1.4.3 Cascaded Motion Loops

A motion loop is nested because the actuator sets a force and the task sets a position. On C-1 an inner loop sets the wheel force from the speed error, $F = K_v(v_r - v)$, and an outer loop sets the speed reference from the position error, $v_r = K_x(x_r - x)$. With $m\dot{v} = F$ the inner loop follows its reference with the time constant $\tau_v = m/K_v$, and the position loop has the characteristic equation (1.3).

$$\tau_v s^2 + s + K_x = 0, \qquad \zeta = \frac{1}{2\sqrt{\tau_v K_x}}, \qquad \tau_v = \frac{m}{K_v} \tag{1.3}$$

With $K_v = 250\,\mathrm{N\,s/m}$ and $K_x = 1.25\,\mathrm{s^{-1}}$ the empty cart has $\tau_v = 0.2\,\mathrm{s}$ and a double root at $-2.5\,\mathrm{s^{-1}}$, critically damped. The inner pole at $5\,\mathrm{s^{-1}}$ is four times the outer gain, which is the separation a cascade needs. Loaded, $\tau_v = 0.4\,\mathrm{s}$, the damping ratio falls to 0.707 and a position step overshoots by 4.3 %. Doubling the mass divides the damping ratio by $\sqrt{2}$. Scheduling K_v in proportion to the believed mass, 500 N s/m loaded, restores the empty cart's roots.

A force the loops are not told about is opposed only through an error. At rest the two proportional loops balance a steady force F_d with the position error of Equation (1.4).

$$e = \frac{F_d}{K_v K_x} \tag{1.4}$$

A 2 % grade pulls the loaded cart back with 19.6 N and leaves 63 mm. An integral term in the inner loop removes it over its integral time. A feedforward of the force removes it at once, if the force is known. A disturbance observer estimates an unknown force from the momentum balance and feeds it forward, so the error decays on the observer's time scale instead [20].

Chapter 3 is this subsection on a multirotor. Its position loop writes a tilt setpoint for a tilt loop that writes a differential thrust, with the two oscillatory modes 5.5 apart [21]. Its altitude gain is scheduled on the believed mass. It compares each loop's error with the proportional-only error of Equation (1.4), and the tilt loop, which has no integral, matches it to 0.01°. Correcting the weight feedforward alone brings a heavy climb from 1.78 m behind to 0.05 m, and an observer with a 0.67 s time constant cuts the drift after a wind step from 1.19 to 0.36 m. Chapter 5 closes its height loop on a heave-rate inner loop, and Chapter 6 feeds its estimated environmental force forward. Chapter 12 has a loop that no controller closes: a constant-voltage arc regrows most of a lost length by itself, with a time constant of 45.5 ms, settling at 3.82 mm against 4.29 mm before the step.

1.4.4 Trajectory Generation

A reference that steps asks the loops for an infinite acceleration and lets the actuator limits decide the path. A trajectory generator gives the loops a reference they can follow. For a move of length D inside a speed limit $v_{\max}$ and an acceleration limit A, the trapezoid, the jerk-limited trapezoid and the quintic polynomial take the durations of Equation (1.5), for a move long enough to reach the speed limit.

$$T_{\mathrm{trap}} = \frac{D}{v_{\max}} + \frac{v_{\max}}{A}, \qquad T_{\mathrm{jerk}} = T_{\mathrm{trap}} + \frac{A}{j_{\max}}, \qquad T_{\mathrm{quin}} = \max\left(\frac{1.875\,D}{v_{\max}}, \sqrt{\frac{5.774\,D}{A}}\right) \tag{1.5}$$

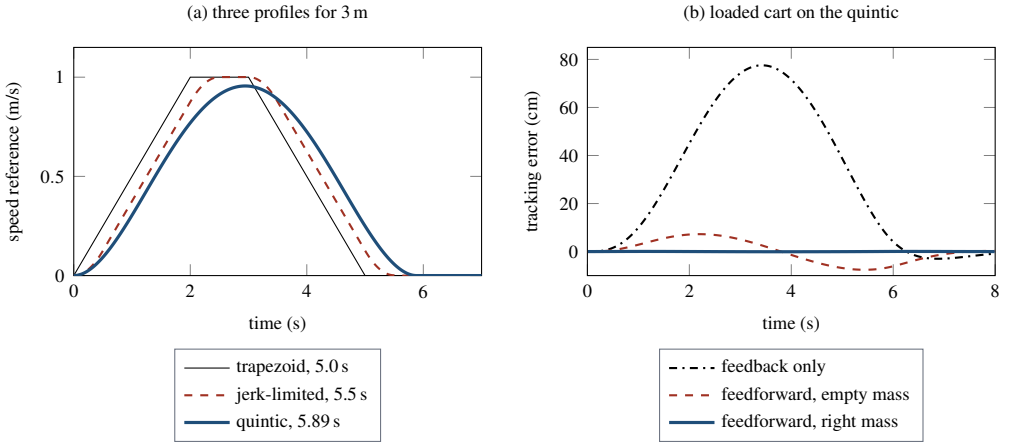

Figure 1.3 A 3 m Move of C-1 at 1 m/s and 0.5 m/s^2. (a) The Trapezoid (Thin), the Jerk-Limited Trapezoid (Dashed) and the Quintic (Thick). (b) Tracking Error of the Loaded Cart on the Quintic Under the Cascade of Section 1.4.3: Feedback Alone (Dash-Dotted) Lags by Up to 78 cm, Feedforward With the Empty Mass (Dashed) Leaves 7.6 cm, and Feedforward With the Right Mass (Thick) Leaves 0.07 cm.

The quintic, $x_r = D(10s^3 - 15s^4 + 6s^5)$ with $s = t/T$, starts and ends at rest with zero acceleration, and it is the minimum-jerk path between two rests [22]. Its peak speed is $1.875\,D/T$, its peak acceleration $5.774\,D/T^2$ and its peak jerk $60\,D/T^3$. The jerk-limited form holds while $v_{\max} \geq A^2/j_{\max}$.

C-1 plans its moves at 1 m/s and 0.5 m/s^2, half its loaded capability, so that feedback keeps the other half. A move of 3 m then takes 5.0 s on a trapezoid, 5.5 s with a jerk limit of 1 m/s^3 and 5.89 s on a quintic, which is limited by its acceleration. Smoothness costs 18 % of the time against the trapezoid, and the quintic's peak jerk is 0.88 m/s^3. Figure 1.3(a) shows the three speed profiles.

A trajectory also carries the feedforward of Section 1.4.3. The outer loop adds the reference speed, $v_r = \dot{x}_r + K_x(x_r - x)$, and the inner loop adds the force the reference acceleration needs, $\hat{m}\ddot{x}_r$. Figure 1.3(b) shows the loaded cart on the quintic. Feedback alone lags by up to 0.78 m, close to the peak speed over K_x. Feedforward with the empty mass leaves 76 mm, because the missing 25 N at the peak acceleration acts like the force of Equation (1.4), which predicts 80 mm. Feedforward with the right mass leaves 0.7 mm.

Chapter 10 plans every move as one quintic per joint and feeds the inverse dynamics of Equation (1.2) forward, and its interception is a fixed point, because the goal depends on the duration and the duration on the goal. Chapter 11 shows the price of smoothness on a stroke: a minimum-jerk move of 0.3 m in 0.25 s needs 27.7 m/s^2 at its peak against 19.2 m/s^2 for a bang-bang move. Chapter 3 generates a jerk-limited reference online, which re-plans from the present state whenever the target moves [23], and Chapter 15 uses its trajectory generator to command a protective stop.

1.4.5 Path Following

When the task is a path rather than a schedule, the error is the lateral distance y to the path and the heading error ψ_e relative to it. On a path of curvature κ, Equation (1.1) gives Equation (1.6).

$$\frac{\mathrm{d}y}{\mathrm{d}t} = v \sin \psi_e, \qquad \frac{\mathrm{d}\psi_e}{\mathrm{d}t} = \omega - v\kappa \tag{1.6}$$

A fixed law $\omega = -k_h \psi_e - k_y y$ gives the linear characteristic polynomial $s^2 + k_h s + k_h k_y v$, whose roots move with speed. The law of Equation (1.7) divides the lateral term by the speed and adds the turn

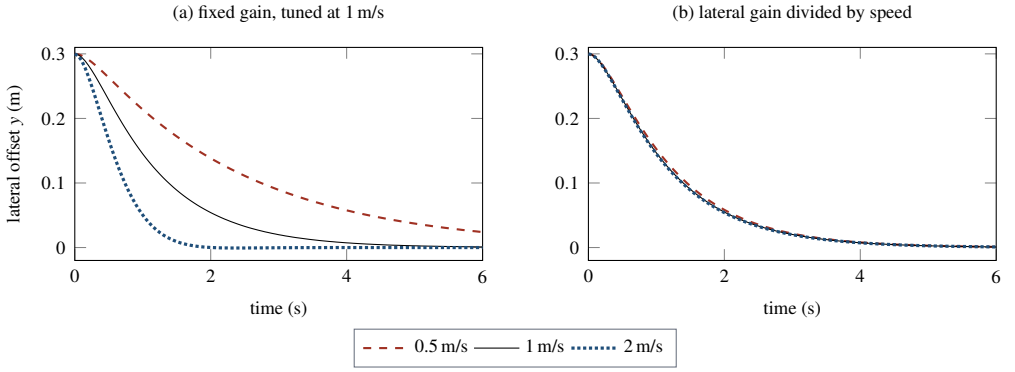

Figure 1.4 C-1 Recovering From a 0.3 m Lateral Offset at 0.5 m/s (Dashed), 1 m/s (Thin) and 2 m/s (Dotted). (a) With Fixed Gains the Response Is Slow at Low Speed and Fast at High Speed. (b) With the Lateral Gain Divided by the Speed the Three Curves Nearly Coincide. The Small Lag at 0.5 m/s Is the Arctangent Limiting a Large Correction.

rate of the path [13, 24].

$$\omega = -k_h \left[\psi_e + \arctan \frac{k_e y}{v} \right] + v\kappa \qquad \Longrightarrow \qquad s^2 + k_h s + k_h k_e = 0 \tag{1.7}$$

The speed in the lateral gain cancels the speed in the kinematics, so the poles stay put. For C-1 with $k_h = 5\,\mathrm{s}^{-1}$ and $k_y = k_e = 0.8$ (in $\mathrm{m}^{-1}\,\mathrm{s}^{-1}$ and s^{-1}), both laws have poles at -1 and $-4\,\mathrm{s}^{-1}$ at 1 m/s. Figure 1.4 shows the recovery from a 0.3 m offset. The fixed law takes 2.52 s to 37 % of the offset at 0.5 m/s and 0.70 s at 2 m/s, and its damping ratio falls to 0.625 at 4 m/s. The speed-divided law takes 1.29 to 1.35 s at all three speeds, and covers twice the distance at twice the speed.

The term $v\kappa$ is a feedforward from the path itself. Without it the steady turn must come from a lateral error, and the balance of Equation (1.7) at $\psi_e = 0$ gives the offset of Equation (1.8).

$$y_\infty = -\frac{v^2 \kappa}{k_h k_e} \tag{1.8}$$

On a 4 m radius C-1 runs 0.063 m outside the path at 1 m/s and 0.25 m at 2 m/s, as the simulation confirms, and the feedforward removes it exactly. What the division by speed does not remove is everything the law leaves out, such as an actuator lag that costs a fixed time and therefore a distance that grows with speed.

Chapter 2 is this subsection on a car, with the Stanley law on a kinematic bicycle. Its cross-track time constant moves only from 0.917 to 0.861 s between 10 and 30 m/s. Its offset without the curvature feedforward is 0.314 m in a 200 m corner at 25 m/s, and a heading error referenced to the wrong point counts the feedforward twice. Its steering rack leaves every cross-track gain stable below $v = L/\tau_\delta = 29\,\mathrm{m/s}$ and sets a finite critical gain above it. Chapter 9 follows A* routes by pure pursuit, closing a heading loop on the bearing to a point ahead on the route [25]. Chapter 12 follows a seam read by a laser ahead of the torch, where a reading belongs to a place along the path and not to a moment.

1.4.6 Navigation and Sensor Fusion

No instrument reports where C-1 is. Its wheel encoders report the distance traveled, accurately over a second and with an error that grows with distance. A beacon reports the position every 0.1 s with a noise of 0.05 m that does not grow. Dead reckoning integrates the encoders alone, and a 2 % error in the wheel radius leaves 0.60 m after 30 m. A complementary filter uses each sensor where it is good,

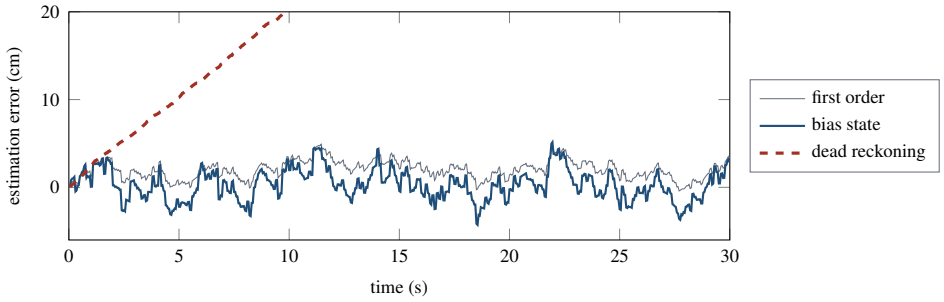

Figure 1.5 Position Estimates of C-1 at 1 m/s With Encoders That Read 2 % High and a Beacon Fix Every 0.1 s With 0.05 m of Noise, Seed 11. Dead Reckoning (Dashed) Leaves the Axis at 10 s and Reaches 60 cm at 30 s. The First-Order Filter (Thin) Settles 1.9 cm Ahead. The Filter With a Bias State (Thick) Has No Steady Error.

as shown in Equation (1.9).

$$\frac{\mathrm{d}\hat{x}}{\mathrm{d}t} = v_m + \omega_c\,(y_m - \hat{x}) \qquad \Longrightarrow \qquad \hat{X} = \frac{\omega_c}{s+\omega_c}\,Y_m + \frac{s}{s+\omega_c}\,\frac{V_m}{s} \tag{1.9}$$

where v_m is the encoder speed, y_m the beacon fix and ω_c the crossover. Below ω_c the estimate follows the beacon and above it the encoders, and the two weights add to one at every frequency. A speed error εv that never stops leaves the steady estimation error $\varepsilon v/\omega_c$. A second state $\hat{b}$ that integrates the innovation estimates that error and removes it, as shown in Equation (1.10).

$$\frac{\mathrm{d}\hat{x}}{\mathrm{d}t} = v_m - \hat{b} + 2\omega_c\,(y_m - \hat{x}), \qquad \frac{\mathrm{d}\hat{b}}{\mathrm{d}t} = -\omega_c^2\,(y_m - \hat{x}) \tag{1.10}$$

Figure 1.5 shows the three estimates of C-1 at 1 m/s with $\omega_c = 1\,\mathrm{rad/s}$. After the first 15 s the first-order filter sits 0.019 m ahead, against $\varepsilon v/\omega_c = 0.020\,\mathrm{m}$, with an rms error of 0.022 m. The filter with the bias state has a mean error under 0.001 m and an rms error of 0.017 m, a third of the beacon's 0.050 m. A Kalman filter has the same structure and chooses the gains from the noise variances instead of a crossover [6, 26].

A report also has an age. A fix taken L seconds ago describes where the machine was, and treating it as current leaves the error of Equation (1.11) at a speed v.

$$e_L = vL \tag{1.11}$$

With fixes 0.2 s late, the first-order filter of C-1 settles 0.181 m behind, which is $\varepsilon v/\omega_c - vL = -0.180\,\mathrm{m}$. Comparing each fix with the estimate stored for its capture time removes the vL part and leaves the 0.019 m of the speed error. A moving target needs the same correction. A product on a belt has moved on by the belt speed times the report's age and the time still to come.

> ✦ DID YOU KNOW? **Two Minutes Without Satellites**
>
> The desert robot of 2005 kept an accurate pose through losses of its satellite fixes lasting up to 2 min, with gyros for its orientation and wheel speeds for its position. The error it gathered was usually of the order of centimeters, and its health monitor cut the top speed to 10 miles an hour while the fixes were lost [13].

The chapters build on these elements. Chapter 5 fuses a height sensor with an accelerometer through three poles at the crossover, with a bias state, and Chapter 9 fuses encoders and a gyro with reflector sightings in an extended Kalman filter that forms each innovation against the pose stored

Table 1.1 Three Ways for C-1 to Cut a Demand of 80 N and 10 N m That Asks One Wheel for More Than Its 50 N

Rule	F_R (N)	F_L (N)	Force F	Moment M
Clip each wheel	50	20	70 N, 88 %	7.5 N m, 75 %
Moment first	50	10	60 N, 75 %	10 N m, 100 %
Force first	50	30	80 N, 100 %	5 N m, 50 %

for the moment of measurement, a reflector delay before the scan arrives. Chapter 7 estimates the gyro bias in a multiplicative extended Kalman filter, and without the bias state its camera drifts 2.35° while the Sun blinds the star tracker. Chapter 10 moves each camera report on by the belt travel since its capture, and an arm aimed at the last report closes behind the product by $64.4 + 148.2L$ in millimeters. Chapter 11 predicts a flight with drag and spin from frames it must stop watching before it is sure, and Chapter 8 flies on mean elements, averages that no instrument reports. It sizes its avoidance burns on linear relative motion, in which a burn moves the satellite along its track by about three times the velocity change for every second that follows [15].

Part of a measured motion can be real and still not an error. A vessel riding over a 4 m sea moves 2.4 m at an eleven-second period without leaving its station, and a foiling boat can follow long waves or let short ones pass under it. The crossover then decides what the loops are shown. Chapter 6 carries a wave resonator in its observer, and the same sea and tuning cost 327.5 MN of thruster command travel with the raw measurement shown to the loops and 72.0 MN with the full observer [2]. Chapter 5 chooses between platforming and contouring by its observer's crossover.

1.4.7 Allocation

The motion loops ask for a total force and a moment, and the actuators must supply them. For C-1 the two wheels give Equation (1.12).

$$F = F_R + F_L, \quad M = \frac{b}{2}\left(F_R - F_L\right) \qquad \Longrightarrow \qquad F_R = \frac{F}{2} + \frac{M}{b}, \quad F_L = \frac{F}{2} - \frac{M}{b} \tag{1.12}$$

A demand of 80 N with 10 N m asks 60 N of the right wheel, which can give 50 N. Table 1.1 compares three ways to cut the demand. Clipping each wheel loses some of both. Serving the moment first keeps the turn and gives up speed, and serving the force first keeps the speed and gives up half the turn. A path follower that loses its turn leaves its path, so a moving machine usually serves its attitude or its turn first.

With more actuators than demands the split is a choice. The weighted least-norm solution of Equation (1.13) spreads a demand τ over the actuator commands $\mathbf{u}$ through the configuration matrix B [27].

$$\mathbf{u} = W^{-1}B^{\top}\left(BW^{-1}B^{\top}\right)^{-1}\tau \tag{1.13}$$

A large weight in W makes an actuator expensive, and a failed one is given a weight so large that the others take its share with no rule written for the case.

Chapter 6 solves Equation (1.13) for four thrusters, six forces and three demands, weighting each machine by the inverse square of its rating. Chapter 3 serves the tilt before the total thrust, so a rotor group at its limit costs altitude and never attitude. Chapter 4 asks the wing first and gives the rotors the rest, and the wing's authority grows with the square of the airspeed. Chapter 7 has an actuator that acts in only two directions: a magnetic coil makes a torque across the Earth's field and none along it, so over an orbit an excess of wheel momentum along the orbit normal decays with a time constant of 15.24 min and one in the orbit plane with 29.5 to 29.6 min.

> ✦ DID YOU KNOW? **Thirty-Eight Millinewtons From Five Hundred Watts**
>
> Electric thrusters generally give less than 75 mN of thrust per kilowatt, so a small spacecraft that can feed one 500 W may get no more than 38 mN. Such a thruster may fire for hundreds or thousands of hours where a chemical one fires for seconds or minutes [9]. Chapter 8 flies one.

1.4.8 Contact and Impedance

Once the tool touches the part, the part fixes some of its coordinates. The controller can then only choose how hard to push along them, and its stiffness turns every position error into force. Impedance control makes that stiffness a design choice, as shown in Equation (1.14) [12].

$$F = K\,(x_r - x) + D\,(\dot{x}_r - \dot{x}), \qquad D = 2\zeta\sqrt{Km} \tag{1.14}$$

Suppose C-1 docks on a charger misplaced by 5 mm. A stiff position servo of 20 kN/m pushes the misplacement with 100 N, the whole drive force. A spring of 500 N/m pushes with 2.5 N, and with $\zeta = 0.7$ it needs 221 N s/m of damping empty. The desired contact force is set by programming the reference past the contact, 40 mm for 20 N.

No controller acts during the landing itself. A mass m meeting a stiff surface k_s at speed v is stopped within a quarter period $(\pi/2)\sqrt{m/k_s}$, and its peak force is Equation (1.15).

$$F_{\text{peak}} \approx v\sqrt{k_s\,m} \tag{1.15}$$

Against a 20 kN/m bumper the empty cart lands with 1000 N per m/s, 20 N at 0.02 m/s and 100 N at 0.1 m/s, and the impedance spring, 40 times softer, adds little in the 79 ms it lasts. The approach speed is the only defense.

A force loop closes on the measured force by moving the reference. With an integral law $\dot{x}_r = k_i(F_{SP} - F_m)$ against a surface much stiffer than the spring, the force follows $F \approx K(x_r - x_c)$, and the loop is an integrator of gain $g = Kk_i$. A measurement delay τ_d takes $g\tau_d$ radians of phase at the crossover g, so the loop holds while Equation (1.16) does.

$$\tau_d < \frac{\pi}{2\,K\,k_i} \tag{1.16}$$

With $k_i = 0.01\,\text{m}/(\text{N s})$ C-1 has $g = 5\,\text{s}^{-1}$, a time constant of 0.2 s and a delay margin of 0.314 s, which a simulation of the delayed loop reproduces.

Chapter 13 is this subsection on a peg. Its stiff servo damages the part beyond about 0.22 mm of fixture offset, and at 0.4 mm a 3 N/mm spring inserts the peg at 13.7 N. Its landings cost 1.35 N per mm/s, against 1.44 N from Equation (1.15), and its force loop, whose gain $K_{c,F}K_z$ is 1.0, loses its phase margin at 18.7 ms of delay. Chapter 14 holds a force on tissue that relaxes and creeps, through an estimate biased by the friction of the port.

1.4.9 Safety Monitoring

A machine that may meet an obstacle must be able to stop before it. After a reaction time t_r and with a braking capability A_b, the stopping distance and the fastest speed that stops within a clear distance d are Equation (1.17).

$$d_s = v\,t_r + \frac{v^2}{2A_b}, \qquad V(d) = A_b\left[\sqrt{t_r^2 + \frac{2d}{A_b}} - t_r\right] \tag{1.17}$$

With $t_r = 0.2\,\text{s}$ the loaded cart needs 0.70 m to stop from 1 m/s and the empty one 0.45 m. The limit is only as good as A_b. A cart that believes its empty braking while carrying the payload is allowed 1.07 m/s with an obstacle 0.5 m ahead. It needs 0.79 m to stop and strikes at 0.76 m/s. Believing the loaded braking allows 0.82 m/s, which stops in 0.50 m. Figure 1.6 shows both limits.

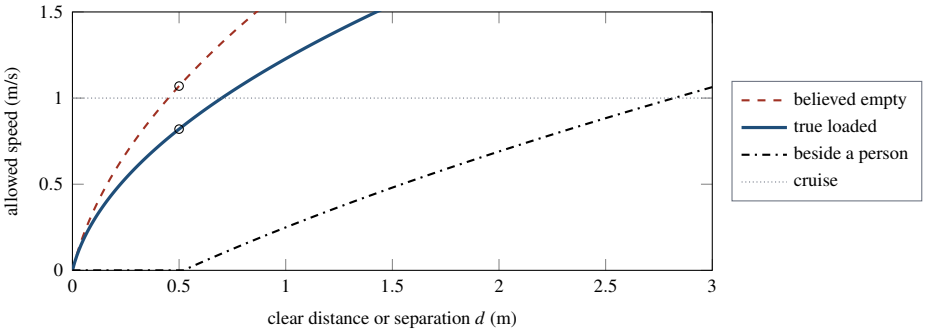

Figure 1.6 Speed Limits of C-1 From Equation (1.17) With the Braking It Believes Empty (Dashed) and Has Loaded (Thick), and From Equation (1.18) for a Person Walking In at 1.6 m/s (Dash-Dotted). The Circles Mark the Two Limits for an Obstacle 0.5 m Ahead. The Dotted Line Is the Cruise Speed.

A person does not stand still. The protective separation distance adds the distance a person walking at v_h covers while the machine reacts and brakes, with the braking time $t_s = v/A_b$ and an intrusion allowance C, as shown in Equation (1.18) [28].

$$S_p = v_h\,(t_r + t_s) + v\,t_r + \frac{v^2}{2A_b} + C \tag{1.18}$$

At 1 m/s, with $v_h = 1.6\,\mathrm{m/s}$ and $C = 0.2\,\mathrm{m}$, the loaded cart needs 2.82 m, and 0.52 m at rest. The braking time enters twice, once as the cart's braking distance and again as time in which the person keeps walking. Leaving out the person's walk during the stop gives 1.22 m, short by the 1.6 m a person covers in that second. Solving Equation (1.18) for v gives the speed law plotted in Figure 1.6, 0.48 m/s at 1.5 m.

A monitor acts through samples. A scanner that reports every 40 ms sees a person at 1.6 m/s come 64 mm closer between reports, so the speed law needs a margin, and a restart needs a hysteresis or the machine stops and starts on every scan. A monitor can also be checked: replaying each decision on the true state and the true braking says whether a person could have been reached, which a monitor that sees only its own estimates cannot say.

> ✦ DID YOU KNOW? **Twenty-Two Meters of Sight**
>
> The desert robot of 2005 found obstacles with its lasers out to about 22 m, enough to avoid them reliably at up to 25 miles an hour. When its camera found no drivable road within 40 m, it slowed to that speed [13]. Its speed was set by the space it could see.

Chapter 9 sets its speed from Equation (1.17) on visible space and believed braking. With a 1000 kg pallet planned on the empty robot's 2.26 m/s^2 it hits a parked cart at 0.74 m/s, while payload-aware braking stops 0.25 m short. Chapter 15 turns Equation (1.18) into a speed law with a margin and a hysteresis that cuts 17 protective stops in a 90 s case to 2 at the cost of one cycle, and its truth audit replays every decision on a copy of the true arm. At 1 m/s with a 3 kg part it needs 1448 mm. Chapter 14 keeps the instrument out of forbidden regions with virtual fixtures, and Chapter 4 earns each change of mode on measured airspeed and lift margin rather than on a clock.

1.4.10 Completing a Task Under Uncertainty

Every subsection above ends at the same point: the task succeeds only if what the controller believes is close enough to what is true. There are four ways to close the gap. A controller can be told beforehand, through a load sheet, a recipe mass or a map, and feed the value forward. It can estimate

Table 1.2 For Each Chapter, the Task, What Is Uncertain and How the Controller Learns of It or Guards Against It

Ch.	Task	Uncertain	Learned or guarded by
2	Hold the lane at any speed	Speed, curvature, linkage misalignment, rack lag	Speed-divided law, curvature feedforward, MPC
3	Fly a route with stops	Load, balance, wind among buildings	Load-sheet feedforward, scheduled gains, disturbance observer, dispatch margin
4	Hand the weight from rotors to wing	Airspeed, lift margin, a failed motor	Wing-first allocation, gates on measured airspeed and lift margin
5	Fly between slam and ventilation	Waves, sensor loss	Complementary observers, platform or contour, heave-rate cascade, fallback
6	Hold station on thrusters	Wave motion, environmental force, a failed thruster	Observer with bias and wave states, weighted allocation, heading supervisor
7	Point the camera through a pass	Gyro bias, stored momentum, a failed wheel	Filter with a bias state, magnetic unloading, eigenaxis limit
8	Keep altitude and avoid debris	Drag, debris covariance, thruster output	Mean elements, predictive planner, collision probability, burn verification
9	Deliver pallets among people	Position, braking with a load, unseen space	Fused estimate, payload-aware envelope, speed for unseen space
10	Grasp products on a belt	Report age, belt speed, payload	Time-stamped prediction, interception at belt speed, payload in the arm model
11	Return the ball	Drag, spin, time left	Filter on a flight model, commit rule, minimum-jerk swing
12	Weld a drifting seam	Torch height, seam ahead of the arc	Arc self-regulation, height through the current, distance-keyed preview
13	Insert a peg	Hole position, friction, late force signal	Impedance, contact detection and search, force loops
14	Work through a port on moving tissue	Tremor, delay, organ motion, tissue creep	Scaling and filtering, virtual fixtures, harmonic predictor, force loop
15	Work beside a person	The person's motion, latency, braking	Separation budget, speed law with margin and hysteresis, truth audit

the unknown online, with an observer, a filter or an adaptive predictor. It can bound the unknown and keep a margin, as a speed law, a transition gate or a probability threshold does. Or it can commit to a prediction before the outcome is known, because the deadline will not wait. Each choice spends something: error while an estimate converges, time when a margin slows the machine, or risk when a prediction is wrong. The chapters measure these costs, often as an error budget that splits a result among its causes.

Table 1.2 lists, for each chapter, the task, what is uncertain and how the controller learns of it or guards against it. Chapter 16 adds the measured results and compares how each system completed its task under uncertainty.

1.5 Instruments and Final Elements in This Volume

Figure 1.7 lists what the controllers of each chapter read and what they move. Nine chapters close their loops on the estimate of a filter or an observer, more than read any single kind of sensor, and a gyro or an accelerometer is the commonest instrument, in six. An electric drive on a wheel, a joint or a rotor is the final element of ten chapters.

An optical encoder chops a light beam with a slotted wheel and sends a train of pulses [29]. With two trains a quarter period apart, the one that leads tells which way the wheel turns, so a counter can add or subtract [29]. Encoders read the joints of four arms in this volume and the wheels of Chapter 9. They count distance, which is why dead reckoning drifts in Section 1.4.6.

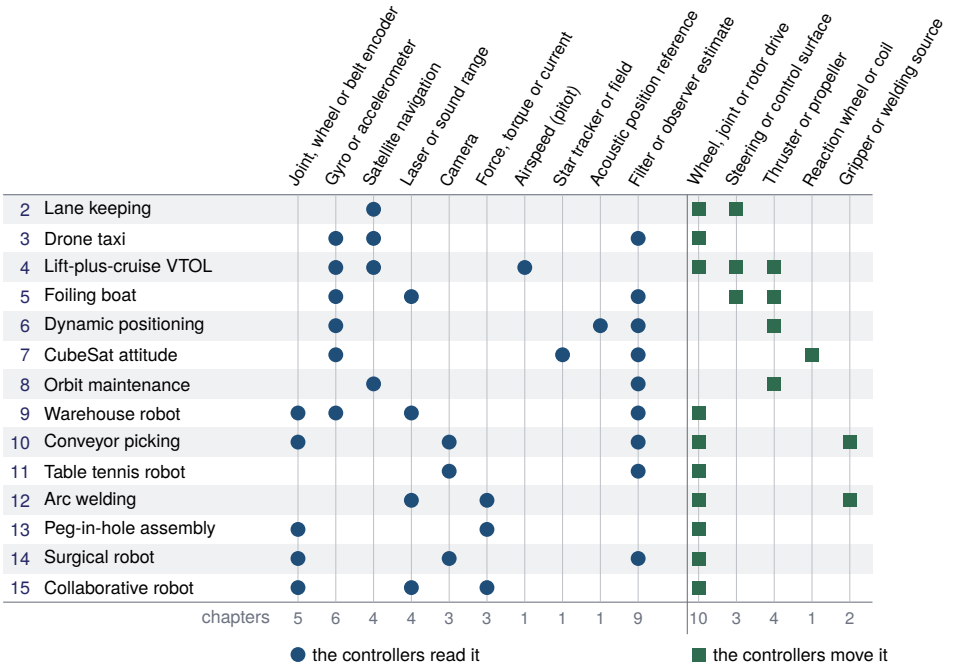

Figure 1.7 Sensors and final elements by chapter: a round mark where the chapter's controllers read the device and a square mark where they move it, as its variables table and its technology section name them. A value computed from other readings is not marked, and the last column of reads marks a loop closed on the estimate of a filter or an observer. Notice that the estimate outnumbers every sensor.

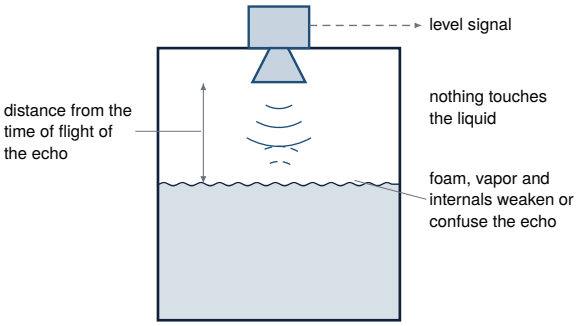

Figure 1.8 Principle of ranging by time of flight, drawn as a radar level gauge. Notice that the distance is half the round trip of the wave, so a timing error counts half, the rule by which the laser scanners of Chapter 9 and Chapter 15 measure the space around them.

A satellite receiver is now the main means of finding the orbit of a spacecraft in low orbit, and it needs at least four navigation satellites in view [9]. On the ground it is fused with gyros, accelerometers and wheel speeds, as on the desert robot of Section 1.2 [13].

A laser scanner finds range from the time of flight of light. A clock time of 1 ns is 150 mm of range, and the accuracy depends on the reflectivity of the target and the angle at which the beam strikes it [30]. Figure 1.8 shows the principle on a radar gauge, which times a radio wave the same way.

A wrist force sensor measures all six components of force and torque in one compliant structure. Strain gauges give it high accuracy but need careful wiring, pick up electrical interference and often

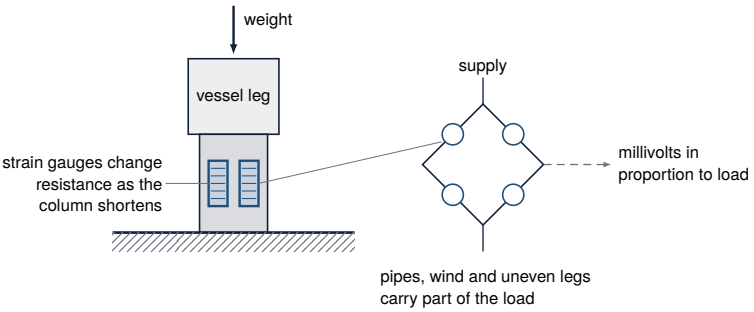

Figure 1.9 Principle of a strain-gauge load cell. Notice that a load that reaches the support by another path is not measured. A wrist sensor puts such a compliant body at the end of the arm, between the arm and its tool (Chapter 13).

Layer	What it does in this volume	Where it runs	How often
People who decide	a driver holds the wheel, a surgeon moves the console, operators screen an orbit	**By hand**, at the wheel, at the console or in operations	as needed; an orbit at least daily
Mission and supervisor	chooses modes and rotor states, plans burns, sets the heading	**Flight software** or the robot's program	10 ms to 6 h in the apps
Safety monitor	stopping envelopes, speed and separation, transition gates	safety-rated sensing, or a computer that sends slow and stop	with each scan; 20 ms to 0.1 s in the apps
Planning and navigation	plans paths and trajectories, fuses the sensors into estimates	**Onboard computer** of the vehicle or the robot	5 ms to 10 s in the apps
Motion loops and allocation	cascades of position, speed, attitude and force, shared among actuators	**Embedded controller** or the robot controller's servos	milliseconds; 0.3 ms to 10 s in the apps
Sensors and drives	encoders, inertial units, receivers, scanners, cameras, motors and thrusters	on wires or the vehicle bus	reports from every 10 s to 200 Hz in the apps

Figure 1.10 The control layers of this volume, the platform that hosts each and how often each acts. Loop periods are from [32] and the app periods from the chapters' technology sections. Notice that the safety monitor sits above the loops it can stop, and a person above it.

couple one axis into another [31]. A strain gauge is bonded to a metal element that stretches and compresses with the force like a very stiff spring [29] (Figure 1.9).

Figure 1.10 places the volume's loops on their platforms. Supervisory and lower loops in industry run with cycle times from milliseconds to minutes [32]. The apps span a wider range, from the 0.3 ms force loops of Chapter 13 to the 6 h planner of Chapter 8, and a person sits above several of them.

1.6 The Simulator Framework

Every chapter of this book has an app that runs in a web browser, collected with the others at https://apmonitor.com/apps/. About 240 apps exist across the series. All but a few are built on one shared framework, so the engine, the PID algorithm, the trends and the data export behave the same way in every chapter. The exceptions are standalone studios, such as the cascade studio and the two Temperature Control Lab apps, which follow the same ideas in their own code. This section describes the framework once and states where the standalone apps differ. Table 1.3 lists what a reader meets.

Table 1.3 What Every App Provides. Standalone Apps Are Described in Their Own Chapters.

Element	What it does
Run, Pause, Reset	Start, hold and restart the run. Reset keeps the reader's settings.
Speed	Simulated time per second of real time. It changes the pace and never the result.
Manual and Auto	Hand each actuator to the reader or to the controller, without a bump.
Scenario list	Numbered cases, each with a briefing, a completion rule and a reload button.
Control panel	Setpoints, manual outputs, tuning, options and disturbances, at a Beginner or an Advanced level.
Trends	SP and PV above the MV for each loop, with limits, event marks and a cursor.
Scorecard	IAE, overshoot, settling time, MV travel and time at a limit for each loop.
Learn card	Predict, Try It, two hints and Explain, opened by the reader in that order.
Show Model	Equations, live variable values, editable parameters and assumptions.
Export CSV	Every logged sample of the run as a text file with a header row.
Instructions	A separate page with background, an operating guide and the model's limitations.

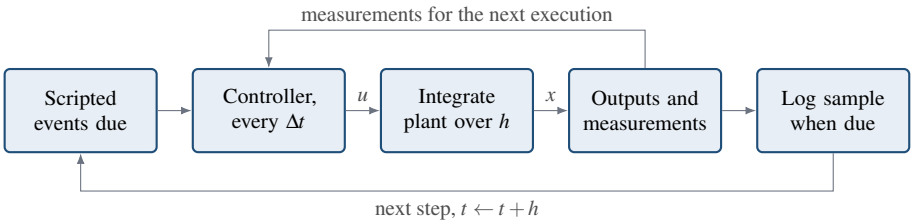

Figure 1.11 Work of the Engine in One Integration Step

1.6.1 The Simulation Engine

Each app carries a plant model written as ordinary differential equations $\frac{dx}{dt} = f(x,u,d,p)$ in the states x, the inputs u, the disturbances d and the parameters p. The engine integrates the model with the classical fourth-order Runge–Kutta method on a fixed step h set for the fastest dynamics of the app. Some apps supply their own step, such as explicit Euler for a discontinuous right-hand side or an exact update for a model driven by events.

Three clocks run at separate rates. The plant advances every step h. The controller executes every control period Δt, equal to h unless the app declares a slower controller, and holds its outputs between executions. The log takes one sample at a fixed interval of simulated time for the trends, the scorecard and the exported file. The order of work in one step is shown in Figure 1.11. The controller also executes once when a scenario is loaded, so every run starts from an initialized controller.

Drawing the screen is kept apart from these clocks. A speed of 5 means five units of the app's time base per second of real time. Speed, frame rate and pauses change how fast a run is shown and never what it computes, so a phone and a workstation produce the same trajectory and the same file.

Noise comes from a deterministic pseudorandom generator. Its seed is declared by the app, or by a scenario that needs its own noise record, and it is restored whenever a scenario is loaded or Reset is pressed. The framework gives the reader no seed setting, so *default seed* in this book means the seed the app ships with. A run repeats bit for bit when the same scenario is loaded and the same settings are made before Run. A change made by hand during a run lands on a different step each time, which is one reason numeric answers carry a tolerance. The standalone simulators draw noise from the browser's own generator, so a noisy run there repeats in its statistics and not sample for sample.

Run and Pause start and hold the clock. Reset restarts the run at $t = 0$, paused, and keeps the reader's settings: every setpoint, tuning value, option, model parameter, disturbance slider and mode the reader changed is put back before the plant and the controller are initialized. Scripted events replay, and a one-shot disturbance button pressed during the run is cleared. The reload button beside

the scenario list restores the scenario's own settings instead.

1.6.2 The Velocity-Form PID Algorithm

One PID implementation serves the PID loops of every framework app. It is the velocity (incremental) form: each execution adds the change in the proportional and derivative terms to a running sum, and the integral is a separate term, as shown in Equations (1.19) to (1.22) for execution k and control period Δt.

$$e_k = a\,(\mathrm{SP}_k - \mathrm{PV}_k), \qquad P_k = K_c\,e_k, \qquad D_k = -a\,K_c\,\tau_D\,\frac{\mathrm{PV}_k^f - \mathrm{PV}_{k-1}^f}{\Delta t} \tag{1.19}$$

$$S_k = S_{k-1} + (P_k - P_{k-1}) + (D_k - D_{k-1}) \tag{1.20}$$

$$v_k = u_b + S_k + I_{k-1}, \qquad u_k = \mathrm{sat}(v_k) \tag{1.21}$$

$$I_k = I_{k-1} + \left[\frac{K_c}{\tau_I}\,e_k + \frac{u_k - v_k}{T_t}\right]\Delta t \tag{1.22}$$

where $a = \pm 1$ is the controller action, PV^f is the measurement after an optional first-order filter, u_b is the output bias, S is the running sum of proportional and derivative increments, v is the output before limiting, u is the output sent to the actuator, I is the integral term and T_t is the tracking time constant, equal to τ_I unless a loop sets another value. The function sat clamps v_k to the actuator range and then applies a rate limit if the loop declares one.

The limit acts on the emitted output u_k only. The sum S_k is never clamped, so it equals $P_k + D_k$ and the output before limiting is the position form $u_b + K_c e_k + D_k + I_{k-1}$. Anti-windup touches the integral alone. While the output is limited, the last term of Equation (1.22) drives I toward the value that reproduces the output applied, which is back-calculation [33]. With no integral action, Equation (1.22) is skipped. Clamping the running sum would instead discard each increment that crosses a limit and keep each one that returns, so symmetric noise would push the output off the limit even in a proportional-only loop.

The derivative acts on the measurement, so a setpoint step gives a proportional kick and no derivative kick. A change of K_c acts on the whole error at the next execution, as in the position form. The action $a = +1$ suits a process whose PV rises with the MV, a reverse-acting controller in the ISA convention, and $a = -1$ gives the direct-acting one.

In Manual the controller keeps executing and sets $I_k = u_{\mathrm{man}} - u_b - S_k$, so the output it would send equals the manual value u_{man}. The first automatic output differs from it by one period's change in P and D. In Auto the engine copies each controller output to the manual station, so a switch to Manual is also bumpless. The standalone studios with PID loops use velocity-form code of their own, with the limit on the output and anti-windup on the integral alone, in which a gain change applies only to later increments.

1.6.3 The Predictive Control and Estimator Options

More than 100 apps offer model predictive control (MPC) as an option beside the structure they open with, which stays the default. Volume 9 develops the theory. In most of them the controller plans with a small linear model in discrete time as shown in Equation (1.23).

$$x_{k+1} = A\,x_k + B\,u_k + E\,w_k + c, \qquad y_k = C\,x_k \tag{1.23}$$

where w is a measured or scheduled disturbance and c is a constant left by linearization. The model comes from a step test at the design point, from the app's balances linearized there or again at every decision, or from the model the reader fits. It never reads the true states or parameters of the run, so a parameter the reader changes in Show Model is model mismatch, as on a real plant.

At each decision the controller minimizes the objective of Equation (1.24) over a horizon of N steps.

$$J = \sum_{k=1}^{N} (y_k - r_k)^{\top} Q (y_k - r_k) + \sum_{j} \Delta u_j^{\top} R \Delta u_j + \sum_{i} \rho_i s_i^2 \tag{1.24}$$

where r is the setpoint, Q weights the tracking error, R is the move suppression on each input change Δu_j and ρ_i is the price of a slack $s_i \geq 0$ that lets an output band be broken when it cannot be held. Actuator limits are hard constraints. The quadratic program is solved in the browser by a dual active-set method [34]. The first move is applied, and the plan is made again at the next decision.

The plan starts from an estimate of the model state. About half of these apps use a steady-state Kalman filter on the model with an added disturbance state, which absorbs what the model did not predict and removes steady offset without an integral term [35]. Others correct the model output with a filtered bias to the same end. About a dozen apps also offer moving horizon estimation (MHE), which solves the estimation problem over a window of recent samples and can hold bounds that a Kalman filter cannot, such as a concentration that is never negative [36].

With the MPC selected, the tuning fields of the loops it replaces leave the Control panel. In most of these apps the Advanced level shows up to three settings, each a multiple of the shipped value: the prediction horizon, the move suppression and the band price. The trend draws the accepted plan ahead of the present time, and a list of constraint prices shows which limits bind. A plan that is not optimal, not finite or outside its limits is rejected, and a fallback, in most apps the shipped controller, holds the actuator until a plan succeeds. The Model Fitting and Tuning Studio builds the same kind of controller from the reader's own data in its tab, MPC from the fit.

1.6.4 Scenarios, Progress and Data

An app teaches through scenarios, the numbered cases at the top of the Control panel, nine to thirteen in most apps. Loading one sets the model parameters, the controller settings, the disturbances, the mode, the initial state and any scripted events. The information button beside the list opens a briefing in four parts: what the case is, why it matters, what to do and one detail for the practicing engineer. The Learn card holds the lesson back, with Predict and Try It before the run, hints on request and Explain afterward.

Each scenario has a completion rule that the app judges from the run on the simulated clock: a run length, the scripted events passed, a setting made by the reader, a state held for a stated time or a target met. Choosing a case earns nothing. A completed case is marked in the list, and a badge row shows the count and a level, Explorer after one case, Practitioner at half and Completed at all. Progress is kept in the reader's browser with no account. A certificate of progress or of completion can be issued from it, as a record of self-paced work and not an assessed credential.

One case, usually the last, is the graded challenge. It states a target, such as an error band to hold through a sequence of upsets. Changing the model parameters or the disturbances voids the attempt in most apps, and a record made with Reveal suggested tuning says so.

Show Model opens the equations, live values of the variables, the editable parameters and the assumptions. The trends plot SP and PV above the MV for each loop, with actuator limits and marks at setpoint moves, disturbances, mode changes and scripted events. The scorecard reports for each loop the integral of the absolute error (IAE), overshoot, settling time, MV travel and share of time at a limit, with the app's own measures, and it does not combine them into one score. IAE and travel accumulate from Reset, so two runs are compared at the same clock reading.

Export CSV saves the log as a text file with a header row of column names that carry their units and one row per logged sample. In most apps the time column is in seconds, minutes or hours according to the run length. The log holds the most recent 60 000 samples unless the app sets another size. Apps with the MPC option begin the file with comment lines that start with #, which a script must skip.

An app exercise in this book names the scenario by its number and name as the app shows them, every setting that differs from the scenario's defaults, the run length in simulated time, the seed

(the default seed unless stated), what to read and where (a trend, a scorecard tile or a column of the exported file) and a tolerance on a numeric answer. Settings are made before Run unless the exercise gives a time.

1.7 Working With AI Agents While Studying

A reader of this book is likely to have an AI assistant within reach, and the apps make its work easy to test. This section gives the editor's view of how to use one while working through a chapter. In short, the assistant may do the algebra and the scripting, and the reader keeps the checking.

What to delegate. Tasks with a clear input and a result that can be tested are good candidates. An assistant can derive the linearization of a chapter's balance at an operating point and the transfer function that follows from it. It can write a short script that reads the exported CSV file and fits a first-order-plus-dead-time model to a step test. It can propose a first tuning from that fit with a named rule, convert units, explain a term in a model or outline a report. The assistant knows nothing about the app until it is told. Give it the equations and parameter values from Show Model, the column names of the exported file and the scenario that produced the data, as described in Section 1.6.4. Ask it to state its assumptions and the units of every result.

What to check, and how. An assistant returns a fluent answer whether or not the answer is right, so each result needs a test that it could fail. Four tests cover most of the errors a chapter will produce.

- *Units.* Every term of an equation carries the same units, and a process gain is in PV units per MV unit. A gain computed with the valve as a fraction is 100 times the gain of a loop whose valve is in percent.
- *Signs.* Decide from the physics which way the PV moves when the MV rises. Compare that with the sign of the fitted gain and with the controller action of Section 1.6.2.
- *Steady state.* Set the derivatives to zero. The model must return the operating point that Show Model reports, and a linearization must give zero deviation there.
- *A run on the app.* Load the named scenario at the default seed, enter the proposed numbers and run. The trend and the scorecard decide. A tuning that reads well on paper and holds the valve at a limit for half the run has failed.

A fitted model gets one more test: compare it with data it was not fitted to, such as a second step in the opposite direction. When a check fails, report the failure to the assistant with the evidence and ask for the cause, then check the correction the same way.

How the exercises train the habit. One exercise in each chapter is marked *Review the agent*. It presents a tuning, a model or a claim as an AI assistant might propose it, with a flaw, and asks the reader to find the flaw in about five minutes. The flaws are ordinary ones: a unit, a sign, a steady state that does not close, a result measured on the wrong scenario or at a different run length, or a limit that was ignored. Each is found with one of the tests above. The purpose is not distrust. An engineer who can say why an answer is right can hand an assistant much larger tasks and still stand behind the result. That larger picture, in which agents do most of the routine work and the engineer supervises and checks, is the subject of Chapter 16.

1.8 Chapter Roadmap

Table 1.4 lists the chapters with the lesson each adds to this one. The level says how much of Section 1.4 a chapter assumes. The time is an estimate for reading the chapter once and running each case of its app once, at about eight minutes a case.

Chapters 3, 9, 13 and 15 form a first course: a cascade told about its load, a speed set by braking, a contact set by stiffness and a stop budgeted against a person. They rest on Sections 1.4.3, 1.4.8

Table 1.4 Roadmap of the Volume

Ch.	Plant	The lesson	Level	Time
2	Car on a kinematic bicycle	Speed enters twice and cancels in the Stanley law. Lag, offset and the rack remain.	intermediate	3 h
3	Passenger multirotor	Tell the loops the load and estimate the wind. The feedforward does the work.	introductory	2.5 h
4	Lift-plus-cruise aircraft	Stop the rotors on measured lift margin, not on a clock.	intermediate	2 h
5	Foiling boat	Choose which waves to reject. A cascade on heave rate contours.	intermediate	2.5 h
6	Vessel on thrusters	Part of a measurement is real and not an error. Structure before tuning.	advanced	3 h
7	CubeSat attitude	Wheels store every torque, and coils act only across the field.	advanced	2.5 h
8	Satellite in low orbit	Flown on predictions and a probability, not on a miss distance.	advanced	2 h
9	Pallet robot	Speed is set by believed braking and visible space.	introductory	2 h
10	Picking arm at a belt	Move old reports forward in time and meet the product at its speed.	intermediate	2 h
11	Table tennis arm	Commit before the prediction is certain, and swing within the time left.	intermediate	2.5 h
12	Arc welding cell	A plant that regulates itself and a sensor that reads places.	intermediate	2 h
13	Peg in a hole	Stiffness turns position error into force. A force loop has a delay margin.	introductory	2 h
14	Surgical instrument	Each layer of the stack has a measurable cost.	advanced	2.5 h
15	Collaborative arm	Stopping in time is a budget, checked against the truth.	introductory	2 h

and 1.4.9. Chapters 2, 4, 5, 10, 11 and 12 add path following, allocation, estimation and timing, and Sections 1.4.4 to 1.4.7 prepare them. Chapters 6, 7, 8 and 14 combine estimation, allocation and supervision on one machine and are best left until Section 1.4.10 is familiar. A reader interested in vehicles can read Chapters 2 to 9 in order, and a reader interested in arms Chapters 10 to 15.

KEY POINTS

- A moving machine reaches its task through geometry and inertia, and what matters is computed from estimates rather than read from one instrument.
- Each inner loop must be faster than the loop it serves, and a force the loops are not told about is opposed only through an error.
- A complementary filter uses each sensor where it is good, and a report taken some time ago describes where the machine was.
- Speed is bounded by the braking a machine believes it has and the space it can see, so the limit is only as good as that belief.

Where else this applies. Feedforward and cascade are the subject of Volume 3, allocation and overrides of Volume 4, estimation of Volume 5 and predictive control of Volume 9.

References

[1] B. Siciliano, L. Sciavicco, L. Villani, and G. Oriolo, *Robotics: Modelling, Planning and Control*, ser. Advanced Textbooks in Control and Signal Processing. London: Springer, 2009.

[2] T. I. Fossen, *Handbook of Marine Craft Hydrodynamics and Motion Control*. Chichester, UK: Wiley, 2011.

[3] F. L. Markley and J. L. Crassidis, *Fundamentals of Spacecraft Attitude Determination and Control*. New York: Springer, 2014.

[4] H. L. Hazen, "Theory of servo-mechanisms," *Journal of the Franklin Institute*, vol. 218, no. 3, pp. 279–331, 1934.

[5] 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.

[6] R. E. Kalman, "A new approach to linear filtering and prediction problems," *Journal of Basic Engineering*, vol. 82, no. 1, pp. 35–45, 1960.

[7] Deep Sea Drilling Project, "Dynamic positioning," Deep Sea Drilling Project, Tech. Rep., 1971.

[8] M. D. Holley, W. L. Swingle, S. L. Bachman, C. J. LeBlanc, H. T. Howard, and H. M. Biggs, "Apollo experience report: Guidance and control systems: Primary guidance, navigation, and control system development," National Aeronautics and Space Administration, Lyndon B. Johnson Space Center, Houston, TX, NASA Technical Note TN D-8227, May 1976. [Online]. Available: https://ntrs.nasa.gov/citations/19760016247

[9] S. V. Weston, C. D. Burkhard, J. M. Stupl, R. L. Ticknor, B. D. Yost, R. A. Austin, P. Galchenko, L. K. Newman, and L. Santos Soto, "State-of-the-art small spacecraft technology," National Aeronautics and Space Administration, Ames Research Center, Moffett Field, CA, Technical Publication NASA/TP-20250000142, Jan. 2025. [Online]. Available: https://ntrs.nasa.gov/citations/20250000142

[10] H. Makino, "Development of the SCARA," *Journal of Robotics and Mechatronics*, vol. 26, no. 1, pp. 5–8, 2014.

[11] M. D. Maisel, D. J. Giulianetti, and D. C. Dugan, *The History of the XV-15 Tilt Rotor Research Aircraft: From Concept to Flight*, ser. Monographs in Aerospace History. Washington, DC: National Aeronautics and Space Administration, NASA History Division, 2000, no. 17, nASA SP-2000-4517, https://ntrs.nasa.gov/citations/20000027499.

[12] N. Hogan, "Impedance control: An approach to manipulation: Part I, theory," *Journal of Dynamic Systems, Measurement, and Control*, vol. 107, no. 1, pp. 1–7, 1985.

[13] S. Thrun, M. Montemerlo, H. Dahlkamp, D. Stavens, A. Aron, J. Diebel, P. Fong, J. Gale, M. Halpenny, G. Hoffmann *et al.*, "Stanley: The robot that won the DARPA Grand Challenge," *Journal of Field Robotics*, vol. 23, no. 9, pp. 661–692, 2006.

[14] S. Haddadin, A. Albu-Schäffer, and G. Hirzinger, "Requirements for safe robots: Measurements, analysis and new insights," *The International Journal of Robotics Research*, vol. 28, no. 11-12, pp. 1507–1527, 2009.

[15] W. H. Clohessy and R. S. Wiltshire, "Terminal guidance system for satellite rendezvous," *Journal of the Aerospace Sciences*, vol. 27, no. 9, pp. 653–658, 1960.

[16] D. B. D'Ambrosio, S. Abeyruwan, L. Graesser, A. Iscen, H. Ben Amor, A. Bewley, B. J. Reed, K. Reymann, L. Takayama, Y. Tassa *et al.*, "Achieving human level competitive robot table tennis," in *2025 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2025, pp. 74–82.

[17] Y. S. Kwoh, J. Hou, E. A. Jonckheere, and S. Hayati, "A robot with improved absolute positioning accuracy for CT guided stereotactic brain surgery," *IEEE Transactions on Biomedical Engineering*, vol. 35, no. 2, pp. 153–160, 1988.

[18] 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.

[19] B. Wie and P. M. Barba, "Quaternion feedback for spacecraft large angle maneuvers," *Journal of Guidance, Control, and Dynamics*, vol. 8, no. 3, pp. 360–365, 1985.

[20] W.-H. Chen, J. Yang, L. Guo, and S. Li, "Disturbance-observer-based control and related methods: An overview," *IEEE Transactions on Industrial Electronics*, vol. 63, no. 2, pp. 1083–1095, 2016.

[21] R. Mahony, V. Kumar, and P. Corke, "Multirotor aerial vehicles: Modeling, estimation, and control of quadrotor," *IEEE Robotics & Automation Magazine*, vol. 19, no. 3, pp. 20–32, 2012.

[22] T. Flash and N. Hogan, "The coordination of arm movements: An experimentally confirmed mathematical model," *The Journal of Neuroscience*, vol. 5, no. 7, pp. 1688–1703, 1985.

[23] T. Kröger and F. M. Wahl, "Online trajectory generation: Basic concepts for instantaneous reactions to unforeseen events," *IEEE Transactions on Robotics*, vol. 26, no. 1, pp. 94–111, 2010.

[24] G. M. Hoffmann, C. J. Tomlin, M. Montemerlo, and S. Thrun, "Autonomous automobile trajectory tracking for off-road driving: Controller design, experimental validation and racing," in *2007 American Control Conference*, 2007, pp. 2296–2301.

[25] B. Paden, M. Čáp, S. Z. Yong, D. Yershov, and E. Frazzoli, "A survey of motion planning and control techniques for self-driving urban vehicles," *IEEE Transactions on Intelligent Vehicles*, vol. 1, no. 1, pp. 33–55, 2016.

[26] W. T. Higgins, "A comparison of complementary and Kalman filtering," *IEEE Transactions on Aerospace and Electronic Systems*, vol. AES-11, no. 3, pp. 321–325, 1975.

[27] T. A. Johansen and T. I. Fossen, "Control allocation: A survey," *Automatica*, vol. 49, no. 5, pp. 1087–1103, 2013.

[28] J. A. Marvel and R. Norcross, "Implementing speed and separation monitoring in collaborative robot workcells," *Robotics and Computer-Integrated Manufacturing*, vol. 44, pp. 144–155, 2017.

[29] T. R. Kuphaldt, *Lessons in Industrial Instrumentation*, version 2.33 ed. Open textbook, Creative Commons Attribution 4.0 license, ibiblio.org/kuphaldt/socratic/sinst, 2024.

[30] 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.

[31] S. Y. Lee, J. Y. Sim, D.-Y. Seok, Y. B. Kim, J. Shim, U. Kim, and H. R. Choi, "A redundant-sensing-based six-axis force/torque sensor enabling compactness and high sensitivity," *Sensors*, vol. 26, no. 3, p. 871, 2026.

[32] 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.

[33] K. J. Åström and L. Rundqwist, "Integrator windup and how to avoid it," in *Proc. 1989 American Control Conference*, 1989, pp. 1693–1698.

[34] D. Goldfarb and A. Idnani, "A numerically stable dual method for solving strictly convex quadratic programs," *Mathematical Programming*, vol. 27, no. 1, pp. 1–33, 1983.

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

[36] C. V. Rao, J. B. Rawlings, and D. Q. Mayne, "Constrained state estimation for nonlinear discrete-time systems: Stability and moving horizon approximations," *IEEE Transactions on Automatic Control*, vol. 48, no. 2, pp. 246–258, 2003.

1.9 Exercises

1.1 Concept check · 2 min

A joint servo forms its error as $e = a\,(\mathrm{SP} - \mathrm{PV})$ with a positive gain, where the PV is the measured joint angle. A positive torque turns the joint in the positive direction, so the torque must rise when the angle falls short of its reference. Which setting and name follow the convention of this book?

(a) $a = -1$, direct acting, because the torque rises with the error

(b) $a = +1$, reverse acting, because the torque falls as the angle rises

(c) $a = +1$, direct acting, because a positive error gives a positive torque

(d) $a = -1$, reverse acting, because the angle lags the reference

1.2 Theory and calculation · 6 min

A mobile robot of 80 kg has an inner speed loop $F = K_v(v_r - v)$ with $K_v = 400\,\mathrm{N\,s/m}$ and an outer position loop $v_r = K_x(x_r - x)$ with $K_x = 1\,\mathrm{s}^{-1}$.

(a) Find the inner time constant and the damping ratio of the position loop. *Checkpoint:* 0.20 s and 1.12.

(b) The robot picks up 80 kg. Find the new damping ratio, and the speed gain that restores the old one. *Checkpoint:* 0.79, and 800 N s/m.

1.3 Theory and calculation · 5 min

A gantry axis moves 0.8 m with a speed limit of 0.5 m/s and an acceleration limit of $1\,\mathrm{m/s^2}$. Find the shortest quintic (minimum-jerk) move and the trapezoidal move, and say which limit sets the quintic.

1.4 Theory and calculation · 6 min

A robot follows a path with the speed-divided law $\omega = -k_h[\psi_e + \arctan(k_e y/v)]$, with $k_h = 4\,\mathrm{s}^{-1}$ and $k_e = 1\,\mathrm{s}^{-1}$, and without the term $v\kappa$.

(a) Find the poles of the linearized lateral loop. *Checkpoint:* a double pole at $-2\,\mathrm{s}^{-1}$.

(b) Find the steady offset on a 10 m radius at 2 m/s and at 4 m/s, and the feedforward that removes it at 2 m/s. *Checkpoint:* 0.10 m, 0.40 m, 0.2 rad/s.

1.5 Theory and calculation · 5 min

A vehicle reacts in 0.3 s and brakes at 1.5 m/s^2. Find the fastest speed at which it can stop within 2 m, check it, and repeat for a braking capability of 0.75 m/s^2.

1.6 Theory and calculation · 6 min

A collaborative arm moves its tool at 0.5 m/s and brakes at 2 m/s^2. Its reaction time is 0.25 s and its intrusion allowance 0.1 m, and a person may walk in at 1.6 m/s. Find each term of the protective separation distance and the total, and the total if the person's walk during the stop is left out.

1.7 Theory and calculation · 5 min

A differential-drive robot with a track of 0.6 m and wheels limited to 40 N each is asked for a total force of 60 N and a moment of 12 N m. Find the wheel forces it asks for, then the force and moment it delivers if it serves the moment first and if it clips each wheel.

1.8 Concept check · 3 min

A first-order complementary filter fuses wheel odometry that reads 3 % high with a position fix, at a crossover of 0.5 rad/s, on a robot moving at a steady 1.5 m/s. Which statement is right?

(a) The estimate lags 0.045 m, and a higher crossover removes it

(b) The estimate has no steady error, because the fix corrects it

(c) The estimate leads 0.09 m, and only a perfect odometer removes it

(d) The estimate leads 0.09 m, and a bias state removes it

1.9 Theory and calculation · 8 min

A 2 kg end effector carries a peg.

(a) It lands at 10 mm/s on a part of stiffness 5 N/µm. Estimate the peak force, and the force a 3 N/mm spring and a 200 N/mm servo need to correct a 0.4 mm misalignment. *Checkpoint:* 31.6 N, 1.2 N, 80 N.

(b) An integral force law $\dot{x}_r = k_i(F_{SP} - F_m)$ with $k_i = 0.002\,\mathrm{m/(N\,s)}$ acts through a spring of 2000 N/m on a much stiffer part. Find the loop gain, the time constant and the largest measurement delay it tolerates. *Checkpoint:* 4 s^{-1}, 0.25 s, 0.39 s.

1.10 Concept check · 3 min

A camera reports a product on a belt moving at 0.4 m/s with a latency of 0.12 s, and the gripper takes 0.08 s to close after it arrives. The arm is aimed at the reported position. What happens, and what corrects it?

(a) It closes 80 mm behind, and predicting from the time stamp corrects it

(b) It closes 48 mm behind, and a faster camera alone corrects it

(c) It closes 32 mm behind, and a faster gripper alone corrects it

(d) It closes on the product, because the arm tracks the camera

CHAPTER 2

Lane Keeping on a Kinematic Bicycle: Speed Dependence of a Stanley Path-Following Loop

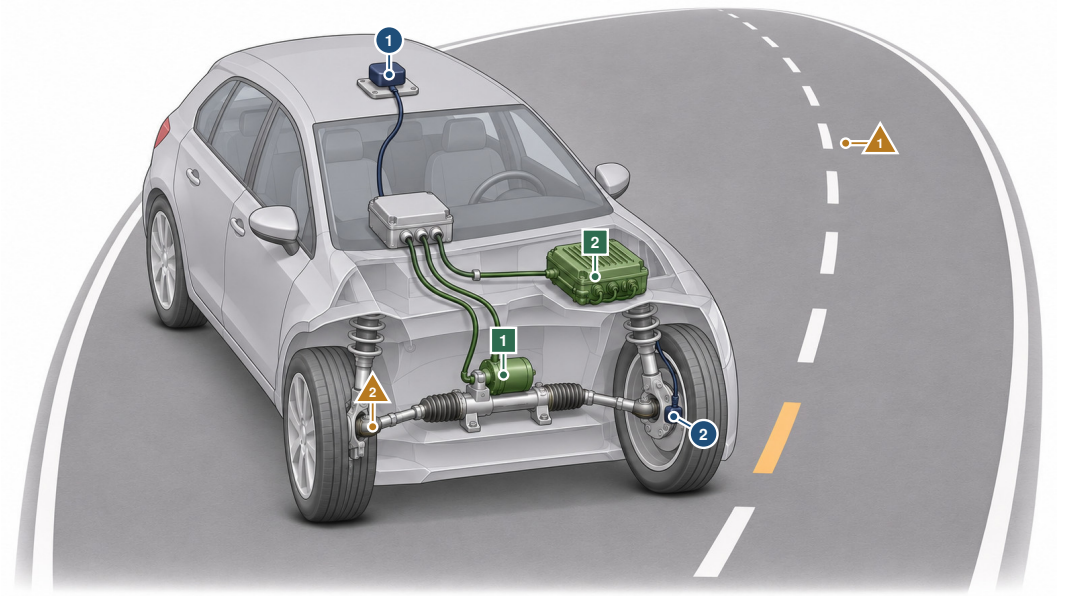

The controller moves 1 the steering motor on the rack, which turns the front wheels
2 the drive and the brakes, which set the speed

The controller reads 1 the position module on the roof: where the car sits across the lane
2 the wheel-speed sensor: how fast the car goes

What disturbs it 1 the bend ahead, known from the road map, which the steering must follow
2 a steering linkage set slightly crooked

Notice the motor on the rack between the front wheels: every correction the controller makes passes through it, with its lag and its end stop. Illustration, not to scale.

IN PLAIN TERMS **Keeping a Car in Its Lane at Any Speed**

The main idea. A lane-keeping controller turns the front wheels to hold a car in the middle of its lane. The faster the car goes, the more a turn of the wheels swings it sideways. So the controller asks for less steering at higher speed.

The challenge. That trick cannot speed up the steering motor, which always takes the same short time to respond. A fast car travels farther in that time. No amount of correction holds a bend tighter than the wheels can turn, and a crooked wheel leaves the car off center.

Why it matters. Systems that help drivers stay in their lane are now fitted to a large share of new cars.

THREE WAYS TO READ THIS CHAPTER

⊘	**Explore**	8 min	the first page, the fact on lane departure warning (Section 2.1), the history of the steering law (Section 2.2), the try-it segment and the fact on one step of gain (Section 2.6), and the key points
🕮	**Study**	60 min	the whole chapter with its three pause points and the twelve exercises at its end
⛰	**Practice**	20 min	the first page, the results (Section 2.6), Technology in Practice (Section 2.7), Limitations (Section 2.8) and the open challenge

The control question A steering angle sets a yaw rate in proportion to the speed. How much of that the Stanley law cancels by construction, and what limits one fixed tuning once it has?

Learning objectives

- Derive the two closed-loop poles of the linearized Stanley loop and explain why only one is free of speed.
- Predict the standing offset left by a missing curvature feedforward and by a misaligned steering linkage.
- Compute the gain at which the rack lag destabilizes the loop, and explain why the sampled controller lowers it.
- Tell an actuator that cannot deliver the required angle from a badly tuned loop.
- Compare the Stanley law with a predictive controller that estimates the misalignment, and name the model error that undoes it.

Prerequisites Sections 1.4.2, 1.4.5, 1.6.2 and 1.6.4. Section 1.6.3 is needed only for the predictive controller.

The app https://apmonitor.com/apps/lane-keeping/

See also Volume 7, Chapter 9 for a vehicle whose stability is set by its speed and Volume 7, Chapter 10 for tire forces at the friction limit. Volume 9, Chapter 4 for offset-free predictive control with an estimated bias.



2.1 Introduction

Most loops in a process dynamics and control course regulate a scalar measurement against a scalar setpoint, and the plant between the actuator and the measurement is a transfer function that a step test can identify [1]. Lane keeping is constructed differently. The error is not read from an instrument, it is computed by projecting a point on the vehicle onto a curve stored in a map, and the actuator does not act on that error directly because a steering angle sets a yaw rate rather than a lateral position. Two integrations separate the manipulated variable from the controlled variable, and the coefficient in front of both of them is the vehicle speed.

That speed dependence is the reason lane keeping merits a chapter of its own. The yaw rate a given road-wheel angle produces is $\dot{\psi} = (v/L)\tan\delta$, so a vehicle traveling at 30 m/s turns three times as fast for the same steering angle as one traveling at 10 m/s, and it covers three times the distance while any fixed-duration transient completes. A controller tuned at one speed therefore operates on a different plant at another, and the standard industrial remedy is a gain schedule indexed by the operating point [1, 2]. Volume 7, Chapter 9 treats a vehicle on which the forward speed decides whether it is stable at all.

The Stanley steering law addresses the same difficulty by construction. It divides the cross-track correction by the speed, so the correction the law requests falls as $1/v$ while the yaw rate that correction produces rises as v, and the two effects cancel. The cancellation is not an interpolated table, it is a model of the plant gain written into the control law, and it holds exactly for the correction term because the tangent of an arctangent is the identity. The cancellation does not extend to the effects the law omits. The steering actuator has a fixed lag in seconds, which at three times the speed costs three times the distance. The lateral acceleration for a given path grows as v^2. The small-angle relations behind the geometry degrade as the lateral acceleration rises.

The lane-keeping loop carries a specific set of process dynamics and control principles. Cross-track regulation on a kinematic bicycle is a double-integrating structure whose loop gain is an explicit

function of an operating variable, which makes it a concrete instance of process order and of gain scheduling. The curvature term is a feedforward from a measured reference rather than from a load, so it separates feedforward action from feedback action in a setting where the required steady input is known in closed form. The steering rack contributes actuator lag, a rate limit and an end stop, which are three distinct actuator limitations with three distinct signatures in the response. The absence of integral action in the geometric law produces a steady-state offset whose size follows from the proportional balance, and the sampled controller executing at 50 Hz places a half-sample of additional lag in the loop. The physical fundamentals are the no-slip rolling constraint at each axle, which fixes the instantaneous center of rotation and therefore the achieved curvature, and the time-scale separation between the fast heading mode and the slower lateral mode. These are the same principles that govern process loops built on material and energy balances [1, 3, 4], applied here to a system whose geometry is directly observable.

The behavior matters industrially because lane departure warning and lane keeping assistance are now fitted to a large fraction of new passenger vehicles, and their measured benefit is a reduction in police-reported crash rates rather than a tracking error in centimeters [5]. The same geometric tracker with the same three actuator limits runs on agricultural guidance, on warehouse vehicles and on mining haul trucks.

> + DID YOU KNOW? **Fewer Crashes With a Lane Warning**
>
> Police-reported crashes from 25 states of the United States, 2009 to 2015, were compared for vehicles with and without lane departure warning. After accounting for the drivers, equipped vehicles had 11 % fewer of the single-vehicle, sideswipe and head-on crashes the warning addresses, and 21 % fewer of those with injuries [5].

Textbook treatments of lateral control state the Stanley law and its stability argument and proceed to a comparison of tracking errors [6]. However, a table of tracking errors does not allow a fixed tuning to be held while the speed is tripled, and it does not separate a loop that is badly tuned from an actuator that cannot deliver the angle the road demands. The purpose of this chapter is to document a lane-keeping simulation on a kinematic bicycle and to quantify what one fixed Stanley tuning does and does not achieve as the vehicle speed, the actuator lag, the steering limit and the lateral measurement quality are each varied in turn. The kinematic bicycle and the Stanley law are both long established, and Section 1.4.5 introduces the speed-divided law on a transport cart, beside a path follower with fixed gains.

Section 2.2 traces automatic steering from buried wires to the Stanley law. Section 2.3 gives the five state equations of the simulator, its road construction, its instrument chain, its hard limits and its sample periods, so that every number in this chapter can be regenerated from the text. Section 2.4 shows that the ideal closed loop carries two real poles, only one of which is independent of speed. It derives in closed form the standing offset left by a missing feedforward, by a heading error referenced to the wrong point on the vehicle and by a misaligned steering linkage, together with the cross-track gain at which the rack lag destabilizes the loop. Section 2.5 sets out the discrete steering algorithm, the speed loop and a model predictive controller offered as a second steering law. Section 2.6 reports twelve configurations, each varying one property of the road, the vehicle or the loop, with outcomes taken from the performance metrics the application computes while it runs. One of them sets the model predictive controller against the fixed Stanley gain, and the comparison is repeated on the others. Section 2.7 names the sensors, the steering actuator and the regulation of a real system, Section 2.8 states what the model leaves out, and Section 2.9 draws the results together.

2.2 Background

Automatic steering was applied on the road well before automatic driving. Agricultural tractors and warehouse vehicles followed buried wires and painted stripes for decades using analog circuits

that closed a loop around a single lateral sensor mounted at the front of the machine. The choice of that mounting point, ahead of the axle that generates the turn, is the essential idea behind every geometric path follower that followed. Highway experiments in the second half of the twentieth century extended the same approach to passenger cars guided by cables laid in the pavement, which removed the sensing problem by building the reference into the road.

Replacing the buried reference with a camera converted lane keeping from a civil engineering problem into a control problem. Recursive estimation of road curvature and vehicle state from an image sequence, formulated as a dynamic model with the road geometry among the estimated quantities, allowed a vehicle to derive its own lateral reference at highway speed on ordinary pavement [7]. The law this chapter teaches comes from the desert race of 2005, as the history box below tells. The urban race two years later moved the emphasis to planning and to behavior around other traffic, while the lateral controller remained a geometric tracker following a curve with a known curvature [8].

> ⧖ HISTORY **A Steering Law From a Desert Race (2005)**
>
> The Defense Advanced Research Projects Agency (DARPA) launched its Grand Challenge in 2003 to spur work on unmanned ground vehicles. In the first race, in March 2004, no robot covered more than 5 % of a 142-mile desert course. The race was repeated on October 8, 2005, and five of the 23 starters finished. The winner, a robot from Stanford University, took 6 h 53 min 58 s. Its steering law pointed the front wheels parallel to the path and added the arctangent of the cross-track error, measured at the center of the front wheels, divided by the speed, so that a small error decays exponentially [9]. The law takes the orientation of the front wheels, not of the body, against the path. On the kinematic equations it was proven globally asymptotically stable, and in the desert it held the cross-track error typically under 0.1 m RMS [10].

Production lane keeping developed along a different route and for a different purpose. Lane departure warning systems measure the same lateral offset but act on the driver rather than on the steering rack, and lane keeping assistance applies a bounded steering torque that the driver can override at any time. Their value is quantified in crash rates rather than in centimeters of tracking error, and the crash rates of equipped vehicles have been compared against otherwise similar unequipped ones in police-reported data [5]. The bounded authority is deliberate. A system that performs most of the steering without visible limits reduces the driver’s supervision of it, and the pattern of use, misuse and disuse that follows from automation whose limits are not visible to its operator is one of the oldest findings in the human factors literature [11].

The control literature on path following has developed alongside both of these. Pure pursuit steers toward a point a fixed distance ahead on the path, which is simple to implement and whose stability depends on that lookahead distance in a way that was characterized early [12] and is still being refined by making the lookahead a function of speed and curvature [13]. Chapter 9 follows its routes by pure pursuit. Surveys of motion planning and control for self-driving vehicles place the geometric trackers, the kinematic model predictive controllers and the dynamic ones on a common footing [14], and direct comparisons of lateral controllers on the same vehicle and the same path report where each one begins to fail [6].

The regulatory structures of the simulator are standard ones. The Stanley law is a proportional geometric feedback with an additive feedforward and no reset action, so the steady-state offset arguments that apply to a proportional-only process loop apply here unchanged [1, 3]. The speed loop is a proportional-integral controller on a pure integrator, which is the same structure as level control on a surge drum. The actuator chain places a first-order lag, a rate limit and a position limit in series with the manipulated variable, and the stability consequence of the lag follows from the Routh criterion applied to the resulting cubic [4].

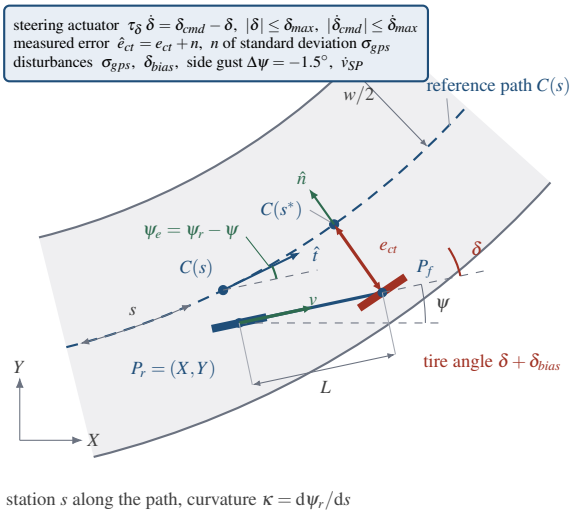

Figure 2.1 Lane keeping geometry, carrying every symbol that enters the state equations and the steering law. The cross-track error is measured at the front axle P_f and the heading error against the road direction at $C(s)$, the projection of the rear axle.

2.3 Vehicle and Road Model

The simulated system is one passenger car on one lane of a flat road. The geometry, the states, the manipulated variable, the disturbances and the measurement points are shown in Figure 2.1.

2.3.1 Signals, Units and Ranges

The signals divide into two controlled variables, two manipulated variables, five states and five disturbances, and their roles, units and permitted ranges are given in Table 2.1. Two loops act on one vehicle. The lane loop regulates a length and the speed loop regulates a speed, and either loop or both may be placed in automatic independently, so the longitudinal input can remain operator-set while the steering is closed. In the action convention of Section 1.3 the speed loop is reverse acting, because the speed rises with the acceleration command. The lane loop has no action setting in the app. Its law raises the steering angle as e_{ct} rises, and a positive steering angle lowers e_{ct}, so on e_{ct} it is a direct-acting loop.

2.3.2 Kinematic Bicycle and Steering Actuator

The two front wheels are collapsed onto the center of the front axle and the two rear wheels onto the center of the rear axle. The wheels are assumed to roll without lateral slip, so the velocity of each axle center lies along the plane of its own wheel, and that rolling constraint places the instantaneous center of rotation at the intersection of the two axle lines. With (X,Y) the rear-axle position, ψ the heading measured counter-clockwise from the world X axis, v the rear-axle speed and δ the road-wheel angle, the vehicle obeys the four relations of Equation (2.1), the kinematic bicycle.

$$\dot X = v\cos\psi, \quad \dot Y = v\sin\psi, \quad \dot\psi = \frac{v}{L}\tan\delta_{\mathrm{eff}}, \quad \dot v = a \tag{2.1}$$

where L is the wheelbase, a is the acceleration delivered to the vehicle and $\delta_{\mathrm{eff}} = \mathrm{clamp}(\delta + \delta_{bias}, -\delta_{max}, \delta_{max})$ is the angle actually taken by the tire. A misaligned linkage adds δ_{bias} between the angle the actuator reports and the angle the wheel takes, and no geometric measurement in the loop observes it. The third relation is the only steering channel in the model and it is the origin of every speed-dependent effect in the application, because the yaw rate it produces is proportional to v.

Table 2.1 Signal Roles, Units and Permitted Ranges

Symbol	Role	Unit	Range or default
e_{ct}	CV, lane loop	m	setpoint 0, band ± 0.10
v	CV, speed loop	m/s	setpoint 5 to 35, default 25
δ	MV, lane loop	rad	$\lvert\delta\rvert \le \delta_{max}$
a_{cmd}	MV, speed loop	m/s^2	-6.0 to $+3.0$
ψ_e	measured input to the law	rad	wrapped to $\pm\pi$
κ	measured input to the law	1/m	from the road table
X, Y	state, rear axle	m	free
ψ	state, heading	rad	free
v	state, speed	m/s	$v \ge 0$
δ	state, road-wheel angle	rad	$\lvert\delta\rvert \le \delta_{max}$
σ_{gps}	DV, measurement	m	0 to 0.6, default 0
δ_{bias}	DV, plant	deg	-2 to $+2$, default 0
$\dot{v}_{SP}$	DV, setpoint ramp	m/s^2	-2 to $+2$, default 0
$\Delta\psi$	DV, side gust event	deg	-1.5
Δy	DV, displacement event	m	$+0.5$ to the right
δ_{man}	operator-set steering	deg	-30 to $+30$
pedal	operator-set longitudinal	%	-100 to $+100$

The steering rack is a first-order lag with an end stop and a rate limit on its command, as shown in Equation (2.2).

$$\tau_\delta \dot{\delta} = \delta_{cmd} - \delta, \qquad |\delta| \le \delta_{max}, \qquad |\dot{\delta}_{cmd}| \le \dot{\delta}_{max} \tag{2.2}$$

where τ_δ is the rack time constant, δ_{max} is the mechanical lock and $\dot{\delta}_{max}$ is the largest rate at which the command may move. These are three separate limits with three distinct effects on the loop. The lag adds phase, the rate limit attenuates fast corrections while leaving slow ones unaltered, and the end stop removes an angle that the geometry requires. The acceleration is clamped to the traction and braking envelope, $a = \mathrm{clamp}(a_{cmd}, a_{min}, a_{max})$, and braking cannot reverse the car, so $\dot{v} = 0$ whenever $v \le 0$ and $a < 0$.

2.3.3 Road Geometry from a Curvature Profile

Every road in the application is defined by a curvature profile $\kappa(s)$ over a sequence of segments, and the four segment kinds that appear are the straight ($\kappa = 0$), the constant-radius arc ($\kappa = 1/R$ with $R > 0$ turning left), the clothoid whose curvature ramps linearly from κ_0 to κ_1 over its length, and the lateral shift, a sinusoidal curvature pulse $\kappa = (2\pi\Delta y_{seg}/\ell^2)\sin(2\pi u/\ell)$ over a segment of length ℓ with u the distance into the segment. Integrating that pulse returns the heading to its original value at $u = \ell$ and, for small heading excursions, displaces the path sideways by Δy_{seg}. Expressing a lane change as geometry rather than as a special case in the controller allows one projection routine to serve every road.

The profile is integrated into an arc-length table at a node spacing of $\Delta s = 0.5\,\mathrm{m}$ using the exact constant-curvature chord over each step, as shown in Equation (2.3).

$$\begin{aligned} \psi_{r,k+1} &= \psi_{r,k} + \bar{\kappa} h, \\ x_{k+1} &= x_k + \frac{\sin\psi_{r,k+1} - \sin\psi_{r,k}}{\bar{\kappa}}, \\ y_{k+1} &= y_k - \frac{\cos\psi_{r,k+1} - \cos\psi_{r,k}}{\bar{\kappa}} \end{aligned} \tag{2.3}$$

where h is the step to the next node or to the end of the current segment, whichever comes first, and $\bar{\kappa}$ is the curvature at the midpoint of that step. The straight-line form $x_{k+1} = x_k + h\cos\psi_{r,k}$,

Table 2.2 The Eight Shipped Roads, Given as Segment Sequences

Road	Segments (lengths in m)
Straight highway	straight 2600
Constant radius	straight 80, arc $R = 200$ for 700, straight
S-curve	straight 50, corner $R = 150$ (140), straight 40, corner $R = -150$ (140), straight 60, corner $R = 180$ (160)
Lane change	straight 70, shift $+3.5$ over 60, straight 140, shift -3.5 over 60, straight
Gentle sweep	straight 120, corner $R = 500$ (500) with 60 m spirals, straight 200, corner $R = -400$ (400) with 60 m spirals
Sudden curvature	straight 120, arc $R = 120$ for 260, straight 120, arc $R = -120$ for 260, straight
Sharp corner	straight 60, corner $R = 25$ (40), straight 120, corner $R = -25$ (40), both with 12 m spirals, straight
Challenge course	straight 40, clothoid to $1/140$ over 30, arc $R = 140$ for 130, clothoid to $-1/160$ over 40, arc $R = -160$ for 200, clothoid to 0 over 40, straight

A *corner* is a clothoid in, a constant arc and a clothoid out, with 30 m spirals unless stated. Every road is padded to 2600 m.

$y_{k+1} = y_k + h \sin \psi_{r,k}$ replaces it when $|\bar{\kappa}| < 10^{-12}$. Every road starts at the origin heading along $+X$ and is padded with a final straight to 2600 m, which is 100 s at 26 m/s. The eight shipped roads are listed in Table 2.2.

2.3.4 Cross-Track and Heading Error Definitions

The vehicle carries two independent projections onto the road table, one for the front axle and one for the rear. Each searches only the window of nodes between $s - 3$ m and $s + 25$ m around its own last station, refines the station along the tangent and reports the normal offset. At reset both trackers start from station zero, where every road begins, and a road changed under a moving vehicle carries the vehicle across at its current station. No global nearest-point search enters the loop, because on a road that doubles back the true nearest point can be a kilometer away along the path and a controller referenced to it would steer toward the wrong lane.

The front-axle position and the two errors are given by Equation (2.4).

$$\begin{aligned} P_f &= \big(X + L\cos\psi,\ Y + L\sin\psi\big), \\ e_{ct} &= \big(C(s^*) - P_f\big) \cdot \hat{n}(s^*), \\ \psi_e &= \mathrm{wrap}\big(\psi_r(s) - \psi\big), \qquad \kappa = \frac{\mathrm{d}\psi_r}{\mathrm{d}s} \end{aligned} \tag{2.4}$$

where s^* is the station of the front-axle projection, s is the station of the rear-axle projection, $\hat{n} = (-\sin\psi_r, \cos\psi_r)$ is the left normal to the path and wrap folds an angle into $[-\pi, \pi]$. The sign convention follows from the minus sign in the second relation. A positive e_{ct} means the vehicle sits to the *right* of the path, which is exactly the case that calls for a positive left steering correction, and a vehicle at $Y = -1$ m on a straight centerline reports $e_{ct} = +1$ m. A positive ψ_e means the road heads to the left of where the car is pointing. Both errors therefore call for a positive correction, which is why the law below can be written with a plus sign in front of every term.

The two reference points differ deliberately. The cross-track error is taken at the front axle, which is the point the steering angle displaces most directly, while the heading error is taken against the road direction under the vehicle. Referencing both to the front-axle projection folds the corner into the heading term, because the road heading at s^* exceeds the road heading at s by approximately $L\kappa$. Section 2.4 derives the effect on the feedforward.

Only the cross-track error is instrumented. The measurement adds a band-limited disturbance, as shown in Equation (2.5).

$$\hat{e}_{ct} = e_{ct} + \sigma_{gps}\, n, \qquad \tau_n \dot{n} = -n + \sqrt{2\tau_n}\, w \tag{2.5}$$

Table 2.3 Constants of the Plant, the Actuator and the Loops

Symbol	Meaning	Value	Adjustable range
L	wheelbase	2.9 m	2.2 to 3.8
$w/2$	lane half-width	1.75 m	1.25 to 2.5
w_{car}	vehicle width	1.86 m	fixed
ℓ_{car}	vehicle length	4.60 m	fixed
δ_{max}	steering end stop	30°	3 to 30
$\dot{\delta}_{max}$	steering rate limit	45 °/s	5 to 120
τ_δ	steering rack lag	0.10 s	0.02 to 0.60
a_{max}	traction limit	3.0 m/s^2	0.5 to 3.0
a_{min}	braking limit	−6.0 m/s^2	−6.0 to −1.0
k_e	cross-track gain	1.2/s	0.1 to 12
k_{soft}	softening constant	1.0 m/s	0.1 to 10
δ_{ff}	curvature feedforward	on	on / off
	steering law	Stanley	Stanley / MPC
T_p	MPC prediction horizon	2 s	0.5 to 4
λ	MPC move suppression	0.01	0.001 to 100
	MPC road preview	on	on / off
K_c	speed loop gain	1.2/s	0.1 to 5
τ_I	speed loop integral time	4 s	0.5 to 30
v_{SP}	speed setpoint	25 m/s	5 to 35
Δt	integration step (RK4)	0.01 s	fixed
Δt_c	control period	0.02 s	fixed
f_{log}	historian rate	20 Hz	fixed
Δs	road node spacing	0.5 m	fixed
τ_n	noise correlation time	0.35 s	fixed

The lane is therefore 3.5 m wide. Random draws are seeded per configuration, at 20260824 for the noisy-measurement configuration and at 20260901 for the challenge, so a noisy run replays identically.

where w is unit white noise, $\tau_n = 0.35\,\text{s}$ and the filter gain is chosen so that $\sigma_{gps} n$ carries standard deviation σ_{gps} once the filter has reached its stationary state. The noise is colored rather than white because a receiver error is correlated over a second or two instead of being independent at each sample. The heading, the road heading, the curvature and the speed reach the controller without error, and the standing offset caused by δ_{bias} is a plant misalignment rather than an instrument fault.

A lane departure is declared, and recorded in the logged data, when Equation (2.6) holds.

$$\max\left(|e_{rear}|, |e_{ct}|\right) + \tfrac{1}{2} w_{car} > \tfrac{1}{2} w \tag{2.6}$$

where w_{car} is the vehicle width and w the lane width. The flag clears again at 0.05 m of hysteresis inside the limit. Both ends of the vehicle are tested, because a front corner can cross the line while the rear axle is still inside it.

2.3.5 Parameters and Numerics

Table 2.3 gives every constant of the plant and of the shipped loops, together with the adjustable settings of the predictive controller, whose fixed weights, prices and variances are given in Section 2.5.5. The first block is the plant, the second is the actuator, the third is the controller and the fourth is the numerics. The vehicle width and length are fixed rather than adjustable, and the width enters the departure test of Equation (2.6) rather than the displayed geometry alone.

A classical fourth-order Runge–Kutta scheme advances the five states on a fixed step of 10 ms. The controller executes every 20 ms, a rate representative of a production lane-keeping electronic control unit, and the data historian logs at 20 Hz. The order of events, control and integration within a step is

the framework's (Section 1.6.1), with the controller on every second step. After each integration the app clamps v at zero and δ at the end stop and tests Equation (2.6). The 20 ms control period places a half-sample of additional lag in the steering loop, which is small against the rack lag at the default setting and is not small against the stability boundary of Section 2.4. The application opens with the vehicle on the centerline at its cruise speed, which is an equilibrium of Equation (2.1) with every derivative that should vanish equal to zero to within 10^{-12}.

2.3.6 Assumptions

No tire model and no sideslip. The kinematic bicycle assumes each wheel rolls in the direction it points, so there is no cornering stiffness, no slip angle and no load transfer, and the path curvature is exactly $\tan\delta/L$. That is a fair representation up to roughly 0.4 g of lateral acceleration on dry asphalt. Above that a real vehicle requires a slip angle to generate the cornering force, the front tires saturate before the rear, and the achieved curvature falls below the commanded one, which is understeer. The application reports the lateral acceleration as $(v^2/L)\tan\delta_{\text{eff}}$. The sharp-corner configuration of Section 2.6 demands $v^2/R = 4.0\,\text{m/s}^2$ and, with the steering pinned at its 6° limit, delivers $3.62\,\text{m/s}^2$, which is the operating region where the kinematic assumption begins to overstate the achievable curvature. The boundary between the kinematic bicycle and a dynamic bicycle carrying tire forces has been examined directly [15, 16], and the lateral dynamics that govern beyond it are standard material [17].

No coupling between the two loops. Longitudinal motion is $\dot{v} = a$ with no drag, no rolling resistance, no grade, no driveline lag and no traction budget shared between turning and braking. A real vehicle cannot brake at $6\,\text{m/s}^2$ and corner at the adhesion limit simultaneously. The speed loop in this application is therefore simpler than a real one by design, so that the configurations isolate the steering.

Perfect localization apart from the injected noise. The heading, the road heading and the curvature are known exactly, and the only measurement error is σ_{gps} on the cross-track term. Constructing the lateral estimate is the principal difficulty in a real lane-keeping system, and that difficulty is absent here.

One flat lane. There is no camber, no superelevation, no width variation, no other traffic and no road grade, and the reference path is the exact centerline rather than something a perception system had to infer.

One first-order actuator. Real electric power steering adds compliance, friction and a torque limit that varies with speed, and the assist map is itself a function of vehicle speed. The model keeps one lag, one rate limit and one end stop, which is enough to produce the three distinct failure modes of Section 2.6 and no more.

2.4 Dynamic Analysis

2.4.1 Steering Law and Ideal Closed-Loop Response

The Stanley law adds three terms, as shown in Equation (2.7).

$$\delta_{law} = \underbrace{\psi_e}_{\text{heading}} + \underbrace{\arctan\left(\frac{k_e\,\hat{e}_{ct}}{v + k_{soft}}\right)}_{\text{cross-track}} + \underbrace{\arctan(L\kappa)}_{\text{feedforward}} \tag{2.7}$$

where k_e is the cross-track gain and k_{soft} keeps the argument finite as v goes to zero. The argument of the arctangent is dimensionless, because k_e carries units of inverse seconds and $\hat{e}_{ct}/(v + k_{soft})$ carries seconds. Neither feedback term regulates the lane position alone. The heading term drives ψ_e to zero, which leaves the vehicle traveling parallel to the lane at whatever offset it holds. The cross-track term must first build a heading error before the vehicle moves laterally and has no mechanism for removing that heading error afterward, so it overshoots and oscillates.

Linearizing Equation (2.1) and Equation (2.4) about a straight road for small angles, with $\psi_r = 0$ so that $\psi_e = -\psi$, gives the pair in Equation (2.8).

$$\dot{e}_{ct} = -v(\delta - \psi_e), \qquad \dot{\psi}_e = -\frac{v}{L}\delta \tag{2.8}$$

Substituting Equation (2.7) with $\kappa = 0$ and an instantaneous actuator, and writing $k' = k_e/(v + k_{soft})$ for the linearized cross-track slope, the heading term cancels identically and what remains is Equation (2.9).

$$\dot{e}_{ct} = -k' v\, e_{ct} = -\frac{k_e\, v}{v + k_{soft}}\, e_{ct} \tag{2.9}$$

The cross-track error decays as a first order lag with time constant $(v + k_{soft})/(k_e v)$, which is 0.917 s at 10 m/s and 0.861 s at 30 m/s for the default $k_e = 1.2/\text{s}$ and $k_{soft} = 1.0\,\text{m/s}$. A threefold change in speed moves the closed-loop time constant by 6 %, and the whole of that 6 % is the softening constant. With $k_{soft} = 0$ the time constant is exactly $1/k_e$ at every speed.

Pause. Try Exercise 2.1 before reading on.

Eliminating δ from Equation (2.8) rather than substituting gives the full ideal characteristic polynomial $(s + k_e^*)(s + v/L)$ with $k_e^* = k_e v/(v + k_{soft})$. The loop has two real poles and only one of them is speed-invariant. The cross-track pole is at $-k_e^*$, close to $-1.2/\text{s}$ at every highway speed, and the heading pole is at $-v/L$, which is –3.45/s at 10 m/s and –10.34/s at 30 m/s. The heading mode is the faster of the two and it accelerates with speed, which is the reason the two-term law produces no overshoot. The time-scale separation between the two modes widens as the speed rises, because speed multiplies the term that removes the heading error at the same time as it is divided out of the term that creates it. The law of Equation (1.7) commands a turn rate through a separate heading gain k_h, and both of its poles are free of speed. The car's law commands a steering angle with a unit heading term, so its heading pole is $-v/L$ and moves with speed, while the cross-track pole stays put in both.

2.4.2 Effect of the Arctangent on Loop Gain

The exact cancellation lives in the cross-track term alone. Because $\tan(\arctan x) = x$, the yaw rate that term produces is $(v/L) \cdot k_e e_{ct}/(v + k_{soft})$, which tends to $k_e e_{ct}/L$ as v grows past k_{soft} with no small-angle approximation anywhere. Both halves of that statement are plotted in Figure 2.2. The correction the law asks for at the lane edge, $e_{ct} = 1.75\,\text{m}$, is 19.29° at 5 m/s and 3.88° at 30 m/s, a factor of five across a factor of six in speed, while the yaw rates those two corrections produce differ by only 16 %.

The arctangent is commonly described as providing large-error saturation, and for this application that description requires qualification. At 25 m/s the argument reaches unity, where the arctangent first departs appreciably from its tangent line, only at $e_{ct} = (v + k_{soft})/k_e = 21.7\,\text{m}$. Inside the lane the argument is at most 0.081 and the arctangent returns 99.8 % of the linear value, so at highway speed the term acts as a proportional gain. The saturation is significant at low speed, where the same crossover error falls to 2.5 m at 2 m/s, and after a large excursion, where it bounds the commanded angle by 90° rather than allowing a large error to demand an angle with no physical meaning.

2.4.3 Curvature Feedforward and the Heading Reference Point

Holding a corner of radius R requires a sustained steering angle $\arctan(L/R)$, and Equation (2.10) confirms that this angle turns the vehicle at exactly the rate the road turns.

$$\delta_{ff} = \arctan(L\kappa) \implies \dot{\psi} = \frac{v}{L}\tan\delta_{ff} = v\kappa \tag{2.10}$$

The demand is set by the reference rather than by a load, and with the feedforward switched off the angle is produced only by allowing a cross-track error to grow until the arctangent term supplies it. In steady cornering the vehicle travels on a circle concentric with the path, so its heading is parallel

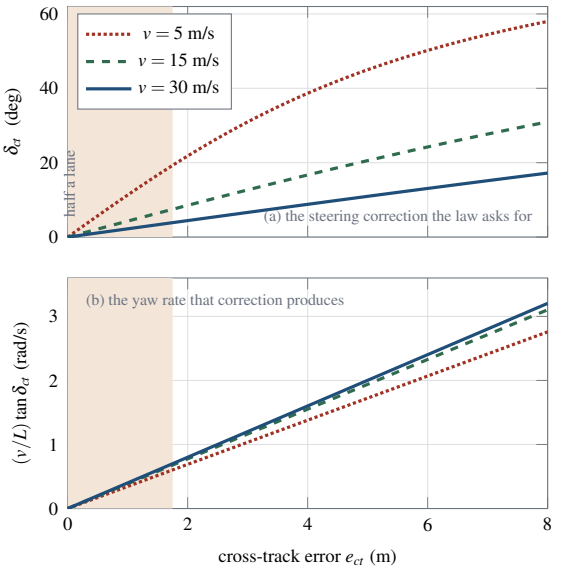

Figure 2.2 Speed entering the loop twice, at $k_e = 1.2/\text{s}$ and $k_{soft} = 1.0\,\text{m/s}$. The correction term $\delta_{ct} = \arctan(k_e e_{ct}/(v + k_{soft}))$ falls with speed while the yaw rate it produces does not, and the shaded strip is half a lane width. The dotted line is $v = 5\,\text{m/s}$, the dashed line 15 m/s and the solid line 30 m/s.

to the road heading beneath it and $\psi_e = 0$. Equating the remaining term to the required angle gives the standing offset of Equation (2.11).

$$e_{ct} = \frac{(v + k_{soft})\,L\,\kappa}{k_e} = \frac{(v + k_{soft})\,L}{k_e R} \tag{2.11}$$

For the shipped constant-radius configuration, $v = 25\,\text{m/s}$, $k_{soft} = 1.0\,\text{m/s}$, $L = 2.9\,\text{m}$, $k_e = 1.2/\text{s}$ and $R = 200\,\text{m}$, the prediction is 0.314 m and the measured steady error inside the corner is 0.314 m, on the outside of the bend. This is the classic proportional-offset argument of a controller without reset action [1], arriving from the reference rather than from a disturbance. Equation (1.8) is its counterpart for the turn-rate law of Chapter 1.

The same algebra determines why the heading error is referenced to the road under the vehicle. Taking it instead against the road direction at the front-axle projection adds $L\kappa$ to ψ_e, because the road heading advances by κ per unit of arc length and the two projections are approximately L apart. The heading term then already contains the feedforward, the vehicle body assumes an attitude offset of $\arctan(L\kappa)$ with the front wheel aligned on the path, and adding $\arctan(L\kappa)$ on top applies the same correction twice. Substituting that variant into the same steady-state balance predicts an offset of $-(v + k_{soft})L\kappa/k_e$, equal in size and opposite in sign. Running the simulator with the heading reference moved to the front-axle projection and the feedforward left on gives –0.3142 m in the same corner where the implemented pairing gives 0.000 m. A feedforward term is additive only when the equivalent term is not already implicit in the feedback path.

🕮 **Pause.** Try Exercise 2.2 before reading on.

2.4.4 Steady-State Offset Without Integral Action

The steering law contains no integral state, so any sustained demand that the feedforward does not supply leaves a standing error. A steering misalignment is the clearest case. The plant delivers

$\delta + \delta_{bias}$, so holding a straight line requires the law to produce $\delta = -\delta_{bias}$, and with $\psi_e = 0$ the offset is given by Equation (2.12).

$$e_{ct} = -\frac{(v + k_{soft}) \tan \delta_{bias}}{k_e} \tag{2.12}$$

At $\delta_{bias} = 0.6°$, $v = 25\,\mathrm{m/s}$ and $k_e = 1.2/\mathrm{s}$ the prediction is −0.227 m and the measured settled value is −0.227 m. Raising k_e to 2.6 reduces it to −0.105 m, and raising the misalignment to 2.0° increases it to −0.757 m. The offset falls as $1/k_e$ and never reaches zero, which is the consequence of a purely proportional geometric law and the reason Section 2.6 treats gain as a trade-off rather than as a free improvement.

2.4.5 Stability Limit Imposed by Actuator Lag

Restoring the actuator lag of Equation (2.2) to the linearized loop of Equation (2.8) turns the second-order characteristic polynomial into the cubic of Equation (2.13).

$$\tau_\delta s^3 + s^2 + \left(\frac{v}{L} + k_e^*\right) s + \frac{k_e^* v}{L} = 0 \tag{2.13}$$

The Routh condition for this cubic is $v/L + k_e^* > \tau_\delta k_e^* v/L$, which is satisfied for every positive gain whenever $\tau_\delta v < L$. At the default rack lag of 0.10 s that inequality holds for all $v < L/\tau_\delta = 29\,\mathrm{m/s}$, so across almost the whole speed range the application allows there is no cross-track gain that destabilizes the linearized loop. Above that speed, and at any speed once the rack is slowed down, the critical gain is finite and given by Equation (2.14).

$$k_{e,crit}^* = \frac{v/L}{\tau_\delta v/L - 1}, \qquad \omega_{crit} = \sqrt{\frac{k_{e,crit}^* v}{L}} \tag{2.14}$$

For the slow-rack configuration, $\tau_\delta = 0.45\,\mathrm{s}$ at 25 m/s, Equation (2.14) gives $k_{e,crit}^* = 2.99$, which corresponds to an entered gain of $k_e = 3.11$ after the k_{soft} correction, and an oscillation period of 1.24 s. A model run outside the settings the application offers, on the straight road and started from a 10 µm displacement so that no limit is active, loses small-signal stability between $k_e = 2.55$ and $k_e = 2.6$ with a measured period of 1.28 s. The 16 % gap is attributable to the sampled-data hold. Reducing the control period from 20 ms to 5 ms in the same run moves the measured boundary to between 2.95 and 3.0, and reducing it to 1 ms moves it to between 3.05 and 3.1, within 1 % of the continuous-time prediction. Tuning this loop at 50 Hz is therefore not tuning the plant of Equation (2.13), and the difference amounts to 0.5/s of gain margin. Sampled measurement and zero-order hold are ordinary features of a digital control loop [1, 4], and this configuration quantifies their cost on a loop whose continuous-time boundary is known in closed form.

2.5 Controller Design

2.5.1 Discrete Steering Algorithm

Under the Stanley law the steering path is a static function of geometry followed by two limits, and it runs to completion every $\Delta t_c = 0.02\,\mathrm{s}$ in both manual and automatic operation. The sequence is

1. project the front axle and read e_{ct} and κ at s^*,
2. project the rear axle and read ψ_r at s, form $\psi_e = \mathrm{wrap}(\psi_r - \psi)$,
3. form $\hat{e}_{ct}$ from Equation (2.5),
4. advance the speed setpoint, $v_{SP} \leftarrow \mathrm{clamp}(v_{SP} + \dot{v}_{SP}\Delta t_c, 2, 40)$,
5. form $\delta_{ff} = \arctan(L\kappa)$ if the feedforward is on and zero otherwise,

6. evaluate Equation (2.7) and clip the result to $\pm\pi/2$ for reporting,
7. select the request, $r = \delta_{law}$ when the steering loop is in automatic and $r = \delta_{man}$ otherwise,
8. clamp, $\delta_c = \mathrm{clamp}(r, -\delta_{max}, \delta_{max})$,
9. rate limit, $\delta_{cmd} \leftarrow \delta_{cmd} + \mathrm{clamp}(\delta_c - \delta_{cmd}, \pm\dot{\delta}_{max}\Delta t_c)$.

The saturation flag is raised when $|\delta_c| \geq \delta_{max}$ and the rate-limit flag when δ_{cmd} differs from δ_c after step nine, so the two failure modes are counted separately in the reported metrics. The clip to $\pm\pi/2$ in step six alters nothing the vehicle experiences, because step eight clamps to $\delta_{max} \leq 30^\circ$ in any case, and it bounds the reported law output for display.

2.5.2 Velocity Form Speed Loop

The longitudinal controller is the velocity-form PID of Section 1.6.2 [2], with the derivative taken on the measurement. In the loop's own symbols the error and the three increments are given by Equations (2.15) and (2.16).

$$
\begin{aligned}
e_{v,k} &= v_{SP,k} - v_k \\
\Delta u_k^P &= K_c \left(e_{v,k} - e_{v,k-1} \right) \\
\Delta u_k^I &= \frac{K_c}{\tau_I} e_{v,k} \Delta t_c \\
\Delta u_k^D &= -K_c \tau_D \frac{v_k - 2v_{k-1} + v_{k-2}}{\Delta t_c}
\end{aligned}
\tag{2.15, 2.16}
$$

where K_c is in 1/s, τ_I and τ_D are in seconds and Δt_c is the control period. The speed loop runs with $\tau_D = 0$, so Δu_k^D vanishes and the loop is proportional-integral. The increments are combined and the demand is clipped to the traction and braking envelope, as shown in Equation (2.17).

$$
\begin{aligned}
z_k &= z_{k-1} + \Delta u_k^P + \Delta u_k^D \\
c_k &= z_k + I_{k-1} \\
a_{cmd,k} &= \mathrm{clamp}\left(c_k,\ a_{min},\ a_{max} \right) \\
I_k &= I_{k-1} + \Delta u_k^I + \frac{\Delta t_c}{\tau_t} \left(a_{cmd,k} - c_k \right)
\end{aligned}
\tag{2.17}
$$

Here z_k accumulates the proportional and derivative increments without limit, c_k is the acceleration the algorithm demands, built on the integral accumulated through the previous scan, and $a_{cmd,k}$ is that demand clipped to the envelope. These are Equations (1.20) to (1.22), with z_k in place of S_k because s is the road station here. Anti-windup is back-calculation with a tracking time $\tau_t = \tau_I = 4\,\mathrm{s}$. There is no bias and no output rate limit, and the plant applies the same envelope the controller clips to. The plant $\dot{v} = a$ is a pure integrator, so proportional action alone already produces zero offset against a constant setpoint, and the integral term is required only on a ramp. In manual operation the pedal is mapped to acceleration by $a = (\mathrm{pedal}/100)\, a_{max}$ for a positive pedal and $a = -(\mathrm{pedal}/100)\, a_{min}$ for a negative one, and I_k is set each scan to the value that makes $z_k + I_k$ equal the operator-set acceleration, so transfer in either direction is bumpless with no additional logic (Section 1.6.2).

2.5.3 Bumpless Transfer of the Steering Loop

The steering law contains no integral state to seed, so continuity across a mode change comes from the rate limit of step nine instead. The command δ_{cmd} is carried across the mode switch and can leave the operator's last value only at $\dot{\delta}_{max}$, which is 45 °/s by default and therefore 0.9° in the first control interval. The law is evaluated and logged in manual operation as well, so the value the loop would command is available before the transfer occurs.

2.5.4 Manual Operation

The application opens with the steering loop in manual on a straight highway at 25 m/s and the road wheels centered. Incremental manual steering is bounded by a constant lateral acceleration envelope, $\min(\delta_{max}, \arctan(La_{lat}/\max(v^2,4)))$ with $a_{lat} = 5.0\,\mathrm{m/s^2}$, which is 1.33° at 25 m/s and 8.25° at 10 m/s, while the direct manual entry reaches the mechanical lock at any speed. The envelope is imposed because full lock at highway speed places the vehicle outside the lane within one second, which leaves no usable range for manual regulation.

Manual operation quantifies the scale of the problem. A steady 1° of steering at 25 m/s bends the path to a radius of $L/\tan 1^\circ = 166\,\mathrm{m}$, and the simulated vehicle, with its rack lag included, moves 1.53 m laterally in the first second after the command reaches the rack and 3.68 m in the first second and a half, against a half lane of 1.75 m. Manual corrections in the logged traces are whole degrees while the controller acts in tenths of a degree, because a manual correction responds to the visible position while the law responds to the position and the heading together.

2.5.5 Model Predictive Steering

A model predictive controller (MPC) is offered as a second steering law, and only configuration 11 selects it. The predictive option of the apps in general is described in Section 1.6.3. This app builds its own model, objective and estimator for the steering, set out below. It replaces δ_{law} in step seven and leaves the clamp, the rate limit and the speed loop as they are. Every 0.1 s it optimizes the steering over the coming $T_p = 2\,\mathrm{s}$, keeps only the opening move and repeats from the newest measurement, a receding-horizon scheme whose constrained form is reviewed in [18]. MPC of front steering on a vehicle model linearized afresh at each sample has been tested at up to 21 m/s on ice [19]. The Stanley law corrects the geometry of Figure 2.1 at the present instant. The MPC plans on that geometry, with the road ahead as a known preview and the steering offset as an estimated disturbance.

The controller's model is not identified. It is Equations (2.1), (2.2) and (2.4) written along the road for small angles, with the rack lag, a ramped command and a constant steering offset d, as shown in Equation (2.18).

$$\begin{gathered} \dot{e}_{ct} = v(\psi_e + \Delta\theta) - v(\delta + d), \qquad \dot{\psi}_e = v\kappa_r - \frac{v}{L}(\delta + d), \\ \tau_{\delta 0}\,\dot{\delta} = u - \delta, \qquad \dot{u} = \rho, \qquad \dot{d} = 0 \end{gathered} \tag{2.18}$$

Here u is the planned command, ρ its rate, d the angle by which the wheels turn past the command, playing the part of δ_{bias}, which the controller never reads, κ_r the curvature at the rear-axle projection and $\Delta\theta = \psi_r(s^*) - \psi_r(s)$ the road heading difference between the two projections, close to $L\kappa$ on an arc. The model is rebuilt and discretized at every decision at the measured speed v. The wheelbase L is read where the feedforward $\arctan(L\kappa)$ reads it, so a changed wheelbase reaches both laws, while $\tau_{\delta 0} = 0.10\,\mathrm{s}$ is the declared rack lag of Table 2.3, which the Stanley law never reads, so a slower rack is an error in the MPC's model alone. Stepped open loop on the commands of runs steered by the Stanley law, with the true offset as d, it errs two seconds ahead by 2.0 mm on configuration 11, 20.7 mm at 6.6° of steering on the $R = 25\,\mathrm{m}$ corner and 2.14 m on the 0.45 s rack of configuration 9.

A zero-order hold over the decision period $T_s = 0.1\,\mathrm{s}$, every fifth control period, gives $N = T_p/T_s = 20$ steps. The decisions are command rates held over nine move blocks, four single steps followed by blocks of 2, 2, 3, 4 and 5. The command therefore ramps in straight segments that the rate limiter passes unchanged, and each control period requests the mean of the first ramp over that period. The objective is Equation (2.19).

$$\begin{aligned} J = \sum_{k=1}^{N-1} \left[\frac{e_k^2}{s_e^2} + \frac{(\psi_k - \psi_k^*)^2}{s_\psi^2} \right] + \lambda \sum_{k=0}^{N-1} \frac{\rho_k^2}{s_\rho^2} \\ + \tilde{z}_N^T P \tilde{z}_N + q_\varepsilon \varepsilon^2 + r_\varepsilon \varepsilon \end{aligned} \tag{2.19}$$

Here e_k and ψ_k are the predicted errors, $\psi_k^* = L\kappa_{r,k} - \Delta\theta_k$ is the heading error that keeps e_{ct} at zero on the road of step k, $s_e = 0.1\,\text{m}$, $s_\psi = 1^\circ$, $s_\rho = 10^\circ/\text{s}$ and $\lambda = 0.01$. The vector $\tilde{z}_N$ is the departure of $(e_{ct}, \psi_e, \delta, u)$ from the rest state of the last step's road and the estimated offset, and P is the stationary Riccati solution with these weights, so the terminal term prices everything after the horizon as the linear quadratic regulator (LQR) would. When no constraint is active and each step carries its own decision, the opening rate of the plan matches the LQR's to 9.5×10^{-9} rad/s, on moves of up to 0.96 rad/s, against an independent construction. The weight λ was chosen once, on configuration 11, as the largest of twelve values within 10 % of the best lane IAE summed over both speeds, among those that kept the lane, rejected no plan and added at most 5 % to the Stanley law's steering travel.

The limits $|\rho_k| \le \dot{\delta}_{max}$ and $|u| \le \delta_{max}$, the angle bound imposed where each block ends and so holding for the whole ramp, are hard and are read as the clamp and the rate limiter read them. The lane is soft, $|e_k| \le w/2 - w_{car}/2 + \varepsilon$ with $\varepsilon \ge 0$, the 0.82 m margin of Equation (2.6) at the defaults, priced at $q_\varepsilon = 10^4$ per m^2 and $r_\varepsilon = 10^3$ per meter. Each plan is a quadratic program in the nine rates and the slack, with 77 inequality rows and a ridge of $10^{-9}(1 + 2\lambda N/s_\rho^2)$, which the dual active-set solver the apps share completes in one iteration on configuration 11 and in 7.5 on average, 36 at most, on configuration 8.

A steady-state Kalman filter on $(e_{ct}, \psi_e, \delta, d)$ at Δt_c, driven by the applied command and by $\Delta\theta$ and κ_r under the vehicle now, supplies the state each plan starts from. It is corrected by the two quantities the Stanley law measures, $\hat{e}_{ct}$ and ψ_e, with variances $(0.1\,\text{m})^2$ and $(0.01^\circ)^2$, against process variances per period of $(1\,\text{mm})^2$, $(0.057^\circ)^2$, $(0.181^\circ)^2$ and $(0.0057^\circ)^2$. Its gain is designed once at 25 m/s with the declared wheelbase and rack lag, and its transition matrix follows the measured speed. The estimate of d settles with a time constant of 2.95 s at 30 m/s and 5.27 s at 10 m/s, and the cross-track channel with 0.64 s. Carrying d as an integrating state lets the plan cancel a constant misalignment without offset [20], the same construction as the offset-free MPC of Volume 9, Chapter 4, whose estimator carries a fuel bias. The filter runs under both laws, so a switch to the MPC starts from a converged estimate.

With the preview on, the default, the plan reads $\Delta\theta_k$ and $\kappa_{r,k}$ from the road table the controller follows, at the stations the two axles reach at the measured speed. With it off the plan holds the front-axle curvature now, the number the feedforward reads, while the filter still reads the road under both axles, so the MPC is not reduced to the information of the Stanley law.

A plan is rejected when its program is not optimal, when its first rate is not finite or exceeds $\dot{\delta}_{max}$, or when its first block ends beyond δ_{max}. The Stanley law, evaluated in every period, then supplies the request through the rate limiter until a plan is accepted. Rejections forced at six instants of configuration 11 move the command by at most 0.58° in one period at the hand-over and 0.80° at the hand-back, inside the 0.90° limiter step, and no shipped configuration rejects a plan. The plant remains Equations (2.1) to (2.6) with the rack lag and misalignment of the run, behind the instrument of Equation (2.5), while the controller reads $\hat{e}_{ct}$, ψ_e, v, the road table and the limits. A decision, which rebuilds the model, solves the Riccati equation and solves the program, takes 0.18 to 0.29 ms on average in Node.js and 0.30 to 0.44 ms in Chrome, at most 1.7 % of a core at five times real time.

2.6 Results

Table 2.4 lists the shipped configurations, the app's numbered scenarios (Section 1.6.4), and the mechanism each one isolates. Each configuration fixes the road it requires and whatever else its lesson depends on and leaves every other setting at the default, so the change under test is always identifiable. Every number in Table 2.5 comes from the performance metrics the application computes while it runs, over a horizon of 40 s, which is long enough for each road to deliver all of its curvature, or 20 s for the offset start and 34 s for the challenge, the time at which its target is judged. Figure 2.3 draws the peak error of every configuration as a bar.

Table 2.4 Shipped Configurations and the Mechanism Each Isolates

Configuration	What it changes and what it shows
1 Straight highway	Manual, 25 m/s. Manual regulation followed by transfer to automatic while running.
2 Constant radius	$R = 200$ m at 25 m/s. Feedforward on against off, and Equation (2.11).
3 S-curve	$R = \pm 150$ m with clothoids. The error appears in the transitions rather than in the constant-curvature corners.
4 Lane change	3.5 m over 60 m. A steering-rate demand that scales with speed.
5 Speed ramp	10 to 30.4 m/s on one sweep with one tuning. The speed cancellation.
6 Sudden curvature	κ steps to $1/120$ with no spiral at 20 m/s. An actuator deficit that no gain removes.
7 Offset start	Starts 1 m right on a straight. First-order recovery at $1/k_e^*$.
8 Steering limit	$R = 25$ m at 10 m/s, $\delta_{max} = 6°$, $k_e = 2.0$. Geometric infeasibility.
9 Slow rack	$\tau_\delta = 0.45$ s on the S-curve. Phase lag and the resulting gain limit.
10 Noisy measurement	$\sigma_{gps} = 0.25$ m on the gentle sweep. Actuator travel against tracking error.
11 MPC at two speeds	Lane change at 30 m/s with 0.5° of misalignment under the MPC, then the Stanley law, then both at 10 m/s. One gain against a model rebuilt at the measured speed.
12 Challenge	Challenge course at 20 m/s with 0.6° of misalignment and 0.07 m of noise, starting from $k_e = 0.5$, $k_{soft} = 4.0$, feedforward off.

Table 2.5 Measured Performance for Every Configuration Over a Fixed Horizon

Configuration	IAE (m s)	max $\lvert e_{ct} \rvert$ (m)	TV(δ) (rad)	sat (%)	out (%)
2 Constant radius, ff on	0.028	0.016	0.04	0.0	0.0
2 Constant radius, ff off	8.797	0.314	0.05	0.0	0.0
3 S-curve	0.124	0.014	0.11	0.0	0.0
3 S-curve, $k_e = 3$	0.050	0.007	0.11	0.0	0.0
4 Lane change, 25 m/s	0.073	0.023	0.15	0.0	0.0
4 Lane change, 12 m/s	0.029	0.004	0.14	0.0	0.0
5 Speed ramp	0.017	0.005	0.02	0.0	0.0
6 Sudden curvature	0.054	0.014	0.11	0.0	0.0
7 Offset start	0.866	1.000	0.07	0.0	1.7
8 Steering limit, $k_e = 2$	19.03	3.096	0.79	28.3	21.8
8 Steering limit, $k_e = 6$	21.20	3.074	1.58	41.4	26.9
8 Steering limit, $k_e = 12$	27.31	3.057	3.54	61.2	35.1
9 Slow rack, $k_e = 1.2$	0.969	0.126	0.17	0.0	0.0
9 Slow rack, $k_e = 3$	3.014	0.456	1.48	0.0	0.0
10 Noisy, $k_e = 1.2$	3.798	0.300	2.13	0.0	0.0
10 Noisy, $k_e = 4$	119.4	27.43	13.51	22.2	26.3
11 MPC, 30 m/s, as loaded	0.326	0.090	0.17	0.0	0.0
11 Stanley, 30 m/s	8.850	0.257	0.17	0.0	0.0
12 Challenge, as loaded	23.14	1.354	0.25	0.0	35.1
12 Challenge, tuned	3.095	0.214	1.27	0.0	0.0

Horizon 40 s except configuration 7 (20 s) and the challenge (34 s). The 12 m/s lane change starts at that speed, because the vehicle is initialized at the speed setpoint. IAE is the integrated absolute cross-track error, max $\lvert e_{ct} \rvert$ is the peak true error, TV(δ) is the total variation of the road-wheel angle in radians accumulated at the 20 Hz logging rate, sat is the fraction of the run with the steering request against its end stop, and out is the fraction with any part of the vehicle outside the lane markings.

2.6.1 Curvature as a Reference Demand

The constant-radius corner separates feedforward action from feedback action with no other effect present. With the feedforward on, the corner costs 0.016 m of peak error and 0.04 of steering travel, and the steering angle holds the geometric value $\arctan(L/R) = 0.831°$ for the whole 28 s of the

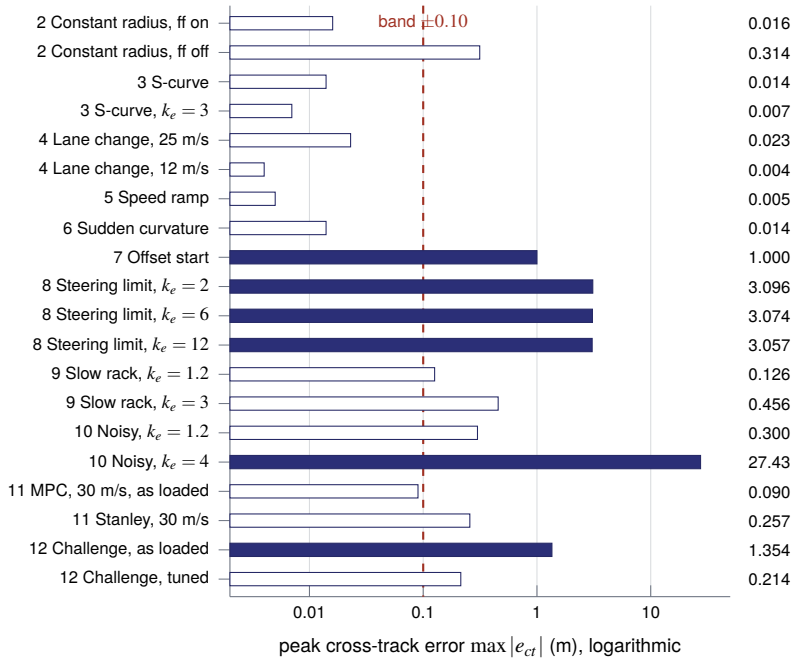

Figure 2.3 The peak cross-track error of Table 2.5, configuration by configuration, on a logarithmic axis, against the band of ±0.10 m (dashed line). Solid bars are runs with some time out of the lane, white bars runs that stayed in it. Notice that the corner the steering cannot reach keeps its peak near 3 m at every gain, that the missing feedforward and the misalignment under the Stanley law leave white bars past the band, and that the noisy road at $k_e = 4$ ends farthest from the lane.

arc. With the feedforward off, the same corner costs 0.314 m of standing error, and the integrated absolute error rises from 0.028 to 8.797, a factor of 317. The steering travel is nearly unchanged at 0.05, because the wheel reaches the same angle in both cases. The feedforward delivers the whole of that error reduction at no cost in actuator activity, which is the argument for supplying a known reference demand directly rather than requiring the feedback loop to develop an error large enough to produce it [1]. Section 1.4.4 makes the same argument for the reference feedforward of a moving cart.

2.6.2 Displacement Recovery at Two Speeds

Figure 2.4 shows the speed-cancellation experiment. The vehicle runs the gentle sweep while its speed setpoint ramps from 10 m/s to 30.4 m/s at 1.2 m/s^2 between $t = 3$ s and 20 s, and a 0.5 m lateral displacement is applied twice, at $t = 6.0$ s with the vehicle at 13.0 m/s and at $t = 27.0$ s at 30.6 m/s, with k_e held at 1.2/s throughout. Evaluated at the 10 ms integration step, the two recoveries decay to $1/\mathrm{e}$ of their initial value in 0.84 s and 0.85 s, against the 0.90 s and 0.86 s that Equation (2.9) predicts at those two speeds, and to 10 % in 2.07 s and 1.98 s respectively. The steering angles that produce those two nearly identical recoveries differ. The peak road-wheel deviation from the angle the road already required is 1.55° at 13.0 m/s, where the vehicle is on the approach straight and the required angle is zero, and 0.60° at 30.6 m/s, where the vehicle is inside the $R = 500$ m arc holding 0.332°. That is a ratio of 2.6 against a speed ratio of 2.35. The regulated variable behaves identically at the two speeds while the actuator excursion behind it differs by a factor of 2.6.

The remainder of the speed ramp shows the same result with no disturbance applied. The cross-track error over the whole 10 to 30.4 m/s ramp never exceeds 0.005 m and the steering angle holds the constant geometric value $\arctan(L/R) = 0.332°$ demanded by the $R = 500$ m sweep, because

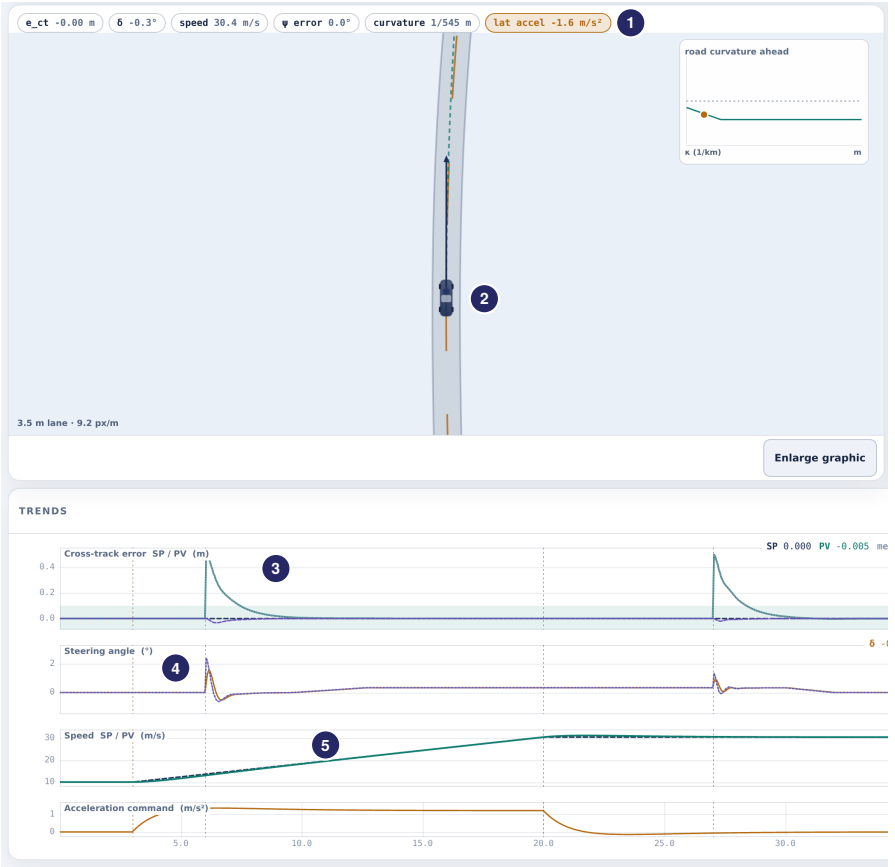

Figure 2.4 One tuning across a three-to-one speed change, shown at $t = 40$ s. The speed setpoint ramps from 10 to 30.4 m/s on the gentle sweep while k_e stays at 1.2/s, and a 0.5 m lateral displacement is applied at $t = 6.0$ s (13.0 m/s) and again at $t = 27.0$ s (30.6 m/s). The two cross-track recoveries have the same shape and the same duration while the steering deviations that produce them differ by a factor of 2.6. The crop keeps the plant view and the trends and cuts the header and the side panels. Callouts: (1) the badges of the lateral error, steering angle, speed and lateral acceleration, (2) the car on the sweep, (3) the two cross-track recoveries, (4) the two steering deviations behind them, (5) the speed ramp.

that angle is set by geometry and carries no speed dependence either. The lateral acceleration does change, reaching 1.85 m/s^2 in the $R = 500$ m arc at 30.4 m/s as v^2/R, and 1.94 m/s^2 while the speed loop overshoots to 31.2 m/s just after the ramp ends.

TRY IT IN THE APP **5 · Increasing Speed (10 → 30 m/s on One Road)**

Load the case and press run. Watch the Speed trend climb from 10 to about 30 m/s while the Cross-track error trend stays flat: over the whole ramp the error never exceeds 0.005 m. The Steering angle trend holds near 0.332°, the angle the 500 m sweep needs at any speed.

2.6.3 Actuator-Limited Responses

Three configurations separate three distinct actuator limitations, and none of the three is a tuning problem.

The sudden-curvature road steps κ from zero to $1/120\,\mathrm{m}$ with no transition spiral, so the feedforward asks for $\arctan(L\kappa) = 1.384°$ instantly and the rack answers with a first-order lag. The resulting peak error is 0.014 m and feedback cannot remove it, because the deficit is in the actuator rather than in the law. Highway geometry standards specify transition spirals so that this step does not occur, which is one of the few cases in which civil engineering design is set explicitly by an actuator rate limit.

The lane change is a curvature pulse rather than a step, and it loads the steering rate rather than the steering angle. The curvature profile $\kappa(s)$ is a property of the road and does not change with speed, so the steering angle profile $\arctan(L\kappa)$ is the same at every speed. What changes is the time available to traverse it. The same 60 m of road is 2.4 s at 25 m/s and 5.0 s at 12 m/s, so the peak steering rate rises from 1.28 °/s at the lower speed to 2.76 °/s at the higher, a factor of 2.2, while both remain far below the 45 °/s limit, which is never reached. The measured peak error is 0.023 m at 25 m/s against 0.004 m at 12 m/s, and the steering travel is 0.15 against 0.14 because the wheel visits the same angles in both cases.

The sharp corner is not a tuning problem. Holding $R = 25\,\mathrm{m}$ needs $\arctan(L/R) = 6.617°$ and the configuration limits the steering to 6°, so the required angle exceeds the available angle and no gain closes the gap. Raising k_e from 2 to 6 to 12 leaves the peak error at 3.096 m, 3.074 m and 3.057 m, an improvement of 1.3 % over a sixfold increase in gain, while the time against the end stop rises from 28.3 % to 61.2 %, the time out of the lane rises from 21.8 % to 35.1 % and the steering travel rises from 0.79 to 3.54. The corner is infeasible for the available steering authority, and distinguishing that condition from a badly tuned loop is the practical result the configuration provides.

Pause. Try Exercise 2.3 before reading on.

2.6.4 Actuator Lag and the Gain Limit

The slow-rack configuration puts $\tau_\delta = 0.45\,\mathrm{s}$ on the S-curve and leaves everything else alone. At $k_e = 1.2/\mathrm{s}$ the loop crosses over near 1.2 rad/s, where the lag contributes $\arctan(1.2 \times 0.45) = 28.4°$ of phase, and the peak error grows from 0.014 m with the standard rack to 0.126 m. Raising k_e to 3 moves the crossover to the frequency where the lag contributes 53.5°, the heading term arrives after the vehicle has already turned, and the peak error over the reporting horizon grows to 0.456 m while the steering travel rises from 0.17 to 1.48. That gain is above the small-signal boundary of Equation (2.14), so the run is unstable rather than merely degraded. The 40 s horizon does not reveal the instability, because the road delivers curvature until 30.7 s and the growth has barely begun. Continued to 80 s on the straight that follows the last corner, the same run leaves the lane at 44.7 s and its envelope grows to 109.8 m over the last 20 s, while the $k_e = 1.2/\mathrm{s}$ run settles to 0.000 m. Reporting a stability conclusion therefore requires a horizon longer than the excitation.

2.6.5 Measurement Noise and Actuator Travel

The noisy-measurement configuration is the case in which raising the gain degrades every reported measure, and its endpoint is not predicted by the linear analysis. Table 2.6 shows the sweep. Between $k_e = 1.2$ and $k_e = 3.8$ the degradation is monotone. The true tracking error grows rather than falls, because the loop cannot distinguish a real displacement from a receiver wander, and the steering

Table 2.6 Cross-Track Gain Against Steering Travel With $\sigma_{gps} = 0.25\,\text{m}$ on the Gentle Sweep

k_e (1/s)	IAE (m s)	max $\lvert e_{ct}\rvert$ (m)	TV(δ) (rad)	rate lim. (%)	max $\lvert\delta_{law}\rvert$ (deg)
1.2	3.798	0.300	2.13	0.0	1.9
2.0	4.918	0.435	3.73	2.5	3.2
3.0	5.878	0.549	5.99	17.7	5.3
3.5	6.326	0.589	7.41	29.7	6.5
3.8	6.781	0.726	8.54	38.35	8.8
4.0	119.4	27.43	13.51	53.3	90.0

travel grows fourfold while the fraction of the run spent against the rate limit rises from zero to 38.35 %.

At $k_e = 4.0$ the loop diverges. The peak error increases from 0.726 m to 27.43 m, the vehicle spends 26.3 % of the run outside the lane and 22.2 % of it against the steering end stop, and the law reaches its own 90° reporting clip. Equation (2.14) shows that this is not a linear instability, because $\tau_\delta v = 2.5\,\text{m}$ is smaller than $L = 2.9\,\text{m}$ at this speed and rack lag, and the same gain on the same road with the noise removed holds the line to 0.001 m. The cause is the rate limiter. The noise-driven command excursions grow with k_e until they demand more than 45 °/s, and a rate limiter that is active for half the run acts as an amplitude-dependent phase lag that no small-signal margin accounts for. The logged traces show the command trailing the law by several degrees for about a second before the first departure at 29.4 s, and at the departure the law requests −23.8° while the command is still passing through −1.6°. The location of the cliff depends on the noise sequence. The table uses the shipped sequence, and across eleven other seeds, on the same gains extended to 4.2 and 4.5, the first gain to leave the lane within 40 s ranges from 3.8 to 4.5, while no seed fails at 3.5.

Actuator travel is therefore read alongside the tracking error rather than the tracking error alone. A tuning that is acceptable on the cross-track record can be cycling the steering rack thousands of times an hour, which is a durability cost in the vehicle and a ride quality cost for the occupants, and it can be one increment of gain below a discontinuous loss of control.

> ✦ DID YOU KNOW? **One Step of Gain Past the Edge**
>
> From the simulator: on the gentle sweep with 0.25 m of measurement noise, a cross-track gain of 3.8 keeps the peak error at 0.726 m. A gain of 4.0 takes it to 27.43 m, with the command on its rate limit for 53.3 % of the run (Table 2.6). Without the noise the same gain holds the line to 0.001 m.

2.6.6 The Graded Challenge

The challenge configuration combines three mechanisms with conflicting gain requirements. The course is a left-hand corner running into a right-hand one at 20 m/s, the linkage is misaligned by 0.6°, the lateral measurement carries 0.07 m of noise, and the starting tuning is deliberately weak at $k_e = 0.5$, $k_{soft} = 4.0$ and the feedforward switched off. The target is judged after the 680 m course, at $t = 34\,\text{s}$. As loaded the configuration fails it, with a peak error of 1.354 m, 5.4 times the band, and 35.1 % of the run outside the lane, without reaching the steering limit. Equation (2.11) accounts for most of that on its own, predicting $24 \times 2.9/(0.5 \times 140) = 0.99\,\text{m}$ in the first corner, and Equation (2.12) adds a further 0.50 m from the misalignment. Enabling the feedforward removes the first contribution completely and at no cost in actuator travel. The second contribution falls as $1/k_e$ while the noise response grows with k_e, and the suggested setting of $k_e = 2.6$, $k_{soft} = 1.0$ with the feedforward on meets the $\pm 0.25\,\text{m}$ target with a peak of 0.214 m, no saturation and no departure, at the cost of raising the steering travel from 0.25 to 1.27. The same settings applied before a Reset give the same result. The 0.036 m margin belongs to the shipped noise sequence. Across eleven other seeds the same tuning peaks between 0.180 m and 0.255 m and exceeds the band on two of them.

Table 2.7 Stanley Law Against MPC With Configuration 11's Settings

Case	Law	max $\lvert e_{cr} \rvert$ (mm)	RMS (mm)	IAE (m s)	TV(δ) (rad)	ψ_e (deg)	rate (%)	sat (%)	out (%)
11, 30 m/s	Stanley	256.6	222.2	8.850	0.168	0.058	0	0	0
	MPC	89.5	20.3	0.326	0.169	0.058	0	0	0
	no prev.	89.5	20.3	0.327	0.181	0.059	0	0	0
11, 10 m/s	Stanley	85.7	78.7	3.136	0.148	0.078	0	0	0
	MPC	33.4	11.1	0.249	0.153	0.063	0	0	0
30 m/s, d 0	Stanley	34.1	6.8	0.090	0.151	0.014	0	0	0
	MPC	5.5	1.1	0.014	0.136	0.035	0	0	0
	no prev.	11.3	1.4	0.017	0.148	0.037	0	0	0
2	Stanley	15.5	2.3	0.028	0.036	0.021	0	0	0
	MPC	8.3	0.8	0.009	0.034	0.024	0	0	0
6	Stanley	14.4	3.1	0.054	0.114	0.060	0.05	0	0
	MPC	16.8	2.4	0.035	0.132	0.066	0	0	0
7	Stanley	1000.0	151.7	0.866	0.065	0.390	0.20	0	1.7
	MPC	1000.0	112.1	0.343	0.190	0.631	1.50	0	1.4
8	Stanley	3096.2	1009.1	19.03	0.790	3.449	0.90	28.3	21.8
	MPC	2605.2	797.5	14.28	0.827	3.054	0.75	31.2	17.9
9	Stanley	125.7	41.8	0.969	0.175	0.127	0	0	0
	MPC	65.5	21.1	0.610	0.244	0.125	0	0	0
10	Stanley	299.9	115.1	3.798	2.127	0.591	0	0	0
	MPC	311.5	126.6	4.151	1.367	0.531	0	0	0
12	Stanley	213.9	104.9	3.095	1.266	0.371	0	0	0
	MPC	124.3	53.7	1.526	0.394	0.189	0	0	0

The first rows are configuration 11 with its 0.5° misalignment and without it (d 0). Stanley is the shipped tuning, $k_e = 2$ on configuration 8 and the suggested tuning on the challenge (12), and *no prev.* the MPC with the preview off. Peak and RMS are the true error at the 10 ms step, ψ_e the RMS heading error, rate the time on the rate limit. Every MPC run plans every 0.1 s and rejects no plan.

2.6.7 Stanley Law Against Model Predictive Steering

Table 2.7 compares the Stanley law, tuned as shipped, with the MPC run at the settings of configuration 11, with no retuning per case and the seeds, noise, instruments and horizons of Table 2.5. The feedforward reads the curvature at the front axle now, while the MPC reads 2 s of road ahead, which the rows marked *no prev.* withhold from its plan.

Configuration 11 drives the lane-change road at 30 m/s with the linkage misaligned by 0.5°, under each law and again at 10 m/s. Without integral action the Stanley law settles where Equation (2.12) puts it, 225.4 mm off the line at 30 m/s and 80.0 mm at 10 m/s, as measured at 20 and 40 s, so one gain resists the offset $31/11 = 2.82$ times less at the higher speed. The MPC's estimate of d is 0.244° at 2 s and 0.499° at 20 s, its error stays under 1 mm after 14.1 s, and its peak of 89.5 mm comes 0.25 s after the start, before the offset is known. Its lane IAE is 0.326 against 8.850 m s at 30 m/s and 0.249 against 3.136 m s at 10 m/s, for steering travel of 0.169 against 0.168 rad and 0.153 against 0.148 rad, as shown in Figure 2.5. At 20 s both laws hold the wheel at −0.500°, the Stanley law through a standing error and the MPC through its estimate.

Without the misalignment the same road separates the laws by less, peaks of 5.5 against 34.1 mm at 30 m/s and 3.0 against 5.8 mm at 10 m/s, which is why configuration 11 carries the offset. The Stanley law turns when the curvature reaches the front axle, and the rack lag then costs 3.0 m of road at 30 m/s, while the MPC turns before the shift arrives. With the preview off it still peaks at only 11.3 mm, but not on the Stanley law's information, because its filter reads the road heading at the front axle and the curvature under the rear axle. A counterfactual outside the application that gives the filter only the front-axle curvature as well peaks at 39.1 mm against 34.1 mm, while the misaligned case keeps its result (lane IAE 0.370 against 8.850 m s). The gain without the offset is therefore the road information's and the gain with it the estimator's. On these clean roads the MPC holds more heading error, 0.035° RMS against 0.014°, because its objective treats 0.1 m of cross-track error and 1° of heading error alike.

The MPC also lowers the peak on configurations 3, 4 and 5 (2.2, 4.0 and 0.6 against 13.8, 22.9 and

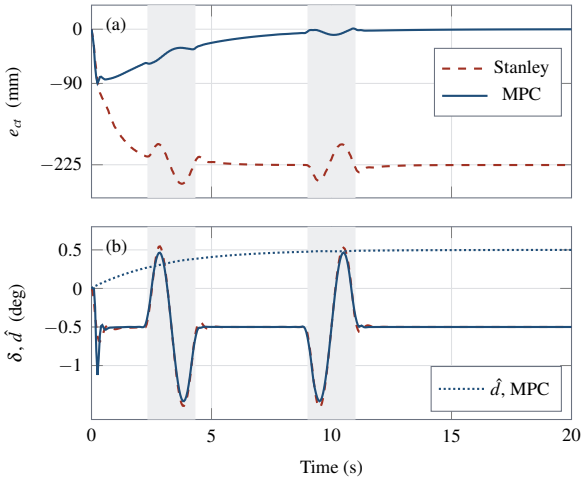

Figure 2.5 Configuration 11 at 30 m/s with 0.5° of misalignment under the MPC as loaded and under the Stanley law. (a) Cross-track error. (b) Road-wheel angle under both laws and the MPC's offset estimate $\hat{d}$ (dotted). The Stanley law is dashed and the MPC solid. Shading marks the two lane shifts.

4.7 mm). It loses the sudden curvature on peak, 16.8 against 14.4 mm, with a lower IAE but more travel. From the offset start it steers harder, with 1.50 % of the run on the rate limit against 0.20 % and heading errors up to 5.58° against 2.60°, because the vehicle starts 0.18 m past the priced lane band. On the steering limit it is out of the lane for 17.9 % of the run against 21.8 % but on the end stop for 31.2 % against 28.3 %, because its preview sees the infeasible corner and its command reaches 6° at 6.71 s against 7.17 s. On the noisy measurement it moves the steering less, 1.367 against 2.127 rad, and tracks less closely, a lane IAE of 4.151 against 3.798 m s, because the filter that estimates the offset also smooths the receiver noise. On the straight highway without noise or misalignment, a −1.5° heading kick at 30 m/s takes it to 133.5 against 118.3 mm, though it wins at 10 m/s (14.8 against 20.9 mm RMS over 8 s), and a 0.5 m displacement is recovered more slowly (147.3 against 124.8 mm RMS at 10 m/s, 127.2 against 117.4 mm at 30 m/s), because the filter admits a measured jump over 0.64 s where the Stanley law acts on it at once.

With the preview off the MPC loses both curvature steps, because a plan that holds the front-axle curvature assumes the whole vehicle is already on the new arc. On the constant radius it commands up to 2.62° where the arc needs 0.831° and peaks at 36.9 against 15.5 mm, and on the sudden curvature it commands 4.05° where 1.384° is needed and peaks at 55.2 against 14.4 mm, with 0.439 against 0.114 rad of travel.

The rack lag is the decisive limitation, because the plan and the filter keep the declared 0.10 s. On the 0.45 s rack of configuration 9 as loaded, at 25 m/s, the MPC halves the peak, 65.5 against 125.7 mm, with more travel, 0.244 against 0.175 rad, but is lightly damped. After a gust on the straight, its largest error over 50 to 70 s is 0.241 of that over 10 to 30 s, against 0.000 for the Stanley law, and bisection on that ratio reaching one puts the MPC's limit at 32.3 m/s on a 0.40 s rack, 26.4 m/s on 0.45 s, 22.2 m/s on 0.50 s and 16.9 m/s on 0.60 s, while the Stanley law at $k_e = 1.2/\mathrm{s}$ is stable to 35 m/s on all four, as Equation (2.14) predicts, because its critical gain exceeds $1/\tau_\delta \geq 1.67/\mathrm{s}$ at every speed. On the S-curve at 30 m/s the MPC's error reaches 40.3 m within 80 s, where the Stanley law peaks at 174 mm. More move suppression lowers the limit, from 27.7 m/s at $\lambda = 0.001$ to 20.4 m/s at $\lambda = 1$, and neither a 0.5 nor a 4 s horizon holds 35 m/s. Told the true rack, in a counterfactual outside the application, the MPC holds the same S-curve to 3.5, 5.1 and 6.9 mm at 25, 30 and 35 m/s, so the loss belongs to the declared actuator and not to the method.

For information, the challenge is met with only the steering law changed. The MPC peaks at 0.124 m, never reaches the end stop or the rate limit, never leaves the lane and moves the steering 0.394 rad against the 1.266 rad of the suggested tuning, over 340 plans with none rejected, and the run breaks no limit that the target leaves unscored. The challenge and its suggested answer are left as shipped.

2.6.8 Implementation

Manual and automatic operation are provided for the steering loop and the speed loop independently, and the road, the tuning, the disturbances and the parameters of Table 2.3 are operator-set inputs. Every logged column, including the true and measured cross-track errors, the law output, the command and the achieved steering angle, is in the exported file of Section 1.6.4.

2.7 Technology in Practice

The car of Section 2.3 measures one lateral offset and is moved by one rack. A real lane-keeping system reads the lane with a camera or a positioning receiver, steers through the power steering and answers to a steering regulation. Figure 2.6 draws the chapter's two loops with such devices in place.

What Measures It

A camera-based system takes the lateral offset from lane lines found in the front-view image and projected onto the ground [21]. A positioning receiver is the other route. On the robot that won the 2005 desert race, receivers on the roof and an inertial measurement unit formed the positioning group that estimated the location and velocity of the vehicle, and the wheel speeds and the steering angle reached its computers over the vehicle bus [9]. In this chapter only $\hat{e}_{ct}$ carries an error, colored noise with a correlation time of 0.35 s, and the heading, the curvature and the speed are exact.

> **TECHNOLOGY** **Sensor: Lane-Detecting Camera**
>
> A camera-based system reads the lane from the front-view image of an onboard camera. A detector finds each lane line in the image, and the camera's intrinsic and extrinsic parameters project those lines onto the ground. The projection assumes flat ground, so an uphill, a downhill or a rough surface gives wrong lane positions even when the detection in the image is right. Lanes may also be hidden by nearby vehicles or by adverse weather. A laser scanner measures depth directly and does not mind the lighting, but its shorter range and high cost have kept cameras the more common choice [21]. The chapter's one imperfect instrument, the colored noise σ_{gps} on $\hat{e}_{ct}$, stands in for such errors without modeling them.

What Moves It

On the desert robot a direct-current motor attached to the steering column provided the electronic steering control [9]. In a production car the steering control may also operate without direct action of the driver, when a driver assistance steering system acts through it [22], and the power steering motor does the work. The regulation bounds that authority from the driver's side, since the effort needed to override the system may not exceed 50 N [22]. The rack of Equation (2.2) stands for the motor, the gear and the linkage together.

> **TECHNOLOGY** **Final Element: Electric Power Steering**
>
> An electric power steering unit adds motor torque to the driver's. Its control unit computes the assist from the torque the driver applies, the steering wheel position and the vehicle speed, with more assist at low speed than at high speed, and the motor turns the steering gear. The motor may sit on the column, on the pinion or on the rack, and the rack form, with its high reduction ratio, suits mid-size to full-size vehicles. Current systems keep a mechanical link from the steering wheel to the gear, and a detected fault in the control unit disconnects power from the motor [23]. The chapter compresses the unit, the gear and the linkage into one lag of 0.10 s, a rate limit and an end stop.

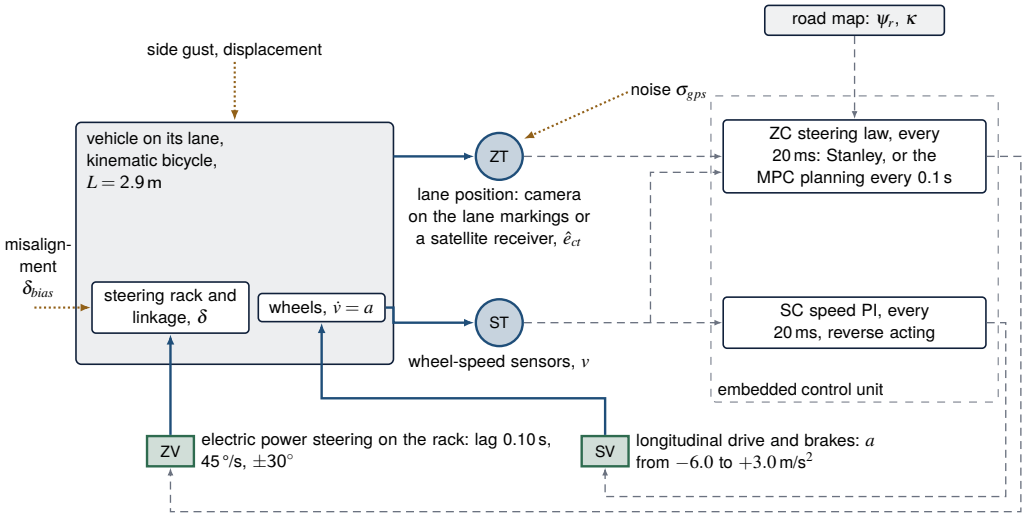

Figure 2.6 The lane loop and the speed loop of the chapter with real devices in place of the schematic's symbols. ZT reads the lane position, ST the wheel speeds, and one embedded control unit (dashed frame) runs the steering law (ZC) and the speed PI (SC) every 20 ms. ZV is the electric power steering on the rack, SV the drive and brakes. Notice that the misalignment enters at the linkage, after the actuator that reports the angle, so no sensor of the loop sees it, while the noise acts on the lane position alone.

Where the Controller Runs

The steering law runs on an embedded control unit in the vehicle. The desert robot's computers polled their sensors at up to 100 Hz and commanded the steering, throttle and brake at up to 20 Hz, the rate of its steering controller [9]. The chapter's controller runs at 50 Hz, a 20 ms period, and its predictive option plans every 0.1 s. The Stanley law is a static function block followed by a clamp and a rate limiter, and the speed loop is the velocity-form block of Section 1.6.2. Above both sits a supervisory layer that the model leaves out. The regulation asks a lane-keeping function to detect whether the driver holds the steering control, to warn after 15 s without hands and to switch itself off at the latest 30 s after its acoustic warning starts [22].

What Is Modernizing It

The first development is lane detection by deep learning on camera images, with networks that predict lanes in three dimensions directly rather than through a flat-ground projection [21]. The second is maps of the lanes estimated online from the vehicle's own sensors and computation, in place of maps built offline at a large cost in manual annotation [21]. The third is steer-by-wire, which removes the mechanical link from the steering wheel to the gear, and whose first production application kept a mechanical back-up [23]. The predictive steering of Section 2.5.5 belongs to the same move toward model-based control, and its loss on a slow rack shows what such a model must get right.

> ⚠ **IN PRACTICE** **What a Lane-Keeping Assist Must Meet**
>
> UN Regulation No. 79 on steering equipment names a function that assists the driver in keeping the vehicle within the chosen lane an automatically commanded steering function of Category B1. For passenger cars the lateral acceleration it is specified to hold may be at most 3 m/s^2, the half-second average of the lateral jerk it generates at most 5 m/s^3, and the driver must be able to override it with at most 50 N at the steering control. Type tests drive a curve at 80 to 90 % of the specified maximum [22]. The sharp corner of configuration 8 demands 4.0 m/s^2, beyond what such an assist may be specified to hold.

2.8 Limitations

The main limitation is the absence of tire forces. The results on geometry, feedforward and actuator limits are unaffected by the omission, and no result about the limit of adhesion can be obtained from this model. A real vehicle at $4\,\mathrm{m/s^2}$ of lateral acceleration generates measurable slip angles, the achieved curvature falls below $\tan\delta/L$, and the gap between the kinematic prediction and the dynamic one widens with speed in a way the Stanley division by v does not model [16]. Volume 7, Chapter 10 carries brush-model tires to the friction limit, where the kinematic picture no longer holds.

Localization is the second limitation and the larger one in practice. The heading, the road heading and the curvature are exact here, and the only imperfect quantity is the lateral offset. A camera-based system estimates all four from the same images with correlated errors, a latency of tens of milliseconds and a curvature estimate that degrades where the road curves most. The noise configuration in this application is therefore a stand-in rather than a model of a real perception chain. A moving horizon estimator built on Equations (2.1) to (2.5) would be the natural way to close that gap within the same model [24].

The steering actuator is one first-order lag with two hard limits. Real electric power steering has a torque limit rather than an angle-rate limit, and production lane-keeping systems bound that torque well below what the rack could deliver so that the driver can override the system at any time. Modeling the authority limit as an angle limit reproduces the saturation behavior but not the handover behavior, and the handover accounts for most of the human factors difficulty [11].

The predictive controller's model omits the actuator as it is. It carries the declared rack lag of 0.10 s, so on the 0.45 s rack of configuration 9 its prediction errs by 2.14 m two seconds ahead and its loop is unstable above 26.4 m/s, a loss the Stanley law, which carries no actuator model, does not suffer at any speed up to 35 m/s. The model also holds the speed over the horizon, linearizes the trigonometric terms and omits the clamp of $\delta + d$ at the end stop and the 20 ms staircase of the command, and the filter that removes the offset delays the answer to a measured jump.

Three smaller omissions are stated for completeness. The two loops do not compete for tire adhesion, so the vehicle can brake at $6\,\mathrm{m/s^2}$ while cornering at the same time. The road is a single lane with no camber, no superelevation and no other traffic. The projection is windowed and monotone, so it cannot jump to another pass of a road that doubles back, but it also cannot recover a vehicle placed far from its last station. The global nearest-point search that would do so is correct only where the road does not double back near the query point, and it appears only in the verification tests for that reason.

> ⚠ **OPEN CHALLENGE** **Seeing the Lane When the Paint Is Hidden**
>
> Both steering laws of this chapter assume a lane they can measure. A camera loses it where nearby vehicles or adverse weather obscure the markings, and a flat-ground projection misplaces it on slopes and rough roads. A recent survey finds notably few dedicated works on detection that pairs the camera with a laser scanner, which lighting does not affect [21]. The regulation concedes the gap: the lateral acceleration an assist is specified for may not be achievable in inclement weather, with different tires or on laterally sloped roads [22]. Progress would be lane estimates that state their own confidence, so that the controller can hand back to the driver before the measurement fails.

2.9 Conclusion

This chapter studied autonomous lane keeping on a five-state kinematic bicycle and a road stored as an arc-length table integrated from a curvature profile. The regulated variable is the cross-track error of the front axle and the heading error is referenced to the road under the vehicle, a pairing that keeps $\arctan(L\kappa)$ out of the feedback path so that the curvature feedforward supplies it once rather than twice. The default steering law is the Stanley construction with no integral action, followed by an end stop and a rate limit, over a velocity form proportional-integral speed loop. The ideal closed loop

carries poles at $-k_e v/(v+k_{soft})$ and $-v/L$, and the standing offsets left by a missing feedforward and by a misaligned linkage are confirmed to three decimal places against the running application.

Twelve configurations quantify the outcome, from a corner that costs 0.314 m without feedforward and 0.016 m with it, to a 25 m corner that no gain holds because 6.617° of steering is not available. One tuning holds a 0.5 m displacement recovery to 0.84 s and 0.85 s at 13.0 m/s and at 30.6 m/s while the steering deviation behind it changes by a factor of 2.6, and the effects the cancellation omits, a fixed actuator time spent covering a growing distance and a lateral acceleration rising as v^2, are what eventually limit it.

A model predictive controller removes the standing error of a 0.5° misalignment at 30 m/s (lane IAE 0.326 against 8.850 m s) but loses stability above 26.4 m/s on a 0.45 s rack that its 0.10 s model misjudges and the Stanley law holds. The same model supports extension to nonlinear model predictive control and to moving horizon estimation of the heading, the curvature and the steering misalignment. A fuller model would add a lateral tire model, a torque-limited rack with driver override and a comparison against pure pursuit on the same roads. The application makes measurable what a geometric path-following law does when the operating point moves beneath it.

KEY POINTS

- The steering law asks for less steering as the car goes faster, so one tuning recovers from a 0.5 m sideways displacement in 0.84 s at 13.0 m/s and 0.85 s at 30.6 m/s.
- Without integral action, a missing curvature feedforward leaves 0.314 m of standing error in the 200 m corner, and a misaligned linkage an offset that never reaches zero.
- A corner that needs 6.617° cannot be held with 6° of steering at any gain: an actuator limit, not a tuning problem.
- A predictive controller that estimates the misalignment removes its offset but loses stability above 26.4 m/s on a 0.45 s rack its model misjudges.

Where else this applies. Chapter 9 follows its routes by pure pursuit, and Volume 9, Chapter 4 removes an offset through an estimated bias, as the predictive steering here does.

References

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

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

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

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

[5] J. B. Cicchino, "Effects of lane departure warning on police-reported crash rates," *Journal of Safety Research*, vol. 66, pp. 61–70, 2018.

[6] M. Rokonuzzaman, N. Mohajer, S. Nahavandi, and S. Mohamed, "Review and performance evaluation of path tracking controllers of autonomous vehicles," *IET Intelligent Transport Systems*, vol. 15, no. 5, pp. 646–670, 2021.

[7] E. D. Dickmanns and B. D. Mysliwetz, "Recursive 3-D road and relative ego-state recognition," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 14, no. 2, pp. 199–213, 1992.

[8] C. Urmson, J. Anhalt, D. Bagnell, C. Baker, R. Bittner, M. N. Clark, J. Dolan, D. Duggins, T. Galatali, C. Geyer, M. Gittleman, S. Harbaugh, M. Hebert, T. M. Howard, S. Kolski, A. Kelly, M. Likhachev, M. McNaughton, N. Miller, K. Peterson, B. Pilnick, R. Rajkumar, P. Rybski, B. Salesky, Y.-W. Seo, S. Singh, J. Snider, A. Stentz, W. Whittaker, Z. Wolkowicki, J. Ziglar, H. Bae, T. Brown, D. Demitrish, B. Litkouhi, J. Nickolaou, V. Sadekar, W. Zhang, J. Struble, M. Taylor, M. Darms, and D. Ferguson, "Autonomous driving in urban environments: Boss and the Urban Challenge," *Journal of Field Robotics*, vol. 25, no. 8, pp. 425–466, 2008.

[9] S. Thrun, M. Montemerlo, H. Dahlkamp, D. Stavens, A. Aron, J. Diebel, P. Fong, J. Gale, M. Halpenny, G. Hoffmann, K. Lau, C. Oakley, M. Palatucci, V. Pratt, P. Stang, S. Strohband, C. Dupont, L.-E. Jendrossek, C. Koelen, C. Markey,

C. Rummel, J. van Niekerk, E. Jensen, P. Alessandrini, G. Bradski, B. Davies, S. Ettinger, A. Kaehler, A. Nefian, and P. Mahoney, "Stanley: The robot that won the DARPA Grand Challenge," *Journal of Field Robotics*, vol. 23, no. 9, pp. 661–692, 2006.

[10] G. M. Hoffmann, C. J. Tomlin, M. Montemerlo, and S. Thrun, "Autonomous automobile trajectory tracking for off-road driving: Controller design, experimental validation and racing," in *2007 American Control Conference*. IEEE, 2007, pp. 2296–2301.

[11] R. Parasuraman and V. Riley, "Humans and automation: Use, misuse, disuse, abuse," *Human Factors*, vol. 39, no. 2, pp. 230–253, 1997.

[12] A. Ollero and G. Heredia, "Stability analysis of mobile robot path tracking," in *Proceedings 1995 IEEE/RSJ International Conference on Intelligent Robots and Systems*, vol. 3. IEEE Computer Society Press, 1995, pp. 461–466.

[13] S. Macenski, S. Singh, F. Martín, and J. Ginés, "Regulated pure pursuit for robot path tracking," *Autonomous Robots*, vol. 47, no. 6, pp. 685–694, 2023.

[14] B. Paden, M. Čáp, S. Z. Yong, D. Yershov, and E. Frazzoli, "A survey of motion planning and control techniques for self-driving urban vehicles," *IEEE Transactions on Intelligent Vehicles*, vol. 1, no. 1, pp. 33–55, 2016.

[15] J. Kong, M. Pfeiffer, G. Schildbach, and F. Borrelli, "Kinematic and dynamic vehicle models for autonomous driving control design," in *2015 IEEE Intelligent Vehicles Symposium (IV)*. IEEE, 2015, pp. 1094–1099.

[16] P. Polack, F. Altché, B. d'Andréa Novel, and A. de La Fortelle, "The kinematic bicycle model: A consistent model for planning feasible trajectories for autonomous vehicles?" in *2017 IEEE Intelligent Vehicles Symposium (IV)*. IEEE, 2017, pp. 812–818.

[17] M. Guiggiani, *The Science of Vehicle Dynamics: Handling, Braking, and Ride of Road and Race Cars*. Cham: Springer, 2023.

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

[19] P. Falcone, F. Borrelli, J. Asgari, H. E. Tseng, and D. Hrovat, "Predictive active steering control for autonomous vehicle systems," *IEEE Transactions on Control Systems Technology*, vol. 15, no. 3, pp. 566–580, 2007.

[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] X. He, H. Guo, K. Zhu, B. Zhu, X. Zhao, J. Fang, and J. Wang, "Monocular lane detection based on deep learning: A survey," arXiv:2411.16316v5, 2024.

[22] United Nations Economic Commission for Europe, "UN Regulation No 79 – Uniform provisions concerning the approval of vehicles with regard to steering equipment [2025/3]," Official Journal of the European Union, L series, 10 January 2025, http://data.europa.eu/eli/reg/2025/3/oj, 2025.

[23] Clemson Vehicular Electronics Laboratory, "Electric power-assisted steering," Clemson University, http://cecas.clemson.edu/cvel/auto/systems/ep_steering.html, accessed 2026-11-20, 2026.

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

2.10 Exercises

2.1 ⌕ Concept check · 2 min

With the default tuning, $k_e = 1.2/\mathrm{s}$ and $k_{soft} = 1.0\,\mathrm{m/s}$, a cross-track error decays with nearly the same time constant at 10 m/s and at 30 m/s, although the yaw rate a given steering angle produces is three times larger at the higher speed. Why?

(a) The law divides the cross-track correction by $v + k_{soft}$, which cancels the v of the yaw-rate equation

(b) The steering rack responds faster at high speed, which offsets the larger yaw rate per degree

(c) The curvature feedforward grows with speed and supplies the part of the correction the gain lacks

(d) The heading pole $-v/L$ stays at one place at every speed and sets the speed of the recovery

2.2 Concept check · 1 min

A programmer measures the heading error against the road direction at the front-axle projection instead of under the rear axle, and leaves the curvature feedforward on. What does the car do in the 200 m corner of scenario 2 at 25 m/s?

(a) It holds the line with zero error, because the heading term and the feedforward then agree

(b) It settles about 0.31 m inside the bend, because $\arctan(L\kappa)$ is counted twice

(c) It settles about 0.31 m outside the bend, as with the feedforward switched off

(d) It settles about 0.63 m outside the bend, because the two errors add

2.3 Concept check · 2 min

In scenario 8, the car meets a 25 m corner at 10 m/s with the steering limited to 6°. Raising the cross-track gain from 2 to 12/s has which effect?

(a) The peak error falls in proportion to $1/k_e$, as the misalignment offset does

(b) The loop becomes unstable, because the rack lag sets a finite critical gain at this speed

(c) The peak error hardly changes, while the time on the end stop and out of the lane grows

(d) The peak error falls into the lane band once the gain passes 6/s

2.4 Theory and calculation · 6 min

The default car has $L = 2.9\,\text{m}$, $k_e = 1.2/\text{s}$ and $k_{soft} = 1.0\,\text{m/s}$, with an instantaneous actuator.

(a) Find the time constant of the cross-track decay at 10 and at 30 m/s. *Checkpoint:* 0.917 s and 0.861 s.

(b) Find the heading pole at both speeds, and the ratio of the heading pole to the cross-track pole at each. *Checkpoint:* -3.45 and -10.34/s; ratios 3.16 and 8.91.

(c) Find the time constant at every speed with $k_{soft} = 0$, and say what the softening constant buys at the price of the speed dependence. *Checkpoint:* 0.833 s at every speed.

2.5 Theory and calculation · 7 min

The challenge course holds a corner of $R = 140\,\text{m}$ at $v = 20\,\text{m/s}$ with a misalignment $\delta_{bias} = 0.6°$. The wheelbase is $L = 2.9\,\text{m}$, and the target band is $\pm 0.25\,\text{m}$.

(a) With the shipped tuning, $k_e = 0.5/\text{s}$ and $k_{soft} = 4.0\,\text{m/s}$ with the feedforward off, find the offset in the corner from the missing feedforward. *Checkpoint:* 0.994 m.

(b) With the same tuning, find the size of the offset from the misalignment. *Checkpoint:* 0.503 m.

(c) With the feedforward on and $k_{soft} = 1.0\,\text{m/s}$, find the smallest k_e at which the misalignment offset alone fits the band, and its size at the suggested $k_e = 2.6/\text{s}$. *Checkpoint:* 0.88/s; 0.085 m.

2.6 Theory and calculation · 7 min

Equation (2.14) gives the stability limit of the linearized loop with the rack lag τ_δ. Take $L = 2.9\,\text{m}$ and $k_{soft} = 1.0\,\text{m/s}$.

(a) Find the speed below which no cross-track gain destabilizes the loop, for the default rack of 0.10 s and for the slow rack of 0.45 s. *Checkpoint:* 29 m/s and 6.44 m/s.

(b) For the slow rack at 25 m/s, find $k^*_{e,crit}$, the gain the reader would enter, and the period of the oscillation at the limit. *Checkpoint:* 2.99, 3.11/s and 1.24 s.

(c) The app, with its 50 Hz controller, loses stability between $k_e = 2.55$ and 2.6. By what fraction is the sampled limit below the continuous one, and what causes the gap? *Checkpoint:* About 16 %.

2.7 App challenge · 7 min

Load scenario 2, Constant-radius turn (R = 200 m), with the default seed.

(a) Run 40 s as loaded and export the CSV file. Read the lane IAE on the scorecard and the `steering_deg` and `cross_track_error_m` columns in the row at 30 s, inside the corner. *Checkpoint:* IAE 0.028 m s, steering 0.831°, error 0.000 m, each within one unit of the last digit.

(b) Set Curvature feedforward to Off, press Reset, run 40 s again and read the same three values. *Checkpoint:* IAE 8.797 m s, steering 0.830°, error 0.314 m, each within one unit of the last digit.

(c) Compare the error of part (b) with Equation (2.11), and explain why the steering angle is the same in both runs. *Checkpoint:* 0.314 m predicted.

2.8 App challenge · 8 min

Load scenario 9, Slow steering actuator ($\tau_\delta = 0.45$ s), with the default seed. A playback speed of 5 suits the long run.

(a) Run 40 s as loaded ($k_e = 1.2$) and read the Peak $|e_{ct}|$ and Out of lane tiles. *Checkpoint:* 0.13 m and 0.0 %.

(b) Press Reset, set Cross-track gain to 3 before Run, and run to 40 s. Read the same tiles. *Checkpoint:* 0.46 m and 0.0 %.

(c) Continue the same run to 60 s and export the CSV file. Find the first row in which `lane_departure` is 1, and read the Peak $|e_{ct}|$ tile. *Checkpoint:* 44.7 s, within 0.1 s; more than 1 m (about 61 m).

2.9 App challenge · 9 min

Load scenario 11, The same gain at 30 m/s, or a model that knows the speed, with the default seed. The exported file begins with comment lines that start with #.

(a) Run 40 s as loaded, with Steering on MPC at 30 m/s. Read the Peak $|e_{ct}|$ tile and the lane IAE on the scorecard, and the `mpc_offset_est_deg` column in the row at 20 s. *Checkpoint:* 0.09 m, 0.326 m s within 0.002, and 0.499°.

(b) Set Steering to Stanley, press Reset and run 40 s. Read the lane IAE and `cross_track_error_m` at 20 s, and compare the error with Equation (2.12). *Checkpoint:* 8.850 m s within 0.002, and –0.225 m against 0.2254 m predicted.

(c) Set Speed setpoint to 10 m/s, press Reset and run 40 s under Stanley. In the row at 40 s, read the column `cross_track_error_m`, and explain the ratio of the two standing errors. *Checkpoint:* –0.080 m; ratio 2.82.

2.10 Application area · 3 min

The kinematic bicycle is a fair model up to roughly 0.4 g of lateral acceleration.

(a) Find $v^2\kappa$ for the 200 m corner at 25 m/s, for the peak curvature $2\pi \cdot 3.5/60^2$ of the lane change at 25 and at 30 m/s, and for the 25 m corner at 10 m/s. Which exceed 0.4 g? *Checkpoint:* 3.13, 3.82, 5.50 and 4.00 m/s^2; the last two.

(b) What would a real car do in those cases, and which results of the chapter are optimistic for that reason?

2.11 Application area · 4 min

Highway designers join straights to arcs with transition spirals.

(a) The sudden-curvature road steps from a straight into $R = 120$ m. Find the steering angle the feedforward then demands at once. *Checkpoint:* 1.384°.

(b) In a 30 m spiral into $R = 150$ m at 25 m/s, find the rate at which the required angle moves, using $\mathrm{d}\delta/\mathrm{d}t \approx vL\,\mathrm{d}\kappa/\mathrm{d}s$, and the lag this leaves behind a rack of 0.1 s. *Checkpoint:* 0.92 °/s; 0.09°.

(c) A production lane-keeping assist limits its steering torque well below what the rack can deliver. Give the reason, and say what modeling that limit as an angle limit leaves out.

2.12 Review the agent · 5 min

An AI assistant is asked to tighten the lane keeping in scenario 10, Noisy cross-track measurement ($\sigma_{gps} = 0.25\,\mathrm{m}$ on the gentle sweep at 25 m/s), and replies as follows. "With the default rack, $\tau_\delta v = 2.5\,\mathrm{m}$ is less than $L = 2.9\,\mathrm{m}$, so by the Routh condition of the cubic no cross-track gain can destabilize this loop. Raise k_e from 1.2 to 4/s. Every offset in the Stanley loop falls as $1/k_e$, so the tracking error will fall by more than three times, and the loop stays stable." Find the flaw and say how you would check it.

CHAPTER 3

Drone Taxi Simulator for Load Feedforward and Wind Estimation Between Buildings

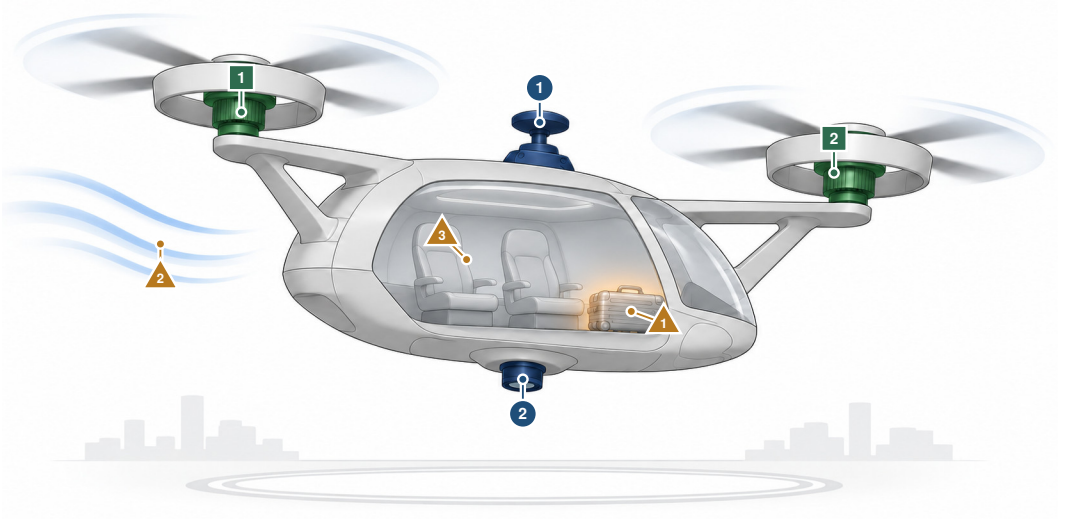

The controller moves 1 the left rotor group, which lifts the taxi
2 the right rotor group: pushing harder on one side tilts the taxi

The controller reads 1 the navigation unit on the roof: position, height and tilt
2 the downward-looking sensor: height above the pad

What disturbs it 1 the suitcase on the sliding tray, which moves the balance point 2 the wind between the buildings
3 passengers boarding at each pad, who change the weight

Notice that the taxi has no wing: the two rotor groups carry its whole weight, and it leans to move sideways or to stand against the wind. Illustration, not to scale.

IN PLAIN TERMS Telling a Flying Taxi What It Carries

The main idea. A flying taxi with no wing is held up by its spinning rotors, and it moves sideways by leaning, pushing harder on one side than the other. A computer adjusts those pushes many times a second to keep it on its path.

The challenge. Every stop brings a new load, and the wind changes behind every building. A computer that only reacts to its own errors lets a heavy taxi lag and drift, so it helps to tell it the load and let it estimate the push of the wind.

Why it matters. A passenger flight has to land near the center of a small rooftop pad.

THREE WAYS TO READ THIS CHAPTER

Explore	9 min	the first page, the fact on the size of research quadrotors (Section 3.1), the history of the quadrotor (Section 3.2), the history of the disturbance observer (Section 3.4), the try-it segment and the fact on the graded route (Section 3.6), and the key points
Study	45 min	the whole chapter with its three pause points and the twelve exercises at its end
Practice	15 min	the first page, the results (Section 3.6), Technology in Practice (Section 3.7), Limitations (Section 3.8) and the open challenge

The control question A passenger multirotor meets a new load at every rooftop and a new wind behind every building. What is it worth to tell the flight controller the load, or to let it estimate the force it was not told about, rather than leave both to the integral terms?

Learning objectives

- Compute the mass, center of mass, trim and thrust margin of a loaded vehicle from its weight and balance and its battery charge.
- Predict from a force and a gain the error a loop keeps when it is not told about the force, and explain why the tilt error never decays.
- Separate the work of a weight feedforward from that of a gain scheduled on the mass.
- Compare an integral and a disturbance observer on a wind that changes along the route, including the cost of a faster observer.
- Explain what the dispatch, gate and go-around rules of the supervisor protect.

Prerequisites Sections 1.4.3, 1.4.4, 1.4.7, 1.6.2 and 1.6.4, with the loop-action convention of Section 1.3.

The app https://apmonitor.com/apps/drone-taxi/



See also Volume 3, Chapter 10 for a position loop that commands an attitude and Volume 3, Chapter 12 for the gravity feedforward of a load that changes along a route. Volume 5, Chapter 13 for an ocean current carried as an estimated state.

3.1 Introduction

A feedback loop learns about a disturbance only from the error that the disturbance has already caused. A heavier load, a shifted center of mass and a steady wind are forces that a controller could be told about or could estimate, yet a first course usually leaves each of them to the integral term. Motion control met the same limitation when a proportional position servo carrying an unknown load torque, such as friction or gravity, was shown to keep a steady and a transient error, and the disturbance observer was the answer [1]. Observer-based rejection has since spread across industrial sectors [2], and flight control was among the first applications of scheduling a controller on the measured operating condition [3].

The class of problem is broad. A crane lifting different containers, an elevator with a varying number of riders, a robot arm picking parts of unknown mass and a ship holding station in changing weather all carry a load that differs from one run to the next and meet an environment that changes with position or time. Each disturbance has a part known before the run, a part that can be inferred from the motion, and a remainder that only feedback can remove. Feedforward, scheduled gains and an observer are the three standard tools for the first two parts [2, 3].

An electric multirotor air taxi carries all three parts on one vehicle. It has no wing, so the rotors hold the whole weight at every instant, and it moves sideways only by tilting that thrust, which puts attitude inside the horizontal position loop of a cascade [4]. Among electric vertical takeoff and landing configurations the multirotor is the most efficient in hover [5], which suits short trips between rooftops, and short trips are where the load changes most often. Buildings shelter the air behind them and shed turbulent wakes [6], gaps between them channel the flow, and in simulated flight of a small unmanned helicopter the building wakes of an 8 kt wind changed its attitude by as much as 10° [7].

> ✦ DID YOU KNOW? **Research Quadrotors Weighed a Few Kilograms**
>
> In 2007 the research groups working on quadrotors consistently chose vehicles of 0.3 to 4.0 kg [8]. A six-passenger quadrotor concept, flown in a piloted motion simulator in 2021, carries a payload of 1,200 pounds [9]. The two-seat vehicle of this chapter is 400 kg empty with a 220 kg payload rating.

The belief this plant is built to dislodge is that a well-tuned feedback loop has no need to know its load. In the application the same gains hold a light climb within 0.09 m of its reference and let

a heavy one fall 1.78 m behind, a suitcase at the end of the baggage tray leaves a standing tilt error of 2.46° that no attitude gain can remove, and an integral that has learned a wind keeps the vehicle leaning after the wind has gone. A transfer function drawn at one operating point shows none of this. Reviews of process control practice and of good practice in control teaching both favor experiments in which the behavior of a realistic plant, rather than algebra, decides the outcome [10, 11].

However, the value of each kind of knowledge is easy to state and hard to see, because it depends on how large the unknown force is against the loop gain and on how fast the force changes along the path. This chapter treats what a flight controller knows about its load and its wind as the experimental variable of a drone taxi simulator, and measures what each level of knowledge is worth on the same route. Cascaded multirotor control, gain scheduling, disturbance observers and jerk-limited references are all standard methods, and Section 1.4.3 sets out the first three on a cart. Section 3.2 follows the domain from helicopter stabilizers to passenger multirotors. Section 3.3 builds a nine-state planar vehicle whose rotor arms are taken about the true center of mass, with weight and balance set at boarding, a battery whose power limit sags at low charge, and a closed-form wind field attached to the buildings that the model and the scene share, each stated with its constants. Section 3.4 describes the three loops, the three levels of knowledge and the supervisor. Section 3.5 derives an error budget that predicts from a force and a gain how far each loop must stray to supply what it was not told, and separates feedforward from scheduling. Section 3.6 gives the measurements of ten cases and a graded route that isolate load, balance, wind, battery and approach one at a time, including the trade-off in observer bandwidth. Section 3.7 names the sensors, the motors and the weight and balance rules of a real multirotor, Section 3.8 states the limitations and Section 3.9 draws the results together.

3.2 Hover, City Wind and the Passenger Multirotor

Vertical flight has always been a control problem as much as a propulsion problem. A helicopter in hover is dynamically unstable, and the treatment of rotorcraft aeromechanics that covers its flight dynamics also supplies the momentum theory, figure of merit and rotor effects used in this chapter [12]. Helicopter control systems grew from simple mechanical stabilizing devices into automatic flight control systems that hold the aircraft for its pilot [13].

Multirotors took a different path. Several fixed-pitch rotors on a rigid frame need no swashplate, and attitude comes from the difference in thrust between rotors, which makes the vehicle mechanically simple and entirely dependent on its feedback loops. The history box below tells how the quadrotor became a research platform. The standard architecture places an attitude loop inside a position loop on inertial estimates [4], and the same vehicles now estimate the external force and moment acting on them from their own dynamics and inertial measurements, for collision detection and for flight in wind [14].

> ⧖ HISTORY **Four Rotors Waiting for Fast Electronics (1920s to 2007)**
>
> The first successful quadrotors flew in the 1920s, yet no practical quadrotor was built for decades, largely because four motors are hard to control at once with enough bandwidth. A crewed quadrotor of 1963 left ground effect, but it had no stability augmentation to lighten the pilot's work, so a steady hover was near impossible, and it stopped at the prototype. Faster microprocessors and micro-electro-mechanical inertial sensors later brought small radio-controlled quadrotors with stability augmentation. By 2007 the quadrotor was a popular research platform, simple to build and maintain and able to hover, with fixed-pitch rotors controlled by motor speed alone. Flight tests then showed that airspeed, angle of attack and airframe design add moments that disturb the attitude and thrust variations that disturb the altitude, effects a design made at hover does not anticipate [8].

The observer used by the third controller mode came from motion control, as the history box in Section 3.4 tells. Gain scheduling, the tool of the second mode, had its early examples in flight control and automotive engine control [3].

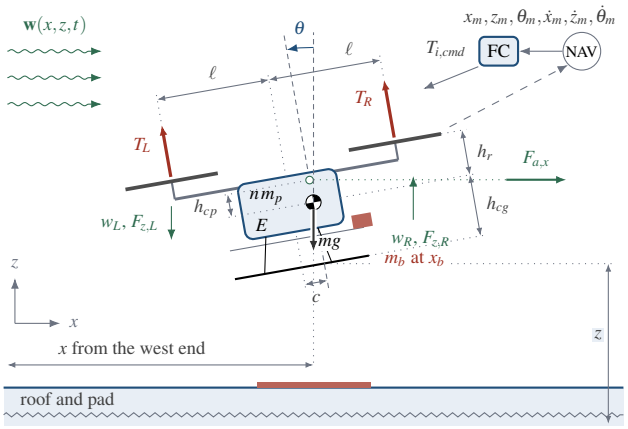

Figure 3.1 The planar air taxi as modeled, with the tilt drawn exaggerated. Two rotor groups at a distance ℓ either side of their centroid produce T_L and T_R along the body axis, a height h_r above the center of mass, which sits a distance c from the centroid set by the baggage m_b at x_b, with the n passengers of mass m_p on the centroid. The wind $\mathbf{w}$ drives a horizontal air force $F_{a,x}$ acting a height h_{cp} over the center of mass and a vertical drag $F_{z,i}$ at each group from its local vertical wind w_i. The skid height z is measured from the street below the break, and the navigation system NAV passes noisy states to the flight computer FC, which commands the group thrusts $T_{i,cmd}$ from the battery energy E.

Wind around buildings has its own literature. The flow around a single surface-mounted block separates and leaves a turbulent wake whose structure depends on the turbulence and shear of the approaching flow [6]. Rotorcraft flying through building wakes see large attitude disturbances [7], and the inability to fly well in adverse weather is among the obstacles listed for wider use of drones in cities [15].

Passenger air taxis carry these problems into commercial service. Of the three main electric configurations the multirotor is the most efficient in hover, the vectored-thrust jet in cruise, and lift-plus-cruise designs sit between [5]. Reviews of urban air mobility list certification and policy, traffic management and ground infrastructure among the questions still open [16], and the battery performance needed to electrify vertical takeoff and landing has been set out as a list of required metrics [17]. Comfort enters through the route. Studies of acceleration and jerk in ground transport vary too widely to fix one acceptable profile [18], so the comfort limits here are design choices, as are all the numerical targets of the application.

3.3 A Planar Air Taxi Among Rooftops

The vehicle belongs to the two-seat multirotor class, 400 kg empty with a 220 kg payload rating, and flies in the vertical plane of its route. The scene draws four rotors on each side, and the model lumps each side into one rotor group with one thrust acting a distance ℓ either side of the rotors' centroid. Horizontal position x points right, the height z of the landing skids above the street points up, and the tilt θ is measured counterclockwise, so a positive tilt leans the thrust toward $-x$. The geometry is shown in Figure 3.1.

Table 3.1 sorts every signal by its role. Nine states are integrated, the position and velocity in x and z, the tilt and its rate, the two group thrusts and the battery energy. The mass, center of mass and inertia change only when passengers or baggage move on a pad or the tray slides in flight, and the wind is a function of position and time.

Table 3.1 Variables, With Nominal Values From the Calm Hover of Case 1

Symbol	Role	Nominal	Range
z	CV, skid height	40 m	0–90 m
x	CV, horizontal position	28 m	0–400 m
θ	CV of the inner loop, tilt	0°	—
T	MV, total thrust	4.80 kN	0–8.4 kN
θ_{SP}	MV of the position loop	0°	±15°
ΔT	MV of the tilt loop, $T_R - T_L$	0 N	±4.05 kN
$\dot{x}, \dot{z}, \dot{\theta}$	states, rates	0	—
T_L, T_R	states, group thrusts	2.40 kN	0–4.2 kN
E	state, battery energy	19.2 kWh	0–24 kWh
n, m_p	DV, passengers, mass each	1, 80 kg	0–2, 50–120
m_b, x_b	DV, baggage and tray position	10 kg, 0	0–40, ±1.5 m
e_m	DV, load-sheet mass error	0 %	±30 %
U_0	DV, background wind at 50 m	0 m/s	±14 m/s
s_t	DV, turbulence scale	1	0–2
—	DV, gust front, push, bag slide	events	—

3.3.1 Rigid Body and Rotor Groups

The thrust points along the body axis, $(-\sin\theta, \cos\theta)$ in the earth frame, which gives the translational balance of Equation (3.1).

$$\begin{aligned} m\ddot{x} &= -T\sin\theta + F_{a,x} + F_p \\ m\ddot{z} &= T\cos\theta - mg + F_{a,z}, \qquad T = T_L + T_R \end{aligned} \tag{3.1}$$

where F_p is a push event of 600 N for 2 s. Tilt buys horizontal force and costs vertical force, and at 5° the vertical share of the thrust falls by 0.38 %. The pitch balance of Equation (3.2) takes each rotor arm from the actual center of mass.

$$I\ddot{\theta} = (\ell - c)\,T_R - (\ell + c)\,T_L + M_a \tag{3.2}$$

where c is the center of mass measured along the body axis from the rotors' centroid, positive toward the right group, and M_a is the aerodynamic moment. Setting $\ddot{\theta} = 0$ in still air gives the trim of a level hover, $T_R - T_L = cT/\ell \approx mgc/\ell$, so a center of mass off the centroid needs a standing difference in thrust. Each group follows its command through a first-order lag, $\tau_T \dot{T}_i = T_{i,cmd} - T_i$ for $i \in \{L, R\}$, and negative thrust is set to zero after each step.

3.3.2 Weight and Balance

The seats sit on the rotors' centroid and the baggage tray slides along the body axis, which gives the mass properties of Equation (3.3).

$$\begin{aligned} m &= m_e + nm_p + m_b, \qquad c = \frac{m_b x_b}{m} \\ I &= I_e + m_e c^2 + nm_p\left(r_p^2 + c^2\right) + m_b\left(x_b - c\right)^2 \end{aligned} \tag{3.3}$$

Passengers change the mass and the inertia, and baggage also moves the center of mass. A 40 kg suitcase at $x_b = -1.5$ m with one passenger gives $m = 520$ kg, $c = -11.5$ cm and a level-hover trim of −294 N. Boarding happens only on a pad. Each passenger walks between the roof door and the cabin in 2.2 s, successive passengers start 0.9 s apart, the baggage takes 1.2 s, and the mass changes as each passenger crosses the cabin door. In flight the mass is constant, while the tray can slide.

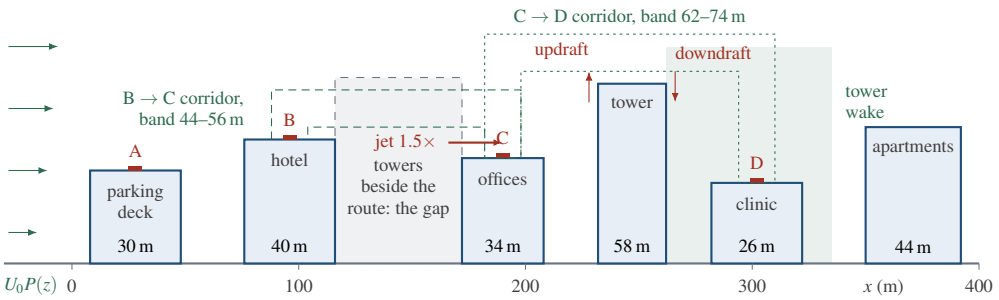

Figure 3.2 The street, with the lengths to one scale and the heights to a larger one, the pads A to D and the buildings' heights. The two towers beside the route over the plaza make the gap, whose channeled flow leaves as a jet over pad C when the wind blows toward $+x$, and the 58 m tower sheds a wake over pad D with an updraft at its windward roof edge and a downdraft behind its leeward one. The dashed outlines are the corridors of the two legs of the graded route, a shaft 16 m wide over each pad and a cruise band 12 m deep centered 10 m above the highest roof between the pads. The arrows at the left show the background profile $U_0P(z)$.

3.3.3 Air Loads

Drag follows the airspeed relative to the local wind, with a different area on each axis, as shown in Equation (3.4).

$$\begin{aligned} v_{r,x} &= \dot{x} - u, \qquad v_{r,i} = \dot{z} + \dot{\theta}\, r_i - w_i \\ F_{a,x} &= -\tfrac{1}{2}\rho\, C_d A_x V_r\, v_{r,x} \\ F_{z,i} &= -\tfrac{1}{2}\rho\, \tfrac{1}{2} C_d A_z \sqrt{v_{r,x}^2 + v_{r,i}^2}\; v_{r,i} \\ M_a &= r_R F_{z,R} + r_L F_{z,L} - h_{cp} F_{a,x} \end{aligned} \tag{3.4}$$

where u is the horizontal wind at the center of mass, w_i the vertical wind at rotor group i, $r_{R,L} = (\pm\ell - c)\cos\theta$ the horizontal lever arms of the groups, $V_r = [v_{r,x}^2 + (\dot{z} - \bar{w})^2]^{1/2}$ with $\bar{w} = (w_L + w_R)/2$, and $F_{a,z} = F_{z,L} + F_{z,R}$. Splitting the vertical area between the groups lets a roof-edge updraft under one group pitch the vehicle, and raising the horizontal drag h_{cp} above the center of mass leans the vehicle downwind. A vehicle drifting with the wind feels no force. The vertical wind is sampled at each group's rotor plane, a height h_r above a center of mass that is itself h_{cg} above the skids.

3.3.4 The City and Its Wind

Six buildings stand along a 400 m street, and four carry a pad of half-width 3 m, as shown in Figure 3.2. They are a parking deck from 8 to 48 m, 30 m tall with pad A at 28 m, a hotel from 76 to 116 m, 40 m tall with pad B at 96 m, offices from 172 to 208 m, 34 m tall with pad C at 190 m, a tower from 232 to 262 m, 58 m tall, a clinic from 282 to 322 m, 26 m tall with pad D at 302 m, and apartments from 350 to 392 m, 44 m tall. Between B and C an open plaza from 116 to 172 m is flanked by two 60 m towers that stand beside the route rather than on it, and they make the gap.

The wind at a point is the background speed $U(t)$ at the reference height of 50 m, shaped by a

Table 3.2 Wind-Field Features for a Wind Toward $+x$. H Is the Building Height, d the Distance Downwind of Its Leeward Face, a the Distance Upwind of Its Windward Face, s the Distance Along the Roof, $S(v) = 1/(1+\mathrm{e}^{-v})$, $B(v) = \mathrm{e}^{-v^2}$, and Lengths Inside S and B Are in Meters

Feature	Where	Contribution
lee shelter	$0 < d \le 6H$	$f_s = 1 - k_s\,\mathrm{e}^{-d/(L_w H)}\,S(h_w - z)\,R$ below the shear layer $h_w = H[1.1 + 0.15\min(d/H, 2)]$, with $k_s = 1.05$ and $L_w = 2.5$
lee turbulence	$0 < d \le 6H$	$f_t^2 \mathrel{+}= \{0.24\,\mathrm{e}^{-d/(3.2H)}\,R\,[0.6\,S(h_w - z) + B((z - h_w)/0.22H)]\}^2$
leeward downdraft	$0 < d \le 6H$	$f_w \mathrel{-}= 0.35\,e_k P(H)\,\mathrm{e}^{-d/(0.6H)}\,B((z - 0.95H)/0.38H)$
windward stagnation	$0 < a \le 1.5H$	$f_s = 1 - 0.55\,\mathrm{e}^{-a/(0.45H)}\,S((0.85H - z)/1.5)$
windward updraft	$0 < a \le 1.5H$	$f_w \mathrel{+}= 0.45\,e_k P(H)\,\mathrm{e}^{-a/(0.35H)}\,B((z - 0.92H)/0.3H)$
roof bubble	on the roof	$f_s = 1 - 0.65\min(1, e_k)\,\mathrm{e}^{-s/(0.8H)}\,b$ and $f_t^2 \mathrel{+}= [0.22\min(1, e_k)\,\mathrm{e}^{-s/(0.8H)}\,b]^2$, with $b = S((1.2H - z)/0.8)\,S((z - H + 0.5)/0.3)$
roof speed-up, roll-over	on the roof	speed $\times\,[1 + 0.12\min(1, e_k)\,B((z - 1.32H)/0.35H)]$, and $f_w \mathrel{+}= 0.3\,e_k P(H)\,\mathrm{e}^{-s/(0.3H)}\,B((z - 1.12H)/0.25H)$
gap	116–172 m, below 60 m	weight $g = S(a_g/2.5)\,E\,S((60 - z)/2.5)$ with $E = 1$ inside and $\mathrm{e}^{-d_g/20}$ past the exit, and $f_t^2 \mathrel{+}= (0.12g)^2$
combination	everywhere	f_s is the smallest shelter factor, $f_u = (1 - g) f_s\,\mathrm{speed} + g\,k_{gap}$, $f_w \leftarrow f_w(1 - 0.7g)$, and f_t^2 starts at 0.06^2

a_g and d_g are the distances past the gap entrance and exit, $R = 1 - S(a_g/2.5)S((60 - z)/2.5)$ removes the wake of a building upwind of the entrance once the channeled flow has replaced it, and e_k is the exposure of building k, the upstream shelter and gap factor evaluated 2 m upwind of its windward roof edge and floored at 0.1, so a low roof in the wake of a tall one has weak edge flows of its own.

power-law profile and by closed-form features attached to the buildings, as shown in Equation (3.5).

$$\begin{aligned} u &= U P(z_c) f_u(x, z_c) + \sigma n_u(t) \\ w_i &= |U| f_w(x_i, z_i) + 0.6\,\sigma n_w(t) \\ \sigma &= s_t |U| P(z_c) f_t(x, z_c) \\ P(z) &= (\max(z, 2\,\mathrm{m})/50\,\mathrm{m})^{\alpha} \\ U &= U_0 \big[1 + 0.08\sin(2\pi t/T_1 + \phi_1) \\ &\qquad + 0.05\sin(2\pi t/T_2 + \phi_2)\big] + U_g \end{aligned} \tag{3.5}$$

where $z_c = z + h_{cg}$ is the height of the center of mass, f_u the horizontal speed as a fraction of the free stream, f_w the vertical speed as a fraction of $|U|$, and f_t the turbulence intensity. The slow variation has periods $T_1 = 23$ s and $T_2 = 9.1$ s, with phases ϕ_1 and ϕ_2 drawn once per run from the scenario seed. The turbulence signals n_u and n_w are first-order colored noise of unit standard deviation with correlation times of 1.2 and 0.6 s, drawn on the controller clock and held over each control step so that the Runge-Kutta stages see a smooth right-hand side. A gust front adds $U_g = \frac{1}{2} A_g [1 - \cos(2\pi(t - t_0)/D_g)]$ for $t_0 \le t \le t_0 + D_g$, with $A_g = 6$ m/s and $D_g = 6$ s unless a case says otherwise. The features that make up f_u, f_w and f_t are listed in Table 3.2. The field mirrors itself when the wind reverses, and the scene draws its arrows and shaded zones from the same functions, so the picture and the flight cannot disagree. The field is illustrative and reduced-order. It is not computational fluid dynamics and represents no particular city.

Table 3.3 Plant Parameters. Every Value Is Adjustable in the Model Dialog Over the Range Shown Unless Marked Hidden

Symbol	Meaning	Value	Range
m_e	empty mass	400 kg	300–500
I_e	empty pitch inertia	700 kg m^2	400–1200
ℓ	centroid to group thrust line	2.0 m	1.4–2.6
r_p	passenger radius of gyration	0.3 m	hidden
h_{cg}	skids to center of mass	1.1 m	hidden
h_r	rotor plane above center of mass	0.8 m	hidden
r_{rot}	rotor radius for strikes	0.75 m	hidden
—	skid half-length	1.2 m	hidden
T_{max}	group thrust rating	4200 N	3000–6000
T_{idle}	idle thrust per group	150 N	0–600
τ_T	group thrust lag	0.12 s	0.03–0.4
A	disk area per group	12 m^2	6–20
FM	overall figure of merit	0.6	0.4–0.8
C_dA_x, C_dA_z	drag areas	8, 24 m^2	2–16, 6–40
h_{cp}	pressure center above mass center	0.4 m	−0.5–1.2
ρ	air density	1.225 kg/m^3	1.0–1.3
g	gravity	9.81 m/s^2	hidden
E_{cap}	usable battery energy	24 kWh	12–40
P_{max}	rated battery power	190 kW	120–300
s_k, p_f	sag knee, power at empty	35, 50 %	10–60, 20–100
P_{av}	avionics and cabin power	1.5 kW	0–5
θ_{max}	tilt command limit	15°	5–30
—	sink that breaks the gear	2.5 m/s	1–5
—	tilt that tips over at touchdown	12°	4–25
σ_x, σ_v	position, velocity noise	0.05 m, 0.03 m/s	0–0.5, 0–0.3
σ_θ, $\sigma_{\dot\theta}$	tilt, gyro noise	0.1°, 0.3 °/s	0–1, 0–3
α	profile exponent	0.28	0.1–0.45
k_{gap}	gap speed-up	1.5	1–2
L_w, k_s	wake length, shelter depth	2.5H, 1.05	1–5, 0.5–1.2

3.3.5 Battery and Available Thrust

Hover momentum theory gives the induced power of an actuator disk as $T^{3/2}/\sqrt{2\rho A}$ [12], and an overall figure of merit FM folds in profile power, motors and inverters, as shown in Equation (3.6).

$$\begin{aligned} \dot{E} &= -\left[\sum_{i=L,R} \frac{T_i^{3/2}}{FM\sqrt{2\rho A}} + P_{av}\right] \\ T_{avail} &= \min\left(T_{max},\ \left[P_g\, FM\sqrt{2\rho A}\right]^{2/3}\right) \\ P_g &= \tfrac{1}{2}\left[P_{max} f(s) - P_{av}\right] \end{aligned} \tag{3.6}$$

where A is the disk area of one group, $s = E/E_{cap}$ the state of charge, and $f(s) = 1$ above the knee s_k and $p_f + (1-p_f)s/s_k$ below it, as the pack voltage sags. At 490 kg the vehicle draws 73.9 kW to hover. The battery limit binds below a charge of 27.2 %, at 17 % a group can deliver 3.72 kN instead of its 4.2 kN rating, and an empty battery delivers nothing. The lift-plus-cruise aircraft of Chapter 4 meets the same limit, a battery that bounds power as well as energy.

3.3.6 Contact, Strikes and Numerics

A touchdown occurs when the skids reach a surface moving down, either z at least 5 mm below the roof height with $\dot{z} < 0$ or a skid end at the surface sinking faster than 0.02 m/s, and the second test is ignored for 0.5 s after liftoff so that a skid rocking off the roof it just left is not a landing. A touchdown in the street, beyond 12° of tilt or faster than 2.5 m/s is a crash, one more than 3 m from a pad center is a building violation, and a rotor tip inside a building, or a skid end more than 0.15 m inside one, is a strike. Otherwise the vehicle is set level with zero velocity and the touchdown speed and pad error

are recorded. Liftoff requires $T\cos\theta > mg$ and is suppressed for 0.3 s after a touchdown, so that a vehicle that touched down carrying almost its weight does not bounce. The parameters are listed in Table 3.3. The states advance by classical fourth-order Runge-Kutta at 2 ms, control executes at 100 Hz, and the log is written at 20 Hz, on the engine of Section 1.6.1. Halving the step, with control still at 100 Hz, moves the touchdowns of the graded route by at most 0.04 mm/s and 0.02 mm with the suggested controller and 0.04 mm/s and 0.06 mm with the shipped one. A scenario starts either on a pad at idle thrust with the load still to board, or in the air trimmed for its load, with the battery at the charge the case sets.

3.3.7 Assumptions

Flight is planar. There is no yaw, roll or sideslip and no wind out of the plane, so the wind of the application is a crosswind to the vertical takeoff and landing paths, and a vehicle flying toward $-x$ flies backward rather than turning.

Each side is one rotor group with a first-order thrust lag. Rotor inflow dynamics, blade flapping and the rotor H-force are absent, and those effects change thrust and moment with airspeed and angle of attack [8]. A real rotor at fixed speed gains or loses thrust in a vertical gust far more strongly than the quadratic drag here, so vertical gusts are understated in calm air.

Drag uses fixed areas in the earth frame, which holds for the small tilts of low-speed flight. Wingborne cruise and rotor tilting are outside the model.

Power follows hover momentum theory with a constant figure of merit, without climb, descent or forward-flight corrections and without vortex-ring state [12]. The battery is an energy store with a sagging power limit and no voltage dynamics, temperature or aging.

The wind field is a set of closed-form patches tuned to the qualitative features of flow around buildings [6]. Turbulence is temporal colored noise scaled by a position-dependent intensity, the same at both groups, and wakes neither meander nor shed vortices.

The skids are rigid, touchdown is instantaneous and sets the vehicle level, and there is no ground effect. Navigation measurements carry small colored noise with no latency, bias or satellite dropouts, which between real buildings are a problem of their own. The rotor radius used for strikes, 0.75 m, is smaller than the 0.98 m implied by the disk area of four 1.95 m rotors per group, so clearance is slightly optimistic.

3.4 Three Loops and Three Levels of Knowledge

The flight controller runs every $\Delta t_c = 10$ ms on sampled navigation measurements and never reads the true state, mass, center of mass or wind. Position and height carry colored noise of standard deviation σ_x with a 0.5 s correlation time, the velocities σ_v with 0.05 s, the tilt σ_θ with 0.05 s and the rate gyro $\sigma_{\dot\theta}$ with 0.02 s. The loop structure is shown in Figure 3.3.

3.4.1 The Shared Loop Algorithm

All three loops use the framework's velocity-form PID of Section 1.6.2, written in this chapter's symbols in Equation (3.7).

$$\begin{aligned} e_k &= a\,(\mathrm{SP}_k - \mathrm{PV}_k), \qquad P_k = K_c e_k \\ S_k &= S_{k-1} + (P_k - P_{k-1}), \qquad c_k = b_k + S_k + I_{k-1} \\ u_k &= \mathrm{clamp}(c_k,\ u_{min,k},\ u_{max,k}) \\ I_k &= I_{k-1} + \left[\frac{K_c}{\tau_I} e_k + \frac{u_k - c_k}{T_t}\right]\Delta t_c \end{aligned} \tag{3.7}$$

where $a = \pm 1$ is the action and b_k is a bias recomputed every step. In the symbols of Equations (1.19) to (1.22), c_k is the output before limiting v_k and b_k takes the place of u_b. The running sum S_k telescopes to $K_c e_k$ for whatever gain is in force, which keeps a scheduled gain meaning the gain, and

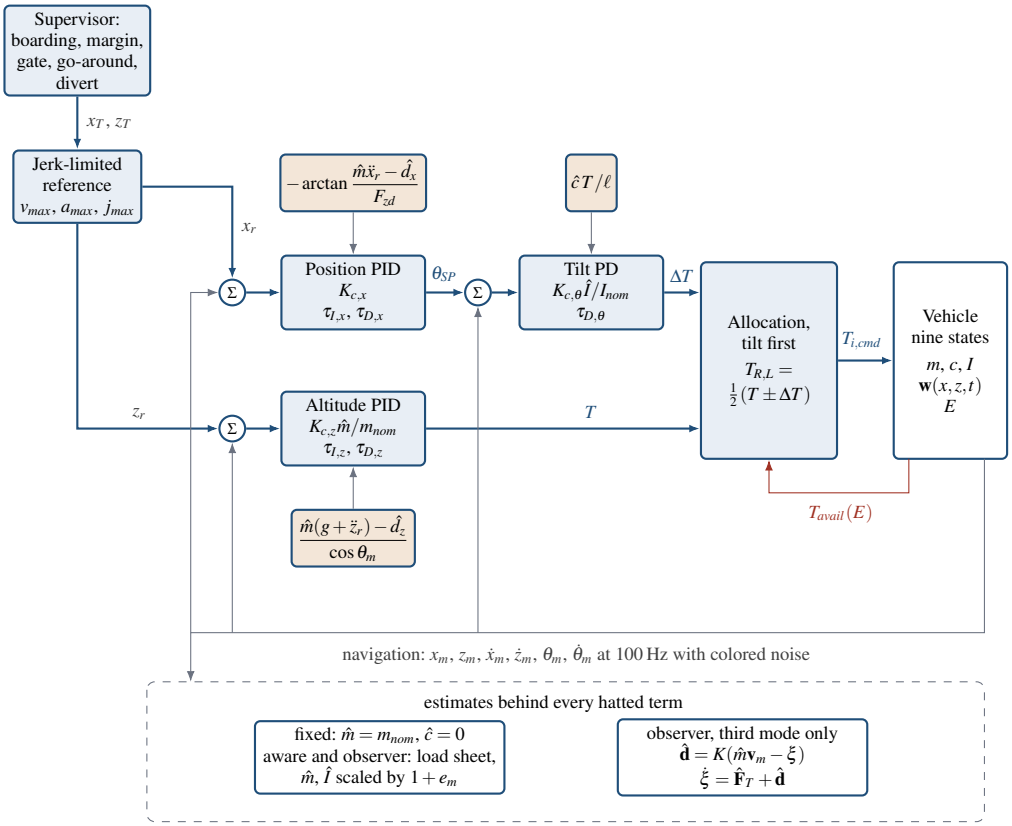

Figure 3.3 Three loops, what they are told and what supervises them. The position loop owns no actuator and writes the tilt setpoint θ_{SP}, the tilt loop writes the differential thrust ΔT, and the altitude loop writes the total thrust T. The shaded blocks are feedforward terms built from the estimates, which are the design point in fixed tuning, the load sheet in the payload-aware mode, and the load sheet with the observer's force estimate $\hat{\mathbf{d}}$ in the third mode. The allocation serves the tilt first within the available thrust, and the supervisor sets the targets of the reference generator.

back-calculation uses $T_t = \tau_I$. No output is rate limited, so the group thrust lag is the only actuator dynamics. The framework's own derivative on the measurement is set to zero in all three loops, and each derivative term is formed from measured rates and carried in b_k instead, so it produces no kick on a setpoint change. Each loop has its own manual and automatic switch, and transfer in either direction is bumpless.

3.4.2 Altitude and Position

The altitude loop has $a = +1$ and is reverse acting in the convention of Section 1.3, because more thrust raises the vehicle, and its bias is the weight feedforward plus the derivative term, as shown in Equation (3.8).

$$
\begin{aligned}
b_z &= \frac{\hat{m}(g+\ddot{z}_r)-\hat{d}_z}{\max(\cos\theta_m,\,0.5)} + K^e_{c,z}\tau_{D,z}\,(\dot{z}_r-\dot{z}_m) \\
T &\in \left[2T'_i+|\Delta T|,\; 2T_{avail}-|\Delta T|\right]
\end{aligned}
\tag{3.8}
$$

with $T'_i = \min(T_{idle}, T_{avail})$ and the upper bound never below the lower, where $\hat{m}$ is the mass the controller believes, $\hat{d}_z$ the observer's vertical force estimate (zero unless the third mode is selected),

z_r the reference and $K^e_{c,z}$ the gain in force. On the pad the bias sits at idle, and the 1.5 s spool-up ramps it from idle to the hover feedforward.

The position loop has $a = -1$ and is direct acting, because a positive tilt accelerates the vehicle toward $-x$, and its output is the tilt setpoint, limited to $\pm\theta_{max}$, with a feedforward for the reference acceleration and the estimated force, as shown in Equation (3.9).

$$\begin{aligned} b_x &= -\arctan\frac{\hat{m}\ddot{x}_r - \hat{d}_x}{F_{zd}} - K_{c,x}\tau_{D,x}\left(\dot{x}_r - \dot{x}_m\right) \\ F_{zd} &= \max\left(0.2\hat{m}g,\ \hat{m}(g+\ddot{z}_r) - \hat{d}_z\right) \end{aligned} \tag{3.9}$$

Both derivative terms act on measured velocity against the reference's own velocity, so the reference supplies its own rate feedforward.

3.4.3 Tilt Loop and Allocation

The tilt loop is proportional-derivative and reverse acting, because more differential thrust raises the tilt, with the rate term on the gyro and the trim of Equation (3.2) fed forward, as shown in Equation (3.10).

$$\Delta T = K^e_{c,\theta}\left[(\theta_{SP} - \theta_m) - \tau_{D,\theta}\dot{\theta}_m\right] + \frac{\hat{c}\,T_{k-1}}{\ell} + I_{k-1} \tag{3.10}$$

with $|\Delta T| \le T_{avail} - T_i'$. It has no reset action, so the term I holds only the value left by a manual transfer or by the trim of an airborne start, and it is multiplied by $\mathrm{e}^{-\Delta t_c/1.5\,\mathrm{s}}$ every step so that the transfer bias bleeds away in automatic. The allocation gives $T_{R,L} = (T \pm \Delta T)/2$. Because the total thrust is limited to what the differential leaves, each group stays between idle and the available thrust, and a group at its limit costs altitude, never attitude. Section 1.4.7 sets this priority beside other allocation rules.

3.4.4 What Each Mode Knows

Fixed tuning believes the design point, $\hat{m} = m_{nom} = 480\,\mathrm{kg}$, $\hat{c} = 0$ and $\hat{I} = I_{nom} = I_e + (m_{nom} - m_e)r_p^2 = 707.2\,\mathrm{kg\,m^2}$, and keeps its gains. The payload-aware mode reads the load sheet written when boarding closes and computes $\hat{m}$, $\hat{c}$ and $\hat{I}$ from Equation (3.3), with $\hat{m}$ and $\hat{I}$ multiplied by $1 + e_m$ to represent an error in the declared weights. It schedules the gains as $K^e_{c,z} = K_{c,z}\hat{m}/m_{nom}$ and $K^e_{c,\theta} = K_{c,\theta}\hat{I}/I_{nom}$. The third mode adds the observer's estimates $\hat{d}_x$ and $\hat{d}_z$ to the two feedforward terms. The observer runs in every mode, so its estimate can be watched, but only the third mode uses it.

3.4.5 The Disturbance Observer

Each translational axis has a generalized-momentum observer, the observer of Section 1.4.3 applied to the two axes, as shown in Equation (3.11).

$$\begin{aligned} \xi_{k+1} &= \xi_k + \left(\hat{F}_{T,k} + \hat{d}_k\right)\Delta t_c \\ \hat{d}_{k+1} &= K\left(\hat{m}v_{m,k} - \xi_{k+1}\right) \\ \hat{\mathbf{F}}_T &= \hat{T}\left(-\sin\theta_m,\ \cos\theta_m\right) - \hat{m}g\,\mathbf{e}_z \end{aligned} \tag{3.11}$$

where $\hat{T}$ is the sum of the controller's own group commands passed through an exact discrete model of the 0.12 s lag, and $K = 1.5\,\mathrm{rad/s}$ is the bandwidth. For the linear model with the thrust known exactly, $\hat{d} = Kd/(s+K)$, so a step force is estimated with a time constant of $1/K = 0.67\,\mathrm{s}$. The observer sees drag, pushes and any weight the load sheet missed as one force, because a mass error $\Delta m = m - \hat{m}$ appears as $-\Delta m(g + \ddot{z})$. On the ground it is reset to zero. The vessel observer of Chapter 6 carries wind and current the same way, as a slowly varying force that it feeds forward.

Table 3.4 Controller, Reference and Supervisor Settings as Shipped

Symbol	Meaning	Value
$K_{c,z}$, $\tau_{I,z}$, $\tau_{D,z}$	altitude PID at m_{nom}	700 N/m, 6, 1.6 s
$K_{c,x}$, $\tau_{I,x}$, $\tau_{D,x}$	position PID	3.5 °/m, 12, 2.0 s
$K_{c,\theta}$, $\tau_{D,\theta}$	tilt PD at I_{nom}	120 N/°, 0.35 s
m_{nom}	design mass of the fixed tuning	480 kg
K	observer bandwidth	1.5 rad/s
v_{max}	horizontal speed limit	10 m/s
a_{max}, j_{max}	acceleration and jerk limits	1.2 m/s^2, 1.0 m/s^3
—	climb, descent, final, flare speeds	3, 2.5, 1.2, 0.35 m/s
μ_{min}	dispatch thrust margin	25 %
—	cruise band	10 m over roofs, ±6 m
—	shaft over each pad	±8 m
—	gate height and window	8 m, 1.0 m, 0.5 m/s
—	gate dwell and longest wait	1, 6 s
—	final go-around, flare height	1.5 m, 3 m
—	go-around wait, divert	3 s, at the third
—	saturation divert	2.5 s, climb or cruise

HISTORY Estimating the Load Torque (Early 1980s)

The disturbance observer was put forward in motion control in the early 1980s, to improve the torque and speed control of a motor by estimating its load torque. Related estimators followed, among them the extended state observer of the 1990s, the core of active disturbance rejection control. All of them lump the external disturbance and the uncertainty of the model together and estimate the total from measured variables, so that the controller can cancel it. Over four decades the idea has reached robot manipulators, motion servos, magnetic levitation, power converters and disk drives. In aeronautics and astronautics its purpose is to cope with disturbance forces and torques caused by external winds, unmodeled dynamics or uncertain aerodynamics [2]. The third controller mode of this chapter applies the idea to the force on each translational axis.

3.4.6 Reference and Supervisor

Each axis follows an online jerk-limited reference, which moves toward its target with the acceleration changing no faster than j_{max} and capped at a_{max} and the speed at the phase limit, and begins braking the moment the distance left equals the exact jerk-limited stopping distance, checking the next substep so that it brakes early rather than overrun that point. It re-plans from its present state whenever the target moves, first stopping any motion away from a new target, which is what a go-around or a changed setpoint needs, and online generation of this kind is established in robotics [19]. Section 1.4.4 gives the durations and peaks of jerk-limited profiles. While it turns back toward a target it builds only the acceleration that the stop to follow leaves room for, and a stop that comes too late, from a moving start, re-plans from where it crosses the target instead of hunting about it. A sweep of 78 000 moves over the slider ranges, from rest and from moving starts, settles every one. It takes ten substeps per control step, and a linear third-order capture filter removes the last few centimeters.

The supervisor boards the load at each stop and then applies the dispatch rule, departing only if the thrust margin $\mu = 2T_{avail}/(\hat{m}g) - 1$ is at least $\mu_{min} = 25\,\%$ and holding on the pad otherwise. After a 1.5 s spool-up it climbs to the leg's cruise band, cruises and descends to an approach gate 8 m above the pad. It continues below the gate only once the vehicle has been within 1.0 m of the pad center and slower than 0.5 m/s sideways for 1 s, and it waits up to 6 s for that. The final descent aims 1 m into the roof at 1.2 m/s, slowing to 0.35 m/s below 3 m, and goes around if the vehicle drifts 1.5 m off center. A go-around climbs back to the band, waits 3 s and approaches again, and the third go-around at one pad diverts the vehicle back to its departure pad. A diverted approach that fails the same way ends in an emergency landing at that pad without the gate test, and the route is marked aborted. Thrust at the available limit for 2.5 s in the climb or the cruise also diverts. In the air the margin is recomputed from $\hat{m}g - \hat{d}_z$ filtered with 1.5 s, for display only. Table 3.4 gathers every setting.

Three models are distinguished throughout. What is simulated is Equations (3.1) to (3.6) together with the wind of Table 3.2. The controller's internal model is algebraic, the believed mass, center of mass and inertia, the hover force balance of the two feedforward terms, and the observer's translational momentum balance with the rotor lag. The analysis model of Section 3.5 is the linearization about level hover, used to interpret the results and present nowhere in the code.

3.5 What a Loop Must Be Told

3.5.1 Linearized Loops and Their Separation

About level hover with $c = 0$ and $T = mg$, Equation (3.2), the thrust lag and the tilt law give the closed tilt loop of Equation (3.12).

$$\frac{\theta}{\theta_{SP}} = \frac{\kappa}{\tau_T I s^3 + I s^2 + \kappa \tau_{D,\theta} s + \kappa}, \qquad \kappa = \ell K_{c,\theta} \tag{3.12}$$

With $I_{nom} = 707.2\,\mathrm{kg\,m^2}$ and $\kappa = 13.75\,\mathrm{kN\,m/rad}$, and without the thrust lag, the loop would have a natural frequency of 4.41 rad/s and a damping ratio of 0.77. The 0.12 s lag turns this into a pair at 6.27 rad/s damped at only 0.34, plus a real root at $4.12\,\mathrm{s^{-1}}$. The linear model predicts 0.417 s to 63 % of a step with 9.4 % overshoot, and the full simulation with the navigation noise switched off gives 0.413 s and 10.1 % for a 5° step in the tilt command, timed from the controller execution that applies it.

For the position loop with an ideal inner loop, $\ddot{x} = -g\theta$ and the PID of Equation (3.9) give the characteristic equation $s^3 + gK_{c,x}\tau_{D,x}s^2 + gK_{c,x}s + gK_{c,x}/\tau_{I,x} = 0$, a pair at 0.70 rad/s with a damping ratio of 0.79 and a real root at $0.10\,\mathrm{s^{-1}}$. Closing both loops with Equation (3.12) inside puts the tilt pair at 5.84 rad/s with a damping ratio of 0.37, the position pair at 1.07 rad/s with a damping ratio of 0.88, and real roots at 2.01 and $0.103\,\mathrm{s^{-1}}$, a separation of 5.5 between the two oscillatory modes, the separation a cascade needs by Section 1.4.3. The rocket of Volume 3, Chapter 10 has the same structure, a position loop that owns no actuator and commands an attitude.

With a correct feedforward the altitude error obeys $m\ddot{e} = -K^e_{c,z}[e + \tau_{D,z}\dot{e} + \int e\,dt/\tau_{I,z}]$ through the thrust lag. Scheduling $K^e_{c,z} \propto \hat{m}$ keeps the roots at 0.27, 0.58, 2.71 and $4.78\,\mathrm{s^{-1}}$ for every load. Fixed gains at 620 kg give a pair at 0.99 rad/s with a damping ratio of 0.98 and real roots at 0.26 and $6.12\,\mathrm{s^{-1}}$, still well damped, so the heavy leg is not a gain problem. Setting the fixed tuning's design mass to 620 kg, which corrects its feedforward and leaves its gain at 700 N/m, holds the heavy climb of case 3 within 0.05 m, and the payload-aware mode with its gain held at the same effective 700 N/m ($K_{c,z} = 540\,\mathrm{N/m}$, scheduled to 698 N/m at 620 kg) gives 0.05 m as well, against 1.78 m for the shipped fixed tuning. The feedforward does the work.

🕮 **Pause.** Try Exercise 3.3 before reading on.

3.5.2 The Error Budget

A loop that is not told about a force F can oppose it only through its error, as Equation (1.4) states for the cart of Chapter 1. The proportional term alone would need $e_P = F/K_P$, where K_P is the force per unit error, and the integral then removes that error over a time of order τ_I. For the altitude loop $e_P = \Delta mg/K^e_{c,z}$, for the position loop $e_P = \arctan(F/mg)/K_{c,x}$, and for the tilt loop, which has no integral, the error is permanent at $e_P = mgc/(\ell K^e_{c,\theta})$. The predictions and the measured errors are compared in Table 3.5.

🕮 **Pause.** Try Exercise 3.1 before reading on.

The tilt loop matches its budget to 0.01°, because nothing else can supply the moment. The loops with integral action reach 90 to 102 % of their proportional-only error, because the integral is slow against the rise of the force. The one exception, the C to D climb of the graded route at 108 %, climbs out of the jet over pad C. The vertical drag of Equation (3.4) grows with the whole airspeed, so climbing through an 11 m/s crossflow adds a downward force of about 350 N that the budget

Table 3.5 Error a Loop Needs to Supply a Force It Was Not Told About. e_P Is the Proportional-Only Value and the Last Column the Measured Peak or Standing Error

Case, mode	Unknown force	e_P	Measured
3, fixed, heavy climb	$\Delta mg = 1.37\,\text{kN}$	1.96 m	1.78 m
4, aware, $e_m = -20\,\%$	1.22 kN	1.68 m	1.54 m
11, fixed, B → C climb	1.18 kN	1.68 m	1.51 m
11, fixed, C → D climb	1.37 kN	1.96 m	2.11 m
6, aware, crosswind	345 N, 4.1°	1.17 m	1.19 m
5, fixed, steady hover	$mgc/\ell = 294\,\text{N}$	2.45°	2.46°

Table 3.6 Scenarios. Wind Is the Background Speed U_0 at 50 m

#	Set-up	Lesson
1	hover 10 m over A, push at 8 s	thrust equals weight, tilt to move
2	A → B, a_{max} 2.5, j_{max} 4, v_{max} 12	comfort, time and energy
3	A → B → C, fixed, heavy second leg	feedforward from a known load
4	B → C, load sheet 20 % light	an estimate repairs the sheet
5	A → B, 40 kg bag at $x_b = -1.5$	trim for the center of mass
6	hover over the street, wind 0 to 8	a steady force and its estimate
7	B → C, wind 9	hovering in the gap's jet
8	C → D, wind 11, turbulence 1.3	an integral that remembers
9	B → C, full load, 17 % charge	dispatch margin
10	B → C, wind 6, gust front at the gate	stabilized approach, go-around
11	B → C → D, wind 11, 36 % charge	graded route

leaves out, and even the payload-aware mode, whose weight is right, lags 0.37 m on that climb. An estimate removes the error on the observer's time scale instead, and the push test in hover confirms the first-order behavior. With the navigation noise switched off, the estimate of a 600 N push reaches 63 % after 0.66 s at $K = 1.5\,\text{rad/s}$, 2.00 s at 0.5 and 0.33 s at 3 rad/s, each within one 10 ms controller period of $1/K$.

A proportional-integral loop that begins and ends with no standing force must also integrate zero error, so an excursion in one direction has to be repaid in the other. After the 600 N push of case 1 with the noise switched off, the vehicle overshoots its setpoint by 0.31 m and the two lobes of the error have equal areas, 3.72 m s each over the 142 s after the push.

3.6 Scenarios and Measured Results

The ten cases and the graded route are listed in Table 3.6. Every value below was measured on the engine at the default seed with the case as shipped unless a change is named, a named change being made before the run starts (Section 1.6.4), and the principal results are collected in Table 3.7.

3.6.1 Hover and Tilt

Before the push at 8 s, case 1 carries 490 kg on 4.80 kN, the weight to within 0.04 %, holds within 0.07 m sideways, and draws 73.9 kW. The 600 N push moves the vehicle 1.46 m, tilting it up to 9.4° on the way back, and it passes back through the setpoint 5.20 s after the push began. It then overshoots by 0.30 m and is not back within 0.1 m until 28.3 s, which is the repayment of Section 3.5 with noise added. With the position loop in manual at −5° from 15 to 22 s, the vehicle accelerates toward $+x$ at $0.66\,\text{m/s}^2$ between four and seven seconds in, below $g\tan 5^\circ = 0.86\,\text{m/s}^2$ because the drag grows with speed, and it covers 18.1 m and reaches 5.0 m/s. Back in automatic, it is within 0.5 m of the setpoint again 18.4 s later.

Table 3.7 Measured Results. Aware Is the Payload-Aware Mode and Observer the Mode With the Wind Estimate

#	Condition	Result
1	before the push	$T/mg - 1 = -0.04\,\%$, 73.9 kW, $\lvert\Delta x\rvert \leq 0.07$ m
1	600 N, 2 s push	peak 1.46 m, back through at 5.20 s, overshoot 0.30 m, tilt $\leq 9.4°$
1	manual $-5°$, 7 s	0.66 m/s^2, 18.1 m, 5.0 m/s
2	as shipped	2.84, 5.36 m/s^3, 6.95 s at limit, lag 5.7 m, 35.3 s, 0.76 kWh
2	a_{max} 1.2, j_{max} 1.0	1.59, 1.75 m/s^3, no saturation, lag 0.66 m, 43.2 s, 0.92 kWh
3	fixed	climbs 0.09 m and 1.78 m behind, heavy liftoff 1.03 s later than aware
3	aware	heavy climb 0.05 m
4	aware	liftoff 1.17 s after climb start, lag 1.54 m, 0.88 m at 5 s
4	observer	lag 0.83 and 0.02 m at 1 and 3 s, then 0.44 m above, $\hat{d}_z = -1.21$ kN
5	fixed	3.15°, 0.76 m in 4 s, 2.27° at gate
5	aware	0.30°, 0.07 m, 0.08° at gate
6	aware, observer	drift 1.19, 0.36 m, back in 19.6, 8.3 s
7	aware, observer	wander 1.90, 0.41 m, gate 1.39, 0.24 m, land 0.71, 0.18 m
8	aware, observer	worst 0.98, 0.33 m, land 0.19, 0.07 m
9	full load	3.72 kN per group, $\mu = 22.4\,\%$, holds
9	one passenger off	$\mu = 45.9\,\%$, departs 3.4 s later
10	aware	go-around 43.2 s, lands 73.9 s at 0.62 m, 1.41 kWh
10	observer	go-around 45.0 s, lands 76.6 s at 0.02 m

3.6.2 Comfort and the Reference

Case 2 lets the reference accelerate at 2.5 m/s^2, which is 95 % of the $g\tan 15° = 2.63\,\mathrm{m/s^2}$ that the tilt limit allows, so the corrections the loop adds push it onto the limit. The tilt command sits there for 6.95 s, the vehicle trails its reference by up to 5.7 m, and the comfort signals peak at 2.84 m/s^2 and 5.36 m/s^3 on the kinematic acceleration filtered at 0.2 s, while liftoff to touchdown takes 35.3 s. At 1.2 m/s^2 and 1.0 m/s^3 the hop takes 43.2 s, peaks at 1.59 m/s^2 and 1.75 m/s^3, never saturates, trails by 0.66 m and lands 0.10 m from the center instead of 0.25 m. Because hover power is paid by the second, the gentle hop uses 0.92 kWh against 0.76 kWh.

3.6.3 The Load and the Load Sheet

Case 3 flies fixed tuning designed for 480 kg, first with one 80 kg passenger and then with two 100 kg passengers and 20 kg of bags. Measured over the last half second of each gate dwell, the hover thrust rises from 4.65 to 6.03 kN, a ratio of 1.298 against a mass ratio of 1.292. The light climb stays within 0.09 m of its reference and the heavy one falls 1.78 m behind, because 1.37 kN of weight has to come from the error, and the heavy liftoff comes 1.03 s later than with the payload-aware mode, whose heavy climb stays within 0.05 m. The climber of Volume 3, Chapter 12 carries a gravity feedforward for the same reason, because its effective weight changes along its route. The error of a wrong feedforward has either sign. With the design mass set to 620 kg, the light climb rises 1.74 m above its reference and lifts off 0.22 s before the climb begins, during the spool-up.

Case 4 gives the payload-aware mode a load sheet 20 % light, 124 kg on a 620 kg vehicle and 1.22 kN of weight it does not know. The spool-up stops short of the weight, so the vehicle sits on the pad until the altitude error has built enough thrust, and it lifts off 1.17 s after the climb begins instead of 0.23 s, falls 1.54 m behind, and is still 0.88 m behind five seconds after liftoff. The observer cannot see a weight the pad carries, so its liftoff is just as late. Once airborne it sees the missing weight as a downward force, and the lag is 0.83 m one second after liftoff and 0.02 m after three. The estimate then arrives on top of a proportional term still correcting the lag, and the vehicle overshoots to 0.44 m above its reference 4.5 s after liftoff. Over the last two seconds before touchdown the estimate averages −1.21 kN of the 1.22 kN missing.

🕮 **Pause.** Try Exercise 3.2 before reading on.

3.6.4 Balance

Case 5 puts a 40 kg suitcase at the left end of the tray. Fixed tuning, which does not know, tilts 3.15° and wanders 0.76 m in the four seconds after liftoff and flies the approach gate with 2.27° between the tilt it is asked for and the tilt it has, while its differential thrust averages −286 N against the −294 N trim. In the steady hover of case 1 with the same load, chosen before Run so that the airborne start carries it, the standing error averages 2.46° from 20 to 40 s, the budget of Table 3.5. The payload-aware mode feeds the trim forward and holds 0.30° and 0.07 m after liftoff and 0.08° at the gate, and in a steady hover its differential thrust matches mgc/ℓ to within 0.1 %.

3.6.5 Steady Wind, the Gap and the Wake

Case 6 steps the background wind from calm to 8 m/s while the vehicle hovers 60 m over the street, where the wind averages 8.36 m/s from 50 to 90 s and the drag settles at 345 N with the vehicle leaning 4.1°. Without an estimate it drifts 1.19 m, 7.2 s after the step, and needs 19.6 s to pass back within 0.1 m, because only the integral can build the standing tilt. With the observer, whose steady estimate is within 1.3 % of the drag, it drifts 0.36 m and is back in 8.3 s.

> **TRY IT IN THE APP 6 · Steady Crosswind**
>
> Load the case and press Run. Watch the Horizontal air force trend: as the wind steps from calm to 8 m/s, the drag settles near 345 N and the dashed estimate follows it. Then watch the Horizontal position trend. With the payload-aware controller the vehicle drifts 1.19 m downwind and needs 19.6 s to pass back within 0.1 m.

The observer bandwidth is a trade. Across K from 0.3 to 4 rad/s in case 6, the peak drift falls from 0.67 to 0.24 m, while the standard deviation of the tilt command in the steady wind rises from 1.08 to 1.35°, against 1.01° without an estimate, both over 50 to 90 s. The estimate error is smallest, 34 to 35 N standard deviation, at 1.0 to 1.5 rad/s, because slower estimates lag the turbulent drag and faster ones pass the velocity noise through the gain $K\hat{m}$, and the shipped 1.5 rad/s sits on that flat minimum.

Case 7 flies B to C in a 9 m/s wind. Cruising through the gap the vehicle moves with the 12.9 m/s channeled flow and averages only 4.5 m/s of airspeed and 112 N of drag, but over pad C, where the jet leaves the gap, it must hover in about 11 m/s against a drag that averages 571 N over the gate dwell. Without an estimate the position loop needs an error and then its integral, and the vehicle wanders up to 1.90 m off the pad line on the way down, 1.39 m at the gate, and lands 0.71 m from the center, outside the 0.5 m target. With the observer it is 0.24 m at the gate and lands 0.18 m from the center. The landing depends on the seed. Over seeds 1 to 10 the payload-aware mode lands 0.21 to 0.91 m from the center, outside the target five times, and wanders 1.12 to 2.04 m, while the observer lands within 0.27 m every time and wanders 0.27 to 0.83 m.

Case 8 crosses the 58 m tower in an 11 m/s wind with the turbulence scaled by 1.3 and descends to the clinic in its wake. At the top of the descent, 68 m up in the tower's shear layer, the mean wind of the field is 5.7 m/s with turbulence of 4.4 m/s standard deviation, and the payload-aware mode's position integral holds 3.3° of lean. By 55 m the mean wind has fallen to 2.1 m/s, yet the integral still holds 3.3° there and 1.5° at the gate, so the vehicle leans on air that is no longer there. It is 0.98 m off the pad line at the top and up to 0.73 m below 60 m, and lands 0.19 m from the center. The observer drops its feedforward as the force goes, keeps the error to 0.33 m and lands 0.07 m from the center. Knowing the force also means asking for all of it, and while braking over the clinic the observer mode holds the tilt command at 15° for 1.55 s, which the payload-aware mode never does. The peak jerks of the two runs, 8.4 and 9.2 m/s^3, both come at the cruise band over the clinic, where the vehicle flies through the tower's shear layer.

3.6.6 Battery and the Approach

Case 9 boards a full load, two 100 kg passengers and 20 kg of bags, on a battery at 17 %. The power limit has sagged to 74 % of its rating, a group can deliver 3.72 kN, and against the 6.08 kN weight

Table 3.8 Graded Route at the Default Seed

	Fixed	Aware	Observer
climb lag B → C, C → D (m)	1.51, 2.11	0.10, 0.37	0.07, 0.15
pad C error (m), sink (m/s)	0.99, 0.45	0.98, 0.45	0.23, 0.41
pad D error (m), sink (m/s)	0.05, 0.33	0.05, 0.34	0.00, 0.38
tilt limit (s)	2.2	2.8	2.9
battery left (%)	22.5	22.4	22.6
route complete (s)	133.9	134.0	132.9
verdict	fail	fail	pass

the margin is 22.4 % when the hold begins, below the 25 % rule, so the supervisor holds on the pad. Taking one passenger off raises the dispatch margin to 45.9 %, and the vehicle departs 3.4 s after the change, the time for the passenger to walk off and the check to pass. What the supervisor knows matters as much as what the loops know. The fixed mode judges the margin on its 480 kg design mass, reads 58 %, and takes off with the full load, and because the true margin of 22 % still covers the flight it lands on C 0.07 m from the center without saturating, while the in-flight margin reads 23 % once the observer has seen the weight and falls to 14 % on the final approach as the battery drains. The rule protects a reserve, not a hard limit.

Case 10 flies B to C in a 6 m/s wind and times a gust front of 8 m/s over 12 s to the vehicle's arrival at the gate, at 38.8 s. The vehicle clears the gate check before the gust has built, starts down at 39.8 s, and is pushed 1.5 m off center at 43.2 s, a go-around. With the gust at its peak the tilt command sits on its limit for 3.0 s while the vehicle climbs back. It starts the second approach at 55.0 s with the gust gone and lands at 73.9 s, 0.62 m from the center, in the jet of case 7 again. The leg costs 1.41 kWh, against 0.94 kWh in a model run with the gust removed, which the case itself does not offer. With the observer the go-around comes at 45.0 s and the second approach lands 0.02 m from the center at 76.6 s. Three 600 N pushes in a row at the start of every final approach show the divert rule. The vehicle goes around at 41.4, 71.8 and 94.2 s, diverts to B at 105.7 s, goes around three more times there and makes an emergency landing on B at 195.2 s, which marks the route aborted.

3.6.7 The Graded Route

The graded route flies hotel, offices and clinic in an 11 m/s wind, boarding two 90 kg passengers and 20 kg of bags at B, a tray position of −1.2 m, and a swap to two 100 kg passengers at C, on a battery that starts at 36 %. It is met with no crash, off-pad touchdown or corridor excursion, every touchdown within 0.5 m of the pad center at under 0.5 m/s, and at least 20 % battery at the end. The three modes are compared in Table 3.8, and Figure 3.5 draws its distances as bars.

Shipped with fixed tuning, the route completes with no violation and 22.5 % left, but the heavy climbs fall 1.51 and 2.11 m behind and the vehicle lands 0.99 m from the center of pad C in the gap's jet, so the scorecard reads one touchdown of two in specification. The payload-aware mode repairs the climbs, to 0.10 and 0.37 m, but not the jet, landing 0.98 m off, and its worst altitude error on the first leg, 0.45 m, comes in the descent into the jet rather than in the climb. The observer repairs both, seeing the load the fixed tuning does not know as a vertical force and the jet as a horizontal one. It lands 0.23 m from the center of C and within 0.01 m of D's, at 0.41 and 0.38 m/s, the route is complete at 132.9 s, and the run ends 1.5 s later at 134.4 s with 22.6 % left. The pass leaves 0.27 m of the touchdown target, 0.09 m/s of the sink limit and 2.6 percentage points of battery to spare. Revealing the suggested setting and then pressing Reset gives the same run. The route is judged on the case as given, so a changed model parameter or disturbance setting voids the attempt. All three spend 2 to 3 s at the tilt limit descending into the jet, which the shaded band marks for the observer run in Figure 3.4.

The verdicts hold beyond the default seed. Over seeds 1 to 20 the observer passes all 20, with its worst touchdown 0.37 m from a pad center and its worst sink 0.491 m/s, while the fixed and

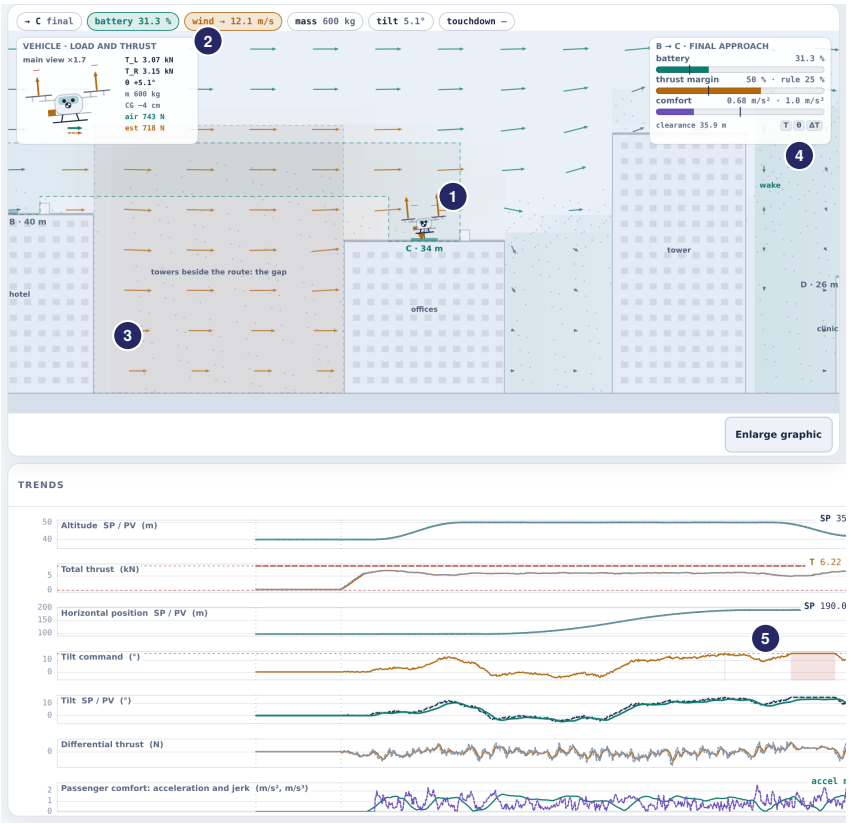

Figure 3.4 The graded route, Challenge · the afternoon route, with the suggested controller that adds the wind estimate, paused at $t = 47.00\,\mathrm{s}$ on the final approach to pad C, cropped to the plant view and the trends. The vehicle (1), drawn 1.7 times its size, carries 600 kg with its center of mass 4 cm left of the centroid, hovers 1.9 m above the offices roof and leans 5.1° into a 12.1 m/s wind, the jet that the amber arrows trace through the gap between the towers beside the route (3). The dashed outline is the corridor of the leg from B, the triangle on the roof predicts a touchdown 0.6 m upwind of the pad center, the two ticks above it mark the approach gate, and the shaded, stippled zone labeled wake behind the tower holds weak, sinking flow over pad D (4). The close-up plate (2) compares the true horizontal air force of 743 N with the observer's 718 N, and the mission plate shows 31.3 % battery and a 50 % thrust margin against the 25 % rule. The trends hold 47 of the 60 s window, from boarding and spool-up on pad B through the climb to the 50 m cruise band, the move from 96 to 190 m and the descent, with the tilt command shaded where it sat at the 15° limit from 34.9 to 37.8 s (5), the tilt following its command, and the differential thrust and comfort signals below. The crop ends the trends a few seconds before the pause and cuts their readouts.

payload-aware modes pass 5 each, and every one of their 30 failures occurs at pad C. The margin on the touchdown speed is thin, and the observer's worst sink leaves 0.009 m/s to spare.

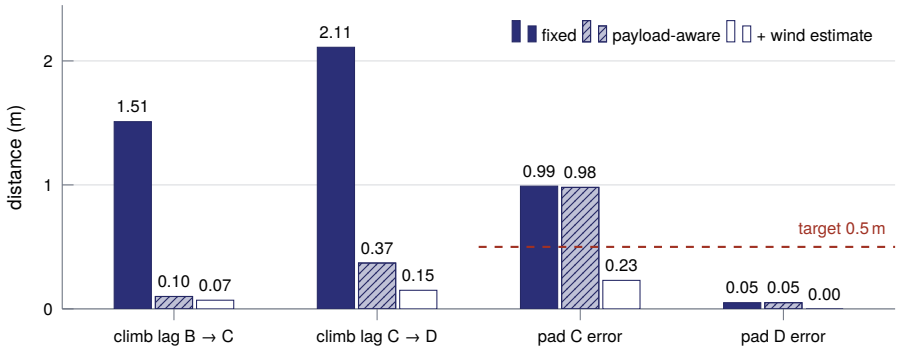

Figure 3.5 The graded route of Table 3.8 at the default seed, drawn as bars: the lag on the two heavy climbs and the touchdown error at pads C and D for the three controller modes. Notice that the load sheet of the payload-aware mode removes most of the climb lag but none of the miss at pad C, which sits in the gap's jet. Only the wind estimate brings that touchdown inside the 0.5 m target, and only it passes the route.

> ✦ DID YOU KNOW? **Knowing the Load Does Not Land in the Jet**
>
> From the simulator: on the graded route the payload-aware mode cuts the heavy climb lags from 1.51 and 2.11 m to 0.10 and 0.37 m, yet it lands 0.98 m from the center of pad C, no closer than fixed tuning at 0.99 m. Only the mode with the wind estimate lands within the target, at 0.23 m (Table 3.8).

3.7 Technology in Practice

The vehicle of Section 3.3 is read by one navigation system and moved by two rotor groups. A real multirotor fuses inertial and satellite measurements, drives its rotors with electric motors and flies under rules for its weight and balance. Figure 3.6 draws the chapter's three loops with such devices in place.

What Measures It

The navigation system NAV of Figure 3.1 stands for an inertial measurement unit, a satellite receiver and a filter that fuses them, as Section 1.4.6 describes. On a research quadrotor of 2007 an extended Kalman filter combined the satellite and inertial measurements into the full state [8]. In this chapter every measured state carries colored noise with no latency, bias or dropout, 0.05 m on position and 0.1° on tilt.

> TECHNOLOGY **Sensor: Inertial and Satellite Navigation**
>
> A multirotor takes its tilt and rates from an inertial measurement unit of three gyroscopes, three accelerometers and three magnetometers. On a research quadrotor of 2007 the attitude estimate was accurate to 2° as long as no acceleration was sustained, a sonic ranger gave the height above ground to 3 to 5 cm, and carrier-phase differential satellite navigation gave the position to 1 to 2 cm [8]. Between buildings the satellites are blocked or reach the receiver only by reflection. Such signals carry range errors of tens or even hundreds of meters, and the reported position can be wrong by the width of the street canyon [20]. The colored noise of a few centimeters in this chapter stands for none of this.

What Moves It

Each rotor group of the model lumps the four rotors on one side into one thrust, which follows its command through a lag of 0.12 s and never exceeds what the battery allows. On a small multirotor that thrust comes from brushless motors whose torque turns the rotors, and the voltage and current of each motor set its speed [8]. The simulator model of a six-passenger concept added each motor's

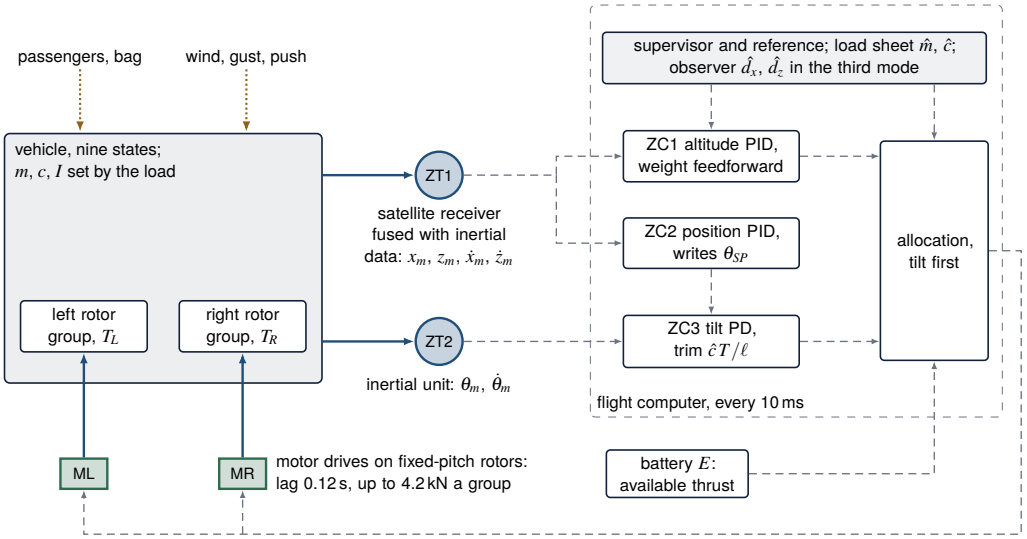

Figure 3.6 The three loops of the chapter with real devices in place of the schematic's symbols. ZT1 reads position and height from a satellite receiver fused with inertial data, ZT2 the tilt and its rate from the inertial unit, and one flight computer (dashed frame) runs the altitude, position and tilt loops (ZC1 to ZC3) every 10 ms. ML and MR are the motor drives of the two rotor groups, held within the thrust the battery allows. Notice that the position loop owns no motor: it writes the tilt setpoint, and only the allocation reaches ML and MR.

armature circuit and the inertia of its shaft and rotor to the flight model, because the power demanded by maneuvering depends on them [9].

> **TECHNOLOGY** **Final Element: Electric Motors on the Rotors**
>
> A small multirotor needs no mechanical linkage at its rotors. Brushless direct-current motors with electronic speed controllers turn fixed-pitch rotors, and the vehicle is controlled by varying motor speed alone [8]. At air-taxi size the choice is open. Rotor speed control gives a lower design weight, but quadrotors sized for it were found not stabilizable with adequate closed-loop performance within the mechanical limits of their drive systems, and a later solution left the motor current unbounded [9]. A six-passenger helicopter concept for air taxis chose blade pitch control for control bandwidth and for autorotation [21]. The chapter folds motors, rotors and either kind of control into one lag of 0.12 s per rotor group.

Where the Controller Runs

The research quadrotor split its computing into two levels. A microcontroller ran the real-time control loops and wrote the motor commands, while a single-board computer above it ran planning, estimation and position control, with satellite fixes arriving at 10 Hz [8]. The flight computer of this chapter runs all three loops every 10 ms on navigation states delivered at 100 Hz. Its altitude and position loops are velocity-form PID blocks (Section 1.6.2), the tilt loop a PD block with the trim fed forward, and the supervisor a layer above them that sets the targets. The six-passenger concept flown in a motion simulator nested a fast inner attitude loop under an outer loop that shaped the commands [9], the same order of loops.

What Is Modernizing It

Three developments bear on the chapter's question of what the controller knows. The first is onboard weight and balance systems, which measure the actual weight and center of gravity instead of

assuming them [22], the measured counterpart of the load sheet that the payload-aware mode reads. The second is statistical models of satellite errors measured in real street canyons, the groundwork for navigation filters that know how far to trust a fix [20]. The third is predictive control: the six-passenger concept's flight control was a discrete-time model predictive controller whose cost weights set its disturbance rejection [9]. The observer of Section 3.4 belongs to a family of methods that estimate a lumped disturbance and cancel it, applied across industrial sectors for four decades [2].

IN PRACTICE **Where the Load Sheet Gets Its Weights**

United States guidance on weight and balance control names three ways to weigh passengers and bags: standard average weights, averages from the operator's own survey, and actual weights. Standard averages add 5 pounds of clothing in summer and 10 in winter to body weights from health surveys. A small sample strays further from the average, and the guidance limits standard averages to multiengine turbine aircraft with five or more passenger seats. An onboard system that measures the actual weight and center of gravity removes the margins kept for assumed passenger weights, though not the load manifest [22]. A two-seat taxi has far fewer seats, and case 4 shows what a sheet 20 % light costs it.

3.8 Limitations

The model is not a flight model of any aircraft, and its numbers carry no claim about a real air taxi's performance, energy or landing precision. Planar flight removes roll, yaw and their coupling, the rotors lack inflow, flapping and H-force, and the wind field is a set of closed-form features rather than a solution of the flow [6, 12]. The payload-aware modes read a load sheet instead of estimating mass online, and the observer estimates forces, not the wind itself. The dispatch rule is applied on the pad and through the saturation divert, not at the approach gate. The touchdown-speed criterion of the graded route is passed narrowly. Where each assumption breaks is stated in Section 3.3.

OPEN CHALLENGE **Holding Height With Rotor Speed Alone**

The first piloted handling-qualities study of an air-taxi quadrotor in a motion simulator, flown in 2021, compared blade pitch control with rotor speed control. Pilot ratings varied strongly with the standard applied, with the tuning and with turbulence, and turbulence amplified the challenge. The authors note that minimal data exist for large rotor speed-controlled vehicles outside of software, and they set out to measure what disturbance rejection costs in motor power [9]. Progress would be flight data and criteria that fix how much motor power a passenger multirotor needs to hold its height in a gust.

3.9 Conclusion

This chapter studied a drone taxi simulator whose experimental variable is what the flight controller knows. Its nine-state planar multirotor has rotor arms about the true center of mass, a sagging battery limit and a closed-form wind field of shelter, a gap jet, roof-edge flows and wakes, and three loops close at 100 Hz under a supervisor. Fixed tuning holds a light climb within 0.09 m and lets a heavy one fall 1.78 m behind, and correcting the feedforward alone brings it to 0.05 m whether or not the gain is scheduled. A suitcase at the end of the tray costs 2.46° of standing tilt error against a predicted 2.45°. An observer with a 0.67 s time constant cuts the drift after a wind step from 1.19 to 0.36 m and the landing error at the gap's jet from 0.71 to 0.18 m, and on the graded route it alone meets the target, on all 20 seeds tested. A battery at 17 % leaves a full load a margin of 22.4 %, and the supervisor holds on the pad. The field is illustrative, flight is planar and the rotors have no inflow dynamics. A fuller model would add an online mass estimate, out-of-plane wind, rotor inflow and a predictive controller that plans the approach within the tilt and thrust limits.

The drone taxi is a study in giving a feedback loop the knowledge it cannot learn fast enough from its own error.

KEY POINTS

- A loop not told its load carries the extra weight only through its error, so fixed tuning lets a heavy climb fall 1.78 m behind.
- Correcting the weight feedforward alone brings that climb to 0.05 m whether or not the gain is scheduled: the feedforward does the work.
- The tilt loop has no integral, so a suitcase at the end of the tray costs 2.46° of standing tilt error against a predicted 2.45°.
- An observer with a 0.67 s time constant cuts the drift after a wind step from 1.19 to 0.36 m, and on the graded route it alone meets the target, on all 20 seeds tested.

Where else this applies. Chapter 6 carries wind and current as a slowly varying force that its observer feeds forward, and the climber of Volume 3, Chapter 12 feeds forward a weight that changes along its route.

References

[1] K. Ohishi, M. Nakao, K. Ohnishi, and K. Miyachi, "Microprocessor-controlled DC motor for load-insensitive position servo system," *IEEE Transactions on Industrial Electronics*, vol. IE-34, no. 1, pp. 44–49, 1987.

[2] W.-H. Chen, J. Yang, L. Guo, and S. Li, "Disturbance-observer-based control and related methods—an overview," *IEEE Transactions on Industrial Electronics*, vol. 63, no. 2, pp. 1083–1095, 2016.

[3] W. J. Rugh and J. S. Shamma, "Research on gain scheduling," *Automatica*, vol. 36, no. 10, pp. 1401–1425, 2000.

[4] R. Mahony, V. Kumar, and P. Corke, "Multirotor aerial vehicles: Modeling, estimation, and control of quadrotor," *IEEE Robotics & Automation Magazine*, vol. 19, no. 3, pp. 20–32, 2012.

[5] A. Bacchini and E. Cestino, "Electric VTOL configurations comparison," *Aerospace*, vol. 6, no. 3, p. 26, 2019.

[6] I. P. Castro and A. G. Robins, "The flow around a surface-mounted cube in uniform and turbulent streams," *Journal of Fluid Mechanics*, vol. 79, no. 2, pp. 307–335, 1977.

[7] D. Galway, J. Etele, and G. Fusina, "Modeling of urban wind field effects on unmanned rotorcraft flight," *Journal of Aircraft*, vol. 48, no. 5, pp. 1613–1620, 2011.

[8] G. Hoffmann, H. Huang, S. Waslander, and C. Tomlin, "Quadrotor helicopter flight dynamics and control: Theory and experiment," in *AIAA Guidance, Navigation and Control Conference and Exhibit*. American Institute of Aeronautics and Astronautics, 2007.

[9] S. Withrow-Maser, J. Aires, A. Ruan, C. Malpica, and S. Schuet, "Evaluation of heave disturbance rejection and control response criteria on the handling qualities evaluation of urban air mobility (UAM) eVTOL quadrotors using the vertical motion simulator," in *Vertical Flight Society Aeromechanics for Advanced Vertical Flight Technical Meeting*, San Jose, CA, 2022, nASA Technical Reports Server 20210026423, https://ntrs.nasa.gov/citations/20210026423.

[10] B. W. Bequette, "Process control practice and education: Past, present and future," *Computers & Chemical Engineering*, vol. 128, pp. 538–556, 2019.

[11] J. A. Rossiter, B. Pasik-Duncan, S. Dormido, L. Vlacic, B. Jones, and R. Murray, "A survey of good practice in control education," *European Journal of Engineering Education*, vol. 43, no. 6, pp. 801–823, 2018.

[12] W. Johnson, *Rotorcraft Aeromechanics*. Cambridge University Press, 2013.

[13] R. W. Prouty and H. C. Curtiss, Jr., "Helicopter control systems: A history," *Journal of Guidance, Control, and Dynamics*, vol. 26, no. 1, pp. 12–18, 2003.

[14] T. Tomic, C. Ott, and S. Haddadin, "External wrench estimation, collision detection, and reflex reaction for flying robots," *IEEE Transactions on Robotics*, vol. 33, no. 6, pp. 1467–1482, 2017.

[15] S. Watkins, J. Burry, A. Mohamed, M. Marino, S. Prudden, A. Fisher, N. Kloet, T. Jakobi, and R. Clothier, "Ten questions concerning the use of drones in urban environments," *Building and Environment*, vol. 167, p. 106458, 2020.

[16] A. Straubinger, R. Rothfeld, M. Shamiyeh, K.-D. Büchter, J. Kaiser, and K. O. Plötner, "An overview of current research and developments in urban air mobility – setting the scene for UAM introduction," *Journal of Air Transport Management*, vol. 87, p. 101852, 2020.

[17] W. L. Fredericks, S. Sripad, G. C. Bower, and V. Viswanathan, "Performance metrics required of next-generation batteries to electrify vertical takeoff and landing (VTOL) aircraft," *ACS Energy Letters*, vol. 3, no. 12, pp. 2989–2994, 2018.

[18] L. L. Hoberock, "A survey of longitudinal acceleration comfort studies in ground transportation vehicles," *Journal of Dynamic Systems, Measurement, and Control*, vol. 99, no. 2, pp. 76–84, 1977.

[19] T. Kroger and F. M. Wahl, "Online trajectory generation: Basic concepts for instantaneous reactions to unforeseen events," *IEEE Transactions on Robotics*, vol. 26, no. 1, pp. 94–111, 2010.

[20] M. Peretic, R. Gilabert, J. Carroll, J. Gutierrez, A. Moore, J. Christie, and E. T. Dill, "Statistical analysis of GNSS multipath errors in urban canyons," in *IEEE/ION Position, Location and Navigation Symposium (PLANS)*, Salt Lake City, UT, 2025, nASA Technical Reports Server 20250000589, https://ntrs.nasa.gov/citations/20250000589.

[21] W. Johnson, C. Silva, and E. Solis, "Concept vehicles for VTOL air taxi operations," in *AHS Technical Conference on Aeromechanics Design for Transformative Vertical Flight*, San Francisco, CA, 2018, nASA Technical Reports Server 20180003381, https://ntrs.nasa.gov/citations/20180003381.

[22] Federal Aviation Administration, "Aircraft weight and balance control," Advisory Circular 120-27F, 6 May 2019, https://www.faa.gov/documentLibrary/media/Advisory_Circular/AC_120-27F.pdf, 2019.

3.10 Exercises

3.1 Concept check · 2 min

In the steady hover of scenario 1 with the 40 kg suitcase of case 5 at the left end of the tray, the fixed tuning holds a tilt error of 2.46° for as long as the vehicle hovers. Why does the error not decay?

(a) The tilt loop has no integral action, so a standing moment needs a standing error
(b) The position loop's integral keeps moving the tilt setpoint away from the tilt
(c) The 0.12 s thrust lag delays every correction beyond the length of the hover
(d) The tilt measurement noise hides an error smaller than its standard deviation

3.2 Concept check · 2 min

In case 4 the load sheet is 20 % light. With the observer added, the vehicle still lifts off 1.17 s after the climb begins, as late as the payload-aware mode without it. Why does the observer not help at liftoff?

(a) Its bandwidth of 1.5 rad/s is too slow to act within 1.17 s
(b) It is reset to zero on the ground, and a weight the pad carries causes no motion to estimate
(c) It estimates only the horizontal force, so a weight error is outside its reach
(d) The dispatch rule holds the vehicle on the pad until the estimate has settled

3.3 Concept check · 2 min

With the fixed tuning, the heavy climb of case 3 falls 1.78 m behind its reference. Which single change brings it within about 0.05 m?

(a) Scheduling the altitude gain on the true mass while keeping the 480 kg weight feedforward
(b) Halving the altitude integral time, so that the integral catches up faster
(c) Setting the design mass to 620 kg, which corrects the weight feedforward at the same gain
(d) Raising the observer bandwidth to 4 rad/s with the fixed controller still selected

3.4 Theory and calculation · 6 min

Case 5 boards one 80 kg passenger and a 40 kg suitcase at $x_b = -1.5$ m on the 400 kg vehicle, with $\ell = 2.0$ m and $g = 9.81\,\mathrm{m/s^2}$.

(a) Use Equation (3.3) to find the mass and the center of mass c. *Checkpoint:* 520 kg and −11.5 cm.

(b) Find the differential thrust $T_R - T_L$ that holds a level hover. *Checkpoint:* −294 N.

(c) The fixed tuning feeds no trim forward and its tilt gain is 120 N/°. Predict the standing tilt error, to two decimals. *Checkpoint:* 2.45°.

3.5 Theory and calculation · 7 min

Use the error budget of Section 3.5 with $g = 9.81\,\mathrm{m/s^2}$.

(a) In case 3 the fixed tuning believes 480 kg and the heavy leg weighs 620 kg, with an altitude gain of 700 N/m. Find the unknown force and the proportional-only error e_P. *Checkpoint:* 1.37 kN and 1.96 m.

(b) In case 4 the payload-aware mode reads a load sheet 20 % light on the same 620 kg and schedules its gain as $K_{c,z}\hat{m}/m_{nom}$ with $K_{c,z} = 700\,\mathrm{N/m}$ and $m_{nom} = 480\,\mathrm{kg}$. Find $\hat{m}$, the unknown force, the scheduled gain and e_P. *Checkpoint:* 496 kg, 1.22 kN, 723 N/m, 1.68 m.

(c) The observer has $K = 1.5\,\mathrm{rad/s}$. How long does its estimate of a step force take to reach 63 % and 95 % of the force? *Checkpoint:* 0.67 s and 2.0 s.

3.6 Theory and calculation · 8 min

Case 9 starts at a charge of $s = 0.17$. The battery has $P_{max} = 190\,\mathrm{kW}$, a sag knee $s_k = 0.35$, power at empty $p_f = 0.5$ and avionics power $P_{av} = 1.5\,\mathrm{kW}$. Each rotor group has $A = 12\,\mathrm{m}^2$ and $FM = 0.6$ in air of $1.225\,\mathrm{kg/m}^3$, and $g = 9.81\,\mathrm{m/s}^2$.

(a) Use Equation (3.6) to find $f(s)$ and the thrust one group can deliver. *Checkpoint:* 0.743 and 3.72 kN.

(b) Find the thrust margin μ for the full load of 620 kg, and the supervisor's decision under the 25 % rule. *Checkpoint:* 22.4 %, hold on the pad.

(c) Find the margin with one 100 kg passenger off, and the margin the fixed mode computes for the full load from its 480 kg design mass. *Checkpoint:* 46.0 % and 58.1 %, each within 0.2 %.

3.7 App challenge · 8 min

Load scenario 3, Light versus full load, with the default seed. A playback speed of 2 suits each run.

(a) Run to 66 s as loaded, with the Fixed controller, and export the CSV file. For each of the two climbs, the rows whose `phase` is `climb`, find the largest value of `z_sp_m` − `z_m`. *Checkpoint:* 0.08 m and 1.75 m, each within 0.03 m.

(b) Press Reset, set Controller to Payload-aware before Run, run to 66 s again and find the same value for the second climb. *Checkpoint:* 0.05 m, within 0.01 m.

(c) Compare part (a) with the proportional-only error of Table 3.5, and say which of the two changes the payload-aware mode makes, the weight feedforward or the scheduled gain, removes the lag. *Checkpoint:* 1.96 m predicted, and the feedforward.

3.8 App challenge · 8 min

Load scenario 6, Steady crosswind, with the default seed. The wind steps from calm to 8 m/s at 5 s. Each run lasts 30 s. Read the `x_m` and `x_sp_m` columns of the exported CSV file.

(a) Run as loaded, with the Payload-aware controller. Find the largest |`x_m` − `x_sp_m`| after 5 s, the time of that row, and the first later row in which the distance is under 0.1 m. *Checkpoint:* 1.19 m at 12.25 s, back at 24.65 s, each time within 0.1 s.

(b) Press Reset, set Controller to + wind estimate before Run, and repeat. *Checkpoint:* 0.36 m at 7.70 s, back at 13.30 s.

(c) Press Reset, also set Observer bandwidth to 4.0 before Run, and read the largest distance again. Name the price of the faster estimate, using the bandwidth sweep of Section 3.6. *Checkpoint:* 0.24 m, within 0.01 m.

3.9 App challenge · 9 min

Load the challenge, the afternoon route, with the default seed, and use a playback speed of 5. Each run ends by itself about 135 s after Run, when the route is complete. Read the Touchdowns in spec and Battery reserve tiles of the scorecard after each run.

(a) Run as loaded, with the Fixed controller. *Checkpoint:* 1 of 2, worst 0.45 m/s and 0.99 m, 22.5 %.

(b) Press Reset, set Controller to Payload-aware before Run, and run again. *Checkpoint:* 1 of 2, worst 0.45 m/s and 0.98 m, 22.4 %.

(c) Press Reset, set Controller to + wind estimate before Run, and run again. Explain why only this run passes. *Checkpoint:* 2 of 2, worst 0.41 m/s and 0.23 m, 22.6 %.

3.10 Application area · 4 min

Case 2 flies the same hop twice. As shipped it takes 35.3 s from liftoff to touchdown and uses 0.76 kWh. With gentle limits of 1.2 m/s^2 and 1.0 m/s^3 it takes 43.2 s and uses 0.92 kWh. The vehicle weighs 490 kg and hovers on 73.9 kW.

(a) Estimate the extra energy of the gentle hop from the hover power and the extra time aloft, and compare it with the measured difference. *Checkpoint:* 0.162 kWh against 0.16 kWh.

(b) An operator wants the gentle limits for passenger comfort on every hop. Name two things the operator must budget for, and why a ground vehicle pays far less for the same choice. *Checkpoint:* Trip time and battery energy. A car does not spend power to stay up.

3.11 Application area · 3 min

Aviation load control often uses standard passenger masses. Case 4 shows a load sheet 20 % light on a vehicle of 620 kg.

(a) How much mass and weight does the sheet miss, with $g = 9.81\,\mathrm{m/s^2}$? *Checkpoint:* 124 kg and 1.22 kN.

(b) Give two reasons an underestimate is the dangerous direction for this vehicle, one in the loops and one in the supervisor. *Checkpoint:* A late liftoff and a lag, and a dispatch margin that reads too high.

3.12 Review the agent · 5 min

A student asks an AI assistant what happens in case 5 when the suitcase slides from the left to the right end of the tray in flight, with Controller set to + wind estimate. The assistant replies: “The disturbance observer estimates whatever the load sheet missed. It will see the shifted center of mass and remove the tilt error within a couple of seconds, about three of its time constants of 0.67 s.” Find the flaw and say how you would check it.

CHAPTER 4

Lift-Plus-Cruise VTOL Simulator for Control Allocation and Measured Transition Gates

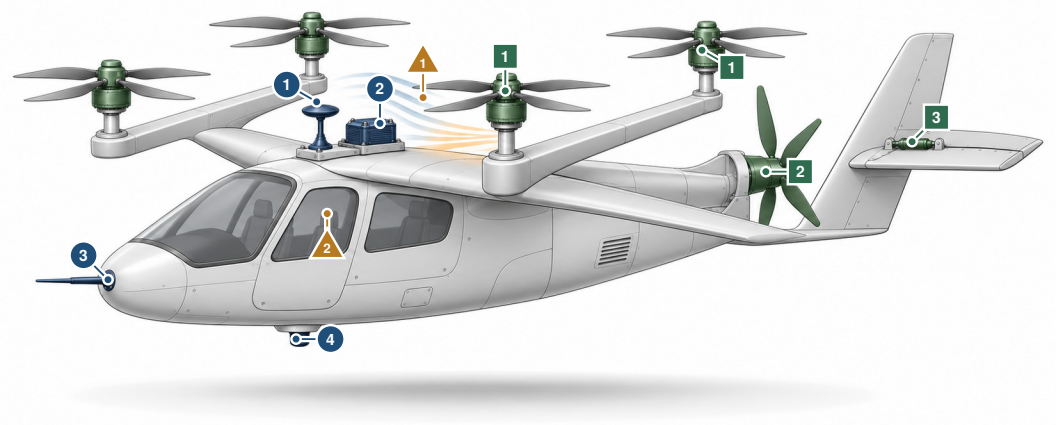

The controller moves 1 front and rear lift rotors 2 pusher propeller 3 elevator on the tail

The controller reads 1 satellite receiver (position and speed) 2 inertial unit (pitch) 3 pitot probe (airspeed) 4 radar altimeter (height)

What disturbs it 1 gust from behind 2 payload in the cabin

Notice the lift rotors, the wing and the pusher on one airframe: in transition the wing takes the weight from the rotors, and the measured airspeed tells when it can carry all of it. Illustration, not to scale.

IN PLAIN TERMS **Handing the Weight From Rotors to a Wing**

The main idea. This electric aircraft lifts off on spinning rotors like a large drone, then flies on its wing like a small plane. The rotors need about three times the power of wing flight, so they stop once the wing can hold the aircraft up.

The challenge. A wing holds the aircraft up only while air flows fast over it, and a gust from behind steals some of that flow. A stop set by the clock can then come too early, while one that waits for the measured airspeed stays safe.

Why it matters. Stopping the rotors too soon makes a loaded aircraft sink toward the trees below it.

THREE WAYS TO READ THIS CHAPTER

Explore	9 min	the first page, the fact on concept vehicles in the Introduction, the history of the first tilt rotor conversion (Section 4.2), the try-it segment, the history of the treetop conversion and the fact on the graded mission (Section 4.6), 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 4.6), Technology in Practice (Section 4.7), Limitations (Section 4.8) and the open challenge

The control question A lift-plus-cruise aircraft hovers on lift rotors that draw about three times the power of wingborne cruise, so every flight hands the weight from the rotors to the wing and back. When may the rotors stop, and should a clock or a measurement decide?

Learning objectives

- Compute the stall, stop and restart speeds of the wing from the weight, and the lift capacity at a rotor stop.
- Explain how a wing-first allocation shares one vertical demand between the wing and two lift-rotor groups.
- Compare a transition timed on the clock with one gated on measurements in a gust and with a heavier load.
- Describe how a lift-derived airspeed exposes a biased pitot.
- Check whether a vertical landing is still feasible after a motor fails.

Prerequisites Sections 1.4.7 (allocation), 1.4.3 (disturbance observer), 1.4.4 (jerk-limited references), 1.4.9 (hysteresis) and 1.6.2 (velocity-form PID). Chapter 3 flies a multirotor on a load sheet.

The app https://apmonitor.com/apps/vtol-transition/

See also Volume 4, Chapter 12 and Volume 4, Chapter 13 for one demand shared among actuators with hard limits. Volume 8, Chapter 11 for a switch on the clock against one on a measured breakthrough, and Volume 8, Chapter 7 for a handover between controllers without a bump.

4.1 Introduction

Many plants have more than one way to produce the effort a loop asks for, and the ways differ in cost, in range and in how their effectiveness changes with the operating point. Volume 4, Chapter 12 shares one push among tires, fans and thrusters at three prices. A ship holding station has thrusters that work at any speed and rudders that work only when it moves, and Chapter 6 weights each thruster of such a vessel by its rating. An aircraft combines control surfaces whose authority grows with dynamic pressure with engines whose thrust does not, and control allocation shares one demand among such effectors in aircraft, marine vessels, road vehicles and spacecraft alike [1]. A process plant has the same shape when a cheap utility carries a load up to its limit and an expensive one supplies the rest. When an effector is switched off rather than merely limited, the plant also becomes a switched system, and the rule that selects the mode belongs to the design, since switching among modes that are each stable can still destabilize the whole [2]. The same choice between a switch on the clock and a switch on a measurement decides when the lead column of Volume 8, Chapter 11 comes off.

An electric lift-plus-cruise aircraft is an extreme case of both. Its lift rotors carry the weight at any airspeed, including none, while its wing carries the weight only above a certain airspeed and below the angle of attack at which it stalls, and the rotors cost several times the power of the wing. NASA's concept vehicles for urban air mobility include a lift-plus-cruise configuration [3], and a physics-based analysis of electric vertical takeoff and landing aircraft found them efficient when cruising but costly in takeoff and climb [4]. In the simulator of this chapter the rotors carry the 19.62 kN weight of a 2000 kg aircraft on 431 kW in hover, while wingborne cruise at 60 m/s with the rotors stopped takes 139 kW. Between the two, the transition hands the weight to the wing, and a supervisor decides when the unloaded rotors may stop and when they must start again. Control design for winged electric aircraft has mostly addressed hover or cruise, although transition needs the rotorcraft and the airplane effectors together [5].

The misconception at stake is that a transition is an event on a schedule, timed once for a nominal aircraft and then replayed. A schedule on the clock and a gate on the measurements make nearly the same decision at the weight and in the air the schedule was tuned for, and in calm air with 250 kg of payload they stop the rotors 42.0 and 43.0 s into the transition, at 49.8 and 51.5 m/s. They part when the load and the air change. With 470 kg on board and a gust from behind at the handover, the schedule stops the rotors at 37.9 m/s, below the 40.4 m/s stall speed of the heavier aircraft, where

no angle of attack can hold the weight, and the aircraft sinks to 9.3 m above a wood. The measured gates keep the rotors turning through the gust and stop them at 55.7 m/s. A transfer function at one operating point represents neither the airspeed on which the wing's effectiveness depends nor the discrete removal of an actuator. Surveys of control teaching and of process control practice recommend laboratory and simulation work in which the plant's behavior, not algebra, carries the lesson [6, 7].

> ✦ **DID YOU KNOW? Eight Lift Rotors and One Pusher**
>
> A lift-plus-cruise concept vehicle described by NASA in 2018, sized for a payload of 1,200 pounds or up to six passengers, has eight lifting rotors and a single pusher propeller. Aircraft of this type have been designed with as few as 4 and as many as 12 lifting rotors [3]. The aircraft of this chapter has eight, lumped into two groups.

However, the two policies can be told apart only in air they were not designed for, and only on a model that carries the physics separating them, lift that follows the air-relative velocity and ends at a stall, rotors that take time to spin up, an allocation that honors every limit and reports what it cannot meet, and sensors that can be wrong. This chapter sets a transition scheduled on the clock against one gated on measurements in a lift-plus-cruise simulator, and shows from the envelope of the wing why the two part company. Lift-plus-cruise transition control, daisy-chain allocation, hysteresis in switching logic and synthetic airspeed are all established methods, and the chapter uses them as they are. Section 4.2 traces the handover from the lift engines of early VTOL aircraft to electric rotors. Section 4.3 builds an eleven-state planar aircraft with air-relative aerodynamics and an explicit stall, momentum-theory power in edgewise flow and a battery power limit, each stated with its constants. Section 4.4 derives closed-form speeds, from the weight and the wing, that locate every gate of the transition, and the lift capacity at the moment of the stop, which predicts the sink that follows. Section 4.5 describes the loops, the constrained allocator and the supervisor with measured gates, feasibility checks and a pitot cross-check. Section 4.6 flies eight cases and a graded mission that isolate the handover, the envelope, the back-transition, a gust, a heavier load, a sensor fault and a motor failure. Section 4.7 names the airspeed sensor, the lift rotors and the flight computer of a real aircraft, Section 4.8 states what the model leaves out, and Section 4.9 draws the lessons together.

4.2 Transition Flight From Lift Engines to Electric Rotors

An aircraft that takes off vertically and cruises on a wing has to hand its weight from one source of lift to the other, and two answers appeared early. One keeps lift and cruise thrust apart. Lift engines used only for vertical flight were evaluated for tactical VTOL aircraft in 1965 [8], and in 1968 a primary fly-by-wire flight control system was designed and developed for the XV-4B V/STOL research aircraft [9]. The other turns one set of thrusters from lifting to propelling, as the XV-15 tilt rotor research aircraft did in a program reported in 1979 [10]. The electric lift-plus-cruise aircraft returns to the first answer with rotors in place of jets, since its lift rotors are stopped in cruise and a separate pusher propels it.

The physics of the handover is older than either. Glauert's momentum theory of the autogyro related the induced velocity of a rotor to its thrust and to the flow through it in forward flight [11], and the same relation, with a figure of merit for the losses, gives the power of the lift rotors in this model. A fixed wing's lift follows its angle of attack up to the stall and then falls toward the normal force of a flat plate, and flight models of small aircraft join the two regimes with a sigmoid in the angle of attack [12].

The distribution of one demand among redundant, limited effectors became a subject of its own with aircraft control surfaces. Constrained control allocation described the moments that limited surfaces can attain [13], and daisy-chain allocation uses one group of effectors until it saturates and passes the remainder to the next [14]. The allocator of this chapter is a daisy chain in that sense, with

the wing first and the rotors second, and it reaches the wing only through the pitch attitude. The general choices of an allocation at its limits are set out in Section 1.4.7.

Electric aircraft revived the problem at a new scale. NASA's concept vehicles for urban air mobility include a lift-plus-cruise design, generic by intent and different in detail from industry arrangements [3, 15]. Transition control laws for such an aircraft have been designed to meet flying-qualities requirements within realistic motor limits, with variable rotor speed sufficient at low speed where the control surfaces lack authority [5]. An airspeed-dependent blending between hover and fixed-wing controllers has been analyzed for stability near the stall, with different responses in the two directions of transition [16], and an automatic flight and transition system for a lift/cruise VTOL aircraft has been verified in flight [17]. The airspeed on which such schedules rest comes from a pitot tube, and a synthetic air data system has been used to detect pitot tube failure on a variable-skew quad plane [18]. The cross-check of this chapter is of that kind, an airspeed inferred from the lift the wing is measured to make.

> ⧖ HISTORY **Converting a Tilt Rotor in Flight (1979)**
>
> A tilt rotor research aircraft flew its first full conversion from helicopter to airplane mode on 24 July 1979. It had taken 15 hours of flight tests over more than three months, widening the envelope in small steps of airspeed and nacelle angle. The wing's control surfaces stayed active in every configuration. As the nacelles tilted toward the airplane position, two mechanical mixing boxes phased the rotor controls out automatically, so the pilot mainly had to hold the attitude while the heavy nacelles moved. The conversion could be flown anywhere in a corridor of airspeeds and nacelle angles, bounded at low airspeed by the stall of the wing [19]. The handover of control followed the nacelle angle, and the wing's stall set how slow it could be flown.

4.3 A Planar Lift-Plus-Cruise Aircraft

The simulated aircraft belongs to the lift-plus-cruise class, 1750 kg empty with 250 kg of payload as shipped, and its flight is confined to the vertical plane through the route. Its eight lift rotors on two booms are lumped into a front and a rear group with one thrust each, a pusher propeller sits at the tail and a wing over the center of mass. Distance x points forward along the route, the altitude z of the landing gear points up and the pitch θ is positive nose up, so the body's forward axis is $(\cos\theta, \sin\theta)$ and its up axis $(-\sin\theta, \cos\theta)$. The values are illustrative and describe no particular aircraft. The geometry, the forces and the instruments are shown in Figure 4.1.

Table 4.1 sorts the signals by role. The eleven integrated states are the distance and the altitude with their rates, the pitch and its rate, the thrusts of the two lift groups and of the pusher, the elevator deflection and the energy left in the battery.

4.3.1 Forces, Moments and Mass

The lift rotors push along the up axis and the pusher along the fuselage, which with the aerodynamic force $\mathbf{F}_a$ gives the translational balance of Equation (4.1).

$$\begin{aligned} m\ddot{x} &= F_{a,x} - T_R \sin\theta + T_p \cos\theta \\ m\ddot{z} &= F_{a,z} + T_R \cos\theta + T_p \sin\theta - mg, \qquad T_R = T_f + T_r \end{aligned} \tag{4.1}$$

The pitch balance of Equation (4.2) takes each group's arm from the center of mass.

$$\begin{aligned} I\ddot{\theta} = \bar{q}S\bar{c}\left[C_m(\alpha) + C_{mq}\frac{\omega\bar{c}}{2V} + C_{m\delta}\delta_e\right] \\ + \ell_f T_f - \ell_r T_r - h_p T_p \end{aligned} \tag{4.2}$$

where $\bar{q} = \rho V^2/2$ is the dynamic pressure of the air-relative speed V, $\omega = \dot{\theta}$ the pitch rate, $\ell_f = \ell - c$ and $\ell_r = \ell + c$, with c the position of the center of mass ahead of the rotors' centroid. The damping term acts only above 0.5 m/s. A level hover in still air needs $\ell_f T_f = \ell_r T_r$, so with $c = -0.15$ m the

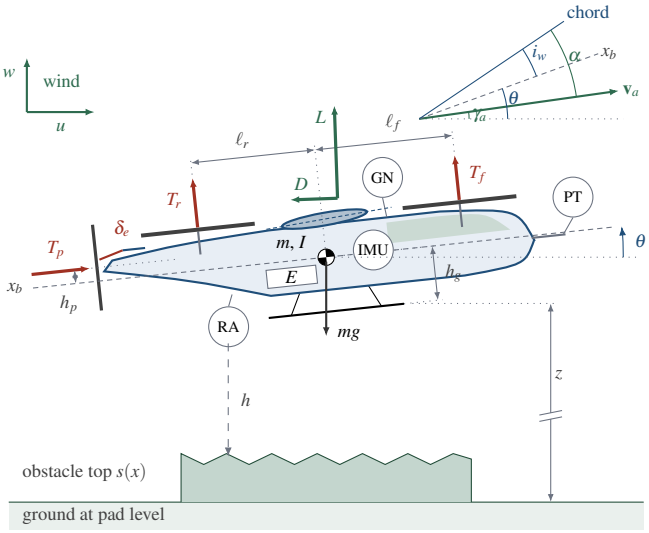

Figure 4.1 The planar lift-plus-cruise aircraft as modeled, drawn with a 6° pitch and not to scale in height. The front and rear lift groups push along the body's up axis with T_f and T_r at arms $\ell_f = \ell - c$ and $\ell_r = \ell + c$ from the center of mass, which sits 0.15 m behind the rotors' centroid ($c = -0.15$ m). The pusher T_p acts along the fuselage h_p above the center of mass, the elevator δ_e at the tail, and the wing at incidence i_w makes lift L normal and drag D parallel to the air-relative velocity $\mathbf{v}_a$. The key defines γ_a, θ, i_w and α with exaggerated angles. GN, IMU, PT and RA are the GNSS receiver (x_m, z_m, $v_{x,m}$, $v_{z,m}$), the inertial unit (θ_m, ω_m), the pitot (V_p) and the radar altimeter (h_m), E is the battery, h_g the height of the center of mass above the gear, z the gear altitude above the pad level and h the height above the surface $s(x)$ below.

Table 4.1 Variables, With Nominal Values in the Hover of Case 1 and in Wingborne Cruise at 60 m/s

Symbol	Role	Hover, cruise	Range
z	CV, gear altitude	50, 110 m	SP 30–200 m
$\dot{x}$	CV, ground speed	0, 60 m/s	SP 35–75 m/s
x	CV, position trim	0, —	0–7 km
θ	CV, pitch loop	0, 0.08°	−8 to 15°
a_z, a_x	MVs, accelerations	0, 0	±5, ±3 m/s^2
M	MV, pitch moment	0, 0	±60 kN m
T_f	state, front group	9.30 kN, 0	0–16 kN
T_r	state, rear group	10.32 kN, 0	0–16 kN
T_p	state, pusher	0, 1.78 kN	−1.5 to 5.5 kN
δ_e	state, elevator	0, −1.59°	−25 to 20°
$\dot{x}$, $\dot{z}$, ω	states, rates	0	—
E	state, battery energy	120 kWh	0–150 kWh
—	MV of the supervisor	rotors on, off	4 states
m_{pay}	DV, payload	250 kg	0–500 kg
u_0, s_t	DV, wind, turbulence	0, 1	±12 m/s, 0–2
b_p	DV, pitot bias	0	±15 m/s
a_r	DV, rear authority	100 %	25–100 %
—	DV, gust, updraft, failure	events	—

rear group carries $\ell_f/(\ell_f+\ell_r) = 52.6\,\%$ of the weight, 10.32 kN against 9.30 kN at 2000 kg. Payload adds mass and inertia, $m = m_e + m_{pay}$ and $I = I_0 + m_{pay} r_{pay}^2$, which gives 2000 kg and 5250 kg m^2 as shipped, and it changes only on a pad, before a run starts or at a Reset, which restarts the case with it.

4.3.2 Air-Relative Aerodynamics and the Stall

The wing's forces follow the velocity relative to the air, $\mathbf{v}_a = (\dot{x}-u, \dot{z}-w)$ with u and w the wind, as shown in Equation (4.3).

$$\begin{aligned} V &= |\mathbf{v}_a|, \quad \gamma_a = \operatorname{atan2}(v_{a,z}, v_{a,x}), \quad \alpha = \theta + i_w - \gamma_a \\ F_{a,x} &= -\frac{D v_{a,x} + L v_{a,z}}{V}, \quad F_{a,z} = \frac{L v_{a,x} - D v_{a,z}}{V} \\ L &= \bar{q} S C_L(\alpha), \quad D = \bar{q} S C_D(\alpha) \end{aligned} \tag{4.3}$$

Lift is normal to $\mathbf{v}_a$ and drag lies along it, so both vanish at zero airspeed, and a gust from behind removes lift at constant ground speed. The coefficients of Equation (4.4) are linear up to the stall and blend into a flat plate past it [12].

$$\begin{aligned} \sigma &= \left[1 + \mathrm{e}^{-M_s(|\alpha| - \alpha_s - \Delta_s)}\right]^{-1}, \qquad C_{L,lin} = C_{L0} + a\alpha \\ C_L &= (1-\sigma) C_{L,lin} + \sigma C_N \sin\alpha \cos\alpha \\ C_D &= (1-\sigma)\left(C_{D0} + k C_{L,lin}^2\right) + \sigma\left(C_{D0} + C_N \sin^2\alpha\right) \\ C_m &= (1-\sigma)\left(C_{m0} + C_{m\alpha}\alpha\right) - \tfrac{1}{2}\sigma \sin\alpha \end{aligned} \tag{4.4}$$

where $M_s = 80\,\mathrm{rad}^{-1}$ sets the sharpness and the offset $\Delta_s = 1.92^\circ$ places the lift peak at the stall angle $\alpha_s = 14^\circ$. There C_L peaks at 1.359, and by 18° it has fallen to 0.621 while C_D has risen from 0.137 to 0.215. The stopped rotors' drag is folded into C_{D0}, and the rotors and the wing do not interfere.

4.3.3 Rotor, Pusher and Battery Power

Each spinning lift group draws the power of Equation (4.5), in which the induced velocity in edgewise flow follows Glauert's momentum theory [11].

$$\begin{aligned} P_i &= \frac{T_i\left[v_i + \max(0, V_c)\right]}{FM} + P_0, \qquad v_h^2 = \frac{T_i}{2\rho A} \\ v_i &= \left[\frac{2 v_h^4}{\sqrt{V_e^4 + 4 v_h^4} + V_e^2}\right]^{1/2} \end{aligned} \tag{4.5}$$

for $i \in \{f, r\}$, where V_e and V_c are the components of $\mathbf{v}_a$ along the fuselage and the up axis, A the disk area of one group, FM the figure of merit on the induced part and P_0 the profile power of a spinning group. A stopped group draws nothing, and a descent earns no credit. In hover $v_i = v_h$, and the two groups draw 420.5 kW at 2000 kg, against an ideal induced power of 277.9 kW, an effective figure of merit of 0.661. The pusher and the battery follow Equation (4.6).

$$\begin{aligned} P_p &= \frac{|T_p|\left(V_f + v_{ip}\right)}{\eta_p}, \qquad v_{ip} = -\frac{V_f}{2} + \sqrt{\frac{V_f^2}{4} + \frac{|T_p|}{2\rho A_p}} \\ \dot{E} &= -\left(P_f + P_r + P_p + P_{av}\right), \qquad P \le P_{max} f(E/E_{cap}) \\ f(s) &= \min\left(1,\ p_f + (1-p_f)\, s/s_k\right) \end{aligned} \tag{4.6}$$

where $V_f = \max(0, V_e)$ for forward thrust and zero in reverse, P_{av} is the avionics and cabin power, and the limit falls linearly below the knee s_k to p_f of P_{max} at empty. The two groups, the pusher and the elevator each lag their commands with a first-order time constant, $\tau \dot{T}_i = T_{i,cmd} - T_i$, and the commands are rate limited before they are sent. A stopped group's thrust below 5 N is zeroed after each step.

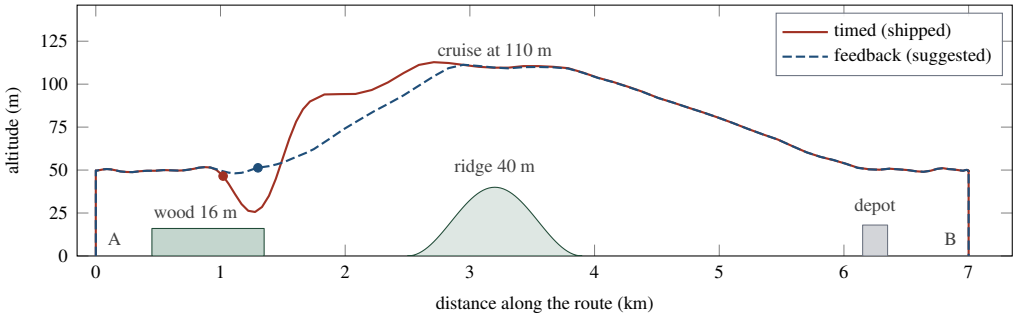

Figure 4.2 The route and the two flights of the challenge, gear altitude against distance, with the timed transition as shipped (solid) and with the suggested feedback transition (dashed). The dot on each line marks its rotor stop command. The timed flight sinks to 25.4 m over the wood, 9.3 m above the trees, and then climbs through its reference, while the feedback flight never falls more than 1.9 m below its reference. Both flights brake to rest over pad B and descend onto it. The vertical scale is exaggerated.

4.3.4 Wind, Route and Contact

The wind of Equation (4.7) is a steady component along the route, one-minus-cosine gusts and first-order colored turbulence, sampled at each control step and held until the next, and the same over the whole airframe.

$$\begin{aligned} u &= u_0 + \textstyle\sum_j \Delta u_j \, g_j(t) + 0.8 \, s_t n_u \\ w &= \textstyle\sum_j \Delta w_j \, g_j(t) + 0.5 \, s_t n_w \\ g_j &= \tfrac{1}{2}\left[1 - \cos\!\big(2\pi(t - t_j)/D_j\big)\right], \quad t_j \le t \le t_j + D_j \end{aligned} \tag{4.7}$$

where n_u and n_w have unit standard deviation and correlation times of 2 and 1 s, the scale s_t is 1 in every case, and u is positive from behind. The gust event adds $\Delta u = 8\,\mathrm{m/s}$ over $D = 8\,\mathrm{s}$, and cases 5 and 9 use 10 m/s with $\Delta w = -1\,\mathrm{m/s}$ over 10 s.

The route runs 7 km from pad A to pad B over a wood of 16 m trees from 450 to 1350 m, a ridge that rises to 40 m as a raised cosine from 2500 to 3900 m, and an 18 m depot from 6150 to 6350 m, as shown in Figure 4.2, and a diversion airfield lies 8 km beyond pad B. A touchdown occurs when the gear reaches the surface while descending, except in the 0.5 s after a liftoff. Within 7.5 m of a pad center it is a landing unless the sink exceeds 3.0 m/s or the pitch 10°, and anywhere else it is a crash, as is any contact of the nose, the tail or a boom tip with the ground or an obstacle. Liftoff requires the rotors and the wing to lift more than $1.001\,mg$.

4.3.5 Parameters, Numerics and Initial States

Table 4.2 gives the parameters. The fourth-order Runge-Kutta method integrates the states on a 2 ms step, the control laws execute every 10 ms, and the log records 20 samples per second, on the engine of Section 1.6.1. Halving the step moves the touchdown of the suggested challenge by 0.012 mm/s and 0.047 mm and the energy drawn up to 238.10 s by 0.014 mWh, and its rotor stop stays on the same control step. In the shipped challenge the timed stop comes one control step later at the half step, at 69.12 s, because the summed clock reaches 42 s after the start of the transition a rounding error later, and the two runs' airspeeds agree to 0.01 mm/s at the same instant. A scenario starts on a pad with the rotors stopped, in a hover trimmed for its load, or wingborne at its speed with the rotors stopped and the pitch, pusher and elevator trimmed.

Table 4.2 Aircraft Parameters. Values Not Marked Hidden Can Be Changed in the Model Dialog Within the Range Given

Symbol	Meaning	Value	Range
m_e	operating empty mass	1750 kg	1400–2100
I_0	pitch inertia, empty	5000 kg m^2	3000–8000
r_{pay}	payload radius of gyration	1.0 m	hidden
ℓ, c	arm, c.m. ahead of centroid	2.9, −0.15 m	2–4, ±0.5
h_p, h_g	pusher line, gear	0.3, 1.2 m	−0.5–1, hidden
T_{max}, T_{idle}	group rating, idle	16 kN, 300 N	10–22, 0–1000
τ_R, —	group lag, rate limit	0.15 s, 25 kN/s	0.05–0.5, 5–60
t_{spin}, P_0	spin-up, profile power	2.0 s, 25 kW	0.5–5, 0–60
A, FM	disk per group, merit	20 m^2, 0.75	10–40, 0.5–0.9
$T_{p,max}$, $T_{p,min}$	pusher limits	5.5, −1.5 kN	3–9, −3–0
τ_p, —	pusher lag, rate limit	0.3 s, 10 kN/s	0.05–1, 2–30
A_p, η_p	pusher disk, efficiency	3.1 m^2, 0.85	hidden, 0.6–0.95
S, $\bar{c}$	wing area, chord	16 m^2, 1.42 m	10–25, 0.8–2.5
i_w, α_s	incidence, stall angle	4°, 14°	0–8, 10–18
C_{L0}, a	lift at zero, slope	0.2, 5.0 rad^{-1}	0–0.5, 3.5–6.3
M_s, C_N	sharpness, flat plate	80 rad^{-1}, 1.9	hidden
C_{D0}	zero-lift drag	0.035	0.02–0.08
k	induced-drag factor	0.05	0.03–0.08
C_{m0}, $C_{m\alpha}$	moment, stiffness	0.03, −0.7 rad^{-1}	±0.1, −2–0
C_{mq}	pitch damping	−10	−30–0
$C_{m\delta}$	elevator effectiveness	−1.1 rad^{-1}	−2 to −0.3
δ_e	elevator travel	−25, 20°	−35 to −10, 10–30
τ_e, —	elevator lag, rate	0.05 s, 60 °/s	0.02–0.2, 20–120
ρ, g	density, gravity	1.225, 9.81	1.0–1.3, hidden
E_{cap}	usable battery energy	150 kWh	60–250
P_{max}	rated battery power	700 kW	400–1000
s_k, p_f	sag knee, power at empty	30, 60 %	10–50, 30–100
P_{av}	avionics and cabin	5 kW	0–20
—	gear sink, pitch limits	3.0 m/s, 10°	1–5, 4–20
σ_x, σ_z	GNSS position, altitude	0.3, 0.15 m	0–2, 0–1
σ_v, σ_q	GNSS velocity, pitot	0.05 m/s, 4 Pa	0–0.5, 0–30
σ_θ, σ_ω	attitude, gyro	0.1°, 0.3 °/s	0–1, 0–3

4.3.6 Assumptions

Flight is planar and longitudinal, with no roll, yaw, sideslip or turns, so a go-around on the wing continues ahead to the airfield beyond pad B instead of circling back.

Each group of four rotors becomes a single thrust, lagged, rate limited and given a 2 s spin-up. Inflow dynamics, flapping and the in-plane H-force of the rotors are absent, and so are the rotors' downwash on the wing and the wing's upwash through the rear rotors, which in a real aircraft change both the lift and the rotor power during the transition.

The aerodynamics are quasi-steady, with no dynamic stall or lift lag. A six-degree-of-freedom analysis of lift-plus-cruise transition locates a near-stall region where the short-period damping falls [16], and the angle-of-attack margin the scorecard reports is not a measure of that damping.

Rotor power is momentum theory corrected by one fixed figure of merit and a profile power, without a vortex-ring state. The battery holds energy and limits power, and it has no voltage, temperature or aging dynamics.

The wind is uniform over the airframe, with no shear and no ground effect, and the sensors carry colored noise without latency or dropouts. The controller's aerodynamic model equals the plant's, so its errors come only from what it cannot measure, the vertical gust, the turbulence and a biased pitot. The margins and limits, a 10 % lift-group margin, 3° of angle-of-attack protection, 25 m of clearance and a 20 % energy reserve, are teaching values, not certification requirements.

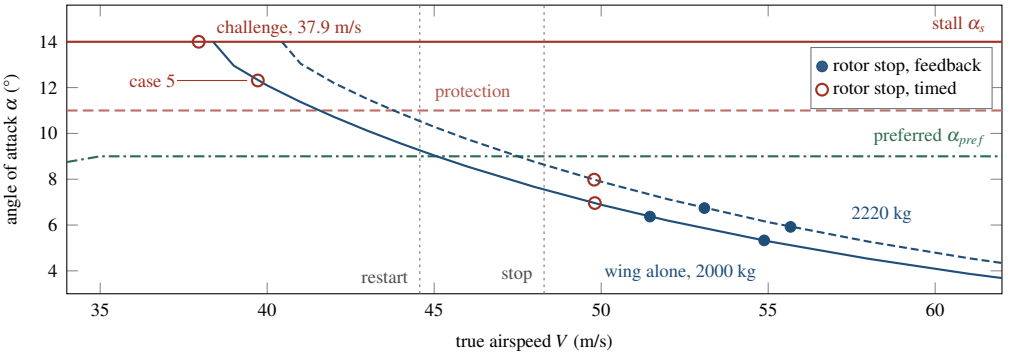

Figure 4.3 The envelope of the handover. The solid and dashed curves give the angle of attack at which the wing alone carries 2000 and 2220 kg in level flight, from the model's own C_L up to the lift peak, and end at the stall speeds of 38.4 and 40.4 m/s. The stall angle is the solid horizontal line, the protection the dashed one and the preferred angle the dash-dotted one. The markers, filled for the feedback transition and open for the timed one, place the true airspeed of each rotor stop command on the curve of its mass. The timed stops of cases 2 and 6 leave the wing margin in hand, the one in the gust of case 5 lies above the protection, and the one in the challenge lies below the stall speed, where no angle holds the weight.

Table 4.3 Transition Speeds From the Weight and the Wing Against the Engine, in m/s

Event	Predicted	Measured
wing takes half the weight, forward (case 2)	32.7	33.5
rotors take half the weight, back (case 4)	32.7	32.7
rotors under 8 % of the weight (case 2)	43.2	42.7
restart gate V_{off}, 2000 kg (case 3)	44.6	44.3
stop gate V_{on}, 2000 kg (case 2)	48.3	48.2
stop after the 2 s dwell (case 2)	50.7	51.5
stop gate V_{on}, 2220 kg (case 6)	50.9	51.0

4.4 Where the Wing Can Take the Weight

Four speeds follow from the weight and the wing alone, and they locate the gates of a transition. In level flight the wing alone carries mg at the angle where $C_L(\alpha) = 2mg/(\rho S V^2)$, so it can carry the weight at all only above the stall speed V_s, and inside the protection only above V_p. The gates of Section 4.5 add the lift margins μ_{on} and μ_{off}, as shown in Equation (4.8).

$$\begin{aligned} V_s &= \sqrt{\frac{2mg}{\rho S C_{L,max}}}, & V_p &= \sqrt{\frac{2mg}{\rho S C_L(\alpha_s - \Delta\alpha_p)}} \\ V_{on} &= V_p\sqrt{1+\mu_{on}}, & V_{off} &= V_p\sqrt{1+\mu_{off}} \end{aligned} \tag{4.8}$$

With $C_{L,max} = 1.359$, $C_L(11^\circ) = 1.159$, $\mu_{on} = 35\,\%$ and $\mu_{off} = 15\,\%$ they are 38.4, 41.6, 48.3 and 44.6 m/s at 2000 kg, and 40.4, 43.8, 50.9 and 47.0 m/s at 2220 kg. Because $V_p = 1.083V_s$, the stop gate lies at $1.258V_s$ and opens after the airspeed gate at $1.2V_s$, which is therefore a backstop at the default settings. The allocation gives the wing what it lifts at a preferred angle α_{pref}, which rises from the fuselage-level 4° below 15 m/s to 9° at 35 m/s, and the rotors the rest. The wing therefore takes half the weight where $\frac{1}{2}\rho V^2 S C_L(\alpha_{pref}(V)) = mg/2$, and the rotors fall under 8 % where it lifts $0.92mg$. The angle the wing alone needs is plotted against the airspeed in Figure 4.3, and the predictions are compared with the engine in Table 4.3.

Table 4.4 Lift Capacity at the Rotor Stop and the Sink That Follows. Capacity Is What the Wing Lifts at 11° at the Stop Airspeed, as a Share of the Weight, and the Loss Is the Largest Distance Below the Altitude Reference in the Next 15 s

Stop	kg	m/s	Capacity	Loss (m)	Drop (m)
feedback, case 2	2000	51.5	153 %	2.5	1.2
feedback, case 5	2000	54.9	174 %	0.9	0
feedback, case 6	2220	53.1	147 %	2.5	1.0
feedback, challenge	2220	55.7	162 %	1.6	0
timed, case 2	2000	49.8	144 %	2.1	0.5
timed, case 6	2220	49.8	129 %	1.8	0.2
timed, case 5	2000	39.7	91 %	6.7	3.2
timed, challenge	2220	37.9	75 %	31.7	21.1

The gates open where Equation (4.8) says, to within 0.1 m/s, because the supervisor evaluates the same expression on an exact load sheet and a pitot that is accurate in calm air. The rotors stop 2.03 s after the gate opens in case 2, while the reference accelerates at 1.2 m/s^2, so the prediction is $48.3 + 2.4 = 50.7\,\mathrm{m/s}$, and the stop comes at a true airspeed of 51.5 m/s with a ground speed of 50.8 m/s, the difference being a 0.7 m/s turbulent headwind. The restart comes 0.3 m/s below its gate because the restart condition must hold for 0.3 s while the aircraft slows. The half-load points fall at and just above the prediction. On the way out the wing takes half the weight at 33.5 m/s, when its true angle of attack trails the commanded 8.65° by 0.54°, and on the way in the rotors carry half at 32.7 m/s.

The capacity at the moment of the stop predicts what follows, as shown in Table 4.4, where the drop is the altitude lost from the stop to the lowest point. Every stop with a capacity of 129 % or more loses less than 2.6 m against a reference that climbs toward the cruise altitude, and loses at most 1.2 m of real height. At 91 % the aircraft drops 3.2 m, and at 75 % it drops 21.1 m. The feedback gates guarantee at least 135 % by construction when the airspeed is right, and delivered 147 to 174 %. The timed schedule guarantees nothing, because its capacity is set by whatever airspeed the clock happens to meet.

4.5 Loops, Allocation and Supervisor

The controller runs every $\Delta t_c = 10\,\mathrm{ms}$ on sampled, noisy measurements, and the true state, airspeed, angle of attack and wind stay hidden from it. Three loops ask for a vertical acceleration, a longitudinal acceleration and a pitch moment, an allocator shares these demands among the wing, the rotors, the pusher and the elevator, and a supervisor runs the mission and decides whether the rotors may turn, as shown in Figure 4.4.

4.5.1 Sensors and Estimators

GNSS gives x_m and z_m with colored noise of 0.3 and 0.15 m and correlation times of 0.5 s, and the two velocities with 0.05 m/s and 0.1 s. The radar altimeter gives the height above what is below with 0.15 m and 0.2 s, and the inertial unit gives θ_m and ω_m with 0.1° and 0.3 °/s at 0.05 and 0.02 s. The pitot measures the dynamic pressure of the airspeed along the fuselage with 4 Pa of noise and a bias b_p, $V_p = \max\big(0, \sqrt{2\max(0, \bar{q}_m)/\rho} + b_p\big)$, which makes it poor below about 10 m/s. The estimators

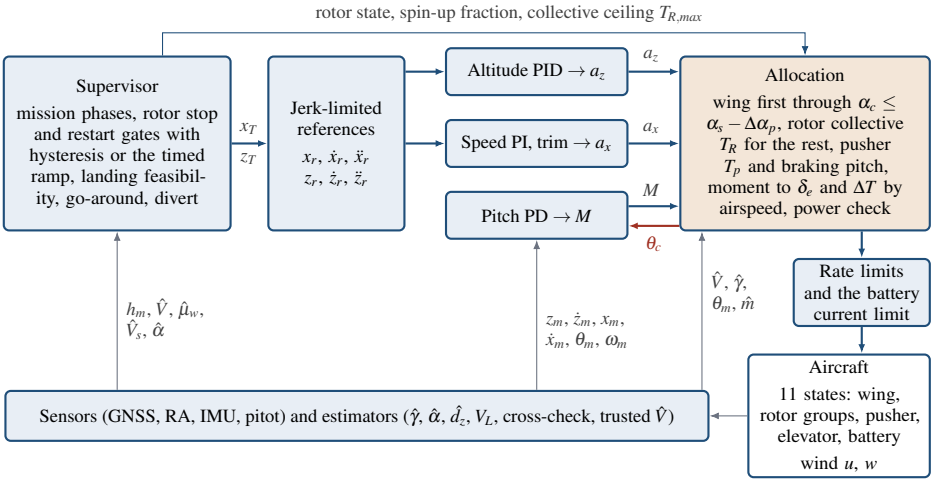

Figure 4.4 The controller as implemented. The allocation writes the pitch loop's setpoint θ_c, which is how it reaches the wing, and the supervisor sets what it may ask of the rotors, through the rotor state and, in the timed transition, a ceiling on the collective. The loops read the measurements, and the allocation and the supervisor also read the estimates, including the trusted airspeed $\hat{V}$.

of Equation (4.9) supply what cannot be measured.

$$\begin{aligned}
\hat{\gamma} &= w_\gamma \arcsin\left(\operatorname{clamp}\left(\frac{v_{z,m}}{\max(\hat{V},10)}, \pm 0.35\right)\right) \\
\hat{\alpha} &= \theta_m + i_w - \hat{\gamma} \\
\dot{\hat{\xi}} &= \hat{F}_z + \hat{d}_z, \qquad \hat{d}_z = K\left(\hat{m} v_{z,m} - \hat{\xi}\right) \\
V_L &= \sqrt{\frac{2\hat{d}_z}{\rho S C_L(\hat{\alpha}) \cos\hat{\gamma}}}
\end{aligned} \tag{4.9}$$

where $w_\gamma = \operatorname{clamp}((\hat{V}-10)/10, 0, 1)$ in m/s, $\hat{F}_z = (\hat{T}_f + \hat{T}_r)\cos\theta_m + \hat{T}_p \sin\theta_m - \hat{m}g$ is built from the controller's own commands through exact models of the actuator lags, $K = 1.5\,\mathrm{s}^{-1}$, and $\hat{m}$ is the load sheet's mass, which is exact. The observer's $\hat{d}_z$ is the vertical aerodynamic force, and V_L, filtered with 0.5 s, is the airspeed at which the wing would make it at $\hat{\alpha}$. The observer is the disturbance observer of Section 1.4.3 applied to the vertical force, and V_L is an analytical redundancy for the pitot. It is formed only above 30 m/s, with $C_L(\hat{\alpha}) > 0.4$ and $\hat{d}_z > 0.3\hat{m}g$. The ratio V_L/V_p is filtered with 1.5 s, a disagreement is a ratio more than 15 % from one, and 2 s of disagreement flags the pitot. The trusted airspeed $\hat{V}$ is the pitot, the lower of the two while they disagree, and V_L once the pitot is flagged, with the pitot's learned error removed below 30 m/s. The estimate $\hat{\alpha}$ cannot see a vertical gust.

4.5.2 The Three Loops

Every loop runs the shared velocity-form PID of Section 1.6.2, written in the symbols of this chapter in Equation (4.10).

$$\begin{aligned}
e_k &= \mathrm{SP}_k - \mathrm{PV}_k, \qquad P_k = K_c e_k \\
S_k &= S_{k-1} + (P_k - P_{k-1}), \qquad c_k = b_k + S_k + I_{k-1} \\
u_k &= \operatorname{clamp}(c_k,\ u_{min},\ u_{max}) \\
I_k &= I_{k-1} + \left[\frac{K_c}{\tau_I} e_k + \frac{u_k - c_k}{T_t}\right] \Delta t_c
\end{aligned} \tag{4.10}$$

It is Equations (1.19) to (1.22) with S_k the running sum of the proportional increments, which is never clamped. Anti-windup is back-calculation with $T_t = \tau_I$, applied only in a loop that integrates, and no loop sets a rate limit. The allocation's own limits, the rotor ceiling, the protected angle and the battery, act after these clamps and are not fed back, so while a lamp reports a demand unmet below the clamp the integral keeps integrating the whole error. The derivative option of the shared PID is left at zero. Each derivative term is built instead from a measured rate and carried in the bias b_k of Equation (4.11), so a setpoint change produces no kick.

$$\begin{aligned} b_z &= \ddot{z}_r + K_{c,z}\tau_{D,z}(\dot{z}_r - v_{z,m}), \qquad b_x = \ddot{x}_r \\ b_\theta &= -\hat{M}_{trim}(\alpha_c) - K_{c,\theta}\hat{I}\tau_{D,\theta}\,\omega_m \\ \mathrm{SP}_x &= \dot{x}_r + \mathrm{clamp}\big(K_x(x_r - x_m), \pm 3\,\mathrm{m/s}\big) \end{aligned} \tag{4.11}$$

with setpoints z_r, SP_x and θ_c against $\mathrm{PV} = z_m$, $v_{x,m}$ and θ_m, and outputs limited to $|a_z| \le 5\,\mathrm{m/s^2}$, $|a_x| \le 3\,\mathrm{m/s^2}$ and $|M| \le 60\,\mathrm{kN\,m}$. Here $\hat{M}_{trim} = \hat{q}S\bar{c}C_m(\alpha_c) - h_p T_p$ is the airframe's moment at the commanded angle of attack and the pusher's, $\hat{q} = \rho\hat{V}^2/2$, and the pitch gain is $K_{c,\theta}\hat{I}$ with $\hat{I}$ the load sheet's inertia. The altitude loop is a PID with $\tau_{I,z} = 12\,\mathrm{s}$, the speed loop a PI with $\tau_{I,v} = 10\,\mathrm{s}$ inside a proportional position trim, and the pitch loop a PD. With $e = \mathrm{SP} - \mathrm{PV}$ and positive gains, an aircraft below, slower than or pitched below its setpoint receives a positive demand in every loop. All three loops are therefore reverse acting, $a = +1$ in the convention of Section 1.3. The pitch loop has no integral action, so its integral term holds only what a manual transfer or a liftoff leaves, multiplied by $\mathrm{e}^{-\Delta t_c/1.5\,\mathrm{s}}$ every step. Manual mode is chosen loop by loop. Individual selections survive before-Run initialization and Reset. The header mode switch sets all three loops together. In manual the operator's demand passes to the allocation and the PID tracks it by choosing I_k, and in the altitude and speed loops the reference follows the aircraft, so the return to automatic is bumpless (Section 1.6.2) and re-plans from where the aircraft is. The references are online jerk-limited profiles of the kind in Section 1.4.4, limited along the route to $1.2\,\mathrm{m/s^2}$ and $0.8\,\mathrm{m/s^3}$ and in height to $0.8\,\mathrm{m/s^2}$, $0.8\,\mathrm{m/s^3}$, 2.5 m/s of climb and 1.5 m/s of descent.

4.5.3 Allocation, Wing First

The allocation works on the estimates. The vertical demand $F_z = \hat{m}(g + a_z)$, less the vertical share of the previous pusher request, is compared with the vertical force the wing makes at the preferred angle, as shown in Equation (4.12).

$$\begin{aligned} \hat{L}_z(\alpha) &= \hat{q}S\left[C_L(\alpha)\cos\hat{\gamma} - C_D(\alpha)\sin\hat{\gamma}\right] \\ \alpha_{pref} &= i_w + (\alpha_{op} - i_w)\,\mathrm{clamp}\big((\hat{V} - 15)/20, 0, 1\big) \\ T_R &= \big[F_z - \hat{L}_z(\alpha_{pref})\big]/\cos\theta_m, \qquad \theta_c = \alpha_c - i_w + \hat{\gamma} \end{aligned} \tag{4.12}$$

with $\alpha_{op} = 9^\circ$ and $\hat{V}$ in m/s. When the wing alone could do more, α_c is lowered to what the demand needs, down to -4°, and the rotors idle at the 570 N collective that holds the more heavily loaded rear group at its 300 N idle thrust. When the rotors cannot supply their share, because of their rating, their spin-up, the timed ramp or the battery, T_R is capped at $T_{R,max}$ and the shortfall is redistributed to the wing by raising α_c, never past $\alpha_s - \Delta\alpha_p = 11^\circ$. What is still missing lights the Z lamp of vertical saturation. Below 12 m/s the wing and the elevator are ignored and the rotors take everything. The pusher then supplies $T_p = (\hat{m}a_x - \hat{F}_{a,x} + T_R\sin\theta_c)/\cos\theta_c$, with $\hat{F}_{a,x}$ the estimated drag force at α_c. A braking demand beyond the pusher's -1.5 kN is redistributed to a nose-up pitch of up to 8° that tilts the rotor thrust aft, never past the protection, and what remains lights the X lamp. If the estimated power exceeds 97 % of the battery's limit, the pusher gives way first and the collective next, which lights the P lamp.

The moment is split by airspeed, as shown in Equation (4.13).

$$\begin{aligned} E_\delta &= \hat{q} S \bar{c} C_{m\delta}, \qquad M_e = w_E M \\ w_E &= \mathrm{clamp}\big((\hat{V} - 20)/20, 0, 1\big) \\ \delta_e &= \mathrm{clamp}(M_e/E_\delta,\ \delta_{min},\ \delta_{max}) \\ \Delta T &= (M - E_\delta \delta_e)/\ell_\Sigma, \qquad \ell_\Sigma = \ell_f + \ell_r \\ T_f &= T_R \ell_r/\ell_\Sigma + \Delta T, \qquad T_r = T_R \ell_f/\ell_\Sigma - \Delta T \end{aligned} \tag{4.13}$$

The elevator's share rises from nothing at 20 m/s to all of the moment at 40 m/s while the rotors turn, and dividing by the estimated effectiveness E_δ schedules its gain with $1/\hat{V}^2$. The collective split in the hover ratio produces no moment, so the differential ΔT carries all of the rotors' share. The differential is served first and limited to the range for which a collective still fits both groups, the collective is then limited to what the differential leaves, and a moment that neither the elevator nor the rotors can make lights the M lamp. A group at its limit therefore costs vertical force, never attitude. This is the moment-first rule of Table 1.1, which Chapter 3 applies to the tilt of a multirotor. The requested commands are rate limited at 25 kN/s for the groups, 10 kN/s for the pusher and 60 °/s for the elevator, and the motors' own current limit, which uses the true airspeed, scales them if their true power would exceed what the battery can give.

4.5.4 The Supervisor

The supervisor flies the mission through the pad, the spool-up, a climb to $h_{tr} = 50$ m, the transition, cruise at $h_{cr} = 110$ m, an approach that begins 3.3 km before pad B, the back-transition, which begins when the along-track reference starts braking, a hover gate over the pad, the vertical descent and the landing. In the feedback transition it stops the rotors only when every condition of Equation (4.14) holds, on measurements and estimates, and the rotors have carried under 8 % of the vertical demand for 2 s.

$$\begin{aligned} &\hat{\mu}_w = \frac{\rho \hat{V}^2 S C_L(\alpha_s - \Delta\alpha_p)}{2\hat{m}g} - 1 \geq \mu_{on}, \qquad \hat{V} \geq 1.2\hat{V}_s \\ &h_m \geq h_{min}, \qquad P_{lim} - P \geq P_{reload}, \qquad \text{airspeeds agree} \end{aligned} \tag{4.14}$$

where $\hat{V}_s$ is Equation (4.8) on the load sheet's mass, $h_{min} = 30$ m, $P_{lim} = P_{max} f(E/E_{cap})$ is the battery's power limit, P the battery's measured power and P_{reload} the power to put the weight back on the rotors at $\hat{V}$. It starts them again after 0.3 s with $\hat{\mu}_w < \mu_{off}$, $\hat{V} < 1.1\hat{V}_s$, $\hat{\alpha}$ within 0.5° of the protection or a disagreement of the airspeeds. On the way down it spins the rotors up once $\hat{V} < V_{spin} = 56$ m/s, or 20 s into the back-transition. The along-track reference does not wait for them, because it is braking into the pad by then and a speed held there could not be braked back within the jerk limit, so V_{spin} alone sets how early they turn. The timed transition caps the collective at $T_{R,max} = 2T_{max}\, r$ with $r = \mathrm{clamp}\big(1 - (t - t_{tr} - t_{u0})/t_u, 0, 1\big)$, where t_{tr} is the start of the transition, $t_{u0} = 28$ s and $t_u = 14$ s, stops the rotors when r reaches zero whatever the airspeed, never restarts them in cruise, and spins them up $t_s = 6$ s after the braking begins.

A vertical landing is feasible when the weaker group can carry its hover share with the margin $\mu_{land} = 10\,\%$, $\min_i(a_i T_{max}/T_{i,trim}) - 1 \geq \mu_{land}$ with a_i the authority left, and the battery can power the hover. The check runs before takeoff, which also needs the charge 3 % above the reserve, and at every step of the flight. In cruise and on the approach a failed check diverts the aircraft on the wing to the airfield, and later a go-around climbs back to h_{tr} on the rotors when that is feasible and fewer than three have been flown, and otherwise diverts if the aircraft is fast enough or above 12 m. The hover gate needs 2 s within 2 m and 0.5 m/s of the pad center within 15 s, and the descent at 1.5 m/s, slowing to 0.5 m/s below 6 m, goes around above 12 m if it drifts 3.5 m off center or its vertical demand stays saturated for 1.5 s. The settings are gathered in Table 4.5.

Three models are kept apart. The simulation model is Section 4.3. The controller's own model is algebraic, made of the same aerodynamic coefficients, the load sheet's mass and inertia, the hover

Table 4.5 Controller, Allocation and Supervisor Settings as Shipped

Symbol	Meaning	Value
$K_{c,z}$, $\tau_{I,z}$, $\tau_{D,z}$	altitude PID	0.3 s^{-2}, 12, 2.2 s
$K_{c,v}$, $\tau_{I,v}$, K_x	speed PI, position trim	0.8 s^{-1}, 10 s, 0.3 s^{-1}
$K_{c,\theta}$, $\tau_{D,\theta}$	pitch PD, times $\hat{I}$	5 s^{-2}, 0.6 s
α_{op}, $\Delta\alpha_p$	preferred angle, protection	9°, 3°
K	vertical-force observer	1.5 s^{-1}
h_{tr}, h_{cr}, V_{cr}	heights, cruise ground speed	50, 110 m, 60 m/s
μ_{on}, μ_{off}	lift margins to stop, restart	35, 15 %
h_{min}, —	stop height, unloaded share	30 m, 8 % for 2 s
V_{spin}	spin-up airspeed	56 m/s
t_{u0}, t_u, t_s	timed schedule	28, 14, 6 s
—	cross-check	15 %, 2 s, above 30 m/s
μ_{land}, —	landing margin, energy reserve	10, 20 %
m_{pay}	payload on the load sheet	250 kg
—	gate window, dwell, longest wait	2 m, 0.5 m/s, 2 s, 15 s

Table 4.6 Scenarios

#	Set-up	Lesson
1	hover 50 m over A, gust at 20 s	the hover split and its cost
2	forward transition from that hover	the load moves, the gate decides
3	wingborne, 42 then 66 m/s at 10 and 70 s	the envelope, hysteresis
4	wingborne 3.2 km out, to pad B	rotors before the wing fades
5	case 2 with a gust at 37 s	measurement against the clock
6	case 2 with 470 kg	gates that move with the load
7	case 2, pitot 12 m/s high at 30 s	an independent airspeed
8	inbound, 470 kg, rear motor fails at 15 s	refusing a landing
9	pad A to B, 470 kg, 40 %, gust at 64 s, timed	graded mission

trim and the power of Equation (4.5), so it errs only through what it cannot measure. The analysis model is Equation (4.8), which interprets the results and enters the code only implicitly, through the gates.

4.6 Measured Cases

Table 4.6 lists the eight cases and the graded mission, which are run as Section 1.6.4 describes. The numbers that follow were produced by the engine run headlessly with the default seed, 20261201, each case as shipped except where a change is stated, and times are simulated seconds from the start of the case. A stated change is made after the case is selected and before Run, and every one quoted here gives the same run when it is carried through a Reset instead. An airspeed at an instant is the true airspeed in force from the control step at that instant, which is how the scorecard records the airspeed of a rotor stop. Table 4.7 gathers the principal outcomes. The exported file carries the true states beside the measured ones, the pitot, lift-derived and trusted airspeeds, the true and estimated angles of attack, the wing and rotor lift beside the weight, the power beside the battery's limit, the phase and rotor state, and the scorecard metrics.

4.6.1 Hover and the Handover

Between 10 and 19.5 s of case 1 the two groups carry the 19.62 kN weight to within 0.02 %, the rear group 10.32 kN and the front 9.30 kN, as the hover trim requires, and the aircraft draws 431 kW, 7.2 kWh a minute. The rotors take 425.7 kW of it, 5.2 kW more than Equation (4.5) gives in still air, because the model charges the rotors for climbing through each updraft of the turbulence and credits nothing on the way down. The pitch loop does not have to hold the unequal split as an error, since the allocation divides the collective in the hover ratio. The gust at 20 s brings the pusher to 300 N of

Table 4.7 Measured Results of Cases 1 to 8

#	Condition	Result
1	hover	$T_R/mg - 1 = 0.02\,\%$, 9.30 and 10.32 kN, 431 kW
1	gust	pusher −300 N, $\lvert\theta\rvert \leq 0.18°$, $x \leq 0.24$ m
2	feedback	492 kW at 2.6 s, half on the wing at 33.5 m/s, gate 48.2 m/s, stop 51.5 m/s at 42.98 s
2	timed	stop 49.8 m/s at 42.00 s
3	60, 42, 66 m/s	4.1°, 139 kW, 9.0° with rotors 13 % and 149 kW, 3.0° and 165 kW
3	speed changes	restart 44.3 m/s at 23.7 s, stop 51.1 m/s at 78.2 s
4	back-transition	spin-up 55.9 m/s, half on the rotors at 32.7 m/s, pitch 5.5°, 464 kW
4	landing	121.6 s, 0.57 m/s, 0.03 m, 8.9 kWh
5	feedback	rotors 9 to 25 %, stop 54.9 m/s at 46.65 s, 1.72 m below
5	timed	stop 39.7 m/s at 42.00 s, 6.7 m below, margin 2.4°
6	feedback	gate 51.0 m/s, stop 53.1 m/s at 44.74 s, timed stop 49.8 m/s
7	cross-check on	lower airspeed 30.73 s, flag 32.72 s, 3.0 m below, stop at 49.3 m/s
7	cross-check off	into the wood at 38.5 s, 856 m along the route
8	470 kg	margin 39.7 to 4.8 %, divert at 15.0 s
8	250 kg	margin 16.3 %, lands at 151.2 s, 0.48 m/s, 0.11 m

reverse thrust, the aircraft moves no more than 0.24 m forward of the pad center, and the pitch stays within 0.18° of level.

Case 2 accelerates from the same hover. Power peaks at 492 kW, 2.6 s in, when the pusher starts while the rotors still carry everything. The rotors carry 97, 86, 42 and 3 % of the weight at 10, 20, 30 and 40 s, at 10.3, 21.7, 34.7 and 47.4 m/s, and the pitch rises to 5.4° as the allocation gives the wing its preferred angle. The altitude stays between 49.3 and 51.9 m through the handover. The estimated rotor share falls under 8 % at 35.55 s, the lift-margin gate opens at 40.95 s, and the rotors stop at 42.98 s at 51.5 m/s, 13.1 m/s above the stall. The timed schedule stops them at 42.00 s at 49.8 m/s, and in this calm air the capacity of 144 % makes the difference invisible.

4.6.2 The Envelope and the Back-Transition

Case 3 starts wingborne at 60 m/s and 110 m. Flown instead at a constant 60, 42 or 66 m/s and averaged from 60 to 150 s, the aircraft draws 139, 149 and 165 kW with the wing at 4.1, 9.0 and 3.0°. At 42 m/s the wing alone would need 10.6°, within 0.4° of the protection, so the allocation keeps it at the preferred 9.0°, and the rotors, restarted, carry 13 % of the weight on 52 kW of which 50 kW is profile power. Slowing from 60 to 42 m/s therefore costs power, 10 kW more, instead of saving it. On the way down the restart comes at 23.7 s at 44.3 m/s, and on the way up the rotors stop again only at 78.2 s at 51.1 m/s. The rotors therefore restart below the 44.6 m/s gate and stop only above the 48.3 m/s one, and between the two they keep whatever state they had, which is the hysteresis of Equation (4.8) and of the safety monitors of Section 1.4.9.

🕮 **Pause.** Try Exercise 4.2 before reading on.

Case 4 starts wingborne at 60 m/s and 50 m, 3.2 km before pad B. Braking starts at 27.6 s, 1.54 km out, and the rotors spin up at 31.5 s at 55.9 m/s, while the wing still carries the whole weight. As the aircraft slows they take what the wing no longer makes and carry half at 51.2 s at 32.7 m/s. With the pusher at its 1.5 kN reverse limit the allocation brakes by pitching up to 5.5°, and the power climbs back to 464 kW by 77.4 s as the aircraft comes to a hover over the pad, where the gate begins at 79.0 s. The 40.5 s descent touches down at 121.6 s at 0.57 m/s, 0.03 m from the center, after 8.9 kWh, and the angle-of-attack margin never falls below 4.1°.

4.6.3 A Gust and a Heavier Load

Case 5 adds a gust from behind at 37 s, 10 m/s along the route and 1 m/s down over 10 s. Before it the rotors carry 9 % of the weight at 42.8 m/s. With the feedback transition the allocation simply gives the rotors back what the wing loses, 25 % of the weight at 42.19 s, the gate stays shut, and the rotors stop at 46.65 s at 54.9 m/s, 3.67 s later than in calm air, with the aircraft never more than 1.72 m below its reference. The timed schedule stops them on the clock at 42.00 s, at 39.7 m/s with a

capacity of 91 %. The vertical demand goes unmet for 3.9 s from 40.9 s, as the ceiling comes down on rotors that the gust is loading again, the aircraft falls 6.7 m below its reference, and the worst angle-of-attack margin shrinks to 2.4°, against 3.8° with the feedback transition.

Pause. Try Exercise 4.1 before reading on.

> **TRY IT IN THE APP 5 · Gust During the Handover**
>
> Load the case and press Run with the Feedback transition. Watch the rotor lift arrows and the Vertical force trend when the gust arrives at 37 s. The rotors take back 25 % of the weight, and they stop only at 46.65 s, at 54.9 m/s, with the aircraft never more than 1.72 m below its reference.

Case 6 carries 470 kg instead of 250 kg, 2220 kg in all. Case 1 flown with that payload hovers on 495 kW, and the estimated stall speed rises from 38.4 to 40.4 m/s. Because the gates are computed from the load sheet, the stop gate opens at 51.0 m/s and the rotors stop at 44.74 s at 53.1 m/s. The timed schedule stops them at 42.00 s at 49.8 m/s, as it would with any payload, and in calm air its capacity of 129 % is still enough. The same schedule meets the gust of case 5 in the challenge.

4.6.4 A Biased Pitot

In case 7 the pitot starts reading 12 m/s high at 30 s, when the true airspeed is 34.7 m/s, so the allocation is told that the wing makes $(46.5/34.7)^2 = 1.8$ times its real lift. With the cross-check on, the lift-derived airspeed disagrees at once. From 30.73 s the allocation trusts the lower of the two, at 32.72 s the pitot is flagged, and from then on the supervisor flies on V_L. The aircraft never falls more than 3.0 m below its reference, and the rotors stop at 41.46 s at a true 49.3 m/s while the pitot reads 61.2 m/s. With the cross-check off, the rotor collective drops from 8.25 to 1.07 kN within half a second of the fault, the aircraft sinks at up to 5.1 m/s, and it comes down on the wood at 38.5 s, 856 m along the route. The Z lamp never lights in that run, because by the allocation's own arithmetic every demand is met.

Pause. Try Exercise 4.3 before reading on.

4.6.5 A Failed Motor

Case 8 flies inbound on the wing with 470 kg. At 15 s one of the four rear motors fails, and the rear group can give 12.0 kN against the 11.45 kN it would carry in a level hover. The lift-group margin falls from 39.7 to 4.8 %, below the 10 % rule, and at 15.00 s the supervisor declares the vertical landing not feasible and diverts to the airfield on the wing, never asking the rear group for thrust it does not have. With 250 kg the rear share is 10.32 kN and the margin 16.3 %, and the aircraft lands on pad B at 151.2 s at 0.48 m/s and 0.11 m from the center, with the rear group working at up to 93 % of what it has left.

4.6.6 The Graded Mission

The challenge lifts off from pad A with 470 kg and the battery at 40 %, begins the transition at 27.11 s and meets the case 5 gust at 64 s, 36.9 s into the transition. It is met with every phase completed, at least 30 s wingborne with the rotors stopped, 25 m of clearance away from the pads, a 2° angle-of-attack margin above 20 m/s, a touchdown slower than 1 m/s no more than 2 m from the pad center, and 20 % of charge or more when it lands, each judged on its worst case. The mission is judged on the case as given, so a changed model parameter or disturbance voids the attempt, and so does an extra disturbance event fired by the learner. The configurations are compared in Table 4.8, and Figure 4.6 draws its distances and stop airspeeds as bars.

As shipped, the timed schedule stops the rotors at 69.11 s, just as the gust has taken the true airspeed down to 37.9 m/s, below the 40.4 m/s stall speed of the loaded aircraft, where even the wing at its lift peak would lift only 88 % of the weight and the protected wing 75 %. The vertical demand had already gone unmet from 67.19 s, as the collective ceiling came down on rotors that the gust

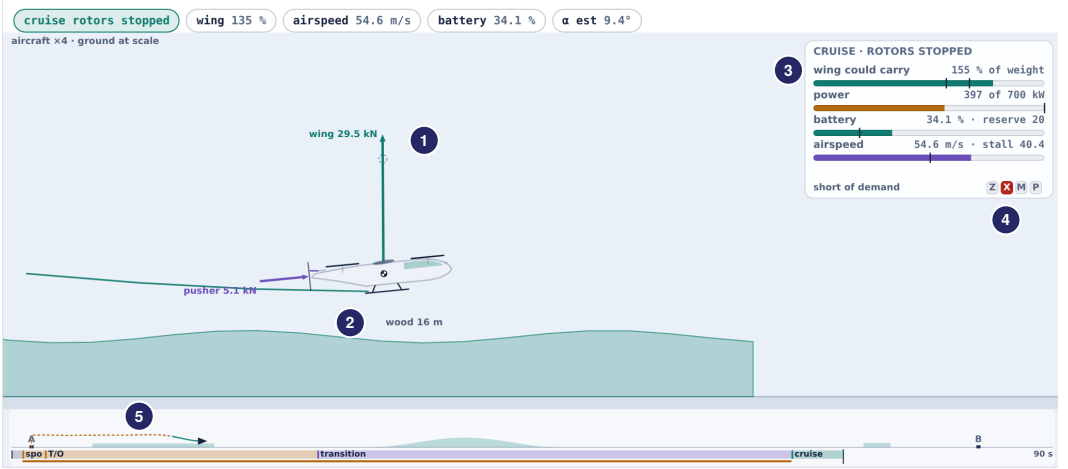

Figure 4.5 The graded mission, Challenge · the full mission, as shipped with the timed transition, paused at $t = 1.23$ min (73.7 s), the lowest point of the flight, 4.6 s after the schedule stopped the rotors, cropped to the plant view. The aircraft (1), drawn four times its size with stopped rotor blades along its booms, is in cruise at 25.4 m, 9.3 m above the wood (2), pitched 5.5° with its wing lifting 29.5 kN, 135 % of the weight, as it pulls out of the sink, and the pusher gives 5.1 kN. The trail in the side view is solid because the rotors were stopped over all of it, and in the route map below it is dotted from pad A to the stop over the wood and solid through the dip (5). The plate reads cruise with the rotors stopped, a wing that could carry 155 % of the weight at the protected angle (3), 397 of 700 kW, 34.1 % of charge and a trusted 54.6 m/s against a 40.4 m/s stall estimate, and the X lamp (4) is lit because the speed loop asks for more than the pusher's 5.5 kN. The trends and the scorecard of the full window are cut.

Table 4.8 Graded Mission at the Default Seed

	Timed	Feedback
rotor stop (s)	69.11	74.42
true airspeed at the stop (m/s)	37.9	55.7
lowest altitude after the gust (m)	25.4	48.1
largest loss below the reference (m)	31.7	1.9
minimum clearance (m)	9.3	32.0
worst angle-of-attack margin (°)	2.2	3.6
longest wingborne interval (s)	80.2	72.8
hover gate over pad B (s)	195.3	195.3
farthest past the pad center (m)	0.21	0.21
touchdown (s)	238.1	238.1
sink (m/s)	0.43	0.43
distance from the center (m)	0.13	0.13
charge at touchdown (%)	25.7	25.6
energy used (kWh)	21.5	21.6
verdict	fail	pass

was loading again. The aircraft sinks at up to 6.2 m/s, from 46.5 m at the stop to 25.4 m at 73.7 s, 9.3 m above the trees, which fails the clearance limit (Table 4.8); Figure 4.5 shows the aircraft at that lowest point. The allocation lowers the nose to keep the angle of attack at its protection as the flight path steepens to −8.4°, and the true angle peaks at 11.8°, so the aircraft sinks rather than stalls. The altitude loop, at its 5 m/s^2 limit and up to 31.7 m behind a reference that climbs toward the cruise



Figure 4.6 The graded mission of Table 4.8 at the default seed, drawn as bars for the timed (solid) and the feedback (hatched) transition. Notice that the timed stop comes below the stall speed of the loaded aircraft, 40.4 m/s (right, dashed line), and that the same flight loses 31.7 m against its reference and passes 9.3 m above the trees, inside the 25 m clearance limit (left, dashed line). The feedback flight keeps 32.0 m of clearance and loses no more than 1.9 m.

altitude, then carries the aircraft back up through that reference to 13.1 m above it at 82.5 s. The overshoot comes from the integral term. Back-calculation with $T_t = \tau_I$ slows it without stopping it, so it grows from 0.09 to 1.33 m/s^2 while the demand is held at the clamp from 69.4 to 76.1 s and peaks at 2.35 m/s^2 at 79.1 s. The rest of the mission is clean, with a 2.2° margin, a touchdown at 0.43 m/s and 0.13 m, and 25.7 % left after 21.5 kWh.

> ⧖ HISTORY **A Conversion Through the Treetops (1979)**
>
> Six days after that first conversion, a test flight called for a maximum acceleration from a low hover over the runway. The nacelles were tilted forward at the maximum rate, at which a conversion from hover to airplane mode can be completed in 12 seconds. The rapid conversion lowered the rotor thrust before the aircraft had the speed for the wing to lift, and the aircraft settled slightly. The crew did not notice the settling until they neared the end of the airfield, and with barely enough speed to climb the aircraft passed through the tops of the trees beyond it. It landed essentially unscathed, with green stains on its blade tips [19]. The timed transition of the graded mission takes the rotor lift away in the same way, before the wing can carry the weight.

The suggested configuration, the feedback transition, keeps the rotors turning through the gust and stops them at 74.42 s at 55.7 m/s. The aircraft never falls more than 1.9 m below its reference, keeps at least 32.0 m of clearance and a 3.6° margin, and flies 72.8 s wingborne. Both flights begin braking into pad B at 143.79 s. The feedback transition spins the rotors up at 147.67 s at 55.9 m/s, and they are up to speed at 53.5 m/s, where the protected wing could still lift 149 % of the weight, while the along-track reference brakes on its jerk-limited profile without waiting for them. The aircraft comes to rest over the pad at most 0.21 m past its center, never drifting back faster than 0.27 m/s, and touches down at 238.1 s at 0.43 m/s, 0.13 m from the center, with 25.6 % left after 21.6 kWh. From 153.2 s on, its position along the route stays within 1 cm of the timed flight's, and the two touchdowns agree to 4 µm. Up to the braking it has used 12.51 kWh against the timed run's 12.49 kWh, and it spends 9.04 kWh from there to the touchdown against 9.01 kWh.

Since the reference never waits, the spin-up airspeed alone decides how much of the wing's margin is left when the rotors take over the weight. Every V_{spin} from 41.5 to 70 m/s, the top of its range, in steps of 0.5 m/s, passes the mission with the same touchdown and a stop within 0.21 m of the pad center, with a worst angle-of-attack margin of 2.2 to 3.6°. The earlier the spin-up, the shorter the wingborne interval, from 84.3 s down to 68.95 s. At 41 m/s the rotors are still spinning up as the

airspeed falls through the 40.4 m/s stall speed of the loaded aircraft, and the margin falls to 1.2°, which fails the mission. Holding the reference's speed while the rotors spin up, as an earlier version of the supervisor did below $1.35\hat{V}_s$, drops its deceleration to zero in the middle of the braking and leaves it to rebuild at the jerk limit. The reference then passes the pad, by 100 m at the default V_{spin} and by 911 m at 52 m/s, and has to fly the aircraft back to it tail first.

Twenty other seeds, 1 to 20, give the same verdicts. The timed schedule stops the rotors at 38.2 to 40.8 m/s and fails 19 of them, 18 on clearance with 10.0 to 24.8 m and 4 on the angle-of-attack margin, and it passes once, with exactly 25.0 m. The feedback transition passes all 20, with 30.4 to 31.9 m of clearance and stops at 55.0 to 57.1 m/s. In every one the aircraft comes to rest over pad B at most 0.43 m past its center and touches down at 0.45 to 0.60 m/s with 25.6 to 25.9 % left.

> ✦ DID YOU KNOW? **Safety Costs Almost No Energy**
>
> From the simulator: in the graded mission the feedback transition flies 72.8 s wingborne against the timed schedule's 80.2 s, and it lands after 21.6 kWh against 21.5 kWh. Only the feedback flight keeps the 25 m clearance, with 32.0 m against 9.3 m (Table 4.8).

4.7 Technology in Practice

The aircraft of Section 4.3 is read by four instruments and moved by two lift-rotor groups, a pusher and an elevator. Figure 4.7 draws the chapter's loops with real devices in their place.

What Measures It

Every gate of Section 4.5 rests on the airspeed, and the airspeed rests on the pitot PT of Figure 4.1. Its 4 Pa of noise and the bias of case 7 stand for the errors of a real probe, and the lift-derived airspeed V_L stands for the independent source a real system compares it with. An unmanned tilt-wing research aircraft listed a pitot-static tube with air data and a combined satellite and inertial unit among its sensors [17], the same set as GN, IMU and PT here.

> TECHNOLOGY **Sensor: Pitot-Static Airspeed**
>
> A pitot tube facing the airflow senses the total pressure, static ports sense the ambient pressure, and the airspeed indicator shows the difference, the dynamic pressure. A wind blowing into the tube makes the same dynamic pressure as motion through still air. Errors almost always mean a blocked tube or port, from moisture, ice, dirt or insects. With the tube blocked and its drain hole open, the indicated airspeed falls to zero, quickly but not at once. With both blocked, it stops following the airspeed and changes with altitude instead [20]. Near hover the output of an airspeed sensor becomes untrustworthy [17], and the chapter's pitot, too, is poor below about 10 m/s.

What Moves It

Each lift group of the model lumps four rotors into one thrust with a lag of 0.15 s, a rate limit of 25 kN/s and a 2 s spin-up, and the supervisor stops and starts the group as a whole. On a real aircraft of this type each lift rotor is driven by its own electric motor, with no mechanical link between the drive systems [3]. A failed motor removes one rotor's thrust, as in case 8, where one of the four rear motors fails.

> TECHNOLOGY **Final Element: Fixed-Pitch Lift Rotors**
>
> Most lift-plus-cruise aircraft proposed have fixed-pitch lift rotors whose thrust is set by their speed [15], with no collective or cyclic pitch. Two-bladed rotors can be indexed to stop aligned with the direction of flight for a small drag penalty, and more than six rotors let controlled flight continue after a failure [3]. With rotor speed control the electric motor sits directly in the control path. A control design to Level 1 handling qualities asked a lift-plus-cruise concept for motor power of 2.15 times its hover power, with heave the most demanding

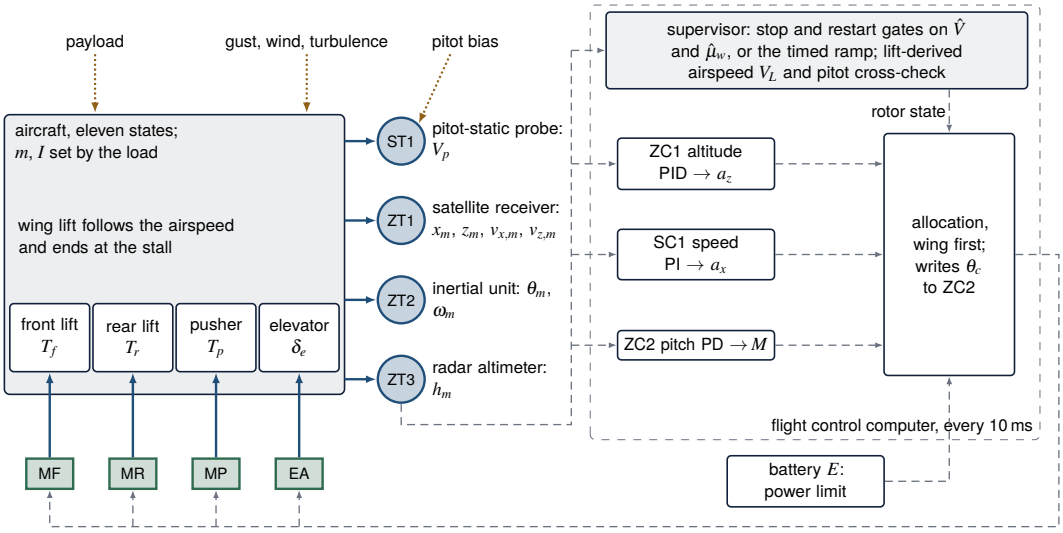

MF, MR: lift motor drives, lag 0.15 s, spin-up 2 s, up to 16 kN a group. MP: pusher drive, −1.5 to 5.5 kN. EA: elevator actuator, −25 to 20°.

Figure 4.7 The loops of the chapter with real devices in place of the schematic's symbols. ST1 reads the airspeed from a pitot-static probe, ZT1 the position and velocities from a satellite receiver, ZT2 the pitch and its rate from the inertial unit and ZT3 the height from a radar altimeter. One flight control computer (dashed frame) runs the supervisor, the altitude, speed and pitch loops (ZC1, SC1, ZC2) and the allocation every 10 ms. MF and MR are the motor drives of the front and rear lift groups, MP drives the pusher and EA moves the elevator. Notice that the drives of the lift rotors take their commands only from the allocation, and the supervisor reaches them only by setting the rotor state.

axis [15].

Where the Controller Runs

The automatic flight and transition system of the tilt-wing research aircraft ran on a flight control computer in layers: a primary flight control system with stability augmentation, whose gains changed with the wing configuration and the airspeed, an attitude loop above it, guidance loops for altitude, rate of climb, speed and heading, and an auto-transition loop that commanded the configuration [17]. The controller of this chapter has layers of the same kind. Its pitch loop is a PD block, the altitude and speed loops are velocity-form PID blocks (Section 1.6.2), the allocation sits between the loops and the effectors, and the supervisor decides the rotor state. All of them run every 10 ms.

What Is Modernizing It

Three developments bear on the chapter's question of when the rotors may stop. The first is synthetic airspeed. Near hover the research aircraft above replaced its untrustworthy air data with the ground speed below 1.0 m/s [17], and the lift-derived airspeed of this chapter is a cross-check of the same family. The second is unified control, which blends optimal control with a generalized control allocation into a single mode from hover to cruise, without effector sets prescribed for each phase [21]. The third is optimized reference commands for the transition. A cost on the rotor thrusts lets the rotors operate only when necessary and leaves the rest to the wing [22], much as the wing-first rule of Section 4.5 does by construction.

> **IN PRACTICE** **Approach Speed Margins Above the Stall**
>
> Pilots read their stall margins off the airspeed indicator. The low end of its white arc marks the stall speed in the landing configuration, and the low end of the green arc the stall speed in a specified configuration, usually clean at maximum takeoff weight [20]. United States flight training guidance flies the approach at the maker's recommended speed, or at 1.3 times the landing stall speed when there is none, and pilots often add half the gust factor in turbulent air [23]. The gates of this chapter keep the same kind of margin: the stop gate opens at 1.258 times the stall speed, behind an airspeed backstop at 1.2 times.

4.8 Limitations

No real aircraft is modeled, so the numbers say nothing about the performance, energy or safety margins of an actual design. Without roll, yaw or sideslip there are no turns and no crosswind, and each lumped rotor group has neither inflow dynamics, flapping and H-force nor any interaction with the wing. The aerodynamics are quasi-steady, and the stability of the aircraft near the stall is not assessed [16]. The energy reserve is a dispatch rule and a scored limit rather than a reason to divert, so a flight that will land below it is still flown to the pad. The controller's aerodynamic model equals the plant's, and the load sheet is exact, so the gates are as accurate as the airspeed they are given. The pitot cross-check works only above 30 m/s, where the wing carries enough to measure the airspeed by its lift. On the way down the deceleration does not wait for the lift rotors, so a spin-up airspeed set too low leaves the wing near its protected angle until they are up to speed, and at 2220 kg a V_{spin} of 41 m/s or less fails the 2° angle-of-attack margin. The point at which each assumption fails is given in Section 4.3.

> **OPEN CHALLENGE** **Proving a Transition Before It Flies**
>
> In transition the effectors of hover and of wing-borne flight act together. Planning the reference commands and tracking them are both open problems, because redundant effectors leave it unclear how to command each one uniquely [22]. Designs have broken the envelope into phases with a control approach for each, and single control laws from hover to cruise are still examined in simulation [21]. With little flight experience of electric multirotor aircraft, new and rigorous requirements for their response to failures are needed [15]. Progress would be transition laws and failure requirements shown to hold in flight, gusts and heavy loads included.

4.9 Conclusion

A lift-plus-cruise aircraft must earn the stop of its lift rotors, and this chapter put that decision under test on a simulator. Its planar aircraft hovers on 431 kW and cruises on the wing on 139 kW, and closed-form speeds from the weight and the wing place the stop and restart gates within 0.3 m/s of the engine. The lift the wing could make at its protected angle at the moment of the stop predicts the sink, from at most 1.2 m of lost height at 129 to 174 % of the weight to 21.1 m at 75 %. A schedule tuned for 250 kg in calm air stops the rotors at 49.8 m/s there and at 37.9 m/s in the graded mission, below the stall speed of the loaded aircraft, which comes down to 9.3 m above the trees, while the measured gates stop them at 55.7 m/s and keep 32.0 m of clearance. On the way down the aircraft brakes to rest over pad B without waiting for the rotors, which are turning again at 53.5 m/s. A lift-derived airspeed catches a pitot reading 12 m/s high in 0.73 s, and a failed motor that leaves a 4.8 % margin turns the landing into a diversion. The model is planar and quasi-steady, and its controller knows the aerodynamics exactly. A fuller model would add turns and crosswind, rotor and wing interference, dynamic stall, an estimated rather than declared mass and a predictive transition that plans the handover within the envelope. The simulator is a browser plant on which a change of mode granted by a clock and one earned by a measurement can be flown over the same wood and compared.

KEY POINTS

- The lift rotors cost far more power than the wing, so they stop once the wing can carry the weight.
- Closed-form speeds from the weight and the wing place the stop and restart gates within 0.3 m/s of the engine.
- The wing's lift capacity at the stop predicts the sink: at most 1.2 m of lost height at 129 to 174 % of the weight, 21.1 m at 75 %.
- In the graded mission a schedule tuned in calm air stops the rotors at 37.9 m/s, below the stall speed of the loaded aircraft, while the measured gates stop them at 55.7 m/s.

Where else this applies. Volume 8, Chapter 11 also ends a step on a measurement rather than on the clock, and Volume 4, Chapter 12 shares one demand among effectors with hard limits.

References

[1] T. A. Johansen and T. I. Fossen, "Control allocation—a survey," *Automatica*, vol. 49, no. 5, pp. 1087–1103, 2013.

[2] D. Liberzon and A. S. Morse, "Basic problems in stability and design of switched systems," *IEEE Control Systems*, vol. 19, no. 5, pp. 59–70, 1999.

[3] C. Silva, W. R. Johnson, E. Solis, M. D. Patterson, and K. R. Antcliff, "VTOL urban air mobility concept vehicles for technology development," in *2018 Aviation Technology, Integration, and Operations Conference*. American Institute of Aeronautics and Astronautics, 2018.

[4] A. Kasliwal, N. J. Furbush, J. H. Gawron, J. R. McBride, T. J. Wallington, R. D. De Kleine, H. C. Kim, and G. A. Keoleian, "Role of flying cars in sustainable mobility," *Nature Communications*, vol. 10, no. 1, p. 1555, 2019.

[5] A. Keller, F. Gandhi, R. Niemiec, and A. Walter, "Full flight regime controller design for a Lift+Cruise eVTOL aircraft," *Journal of the American Helicopter Society*, vol. 71, no. 3, pp. 1–22, 2026.

[6] B. W. Bequette, "Process control practice and education: Past, present and future," *Computers & Chemical Engineering*, vol. 128, pp. 538–556, 2019.

[7] J. A. Rossiter, B. Pasik-Duncan, S. Dormido, L. Vlacic, B. Jones, and R. Murray, "A survey of good practice in control education," *European Journal of Engineering Education*, vol. 43, no. 6, pp. 801–823, 2018.

[8] D. H. Jacobson, "An evaluation of lift engines in tactical VTOL aircraft," *Journal of Aircraft*, vol. 2, no. 5, pp. 424–429, 1965.

[9] V. Jackson and D. Patch, "Design and development of the primary fly-by-wire flight control system for the XV-4B V/STOL research aircraft," in *Control and Flight Dynamics Conference*. American Institute of Aeronautics and Astronautics, 1968.

[10] J. Magee and K. Wernicke, "XV-15 tilt rotor research aircraft - program report," in *Atlantic Aeronautical Conference*. American Institute of Aeronautics and Astronautics, 1979.

[11] H. Glauert, "The theory of the autogyro," *The Journal of the Royal Aeronautical Society*, vol. 31, no. 198, pp. 483–508, 1927.

[12] R. W. Beard and T. W. McLain, *Small Unmanned Aircraft: Theory and Practice*. Princeton: Princeton University Press, 2012.

[13] W. C. Durham, "Constrained control allocation," *Journal of Guidance, Control, and Dynamics*, vol. 16, no. 4, pp. 717–725, 1993.

[14] J. M. Buffington and D. F. Enns, "Lyapunov stability analysis of daisy chain control allocation," *Journal of Guidance, Control, and Dynamics*, vol. 19, no. 6, pp. 1226–1230, 1996.

[15] W. Johnson and C. Silva, "NASA concept vehicles and the engineering of advanced air mobility aircraft," *The Aeronautical Journal*, vol. 126, no. 1295, pp. 59–91, 2022.

[16] J. P. Spadão, R. M. G. Vasconcellos, M. Sartorato, and W. M. dos Santos, "Dynamic stability and control authority blending in lift-plus-cruise eVTOL transition flight," *Dynamics*, vol. 6, no. 2, p. 21, 2026.

[17] K. Muraoka, T. Yokoyama, N. Yokoyama, and M. Sato, "Flight verification of automatic flight and transition system for lift/cruise thrust type VTOL aircraft," *Transactions of the Japan Society for Aeronautical and Space Sciences*, vol. 67, no. 5, pp. 285–294, 2024.

[18] F. Larocque, T. De Ponti, B. Remes, and E. Smeur, "Synthetic air data system for pitot tube failure detection on the variable skew quad plane," in *AIAA SCITECH 2025 Forum*. American Institute of Aeronautics and Astronautics, 2025.

[19] M. D. Maisel, D. J. Giulianetti, and D. C. Dugan, *The History of the XV-15 Tilt Rotor Research Aircraft: From Concept to Flight*, ser. Monographs in Aerospace History. Washington, DC: National Aeronautics and Space Administration, NASA History Division, 2000, no. 17, nASA SP-2000-4517, https://ntrs.nasa.gov/citations/20000027499.

[20] Federal Aviation Administration, "Pilot's handbook of aeronautical knowledge," FAA-H-8083-25B, Chapter 8, Flight Instruments, https://www.faa.gov/sites/faa.gov/files/regulations_policies/handbooks_manuals/aviation/phak/10_phak_ch8.pdf, 2016.

[21] M. J. Acheson, I. M. Gregory, and J. Cook, "Examination of unified control incorporating generalized control allocation," in *AIAA Scitech 2021 Forum*. American Institute of Aeronautics and Astronautics, 2021.

[22] J. L. Bullock, S. Cheng, A. Patterson, M. J. Acheson, N. Hovakimyan, and I. M. Gregory, "Reference command optimization for the transition flight mode of a lift plus cruise vehicle," in *AIAA SCITECH 2024 Forum*. American Institute of Aeronautics and Astronautics, 2024.

[23] Federal Aviation Administration, "Airplane flying handbook," FAA-H-8083-3C, Chapter 9, Approaches and Landings, https://www.faa.gov/sites/faa.gov/files/regulations_policies/handbooks_manuals/aviation/airplane_handbook/10_afh_ch9.pdf, 2021.

4.10 Exercises

4.1 ⌕ Concept check · 2 min

In scenario 5 a gust from behind arrives at 37 s, during the handover. With the timed transition the rotors stop at 42.00 s. At that instant the ground speed is 49.6 m/s, close to the 49.7 m/s of scenario 2 in calm air at the same instant, but the true airspeed is 39.8 m/s. Why has the wing lost lift?

(a) The gust slows the aircraft over the ground, and the lift follows the ground speed

(b) The downdraft of the gust tilts the lift backward without changing its size

(c) The pitot reads low in the gust, so the allocation lowers the angle of attack

(d) Lift follows the velocity relative to the air, and wind from behind subtracts from it

4.2 ⌕ Concept check · 1 min

In scenario 3 the cruise speed is cut to 42 m/s and later raised to 66 m/s. The supervisor restarts the rotors at 44.3 m/s on the way down and stops them only at 51.1 m/s on the way up. What is the gap between the restart margin $\mu_{off} = 15\,\%$ and the stop margin $\mu_{on} = 35\,\%$ for?

(a) It keeps turbulence near one airspeed from starting and stopping the rotors again and again

(b) It gives the rotors their 2 s of spin-up before the wing runs short of lift

(c) It corrects for the pitot, which reads differently while slowing and while accelerating

(d) It keeps the rotors stopped for as long as possible to save their profile power

4.3 ⌕ Concept check · 2 min

In scenario 7 with the airspeed cross-check off, the pitot starts reading 12 m/s high at 30 s. The aircraft sinks into the wood at 38.5 s, yet the Z lamp of vertical saturation never lights. Why?

(a) The lamp is switched off together with the cross-check, so no fault can be shown

(b) The allocation sizes the wing's share with the biased airspeed, so by its own arithmetic every demand is met

(c) The altitude error stays too small for the loop to ask for more lift than the wing makes

(d) The bias stops the rotors, so neither group can reach a limit that would light the lamp

4.4 ⌨ Theory and calculation · 7 min

Each lift group sits $\ell = 2.9\,\mathrm{m}$ from the rotors' centroid, and the center of mass lies 0.15 m behind the centroid ($c = -0.15\,\mathrm{m}$). The aircraft has a mass of 2000 kg with $g = 9.81\,\mathrm{m/s^2}$. Each group has a disk area $A = 20\,\mathrm{m^2}$, a figure of merit $FM = 0.75$ and a profile power $P_0 = 25\,\mathrm{kW}$, and $\rho = 1.225\,\mathrm{kg/m^3}$.

(a) Find the thrust of each group in a level hover in still air, and the share of the weight the rear group carries. *Checkpoint:* $T_f = 9.30\,\text{kN}$, $T_r = 10.32\,\text{kN}$, 52.6 %.

(b) Find the power the two groups draw in that hover from Equation (4.5). *Checkpoint:* 420.5 kW.

(c) Find the ideal induced power $\sum T_i v_h$ and the effective figure of merit, the ideal power divided by the power of part (b). *Checkpoint:* 277.9 kW and 0.661.

4.5 Theory and calculation · 7 min

Use $C_{L,max} = 1.359$ at the lift peak, $C_L = 1.159$ at the protected angle of 11°, $S = 16\,\text{m}^2$, $\rho = 1.225\,\text{kg/m}^3$ and $g = 9.81\,\text{m/s}^2$. The aircraft carries 470 kg of payload, 2220 kg in all.

(a) Find the stall speed V_s and the protected speed V_p of Equation (4.8). *Checkpoint:* 40.4 m/s and 43.8 m/s.

(b) Find the stop gate V_{on} and the restart gate V_{off} for $\mu_{on} = 35\,\%$ and $\mu_{off} = 15\,\%$. *Checkpoint:* 50.9 m/s and 47.0 m/s.

(c) The timed schedule of the graded mission stops the rotors at a true airspeed of 37.9 m/s. Find the share of the weight the wing can lift at that speed at the protected angle and at the lift peak. *Checkpoint:* 75 % and 88 %.

4.6 Theory and calculation · 5 min

A lift group is rated at 16 kN with four motors. With 2220 kg on board, one of the four rear motors fails, so the rear group can give three quarters of its rating. Use the hover shares of the arms $\ell_f = 3.05\,\text{m}$ and $\ell_r = 2.75\,\text{m}$, $g = 9.81\,\text{m/s}^2$, and the landing rule $\min_i(a_i T_{max}/T_{i,trim}) - 1 \geq 10\,\%$.

(a) Find the margin of the weaker group after the failure. *Checkpoint:* 4.8 %, the rear group.

(b) Find the largest total mass that still meets the rule, and the largest payload on the 10 kg grid of the Payload control for an empty mass of 1750 kg. *Checkpoint:* 2114.7 kg, so 360 kg.

4.7 App challenge · 8 min

Load scenario 5, Gust during the handover, with the default seed, and run at speed 5. Use the scorecard tiles Rotors stopped at, Peak altitude loss and Worst angle-of-attack margin.

(a) Run as shipped, with the Feedback transition, to 60 s, and read the three tiles. *Checkpoint:* 54.9 m/s, 1.7 m and 3.8°.

(b) Press Reset, set Transition to Timed before Run, run to 60 s and read the same tiles. *Checkpoint:* 39.7 m/s, 6.7 m and 2.4°.

(c) In the exported file of the second run, read `speed_mps` and `airspeed_mps` at 42.00 s, within 0.2 m/s. Then find the share of the weight the wing can lift at the protected angle at the stop airspeed, with $V_p = 41.6\,\text{m/s}$ at 2000 kg. *Checkpoint:* 49.6 m/s and 39.8 m/s, a capacity of 91 %.

4.8 App challenge · 8 min

Load scenario 7, Airspeed bias, with the default seed, and run at speed 5.

(a) Run as shipped, with the cross-check On, until the Transition logic card reports ROTORS STOPPED. Read the Rotors stopped at and Peak altitude loss tiles, and find in the exported file the first row in which `airspeed_disagree` is 1. *Checkpoint:* 49.3 m/s with the pitot at 61.2, 3.0 m, and the row at 32.75 s.

(b) Press Reset, set Airspeed cross-check to Off before Run, and run until the run stops by itself. Read the Mission tile, and find in the file the first row in which `phase` reads `crashed`, with its time and `x_m` within 1 m. *Checkpoint:* Crashed, at 38.55 s and 856 m.

(c) In the same file, read `rotor_lift_z_kN` at 30.00 and 30.50 s, and explain the change. *Checkpoint:* 8.23 and 1.07 kN.

4.9 App challenge · 8 min

Load the challenge, the full mission, with the default seed, and run at speed 10. Each flight lasts about 240 s and pauses itself after the landing.

(a) Fly it as shipped, with the Timed transition. Read the Min clearance, Rotors stopped at, Peak altitude loss and Worst angle-of-attack margin tiles, and say which part of the target fails. *Checkpoint:* 9.3 m, 37.9 m/s, 31.7 m and 2.2°. The clearance fails.

(b) Press Reset, set Transition to Feedback before Run, and fly the whole mission. Read Min clearance, Worst angle-of-attack margin, Wingborne, rotors stopped, Touchdown and Battery at the end. *Checkpoint:* 32.0 m, 3.6°, 73 s longest, 0.43 m/s at 0.13 m, 25.6 %.

4.10 Application area · 4 min

The engine measures 431 kW in hover at 2000 kg, 139 kW in wingborne cruise at 60 m/s, and 149 kW at 42 m/s, where the restarted rotors carry 13 % of the weight on 52 kW, of which 50 kW is profile power.

(a) Find the energy of one minute of hover and of one minute of cruise at 60 m/s, and the energy per kilometer of cruise at 60 and at 42 m/s. *Checkpoint:* 7.2 and 2.3 kWh, then 0.64 and 0.99 kWh per kilometer.

(b) A dispatcher proposes flying the approach at 42 m/s to save energy, because the airframe drag falls with the square of the speed. Say why the power rises instead. *Checkpoint:* The wing needs more angle than the allocation allows, so the rotors restart and their profile power outweighs the drag saved.

4.11 Application area · 4 min

The chapter's controller knows the payload exactly from the load sheet, and its pitot cross-check works only above 30 m/s.

(a) Suppose the load sheet says 250 kg while 470 kg is on board. The stop gate is computed from the sheet. Where does it open, and what share of the true weight can the wing lift there at the protected angle? Use the speeds of Equation (4.8) at 2000 and 2220 kg. *Checkpoint:* At 48.3 m/s, 122 % instead of the intended 135 %.

(b) Why can the cross-check not work at low airspeed, and what does a practicing engineer have to add for that part of the flight? *Checkpoint:* The wing makes too little lift to infer an airspeed from it.

4.12 Review the agent · 6 min

An AI assistant reports: "Raising the altitude loop's Integral time from 12 to 60 s cures the challenge. The sink over the wood was integral windup in the altitude loop, so the shipped Timed transition is safe and the measured gates are not needed." The challenge fires its gust 36.9 s after the transition starts. Test the claim with the default seed at speed 10, then find the flaw. Each flight pauses itself after the landing, near 256 s.

(a) Load the challenge, set Integral time to 60 s before Run, keep the Timed transition and fly the mission. Read the Min clearance, Rotors stopped at and Worst angle-of-attack margin tiles. *Checkpoint:* 19.4 m, 38.9 m/s and 2.6°. The clearance fails.

(b) Press Reset, keep Integral time at 60 s, set Transition to Feedback before Run and fly it again. Read Min clearance and the verdict. *Checkpoint:* 31.3 m, and every part passes.

CHAPTER 5

Flying a Foiling Boat Between a Slamming Hull and a Ventilating Foil

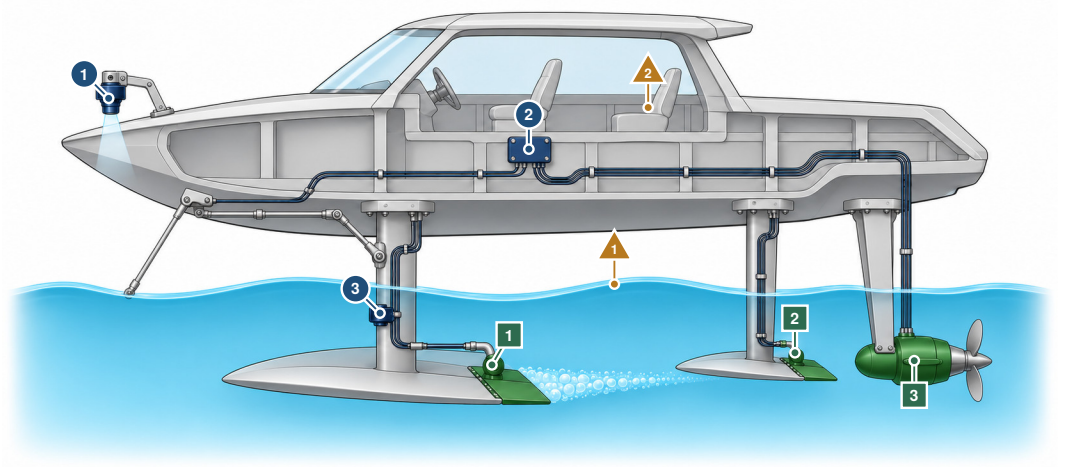

The controller moves 1 flap on the main foil 2 flap on the rear foil 3 pod motor and propeller

The controller reads 1 ultrasonic height sensor at the bow 2 inertial unit 3 speed log

What disturbs it 1 waves 2 a passenger walking aft

Notice the narrow band between the hull above the waves and the main foil below them, and the bubbles behind that foil, a hint of the air it draws when it flies too high. Illustration, not to scale.

IN PLAIN TERMS **Flying a Boat on Underwater Wings**

The main idea. A foiling boat rides on small wings under the water, so its hull flies above the waves. A controller tilts a flap on the main wing to hold the boat at the right height.

The challenge. Too low, the hull smacks into a wave, and too high, air reaches the wing and takes most of its lift. The controller must choose which waves to follow and which to ride over, and big waves are a reason to stop flying.

Why it matters. A wrong choice ends the flight in a crash, and a right one gives a smooth ride.

THREE WAYS TO READ THIS CHAPTER

Explore	12 min	the first page, the fact on a racing dinghy in the Introduction, the two history segments (Section 5.2), the try-it segments of cases 6 and 9 and the fact on the challenge (Section 5.6), 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 measured cases (Section 5.6), Technology in Practice (Section 5.7), Limitations (Section 5.8) and the open challenge

The control question A foiling boat flies in a band fixed by its struts, with a slamming hull below and a ventilating foil above, while the waves it meets are not fixed at all. Which waves should its height loop ride over level and which should it follow, and when is the right decision to stop flying?

Learning objectives

- Derive the free height stiffness from the surface factor and explain why a deep foil leaves the height an integrator.
- Compute the encounter frequency of a sea met from astern and from ahead, and use it with the band to choose between platforming and contouring.
- Explain a mechanical wand as a proportional loop whose linkage ratio is its gain and whose length is its setpoint.
- Compare a PID, a heave-rate cascade and an MPC with a preview on the same sea, and judge what the preview and the retuning are each worth.
- State the loop action of the height and pitch loops in the convention of Chapter 1.

Prerequisites Sections 1.3 (loop action), 1.4.3 (nested loops), 1.4.6 (complementary filter with a bias state, wave motion), 1.6.2 (velocity-form PID) and 1.6.4 (scenarios and the challenge). Section 1.6.3 is needed only for the MPC.

The app https://apmonitor.com/apps/hydrofoil-boat/

See also Volume 3, Chapter 2 for choosing an inner loop by where a disturbance enters and Volume 3, Chapter 3 for how fast an inner loop must be. Volume 5, Chapter 13 for inertial navigation through a lost measurement. Chapter 6 for wave motion that is real and not an error.

5.1 Introduction

At 20 kn the simulated launch flies with the keel at its bow 0.50 m above the water and its main foil 0.40 m under the surface, on a flap of 0.50° and 17.7 kW from the motor. Let the foil rise 0.30 m and it is within 0.3 chords of the surface, where air reaches it and it keeps 30 % of its lift. Let the boat sink 0.50 m and the keel is in the water at flying speed. The usable band is 0.80 m tall. In the 0.5 m sea of case 6, met from astern, a loop that holds the boat level over the waves lets a trough uncover the main foil, which ventilates at 11.20 s, and the hull slams into the next crest at 11.65 s. The same loop told to follow the surface flies the case's 40 s without either event, at a vertical acceleration of 2.67 m/s^2 at the passenger's seat.

Three beliefs about flying on foils are contradicted by the cases. The first is that a submerged wing has no height of its own. Near the surface it does, because lift falls as the foil rises: with the main foil 0.15 m deep the boat answers a lost 0.8 kN of lift by sinking 4.2 cm and staying there. The free height stiffness at that depth is 24.0 kN/m, while with the foil 0.70 m deep it is 0.39 kN/m and the hull is back in the water within 5 s. The second is that a good height loop rejects the waves. Whether it should depends on how fast the boat meets them, which sets whether a wave can be followed at all, and on how tall they are against the band. The third is that enough tuning flies any sea. In a 1.4 m sea from ahead, 33 795 settings tried at 20 kn flew no longer than 12.49 s of the minute asked.

The plant belongs to a wider class, lift surfaces that work near a boundary which removes their lift. Wing-in-ground-effect craft fly close over the water on a lift that depends on their height above it [1], and surface-piercing and fully submerged hydrofoils trade the free stiffness of the boundary against exposure to the waves [2, 3]. Ventilation is a change of flow regime with a region where both regimes are stable [4], and it bears on the safety and controllability of any lifting surface near a free surface [5]. The choice between holding a body still and letting it follow the terrain is a central trade of active vehicle suspension [6], and blending a position sensor with an inertial one through complementary filters is a long-established estimator [7].

DID YOU KNOW? A Racing Dinghy on a Wand

A class of one-person racing dinghy that flies on hydrofoils is described as the fastest one-person dinghy in the world, yet its rules cap the boat at 3.355 m in length, 2.25 m in beam and 8.25 m^2 of sail. A mechanical wand at its bow holds the ride height [8], as in case 4.

However, the hydrodynamics, the ventilation physics and the motion control of hydrofoil craft are reported in separate literatures [4, 9–11], and none of them lets a student fly the boat. This chapter measures what a mechanical wand, a proportional-integral-derivative (PID) loop and a heave-rate cascade can and cannot do for a foiling boat whose lift falls near the surface and collapses at it, on a browser simulator in which the boat, the sea, the instruments and the controllers share one model. The rigid-body equations, the Pierson–Moskowitz sea, the complementary filter and the PID law are established methods.

Section 5.2 follows the domain. Section 5.3 gives the boat and the sea as a nine-state model whose hard events are located inside the integration step, with a split between plant and instruments under which no law sees the true foil depths, and with the full set of equations and defaults needed to rebuild the boat. Section 5.4 shows that the free height stiffness falls by a factor of 62 between 0.43 and 2.0 chords of depth, and that the band between slam and ventilation decides the platform-or-contour choice only together with the encounter frequency. Section 5.5 describes the wand, the observers, the PID, the cascade, the fallback, the takeoff sequencer and an optional model predictive controller (MPC) with a preview. Section 5.6 reports ten cases and a challenge, each credited in the app by a completion rule. There the same contouring structure is clean on 7 of 17 draws of a sea under the PID and on all 17 under the cascade, and in a sea 1.75 times the band no setting in two searches of 33 795 each flies the minute. A cascade retuned by the MPC's own rule flies the 0.65, 0.8 and 1.0 m following seas clean at 0.93, 1.01 and 1.04 m/s^2, where the MPC rides 1.00 and 1.16 and ventilates. Section 5.7 names the height sensor, the flap servos and the ride control computer of real craft, and Sections 5.8 and 5.9 state the limitations and the conclusions.

5.2 Hydrofoils, Their Surface and Their Controllers

Hydrofoil craft were built and flown early in the twentieth century, and a survey of their development published by the American Society of Naval Engineers in 1985 runs to 58 pages [12]. Two families emerged. Surface-piercing foils cross the surface at an angle, so a boat that rises wets less foil and loses lift, and the height control is in the geometry. Fully submerged foils run below the surface, where the waves disturb them less, but they have no such geometric restoring force and fly only under an automatic control system [2, 3]. The simulated launch belongs to the second family, with a remnant of the first: the proximity loss that the surface imposes on a submerged foil.

HISTORY Measuring the Surface Loss (1955 and 1959)

A National Advisory Committee for Aeronautics report of 1955 investigated, in theory and experiment, how the free water surface changes the lift and drag of hydrofoils of aspect ratio 10 and 4. It found the agreement between theory and experiment satisfactory for engineering estimates of hydrofoil characteristics from aerodynamic data [9]. Ventilation was studied at the same time on surface-piercing struts, where cavities open to the atmosphere develop, and the conditions for their existence could be derived from the measurements. The same study concluded that further experiments were needed to define the inception of ventilation as a function of the boundary-layer state or the Reynolds number [13].

That loss, and the ventilation of struts near the surface, were first measured in the 1950s, as the history box above tells. Modern towing-tank experiments on a surface-piercing hydrofoil identify fully wetted, partially ventilated and fully ventilated regimes whose stability regions overlap, so that the flow shows hysteresis: ventilation forms when air reaches separated flow at sub-atmospheric pressure, and it is washed out by a re-entrant jet [4]. A review of ventilating lifting bodies sets out

the scaling that model tests must respect, among them the Froude and cavitation numbers [5]. The simulator's latch is a teaching caricature of that hysteresis, not a fit to it.

Control of fully submerged foil craft became a design subject of its own. A model of a fully submerged hydrofoil craft was flown under longitudinal motion control in following seas in experiments [10], the same authors designed a longitudinal controller around an optimal preview servo system [14], and the control of the vertical motion of a hydrofoil vessel has been treated as a design problem of its own [11]. The seakeeping, slamming and hydrofoil chapters of a monograph on high-speed marine vehicles connect the loads to automatic control [3]. A different line runs through small sailing craft. A racing dinghy class adopted hydrofoils and renewed interest in foiling, with full-scale tow tests of its T-foils and hull [15], and its hydrofoil control system is now tested on a full-scale rig [8]. The simulator's wand is a mechanical height control in the spirit of those small craft, a stick on the water linked to the flap. The irregular sea is the Pierson–Moskowitz spectrum of a fully developed wind sea [16], and passenger comfort is judged against the literature on whole-body vibration [17].

> ⧖ HISTORY **Feelers on the Water (1930s to 1960)**
>
> Before electronics, fully submerged foils were held at their depth by mechanical feelers on the water. One pioneer chose submerged foils to get away from the disturbing influence of the surface in waves, and tried several mechanical depth feelers on six test boats over eight years. They worked in calm water but not in slight seas, and he turned to surface-piercing foils instead. Another inventor showed a small craft at a boat show in New York in 1951. The forward feeler of his system sensed the oncoming waves ahead of the foils. Its spring-loaded heel smoothed out small waves while its buoyant float answered the larger ones. The United States Navy completed its evaluation of the most sophisticated mechanically controlled hydrofoil in 1960, by which time electronic control was well under way [18].

5.3 The Boat and the Sea as Modeled

The launch is a generic design with teaching values throughout, as shown in Figure 5.1. A main foil of 0.8 m^2 and chord 0.35 m sits under the center of mass and carries the weight. A rear foil of 0.25 m^2 and chord 0.25 m, 2.6 m aft, is the stabilizer and carries almost nothing in trim, so the main foil alone fixes the takeoff speed. Both hang 0.90 m below the keel on struts, and each has a flap. An electric pod on the rear strut drives the boat, an ultrasonic sensor at the bow measures the ride height, an inertial measurement unit (IMU) measures pitch, pitch rate and vertical acceleration, a log measures speed, and a wand pivoted at the bow trails forward to the water. The frame has x forward and z up from the mean water line, and the pitch θ is positive bow up.

Each variable carries its role in Table 5.1. Nine states are integrated: the position along the course x and the speed V, the height z of the center of mass and its rate w, the pitch θ and its rate q, the two flaps δ_f and δ_r, and the delivered thrust T. The ventilation latches, the hull's contact with the water and the slam count are discrete states kept between steps.

5.3.1 Rigid-Body Motion

The boat moves in surge, heave and pitch under the foil forces, the hull's contact forces, the thrust and the weight, as shown in Equation (5.1).

$$\begin{aligned} m\dot{V} &= T - D_f - D_r - R_{hull} - D_{air} \\ m\dot{w} &= L_f + L_r + F_{hull} - mg \\ I_{yy}\dot{q} &= \textstyle\sum_j \left(r_{x,j} F_{z,j} - r_{z,j} F_{x,j} \right) \\ r_x &= x_b \cos\theta - z_b \sin\theta, \quad r_z = x_b \sin\theta + z_b \cos\theta \end{aligned} \tag{5.1}$$

where $\dot{x} = V$, $\dot{z} = w$, $\dot{\theta} = q$, and the sum runs over the two foils, their struts, four hull stations and the pod, each at its body offset (x_b, z_b) from the center of mass rotated into the world frame. The center

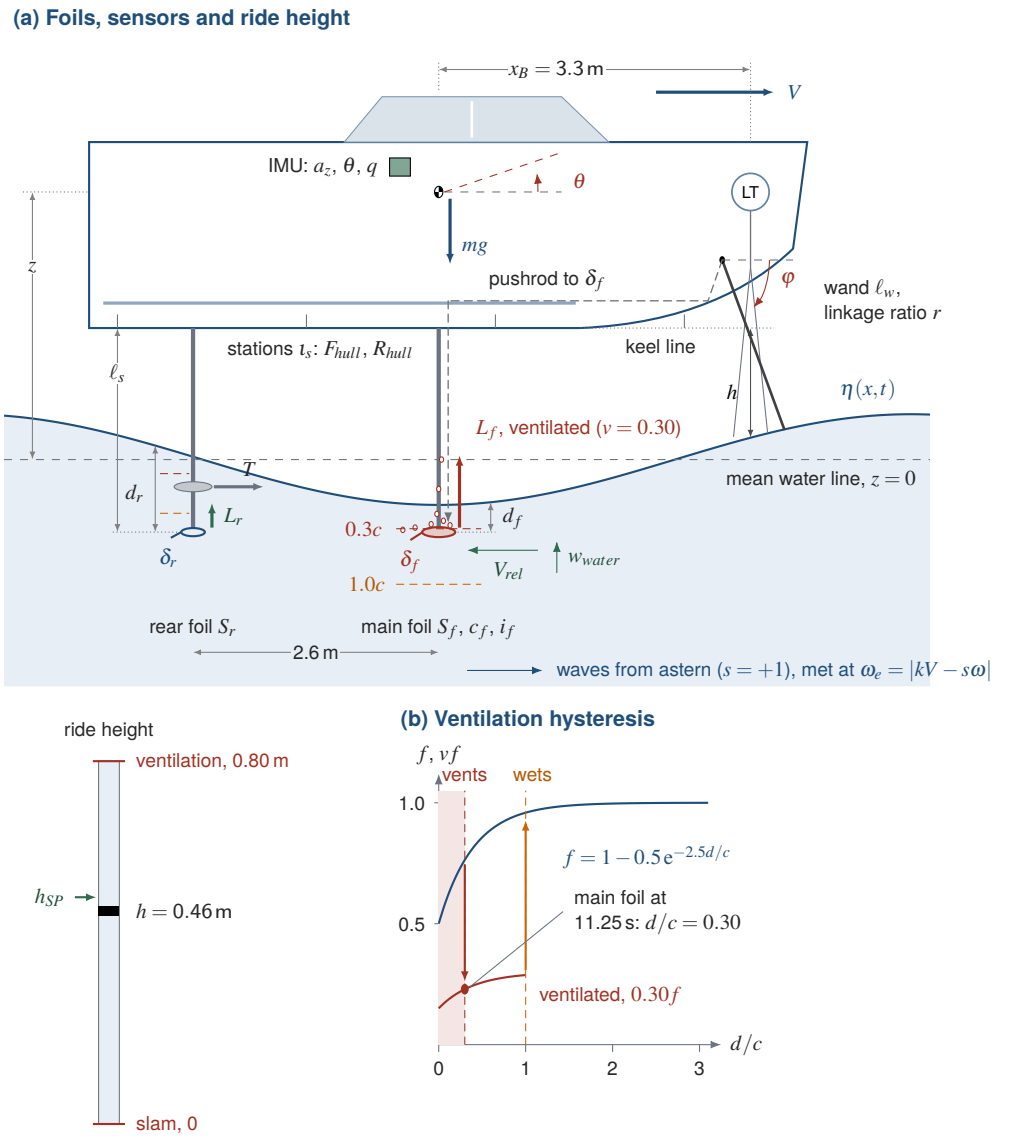

Figure 5.1 (a) The launch as the Sea view draws it at the ventilation of case 6 (platforming, $t = 11.25$ s). The center of mass at height z and pitch θ carries the weight mg and the IMU. The main foil, in a trough at depth d_f under the local surface $\eta(x,t)$, has passed its ventilation depth $0.3c$ and keeps $v = 0.30$ of its lift L_f until it is deeper than $1.0c$, while the rear foil at depth d_r makes L_r. Each foil sees the water at V_{rel} with the orbital velocity w_{water} and carries a flap, δ_f or δ_r. The pod gives the thrust T at speed V. The bow sensor LT measures the ride height h of the keel line above the water, shown on the gauge at lower left between a slam (0) and a ventilation (0.80 m) with its setpoint h_{SP}. The wand of length ℓ_w at angle φ moves the main flap through a pushrod of ratio r. (b) The declared surface factor $f(d/c)$ with its ventilated branch and the two hysteresis arrows. Vertical scale exaggerated and the boat drawn level. Not to scale.

of mass is 0.60 m above the keel, the foils $\ell_s + 0.60\,\text{m} = 1.50\,\text{m}$ below it and the pod 1.45 m below it. A passenger of 85 kg who walks x_p aft moves every body offset by $85x_p/m$, 5.3 cm per meter walked.

Table 5.1 Variables by Role, With Their Values in Trim at 20 kn and 0.5 m

Variable	Role	Trim	Range
h	CV, ride height at the bow	0.50 m	SP 0.1–0.9 m
θ	CV, pitch	0°	SP −3 to 5°
V	CV, speed	20 kn	SP 0–30 kn
δ_f	MV, main flap	0.50°	±15°
δ_r	MV, rear flap	0.12°	±10°
u_T	MV, throttle	29.6 %	0–100 %
H_s	DV, significant wave height	0	0–1.4 m
—	DV, wake, spray, gust, slowdown	events	—
x_p, w_d	DV, passenger aft, weed	0, 0	0–4 m, 0–40 %
z, w	states, heave	1.10 m, 0	—
θ, q	states, pitch	0, 0	—
x, V	states, surge	—, 10.29 m/s	—
T	state, thrust	1035 N	$\le T_{avail}(V)$
d_f, d_r	outputs, foil depths	0.40 m, 0.40 m	truth only
a_{seat}	output, seat acceleration	0	rms scored

5.3.2 Foil Lift Near the Surface

Each foil makes a lift from the speed of the water past it, its effective angle, the surface factor and the ventilation factor, as shown in Equation (5.2).

$$\begin{aligned} L_i &= \tfrac{1}{2}\rho V_{rel}^2 S_i C_{L\alpha}\,\mathrm{sat}_{\alpha_s}(\alpha_{e,i})\, f(d_i/c_i)\, v_i \\ \alpha_{e,i} &= \theta + i_i + k_\delta \delta_i + (w_{water} - \dot{z}_i)/V_{rel} \\ f(d/c) &= 1 - A_s\, \mathrm{e}^{-B_s d/c} \\ D_i &= \tfrac{1}{2}\rho V_{rel}^2 S_i \left(C_{D0} + \frac{(C_L f v)^2}{\pi AR_i\, e} \right) \end{aligned} \tag{5.2}$$

where d_i is the depth of foil i under the local surface, $V_{rel} = \max(0.5,\ \dot{x}_i - u_{water})$ is the speed of the water past it with the orbital velocities u_{water} and w_{water} taken at its depth, $\dot{z}_i$ is the foil's own vertical velocity, sat_{α_s} caps the angle at the stall $\alpha_s = 8°$, $AR_i = S_i/c_i^2$ and $e = 0.85$. The heave rate in α_e is the foil's natural damping, and each strut adds a drag of $\frac{1}{2}\rho V_{rel}^2 (0.2\,\mathrm{m})\, 0.012$ per meter wetted. The surface factor, with $A_s = 0.5$ and $B_s = 2.5$, is 0.697 at 0.2 chords, 0.857 at 0.5, 0.959 at 1 and 1.000 at 3. It is a declared teaching form with the shape of the measured free-surface loss [9] and has not been fitted to any data. Weed on the main foil multiplies its lift coefficient by $1 - w_d$ and adds $0.03\, w_d$ to its profile drag.

5.3.3 Ventilation as a Latch

The ventilation factor is a latch with hysteresis, as shown in Equation (5.3).

$$v_i = \begin{cases} 0.30 & \text{latched: set if } d_i < 0.3c_i,\ \alpha_{e,i} > \alpha_v \\ & \text{and cleared if } d_i > 1.0c_i \\ 0 & d_i \le 0 \text{ (out of the water)} \\ 1 & \text{otherwise} \end{cases} \tag{5.3}$$

where $\alpha_v = 2°$. For the main foil the latch sets shallower than 0.105 m and clears deeper than 0.35 m. The latch is evaluated after each integration step, with the crossing instant interpolated inside the step, so the right-hand side stays smooth within a step. The thresholds and the retained fraction are teaching values. Real ventilation depends on the strut, the speed, the cavitation number and the history of the flow [5, 13].

5.3.4 The Hull, the Hump and the Slam

The hull is a V-bottom of 20° deadrise represented by four buoyant stations, and its resistance scales with the weight it still carries, as shown in Equation (5.4).

$$\begin{aligned} F_{hull} &= \textstyle\sum_{s=1}^{4} \max\left(0,\ K_b \iota_s^2 + C_b\, \iota_s \dot{\iota}_s\right) \\ R_{hull} &= F_{hull}\left(0.25\, Fn^2 + r_h\, \mathrm{e}^{-((Fn-0.45)/0.15)^2}\right) \\ Fn &= V/\sqrt{g\, L_h} \\ D_{air} &= \tfrac{1}{2}\rho_a C_a A_a\, V|V| \end{aligned} \tag{5.4}$$

where ι_s is the immersion of station s below the local surface, $K_b = 55\,300\,\mathrm{N/m^2}$, $C_b = 10\,300\,\mathrm{N\,s/m^2}$, $L_h = 7.4\,\mathrm{m}$, $r_h = 0.10$, and $\rho_a C_a A_a = 1.2 \times 2.4\,\mathrm{m^2}$. Because the resistance falls as the foils unload the hull, the sum of hull and foil drag peaks before takeoff. A run is foilborne once the hull has been clear of the water for 0.5 s. Five keel points are watched, and one that enters the local surface while the boat is foilborne at a closing speed above $v_{slam} = 0.5\,\mathrm{m/s}$ is a slam, located inside the step, which ends the clean run and pauses the simulation. A slower contact is a touchdown. The slam speed is a teaching value. Real slamming loads come from the high-speed craft rules of the classification societies and from the analyses that underlie them [3].

5.3.5 Actuators

The flaps and the thrust follow their commands through lags and limits, as shown in Equation (5.5).

$$\begin{aligned} \dot{\delta}_i &= \mathrm{clip}\Big(\frac{\mathrm{clip}(\delta_{i,cmd}, \pm\delta_{i,max}) - \delta_i}{\tau_s},\ \pm\dot{\delta}_{max}\Big) \\ \tau_T \dot{T} &= u_T\, T_{avail}(V) - T \\ T_{avail} &= \min\left(T_{max},\ \eta P_{max}/\max(1, V)\right) \end{aligned} \tag{5.5}$$

where $\tau_s = 8\,\mathrm{ms}$ is a 20 Hz servo, $\dot{\delta}_{max} = 60^\circ/\mathrm{s}$, $\tau_T = 0.3\,\mathrm{s}$, and the delivered thrust is $\min(T, T_{avail})$. The thrust limit of 4.5 kN binds below 8.00 m/s and the power limit of 60 kW with $\eta = 0.6$ above it.

5.3.6 The Sea

The sea is a long-crested sum of sixteen linear waves on deep water, from astern ($s = +1$) or ahead ($s = -1$), as shown in Equation (5.6).

$$\begin{aligned} \eta(x,t) &= \textstyle\sum_k a_k \cos\psi_k, \quad \psi_k = k_k x - s\,\omega_k t + \phi_k \\ S(\omega) &= \tfrac{5}{16} H_s^2\, \omega_p^4\, \omega^{-5}\, \mathrm{e}^{-1.25(\omega_p/\omega)^4} \\ a_k &= \sqrt{2 S(\omega_k)\Delta\omega_k}, \qquad k_k = \omega_k^2/g \\ (u,w)_{water} &= \textstyle\sum_k s\, a_k \omega_k\, \mathrm{e}^{k_k z}(\cos\psi_k,\ \sin\psi_k) \\ \omega_e &= |kV - s\,\omega| \end{aligned} \tag{5.6}$$

where $\omega_p = 2\pi/T_p$ with $T_p = c_T\sqrt{H_s}$ and $c_T = 5.0\,\mathrm{s/m^{1/2}}$, the orbital velocities are taken at depth z, and ω_e is the frequency at which a boat at speed V meets a wave of frequency ω. The frequencies lie on a geometric grid from 0.9 to 5 rad/s, jittered inside each band, and the phases come from the scenario's seed. The spectrum is the Pierson–Moskowitz form [16], and the amplitudes are rescaled so that the variance is exactly $(H_s/4)^2$. A boat at 20 kn meets the peak of a 0.5 m sea from astern at 1.54 rad/s and from ahead at 5.09 rad/s, the peak of a 0.2 m sea from astern at 5.47 rad/s and that of a 1.4 m sea from ahead at 2.25 rad/s. A new sea state is approached at 0.3 m per minute with the phases continuous, a run that starts foilborne fades its sea in over 3 s, and a wake is one extra wave train of period 4 s under a Gaussian envelope of 2 s.

Table 5.2 Model Parameters. H Marks the Hidden Ones, and the Rest Are Adjustable

Symbol	Meaning	Value	Unit
m, I_{yy}	mass, pitch inertia	1600, 4500	kg, kg m^2
z_K	center of mass above the keel (H)	0.60	m
ℓ_s	strut, keel to foil	0.90	m
x_B	bow sensor ahead of mass center (H)	3.3	m
ρ, g	density, gravity (H)	1025, 9.81	kg/m^3, m/s^2
S_f, c_f, i_f	main foil area, chord, incidence	0.80, 0.35, 4	m^2, m, °
x_f	main foil ahead of mass center (H)	0	m
S_r, i_r	rear foil area, incidence	0.25, 0	m^2, °
c_r, x_r	rear foil chord, position (H)	0.25, −2.6	m
$C_{L\alpha}$, α_s	lift slope, stall angle	5.0, 8	rad^{-1}, °
k_δ, C_{D0}	flap effectiveness, profile drag	0.5, 0.008	—
A_s, B_s	surface loss at the surface, decay	0.5, 2.5	—, per chord
d_v, d_w	ventilates, wets again	0.3, 1.0	chords
α_v, v	ventilation angle, lift kept	2, 0.30	°, —
$\delta_{f,max}$, $\delta_{r,max}$	flap limits	15, 10	°
τ_s, $\dot{\delta}_{max}$	servo lag, rate limit	8, 60	ms, °/s
T_{max}, P_{max}, η	thrust, power, efficiency	4.5, 60, 0.6	kN, kW, —
τ_T	thrust lag	0.3	s
K_b, C_b	station buoyancy, damping (H)	55300, 10300	N/m^2, N s/m^2
r_h, v_{slam}	hump height, slam speed	0.10, 0.5	—, m/s
k_n	instrument noise multiplier	1	—
c_T	peak period per $\sqrt{H_s}$	5.0	s/m$^{1/2}$
—	sea build rate	0.3	m/min

5.3.7 Instruments

The controllers read a packet and nothing else. The bow sensor reports the height of the keel line above the water under the bow, $h_m = h + n_h$, with colored noise of 1 cm standard deviation, and reports nothing for the spray time after a spray event. The IMU reports the pitch ($\pm 0.05°$), the pitch rate ($\pm 0.1\,°/\mathrm{s}$) and the vertical acceleration ($\pm 0.02\,\mathrm{m/s^2}$ plus a bias of $0.02\,\mathrm{m/s^2}$), and the log reports the speed ($\pm 0.05\,\mathrm{m/s}$). Each noise is a first-order filtered Gaussian from the seeded generator, so a run repeats exactly after Reset. The wand is a mechanism, not an instrument: its angle below the horizontal is the one at which a stick of length ℓ_w from a pivot 3.0 m ahead of the center of mass and 0.30 m above the keel reaches the local surface, without noise. The true foil depths, the ventilation state, the sea, the weed and the passenger never reach a controller.

5.3.8 Parameters and Numerics

Every parameter appears in Table 5.2. Hidden parameters fix the geometry and are not shown, while the others are adjustable in the Parameters panel. Integration uses fourth-order Runge–Kutta on a fixed step of 2 ms, the controllers execute every 20 ms, and samples are logged at 50 Hz. At half the step the first ventilation of case 6 moves from 11.1985 s to 11.1987 s and its slam from 11.6470 s to 11.6444 s, and the slam of case 9 from 6.1421 s to 6.1418 s. A foilborne start is trimmed in calm water by Newton iteration on the two flaps and the thrust, and a start at rest floats on the hull.

5.3.9 Assumptions

The boat moves in surge, heave and pitch only, running straight into or away from a long-crested sea, so roll, yaw, leeway and oblique waves are absent. The lift is quasi-steady, with no unsteady lift lag, no added mass and no downwash from the main foil onto the rear one, which is least defensible at the 2.6 Hz wand oscillation of case 4. The surface factor and the ventilation latch are declared forms, and the ventilated flow has no partial regime and no dependence on speed. The waves are linear and do not break, and the boat does not change them. The hull has buoyancy and damping but no planing lift, so its contribution near takeoff is approximate, and its resistance curve is declared. The slam speed and the comfort limit are teaching values.

Table 5.3 Free Height Stiffness at 20 kn in Calm Water (Case 3 Geometry)

h (m)	d_f (m)	d_f/c_f	Stiffness (kN/m)	Damping (kN s/m)	τ (s)
0.75	0.15	0.43	24.00	23.3	1.0
0.50	0.40	1.14	3.43	27.0	7.9
0.20	0.70	2.00	0.39	27.6	70.5

5.4 Dynamic Character of the Flying Boat

5.4.1 Takeoff Speed and the Hump

The main foil at its stall angle, 1.5 chords deep, carries the 15.70 kN weight from the takeoff speed of Equation (5.7), which is 14.48 kn.

$$V_{to} = \sqrt{\frac{mg}{\frac{1}{2}\rho S_f C_{L\alpha} \alpha_s f(1.5)}} = 7.449\,\mathrm{m/s} \tag{5.7}$$

At the stall it can lift 10.2 kN at 6 m/s, 18.1 kN at 8 m/s and 40.7 kN at 12 m/s, where the boat trimmed with the foil 1.5 chords deep flies at 3.06° of effective angle. In trim at 0.5 m the main flap moves from 5.30° at 16 kn to 0.50° at 20 kn and −2.11° at 24 kn, while the power rises from 13.2 to 17.7 and 25.7 kW, and the ratio of weight to thrust falls from 16.3 to 12.6. With the main flap held at 15°, as the takeoff sequencer holds it, the steady resistance peaks at 2253 N at 8.0 kn (a Froude number of 0.48), falls to 1804 N at 11.0 kn, and the hull clears the water at 12.5 kn. Below 8 m/s the pod delivers 4.5 kN at full throttle, so 50 %, 2250 N, stops short of the hump and 55 % passes it.

5.4.2 A Stiffness the Surface Gives and Takes Away

A small rise of the boat brings the main foil closer to the surface and lowers f, so the lift falls and the boat sinks back. The slope of the lift against height is the free stiffness, and the heave rate in α_e supplies a damping. Both are measured at 20 kn in calm water by perturbing the trimmed state, as listed in Table 5.3. The stiffness falls from 24.0 to 0.39 kN/m between ride heights of 0.75 and 0.2 m while the damping stays between 23 and 28 kN s/m, so the time constant of the free height grows from 1.0 s to 70.5 s. Deep, the height is in effect an integrator of the lift error.

🕮 **Pause.** Try Exercise 5.2 before reading on.

The stiffness does not make the boat stable on its own. With both flap commands and the throttle fixed, the linearization at 20 kn and 0.5 m has eigenvalues at $0.085\,\mathrm{s}^{-1}$ and $0.148 \pm 0.437j\,\mathrm{s}^{-1}$, and at 16 kn at $0.126\,\mathrm{s}^{-1}$ and $0.209 \pm 0.357j\,\mathrm{s}^{-1}$, because height, drag and speed are coupled through the hull and the thrust. Released from the 16 kn trim with the flap moved by 0.05°, the boat reaches 0.616 m at 5 s and is back on its hull by 20 s. Every case therefore runs with some loop closed, if only the pitch and speed loops.

5.4.3 The Band and the Encounter Frequency

The ride height is defined at the bow sensor as the height of the keel line above the water there. In calm level flight the keel touches at $h = 0$ and the main foil reaches 0.3 chords at $h = \ell_s - 0.3c_f = 0.795$ m, so the band is 0.80 m. In waves the surface under the bow and under the foil differ, and the band is a statement about the wave height measured from crest to trough against 0.80 m. Platforming holds the boat level and lets the waves pass between hull and foil, which works while they are smaller than the band. Contouring follows the surface, which works while they come slower than the loop can move the boat. From astern the boat overtakes the long waves of a large sea and meets them slowly, 1.54 rad/s for the 0.5 m sea of case 6, while the short waves of a small sea arrive fast, 5.47 rad/s for a 0.2 m sea. From ahead every wave is met faster.

5.5 Implemented Control System

Every law runs every $\Delta t_c = 20\,\text{ms}$ on the packet of the previous section. The height loop drives the main flap, the pitch loop the rear flap and the speed loop the throttle. The header's Manual mode gives all three stations to the student, and the Hand height controller gives only the main flap while the pitch and speed loops keep running. Every law tracks its station when it is not in command, so Manual to Auto and a change of structure are bumpless.

5.5.1 The Wand

The wand is a proportional controller made of a stick and a pushrod, as shown in Equation (5.8).

$$\delta_f = \delta_n - r(\varphi - \varphi_n), \qquad \frac{\partial \delta_f}{\partial h} \approx -\frac{r}{\ell_w \cos\varphi} \tag{5.8}$$

where φ is the wand's angle below the horizontal, $\delta_n = 0.5^\circ$ is the trim flap at 20 kn, $\varphi_n = 53^\circ$ and the linkage ratio r (default 0.6) is the gain per degree. The boat settles where the flap holds the weight, which at 20 kn is near φ_n, so the pivot flies about $\ell_w \sin\varphi_n$ above the water and the ride height about $\ell_w \sin\varphi_n - 0.30\,\text{m}$: 0.50 m for the default $\ell_w = 1.0\,\text{m}$ and 0.66 m for 1.2 m. The length is the setpoint. The gain per meter at the default is 57 °/m of flap.

5.5.2 Height Estimates

The electronic loops read complementary observers of $H = h_m - x_B \sin\theta$, the keel line's height at the station of the mass center, given in Equation (5.9). Each is the filter with a bias state of Section 1.4.6 with a third pole at the crossover.

$$\begin{aligned} \dot{\hat{H}} &= \hat{v} + 3\omega_c e, \qquad \dot{\hat{v}} = a_z - \hat{b} + 3\omega_c^2 e, \qquad \dot{\hat{b}} = -\omega_c^3 e \\ e &= H - \hat{H}, \qquad \hat{h} = \hat{H} + x_B \sin\theta_m \end{aligned} \tag{5.9}$$

where $\hat{b}$ estimates the accelerometer's bias and all three poles sit at $-\omega_c$, so a constant bias leaves no steady error. The observers are advanced by a semi-implicit Euler step every Δt_c. Below ω_c the estimate follows the sensor, and above it the IMU [7]. The platform observer runs at $\omega_c = 0.6\,\text{rad/s}$ (adjustable), the contour observer at 8 rad/s, and a third at 0.15 rad/s whose residual is the surface under the bow. Its mean square, averaged over 8 s, gives the wave-height estimate $\hat{H}_s = 4\sqrt{\overline{\eta^2}}$. The heave rate used by the cascade is $\hat{v} + x_B q_m \cos\theta_m$ from the platform observer. With the sensor blind, every observer coasts on the IMU and its last bias estimate. Volume 5, Chapter 13 meets the same loss under water, where its Doppler log loses the seabed.

5.5.3 PID and Cascade

The PID is the shared velocity-form law of Section 1.6.2, Equations (1.19) to (1.22). Its error is $e_k = h_{SP,k} - \hat{h}_k$ on the height estimate, its derivative acts on the measurement through a first-order filter of 0.1 s, its output range is the main flap's $\pm\delta_{f,max}$ and its tracking time is τ_I. Its output bias is $\delta_{map}(V_m)$, the flap that holds 0.5 m in calm water at the measured speed, tabulated once from the declared model from 14 to 30 kn. The defaults are $K_c = 20^\circ/\text{m}$, $\tau_I = 2\,\text{s}$ and $\tau_D = 0.2\,\text{s}$. A larger flap means more lift and a higher boat, so the loop is reverse acting, $a = +1$ in the convention of Section 1.3. The cascade replaces the single loop by two, as shown in Equation (5.10).

$$\begin{aligned} v_{SP} &= \text{clip}\big(K_h(h_{SP} - \hat{h}),\ \pm 1.5\,\text{m/s}\big) \\ \delta_{f,cmd} &= \delta_{map}(V_m) + \text{PI}_w\big(v_{SP} - \hat{\dot{h}}\big) \end{aligned} \tag{5.10}$$

where $K_h = 4\,\text{s}^{-1}$ and the inner PI has a gain of 8° per m/s and $\tau_{I,w} = 1\,\text{s}$, in the same velocity form and also reverse acting. The inner loop works on the IMU's heave rate, which the waves do not disturb directly and which survives the loss of the height sensor. Choosing an inner loop by where the disturbances enter is the subject of Volume 3, Chapter 2.

5.5.4 Wave Handling, Fallback and the Other Loops

The height estimate is the platform one, the contour one, or the one chosen by the schedule (Auto), which contours once $\hat{H}_s$ exceeds 0.37 m and platforms again below 0.33 m, a switch height of 0.35 m with 0.02 m of hysteresis. With the sensor lost, Hold flap freezes the last command, and Fly on IMU keeps the PID or cascade closed on the coasting platform estimate at the ride-height setpoint. The pitch loop is the same PID on the IMU's pitch to the rear flap. More rear flap lifts the stern and lowers the bow, so a bow-high error asks the rear foil for more lift and the loop is direct acting, $a = -1$. Its settings are $K_c = 2$, $\tau_I = 2\,\text{s}$, $\tau_D = 0.1\,\text{s}$ and a derivative filter of 0.05 s. The speed loop is a PI on the log to the throttle, at 7.7 % per knot and $\tau_I = 5\,\text{s}$.

🕮 **Pause.** Try Exercise 5.3 before reading on.

The takeoff sequencer is a phase register with minimum dwell times. Above 60 % of the takeoff speed it enters HUMP and holds the main flap at 15° to unload the hull. At 95 % of the takeoff speed, before the hull unsticks, it enters FOILBORNE and hands the flap to the height loop, whose setpoint ramps from the measured height to the target at 0.25 m/s. It returns to DISPLACEMENT below 85 % of the takeoff speed with the keel down. With the speed setpoint below 95 % of the takeoff speed it lands the boat instead, ramping the setpoint at 0.15 m/s to –0.05 m and centering the flap once the hull is in. The takeoff speed is the declared boat's 14.48 kn, so a lighter boat lifts off with the flap still held at 15°. At 1200 kg, case 2's PID at 75 % throttle climbs to a 1.38 m ride height and ventilates the main foil at 4.2 s.

5.5.5 Model Predictive Control

The fifth height controller is an MPC on the main flap, built on the general setup of Section 1.6.3 [19, 20]. Every Δt_c it solves a quadratic program over $N = 105$ steps of $T_p = 40\,\text{ms}$, a horizon of 4.2 s, and applies the first move. The flap is constant over each of eight move blocks, whose lengths in steps run 1, 2, 3, 5, 8, 13, 21 and 52. The pitch and speed loops, the sequencer and Hold flap act as shipped and take precedence over the plan, and the PID's and the cascade's tuning and wave handling leave the panel, because both loops then only track the applied flap and the MPC has its own observer.

The controller's model is the declared boat. The forces of Equations (5.1) to (5.5) are differentiated numerically at the declared parameters, in calm water with no passenger and no weed, about the foilborne trim at $h = 0.5\,\text{m}$ and $\theta = 0$, at each speed from 15 to 30 kn, once at load, and interpolated at the measured speed, which gives Equation (5.11).

$$\begin{aligned} x &= [z,\ w,\ \theta,\ q,\ \delta_f,\ \delta_r,\ d,\ g,\ m]^\top \\ \dot{w} &= \mathbf{a}^\top \xi + d + g - a_z \eta_f, \quad \dot{q} = \mathbf{c}^\top \xi + m - c_z \eta_f \\ \tau_s \dot{\delta}_f &= u - \delta_f, \qquad \tau_s \dot{\delta}_r = K_\theta(\theta + T_\theta q) - \delta_r \\ \dot{d} &= 0, \qquad \dot{g} = -g/\tau_g, \qquad \dot{m} = 0 \end{aligned} \tag{5.11}$$

where every variable is a deviation from the trim at the measured speed, $\xi = [z, w, \theta, q, \delta_f, \delta_r]^\top$, $\mathbf{a}$ and $\mathbf{c}$ hold the derivatives of the heave and pitch accelerations with respect to ξ (first entries a_z, c_z), u is the flap command and η_f the surface at the main foil, which enters as a negative heave. The rear flap follows the shipped pitch loop's proportional and derivative action ($K_\theta = 2$, $T_\theta = 0.1\,\text{s}$, integral left out), d and m are constant heave and pitch disturbances that make the controller offset-free [21], and g is a heave disturbance that fades with $\tau_g = 0.1\,\text{s}$, the orbital push of the waves, which reverses within half an encounter period. The model is sampled exactly with a zero-order hold at T_p for the plan and at Δt_c for the estimator. At 20 kn its derivatives agree with hand formulas from the declared parameters within 0.9 % (65.87 m/s^2 of heave acceleration per radian of flap). After a 0.3° flap step in calm water its heave matches the engine's within 1 % to 0.5 s and 2 % at 1 s, then runs ahead, 8.2 cm against 7.1 cm at 2 s and 14.1 cm against 9.1 cm at 4 s, because the pitch integral and the change of the surface factor with depth are left out.

The plan judges the keel's height above the mean water at the bow, $\bar{h}$, at the main foil's station, k_f, and at the stern, k_s, 3.6 m aft of the center of mass, and the heave rate w, with the objective and limits of Equation (5.12).

$$\begin{aligned}
\min\ & \textstyle\sum_{k=1}^{N} \big[Q_h(\bar{h}_k - h_{SP})^2 + Q_w w_k^2 \\
& + Q_k(\bar{h}_k - h_{SP} - \hat{\eta}_{f,k})^2 \big] + \textstyle\sum_{j=1}^{8} S\,(u_j - u_{j-1})^2 \\
& + \textstyle\sum_{b=f,s} \big(\rho\, \varepsilon_b^2 + \sigma \varepsilon_b\big) \\
\text{s.t. } & \hat{\eta}_{f,k} + k_{lo} - \varepsilon_f \le k_{f,k} \le \hat{\eta}_{f,k} + k_{hi} + \varepsilon_f \\
& k_{s,k} \ge \hat{\eta}_{s,k} + k_{lo} - \varepsilon_s, \quad \varepsilon_b \ge 0, \quad |u_j| \le \delta_{f,max} \\
& |u_1 - u_0| \le \dot{\delta}_{max} \Delta t_c \\
& |u_j - u_{j-1}| \le \tfrac{1}{2} \dot{\delta}_{max} T_p (n_{j-1} + n_j), \quad j \ge 2
\end{aligned} \tag{5.12}$$

where $\hat{\eta}_{f,k}$ and $\hat{\eta}_{s,k}$ are the water the plan expects under the foil's station and the stern at step k, $k_{lo} = 0.10$ m is a slam margin, $k_{hi} = 0.70$ m keeps the main foil 0.20 m deep, above its 0.105 m ventilation depth, ε_f and ε_s are single slacks, each covering its band over all N steps, u_j is the flap of block j with n_j steps and u_0 the command applied now. The weights are $Q_h = 100$, $Q_k = 100$, $Q_w = 30$, $S = 20$, $\rho = 10^5$ and $\sigma = 10^3$. The Q_k term is a partial contour, which leans the plan toward the expected water by $Q_k/(Q_h + Q_k)$, one half, and leaves it exactly at the setpoint in calm water. The flap's box is hard and read as the shipped loops read it, and the servo's 60 °/s is hard as a ramp, 1.2° for the first move and the travel between two block centers for the later ones.

The controller reads the packet only. Its own observer of the form of Equation (5.9) lowers its crossover from 2 to 0.3 rad/s over 3 s, so it learns the accelerometer's bias early and is inertial over the sea's 0.9 to 5 rad/s, and the bow reading against its height is the surface under the bow. A steady-state Kalman gain computed at 20 kn corrects the state from that height and rate and the IMU's pitch and rate.

The preview is the water the bow sensor has already measured, logged against the distance run and read back 3.3 m behind the sensor for the main foil and 6.9 m behind it for the stern, as if the surface stood still, as in preview control on a measured road [14, 22]. Chapter 12 reads its seam preview by place in the same way. At 20 kn the foil's water is known 0.32 s ahead, and past that the latest reading is held. Waves travel, and that is the preview's declared error: on the predicted change of the water at the foil it is 0.104 m rms 0.16 s ahead on case 6's sea against 0.145 m for assuming no change, equal to it 0.32 s ahead (0.087 and 0.088 m), and worse than holding the reading from ahead (case 9, 0.184 against 0.093 m). Switching the preview Off leaves the plan the water under the bow now, held over the horizon, which is what a contouring loop knows.

The controller accepts a plan whose solver status is optimal, whose numbers are all finite, whose planned flaps sit in the box and whose opening entry equals the move. When a plan fails these tests the cascade, whose inner PI has tracked the applied flap exactly, flies the flap and the historian records the event until a later plan passes. On every entry to MPC control the command blends from the cascade's law to the plan's over 0.3 s, never faster than the servo. Over seven forced fallbacks on case 6's sea and four in the 0.8 m sea the blend lowers the mean peak seat acceleration in the 1.5 s after resumption from 2.61 to 2.01 m/s^2 and from 4.87 to 3.87 m/s^2.

A 20 ms grid of 75 steps ventilates in the 0.8 m following sea at 14.4 s where the 40 ms grid flies clean, and 38 or 105 steps of 40 ms give the same seat accelerations within 0.01 m/s^2, so the horizon covers the longest encounter period compared, 4.09 s. The weights were picked once on case 6's seed from 108 settings, 37 passing, by a rule: no slam or ventilation at 0.5, 0.65 and 0.8 m, then the least seat acceleration at 0.5 m. The margin entered the rule only after five grids on 0.5 m alone had picked weights that ventilate at 0.65 m, so the results at 0.65 and 0.8 m are in-sample.

In Chrome 157 the solves of case 10 take 1.7 % of a core at real time and 5.4 % at the top speed of five times, the largest 3.3 ms, and the page keeps 60 frames per second.

Table 5.4 The Eleven Cases

#	Case as shipped	Teaching point
1	At rest, Hand height controller, throttle 75 %, flap 6°, 0.1 m sea	takeoff by hand
2	At rest, PID with the sequencer, hand throttle 40 %	hump, handover
3	Flying at 0.75 m, flap fixed, 0.5° less flap at 3 s	free stiffness
4	Wand, calm, gust of 1.5 m/s at 5 s	a mechanical P loop
5	PID, 0.3 m sea astern, spray 2 s at 10 s, Hold flap	sensor loss
6	PID platforming, 0.5 m sea astern, 40 s	platform or contour
7	Cascade, calm, pitch loop P only, pitch SP 1.5° at 3 s	pitch and heave
8	PID, no integral, passenger 3 m aft at 5, weed 20 % at 15 s	integral action
9	Suggested design, 1.4 m sea ahead, seed 2, 60 s	a sea to leave
10	MPC with the preview On, case 6's sea, 40 s	a look-ahead
11	Challenge: sea astern 0.2 to 0.5 m, wake at 30, spray at 45 s	the whole problem

5.5.6 The Three Models

The plant is everything in Section 5.3. For the wand, the PID and the cascade the control model is implicit in the tuning, with the static flap map as the only part of the plant they hold, computed from the declared parameters at load with no passenger and no weed. The MPC holds the declared boat of Equation (5.11), and the sequencer that serves every electronic structure holds the declared takeoff speed. A learner's change to the boat is therefore model error for every law. Set before Run, a mass of 1900 kg or the passenger 3 m aft leaves the PID's and the cascade's events and seat rms in the seas of cases 6 and 10 as they were with the map rebuilt for that boat, and moves their flap travel by under 2 %, because the integral carries what the map misses. The observers carry a double integrator with a bias, which knows nothing of the foils. No law is told the foil depths or the ventilation state, so a ventilated foil reaches a controller only through the IMU and the height sensor, as a sudden fall.

5.6 Eleven Cases, Measured

The cases appear in Table 5.4. Each quoted value comes from the engine of Section 1.6.1 driven from a script, stepped exactly as the page steps it, on the case's own seed, with any setting a sentence names entered before the run. Vertical accelerations are root mean square values at the seat, 1 m ahead of the center of mass, while foilborne.

5.6.1 By Hand, Over the Hump and Onto the Foils

In case 1 the hull is clear at 8.53 s at 17.1 kn with the throttle still at 75 %, and the flap that lifted it out is too much for the rising speed. The main foil ventilates at 8.83 s, and the boat porpoises in and out for the rest of the minute at 7.60 m/s^2 while the speed runs to 32.6 kn. With the speed loop holding 18 kn and the flap at 3° the boat is foilborne from 7.08 s and never slams, although it still ventilates at 7.35 s, and with 5° it slams at 21.41 s. In case 2, at 40 % throttle the boat is doing 6.28 kn at 15 s with 1796 N of drag and the hull carrying 13.4 kN, 50 % stops at 7.5 kn, and 55 % is foilborne at 18.20 s. With 75 % from the start the sequencer takes the boat off at 7.63 s without a ventilation, between 0.42 and 0.85 m. Without the sequencer the PID meets the unsticking hull with its whole error and overshoots to 1.04 m, ventilating at 7.75 s, while the cascade without it stays between 0.42 and 0.77 m.

5.6.2 Stiffness, the Wand and the Pitch Coupling

In case 3 the 0.5° flap step at 0.75 m takes 0.785 kN off the main foil. The boat drops from 0.750 to 0.708 m within a second and is still at 0.709 m at 14.9 s. From 0.5 m it sinks to 0.245 m by 8 s and to the water by 15 s, and from 0.2 m it is in the water, at –0.211 m, by 8 s. In case 4 the wand flies at 0.50 m with a length of 1.0 m, at 0.34, 0.65 and 0.72 m with 0.8, 1.2 and 1.3 m. Up to a ratio of 2.6 it rides out the gust, at 0.64 m/s^2. At 2.7 an oscillation of 2.55 to 2.74 Hz grows into a porpoise of 0.69 Hz and 1.27 m, which ventilates at 17.73 s and slams at 18.17 s. The instant is sensitive because 2.7 lies just past the threshold: at half the step the boat slams at 16.76 s without ventilating first. In

case 7 a proportional pitch loop of gain 0.5 reaches 0.31° of the 1.5° step by 12 s, because the rear foil resists being pitched. A 0.5 s integral reaches 1.478° with the rear flap at −2.50°, while the height loop moves the main flap from 0.41° to −2.77° and holds the ride height between 0.473 and 0.511 m. Each degree of pitch adds 3.68 kN at the main foil and each degree of main flap 1.84 kN, so the pairing costs about two degrees of main flap per degree of pitch.

5.6.3 Sensor Loss and Steady Disturbances

In case 5 the bow sensor is blinded for 2 s at 10 s. Holding the flap, the boat climbs on a wave with nothing to stop it: the main foil ventilates at 11.17 s and the hull slams at 11.80 s. Flying on the IMU, the PID still ventilates at 11.18 s but does not slam, and the cascade on the IMU flies the 25 s clean at 0.59 m/s^2. A shadow platform observer blinded at 10 s wanders 3.2 cm from the true height in the case's sea in those 2 s, and 2.2 cm in calm water. In case 8, with no integral action, the passenger leaves the boat 0.563 m high with the bow up 0.41° and the weed leaves it at 0.478 m. A 2 s height integral brings it to 0.500 m by 29.9 s, and with a pitch integral too the pitch is back to 0.00°, with the rear flap at 1.40° and the main flap at 2.43°. Integral action stops at the flap's authority: with 40 % weed at 18 kn the flap reaches its 15° stop and the hull slams 0.74 s after the weed arrives, while at 20 kn the flap holds at 7.1°.

5.6.4 Platforming Against Contouring

Case 6 shows the platform-or-contour choice on one sea, as shown in Figure 5.2. Platforming with the 0.6 rad/s cutoff, the PID holds the boat near its height over the mean water while a trough, with the surface 0.200 m below the mean over the main foil, uncovers the foil to 0.104 m. It ventilates at 11.20 s, the boat drops, and the stem meets the next crest, 0.280 m above the mean, at 11.65 s and 0.77 m/s. Figure 5.3 shows the app at that ventilation. Contouring, the same PID flies the 40 s clean between 0.19 and 0.79 m at 2.67 m/s^2. The cutoff is the same choice on a dial: at 1, 1.5 and 3 rad/s the boat flies clean at 1.15, 1.29 and 1.93 m/s^2. The cascade contours at 1.38 m/s^2, while its platforming ventilates later, at 21.69 s. In a 0.2 m sea the comfort ranking reverses, 0.87 m/s^2 platforming against 2.31 contouring, because the small waves arrive at 5.47 rad/s and pass under the hull anyway.

Pause. Try Exercise 5.1 before reading on.

The lesson holds on the case's sea, and the structure decides how far it generalizes. Over 17 draws of the same sea state (the case's seed and seeds 1 to 16), PID platforming ventilates on all 17 and slams on 8, PID contouring is clean for 40 s on 7 and ventilates on 10, and cascade contouring is clean on all 17. From ahead the 0.5 m sea is met at 5.09 rad/s, and the PID's contouring slams at 7.35 s, sooner than its platforming ventilates (34.82 s), while the cascade's contouring flies the 40 s clean at 2.94 m/s^2. The app therefore teaches the choice from astern, where the long waves are slow enough to follow, and states the head-sea result in its limitations. Chapter 6 meets the same question on a vessel whose wave motion is real and not an error.

> **TRY IT IN THE APP 6 · Waves: Hold the Height, or Follow the Sea**
>
> Load the case and press Run. Watch the main foil on the Sea view as a trough passes over it: at 11.20 s it ventilates and turns red, and at 11.65 s the hull slams into the next crest. Press Reset, set In waves to Contour and run again. The boat flies the 40 s clean, at a seat acceleration of 2.67 m/s^2.

5.6.5 A Sea Nothing Tried Could Fly

Case 9 sends the challenge's suggested design into a 1.4 m sea from ahead, 1.75 times the band, on draw seed 2, and asks for 60 s clean at 20 kn. At 20 kn the boat meets 44 waves in the minute at the bow, 19 of them taller than the band from crest to trough and the tallest 1.83 m, and the part of the sea met faster than 3 rad/s has a significant height of 1.03 m on its own. Platforming cannot hold the boat level over waves that tall, and contouring can follow the long waves but not the fast remainder. The suggested design slams at 6.14 s at 2.0 m/s, as shown in Figure 5.4.

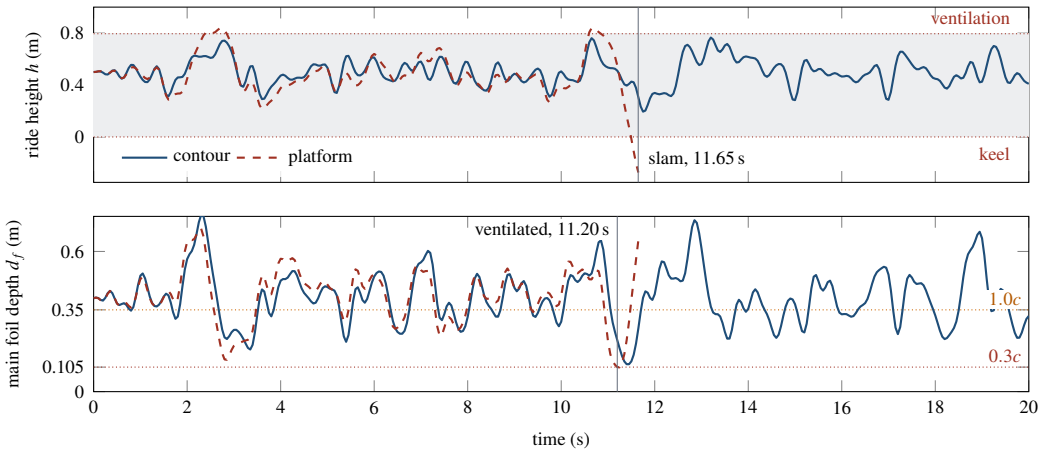

Figure 5.2 Platforming against contouring on the 0.5 m sea of case 6, PID, the same seed. Platforming (red, dashed) lets a trough uncover the main foil past $0.3c$ at 11.20 s and slams on the next crest at 11.65 s, where the run stops. Contouring (blue, solid) keeps the foil above $0.3c$ and the ride height inside the band. The bow can pass the calm-water band without an event, because the surface under the foil differs from that under the bow. The limit lines are dotted. Data from the engine every 0.05 s.

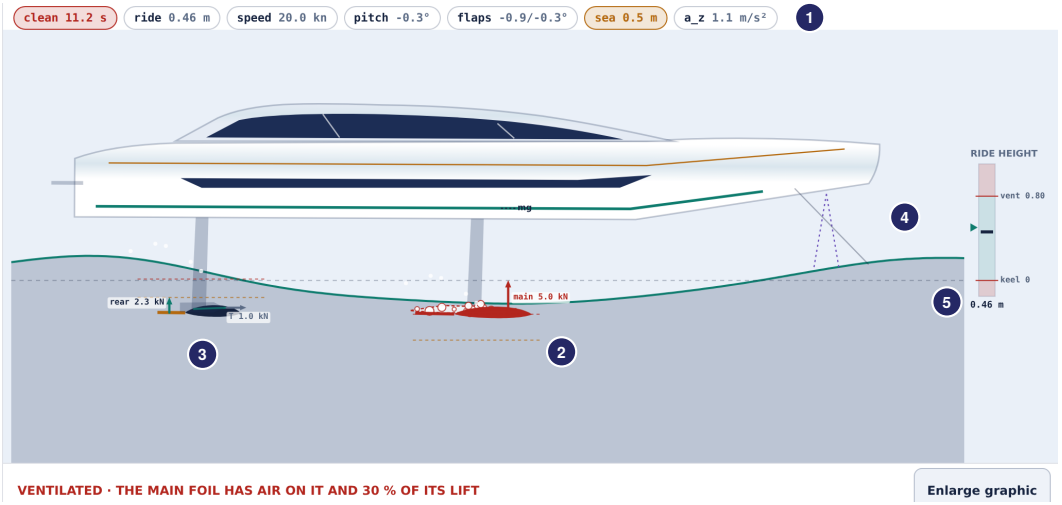

Figure 5.3 Case 6 at its ventilation, $t = 11.25$ s, on the Sea view, cropped to the plant view. The PID platforms on a 0.5 m sea from astern at 20 kn. (1) The badges show 11.2 s clean, flaps at −0.9 and −0.3° and 1.1 m/s^2. (2) The main foil is 0.10 m under a trough, drawn red under an air sheet with bubbles up its strut, and its lift arrow shows 5.0 kN against (3) the 2.3 kN of the rear foil and the mg tick at 15.7 kN. (4) The bow sensor's beam and the wand reach down to the water at the bow. (5) The gauge shows the true ride height at 0.46 m inside the band from a slam at 0 to a ventilation at 0.80 m, and the status line below reads VENTILATED. The trends and the scorecard of the full window are cut.

The infeasibility is a statement about a search, not a proof. A first sweep held the speed at 20 kn and drew 27 795 settings of every control in reach (the structure including the hand and the wand, the wave handling, the cutoff, the ride height, the pitch setpoint, every gain including zero, the hand flap

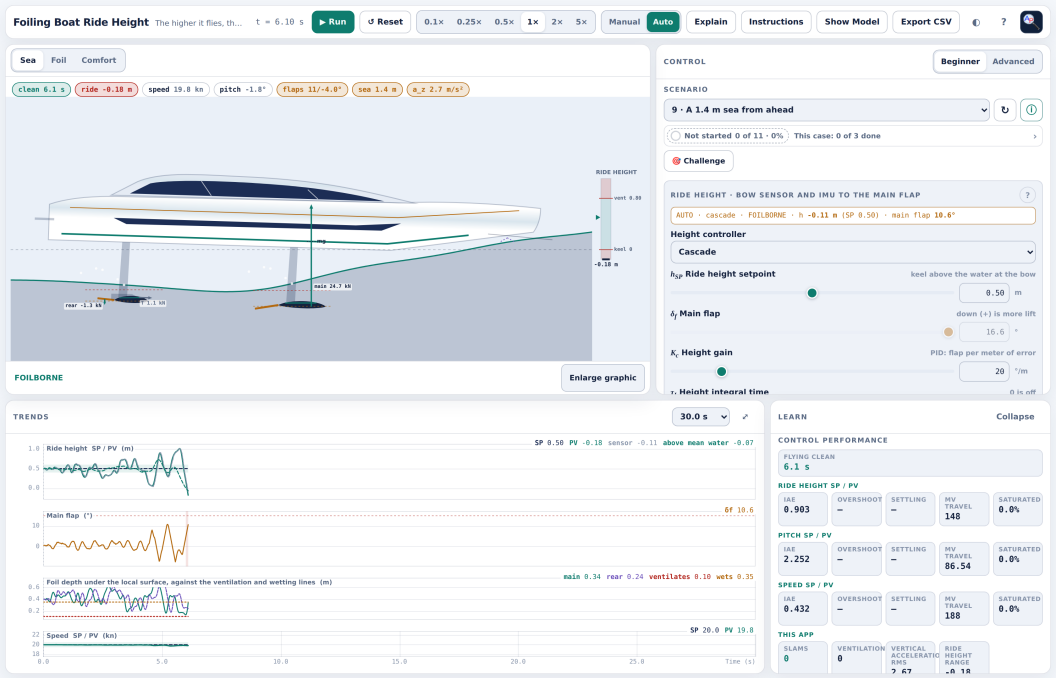

Figure 5.4 Case 9 at $t = 6.10$ s, 0.04 s before its slam, on the Sea view. The suggested cascade with automatic contouring runs into a 1.4 m sea from ahead at 19.8 kn. The bow is over a rising crest: the keel line at the sensor is 0.18 m below the water there while the stem, whose keel rises toward the bow, is still 0.10 m clear, so the gauge reads below its slam mark before any contact. The main flap stands at 10.6° on its way to the 15° command, and the main foil, 0.34 m deep, makes 24.7 kN against the 15.7 kN weight while the rear foil pushes down with 1.3 kN and the bow pitches 1.8° down. The badges show 6.1 s clean and 2.7 m/s^2, and the trends cover the 6.1 s flown.

and the passenger's place), followed by ten rounds of hill-climbing from the best: none of the 33 795 runs flew past 12.49 s. A second sweep of the same size freed the speed from 14.5 to 30 kn and the hand throttle. A hill-climb tuned to this one draw found 59.07 s with a soft PID at about 26 kn, but the ten best settings of that search fly a median of 4.2 to 5.8 s on fifteen other draws of the same sea and none flies the minute on any of them: the search found a passage through these particular waves, not a way to fly the sea. The case rewards the seaman's decision. With the speed setpoint at 12 kn before the run, the boat never leaves its hull and the longest moment with the hull clear is 0.07 s, with no slam in the minute, while slowing from flight at 3 s brings the hull down onto a crest at 4.61 s. The same suggested design flies the challenge's 0.5 m sea from astern clean.

> **TRY IT IN THE APP** **9 · A 1.4 m Sea From Ahead**
>
> Load the case and press Run. The suggested design slams at 6.14 s. Press Reset, set the Speed setpoint to 12 kn and run again. The boat never leaves its hull, and the minute passes with no slam. Slowing below the takeoff speed is the decision this case rewards.

5.6.6 Predictive Control Against the Cascade

Case 10 flies case 6's sea under the MPC, then without the look-ahead, then on the cascade. Table 5.5 sets the MPC at case 10's settings against the shipped cascade and three classical rows, one fixed setting each, with the same seed, instruments and noise. The retuned cascade was chosen from 192

Table 5.5 Seat Acceleration (m/s^2 rms) of a Clean Run, Otherwise the Time in Seconds of the First Slam (s) or Ventilation (v), 20 kn

	Shipped		Re-	Plat.	Cont.	MPC		
Sea	Auto	Cont.	tuned	$+\hat{\eta}$	$+\Delta\hat{\eta}$	On	Off	Off_0
case 6's, 0.5 m	1.30	1.38	0.90	0.71	0.81	0.74	0.98	14.4v
0.65 m astern	1.36	1.42	0.93	11.0v	0.88	1.00	14.5v	13.9v
0.8 m astern	2.7v	1.45[a]	1.01	11.0v	0.94	1.16	14.4v	2.8v
1.0 m astern	2.7v	2.7s	1.04	2.7v	2.7v	14.4v	14.4v	2.7v
1.2 m astern	2.7v	2.7s	37.5s	2.7v	2.7v	2.7s	2.7s	2.7v
0.5 m ahead	2.83	2.94	2.02	18.8v	34.8v	2.03	2.51	7.1v
case 9	6.1s	6.1s	5.9v	5.1v	5.1v	6.1s	6.2s	5.7v
challenge[b]	1.32	1.35	0.91	0.61	0.78	0.74	1.05	33.9v
1.2 m wake[b]	1.33	1.37	0.92	34.9s	0.80	34.0v	34.0v	31.4s

Case 6's seed, 40 s, except case 9 and the challenge (seed 2, 60 s). Plat. $+\hat{\eta}$ and Cont. $+\Delta\hat{\eta}$: the cascade with the previewed water or its change on its setpoint (script only). Off: the water under the bow now held. Off_0: the mean water (the first Off). The MPC planned every 20 ms, 2,000 to 3,000 plans a run, no fallback, 0.19 to 0.22 ms a plan headless. [a] One touchdown below the slam speed. [b] Classical rows with Fly on IMU, the MPC with Hold flap. Last row: the challenge with a 1.2 m wake.

settings, 50 passing, by the MPC's rule: $K_h = 2\,\text{s}^{-1}$, an inner gain of 16° per m/s, $\tau_{I,w} = 1$ s, Auto and $\omega_c = 0.3$ rad/s, within the panel's ranges and with no preview. The two feedforward rows exist only in the measurement script and add the MPC's own preview to the cascade's setpoint: the previewed water at the foil at a gain of 0.5 while platforming, and the best of 72 forms under the rule, 0.75 times its previewed change while contouring.

On case 6's sea the MPC flies clean at 0.74 m/s^2 with 556° of flap travel, against 1.30 m/s^2 and 729° for the shipped cascade on Auto, and its plan presses on a band at 67 of its 2,000 decisions, by at most 3.6 cm. Without the look-ahead it still flies clean, at 0.98 m/s^2 with 781°, so the preview accounts for 0.24 of the 0.56 m/s^2 by which the MPC beats the shipped cascade there. At 0.65 m Off ventilates at 14.5 s where On is clean. A first Off, against the mean water, also dropped the water under the bow now and so ventilated on case 6's sea at 14.4 s for a reason other than the missing look-ahead.

The MPC is not the only structure that flies the taller following seas. The retuned cascade, with no preview, flies 0.65, 0.8 and 1.0 m astern clean at 0.93, 1.01 and 1.04 m/s^2, where the MPC rides 1.00 and 1.16 and ventilates at 14.4 s, and in the 1.2 m sea it lasts to 37.5 s where the MPC slams at 2.7 s. The contouring cascade with the previewed change rides 0.65 and 0.8 m at 0.88 and 0.94 m/s^2. The platforming feedforward matches the MPC on case 6's sea, 0.71 against 0.74 m/s^2, and beats it on the challenge, 0.61 against 0.74, but ventilates at 11.0 s at 0.65 m. Figure 5.5 shows two of these seas. By seat acceleration the MPC is the smoothest of the panel's settings only on case 6's sea and the challenge, and the script's platforming feedforward is smoother on both. What the MPC brings is the flap's box, the servo's rate and both keel bands as limits of one plan, a slack that states by how much a band cannot be met, and, against the shipped Auto, the 0.8 m sea flown clean and the 1.0 m sea to 14.4 s where Auto ventilates at 2.7 s in both. Its peak seat acceleration is higher than the shipped cascade's at 0.65 m, 5.8 against 4.5 m/s^2.

Case 9 and following seas of 1.2 m and more defeat every row. In case 9 the MPC slams at 6.1 s as the shipped cascade does, with 18 slams and 21 ventilations in the minute against 15 and 16 if the run continues and a slack of 0.78 m. With the challenge's wake at its 1.2 m maximum the MPC ventilates at 34.0 s where the shipped Auto flies clean, since an approaching wave train is where the frozen surface errs most. No fix measured keeps the other seas within the budget: a slack every second step flies that wake but slams at 1.0 m and costs 6.3 ms a solve, and previews blended or fading ventilate at 0.65 or 0.8 m. In case 8 the MPC is still 3.4 cm low at 29 s, where the PID with a 2 s integral is within 0.5 cm by 10 s.

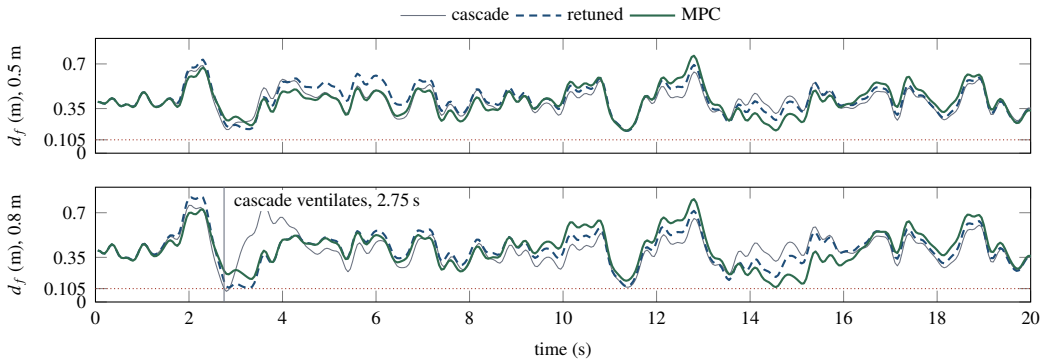

Figure 5.5 Main foil depth under the local surface on case 6's 0.5 m sea (top) and on the same seed at 0.8 m (bottom), from astern at 20 kn, under the shipped cascade on Auto (gray, thin), the cascade retuned by the MPC's rule (blue, dashed) and the MPC with its preview (green, thick solid), with the ventilation depth $0.3c$ dotted. In the taller sea the shipped cascade ventilates at 2.75 s, while the retuned cascade and the MPC keep the foil deeper than $0.3c$, with least depths of 0.110 and 0.111 m over the 40 s. Data from the engine every 0.05 s.

Table 5.6 The Challenge on Its Sea (Seed 2) and Over Sixteen Seas

Configuration	Clean (s)	a_{seat} (m/s^2)	Verdict	Pass of 16
PID, platform, Hold flap (shipped)	40.9	1.32	fails	1
cascade, Auto, Fly on IMU (suggested)	60.0	1.32	passes	12
cascade, Auto, Hold flap	46.3	1.40	fails	—
cascade, platform, Fly on IMU	40.9	0.87	fails	7
cascade, contour, Fly on IMU	60.0	1.35	passes	12
PID, Auto, Fly on IMU	57.8	2.60	fails	—
wand	60.0	4.18	fails (comfort)	—
MPC, preview On, Hold flap	60.0	0.74	passes	—

5.6.7 The Challenge

The challenge flies 60 s at 20 kn on a sea from astern that builds from 0.2 toward 0.5 m, through a 0.5 m wake at 30 s and 2 s of spray on the sensor at 45 s. Its target is the whole minute foilborne with no slam and no ventilation and a seat acceleration of at most 1.5 m/s^2, and its score is the seconds flown clean. The comfort limit is a teaching value. Ratings of discomfort under whole-body vibration come from the literature and standards on human vibration [17]. The challenge is void if a model parameter, the sea, the wake or the spray is changed. The sea is one fixed draw, seed 2, chosen because it shows the result that most draws show, and Table 5.6 lists eight configurations on it with the pass count over sixteen seas where it was measured. As shipped, the PID platforms and freezes the flap in the spray, and it ventilates at 40.91 s as the sea passes 0.405 m. The suggested design, the cascade with automatic contouring and the IMU fallback, flies the minute clean at 1.32 m/s^2. Each part matters on this sea, and contouring throughout also passes, a result the application states, and so does the MPC. Figure 5.6 draws the table as bars.

> ✦ **DID YOU KNOW? A Stick That Flies but Shakes**
>
> From the simulator: on the challenge's sea the wand, a stick and a pushrod with no electronics, flies all 60.0 s with no slam and no ventilation. It still fails, at a seat acceleration of 4.18 m/s^2 against the limit of 1.5, which the suggested cascade meets at 1.32 (Table 5.6).

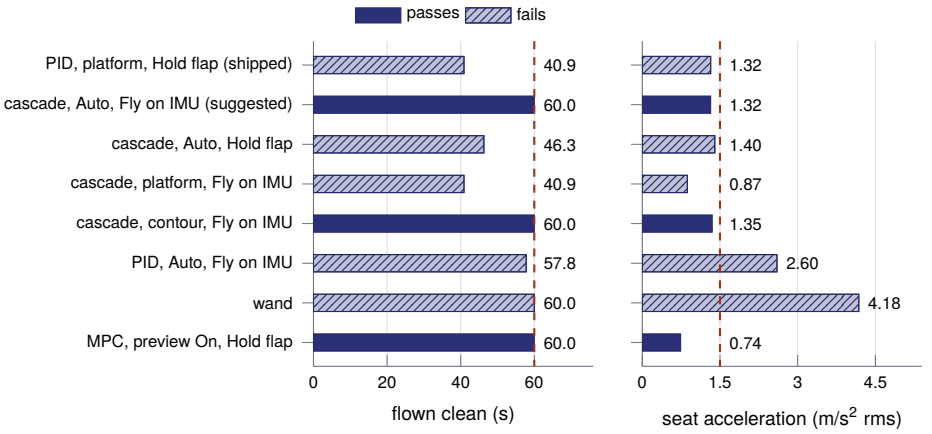

Figure 5.6 The eight configurations of Table 5.6 on the challenge's sea, drawn as bars, solid for a pass and hatched for a fail. Left: the seconds flown clean, against the 60 s the challenge asks (dashed line). Right: the seat acceleration, against the comfort limit of 1.5 m/s^2 (dashed line). A value left of a dashed line is printed just right of it. Notice that the wand flies the whole minute and fails on comfort alone, and that the shipped and the platforming configurations stop at the same 40.9 s.

5.7 Technology in Practice

The launch of Section 5.3 is read by a bow height sensor, an inertial unit and a speed log, and moved by two foil flaps and an electric pod. Figure 5.7 draws its loops with real devices in their place.

What Measures It

Every height law of Section 5.5 rests on the bow sensor LT of Figure 5.1. Its 1 cm of colored noise and the spray event of case 5 stand for the noise and the lost echoes of a real ultrasonic sensor. Naval hydrofoils sensed pitch and roll with a vertical gyroscope and heave with a servoed accelerometer beside the height sensor [18], the same set as the IMU and LT here.

> **TECHNOLOGY** **Sensor: Ultrasonic Ride-Height Sensor**
>
> An ultrasonic sensor times a sound pulse from a piezoelectric transducer to the surface and back, and the time of flight gives the distance. Its calibration rests on the speed of sound in the air, which changes with the air's temperature and pressure. Foam can fool it, and every echo instrument has a dead zone close to its face [23]. Sonic height sensors flew on the early electronically controlled hydrofoils, and one naval craft changed to a radar height sensor to eliminate noise and heavy rain interference [18]. On a trend such interference shows as noise or as a lost reading, which is how the spray of case 5 blinds the bow sensor for 2 s.

What Moves It

The model drives each flap through an 8 ms lag, a rate limit of 60 °/s and a stop at 15° or 10°, and the pod through a 0.3 s thrust lag under its thrust and power limits. On a small foiling dinghy a mechanism does the height loop's job. At rest its wand at the bow lies almost horizontal and pushes the front foil's flap fully down for the most lift, and as the boat rises the wand turns toward the vertical and reduces the flap [8]. The wand of Equation (5.8) is that linkage.

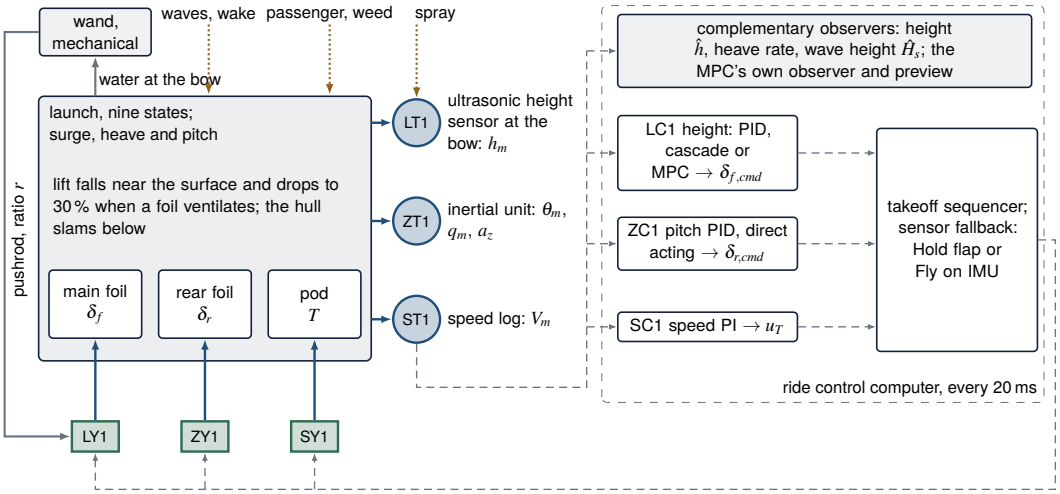

LY1, ZY1: main and rear flap servos, lag 8 ms, 60 °/s, stops at ±15° and ±10°. SY1: pod motor drive, thrust lag 0.3 s, 4.5 kN and 60 kW. Wand: a stick from the bow to the water, linked to the main flap with ratio r, in place of LC1.

Figure 5.7 The loops of the chapter with real devices in place of the schematic's symbols. LT1 reads the ride height at the bow with an ultrasonic sensor, ZT1 the pitch, pitch rate and vertical acceleration with an inertial unit and ST1 the speed with a log. One ride control computer (dashed frame) runs the observers, the height, pitch and speed loops (LC1, ZC1, SC1), the takeoff sequencer and the sensor fallback every 20 ms. LY1 and ZY1 are the servos of the main and rear flaps and SY1 the pod's motor drive. Notice that the wand reaches the main flap through a pushrod and no computer, and that spray on the bow sensor is the only disturbance that enters at an instrument.

TECHNOLOGY Final Element: Foil Flap on a Hydraulic Servo

The naval hydrofoils of a 1974 survey made their lift control with trailing-edge flaps or with whole foils pivoted in incidence. An electrohydraulic servo valve drove each control surface from a hydraulic supply at 3000 psi with redundant pumps. In rough water the outboard flaps of one craft were bottoming quite frequently, the sign that they lacked authority [18]. A flap held at its stop shows on a trend as a flat line while the height error grows, as when 40 % weed at 18 kn drives the main flap of case 8 to its 15° stop.

Where the Controller Runs

The autopilots of naval hydrofoils in 1974 were analog, hard-wired for each ship from plug-in modules, and four servo amplifiers drove the control surfaces, one on each foil and the rudder, through electrohydraulic servo valves [18]. An electronic ride control built for a small foiling dinghy runs on a single-board computer: one software node reads each sensor, a PID node computes the flap, and knobs on the boat set the gains [8]. No open source consulted gives the cycle time of a ride control computer. The PID, the cascade, the MPC, the pitch and speed loops and the sequencer of this chapter all run every 20 ms, with the observers in the same cycle.

What Is Modernizing It

Three developments bear on the chapter's question. The first is radar height sensing in place of sonic, chosen on one naval craft to remove noise and rain interference [18]. The second is the digital autopilot. A universal digital autopilot program for naval hydrofoils, started in 1972, ran a background program that checked every element ten times a second and switched to a backup circuit when it found a fault [18]. The third is electronic ride control on small sailing craft. An ultrasonic sensor and an attitude sensor replace the wand, damping the height signal makes the boat platform and

amplifying it makes it contour, and the sailor can change both the ride height and the response to the sea while sailing [8]. The wave handling of Section 5.5 is a choice of the same kind, made by the controller instead of the sailor.

> **IN PRACTICE** **Design Rules for a Foil Control System**
>
> A 1977 specification for naval hydrofoil control systems asked each loop for a gain margin of 2 to 1 upward and 4 to 1 downward and a phase margin of 30 degrees. The effective strut length had to be at least 1.4 times the significant wave height that 90 percent of the seas stay below. In contouring, strut immersion could vary at most 25 percent more than the wave height for waves met slower than 5 rad/s. The document defines a foil broach as a foil that comes near or clear of the surface and loses its lift, and a hull slam as the impact on the next wave that follows it [24], the sequence of case 6.

5.8 Limitations

The boat is a consistent teaching design, not a model of any craft, and its dimensions, foil sizes, struts and limits were chosen so that the lessons appear. The surface factor and the ventilation latch are declared forms that have not been checked against measured data, so the depth at which the lift starts to fall, the retained 30 % and the wetting depth are plausible rather than predicted, and the flow has no partially ventilated regime. The lift is quasi-steady, the waves are linear and the boat has three degrees of freedom, so the wand's fast oscillation, the steep crests of case 9 and any oblique sea are where the model is weakest.

The cases rest on fixed draws of a random sea, and two of their lessons are draw-dependent and stated so: the PID's contouring is clean on 7 of 17 draws of the sea of case 6, and the challenge's suggested design passes on 12 of sixteen. The infeasibility of case 9 is the result of two searches of 33 795 settings, not a proof. The ride height is measured at one point, so in steep waves it can read below the slam mark with no contact, as in Figure 5.4. Several true quantities in the CSV export, the foil depths and lifts among them, carry no `truth_` prefix, although no controller reads them.

The MPC's model is linear about one trim and leaves out the surge, the orbital angle of attack except through a disturbance fading in 0.1 s, the surface factor's change with depth, ventilation and the stall, so it predicts no ventilation. Its frozen-surface preview errs by 0.104 m rms 0.16 s ahead on case 6's sea, and its results at 0.65 and 0.8 m are in-sample.

> **OPEN CHALLENGE** **Predicting When a Foil Ventilates**
>
> The triggers of natural ventilation are diverse and seemingly stochastic, which makes its onset hard to predict. The most complete map of the flow regimes holds only for the foil and aspect ratio it was measured on. A recent towing-tank study found that two paths to the same steady condition can end in different regimes, and that the onset depends on small disturbances of the free surface [25]. Progress would be ventilation maps that carry over to new foils, to waves and to a moving craft, so that a latch like the one of this chapter could be fitted rather than declared.

5.9 Conclusion

A foiling launch flies in a band 0.80 m tall between a slam and a ventilation, and this simulator lets a student find the band by losing it. Near the surface the foil gives the boat a free stiffness of 24.0 kN/m, which falls to 0.39 kN/m with the foil two chords deep, so a deep foil needs a controller. A wand makes that controller from a stick, with its ratio as the gain and its length as the setpoint, until a ratio of 2.7 porpoises the boat out of the water. In a 0.5 m sea from astern platforming ventilates at 11.20 s and slams at 11.65 s while contouring flies clean, and the cascade on the IMU's heave rate contours that sea clean on all 17 draws against 7 for the PID. In a 1.4 m sea from ahead no setting found flies the minute, and the right decision is to slow below the takeoff speed. An optional MPC rides case 6's sea at 0.74 m/s^2 against the shipped cascade's 1.30, but a cascade retuned by the same rule rides the 0.65 to 1.0 m following seas more smoothly, so the MPC adds limits and a slack in one plan rather

than comfort. Further work would add roll and oblique seas, an unsteady lift lag, a partially ventilated regime fitted to towing-tank data, and the planing lift of the hull. The simulator shows that a height loop on a lift surface near a boundary has to choose which disturbances to reject.

KEY POINTS

- A foiling boat flies in a narrow band: too low and the hull slams, too high and the main foil draws air and loses most of its lift.
- Near the surface the foil gives a free stiffness of 24.0 kN/m, which falls to 0.39 kN/m two chords deep, so a deep foil needs a controller.
- In a 0.5 m sea from astern contouring flies clean while platforming ventilates and slams, and the cascade is clean on all 17 draws against 7 for the PID.
- In a 1.4 m sea from ahead nothing tried flies the minute, and the right decision is to slow below the takeoff speed.

Where else this applies. Volume 7, Chapter 11 trades a still body against a wheel that follows the road, also with a preview, and Chapter 6 chooses which wave motion a vessel should leave alone.

References

[1] K. V. Rozhdestvensky, "Wing-in-ground effect vehicles," *Progress in Aerospace Sciences*, vol. 42, no. 3, pp. 211–283, 2006.

[2] A. J. Acosta, "Hydrofoils and hydrofoil craft," *Annual Review of Fluid Mechanics*, vol. 5, no. 1, pp. 161–184, 1973.

[3] O. M. Faltinsen, *Hydrodynamics of High-Speed Marine Vehicles*. Cambridge: Cambridge University Press, 2005.

[4] C. M. Harwood, Y. L. Young, and S. L. Ceccio, "Ventilated cavities on a surface-piercing hydrofoil at moderate Froude numbers: cavity formation, elimination and stability," *Journal of Fluid Mechanics*, vol. 800, pp. 5–56, 2016.

[5] Y. L. Young, C. M. Harwood, F. Miguel Montero, J. C. Ward, and S. L. Ceccio, "Ventilation of lifting bodies: Review of the physics and discussion of scaling effects," *Applied Mechanics Reviews*, vol. 69, no. 1, p. 010801, 2017.

[6] D. Hrovat, "Survey of advanced suspension developments and related optimal control applications," *Automatica*, vol. 33, no. 10, pp. 1781–1817, 1997.

[7] W. Higgins, "A comparison of complementary and Kalman filtering," *IEEE Transactions on Aerospace and Electronic Systems*, vol. AES-11, no. 3, pp. 321–325, 1975.

[8] S. Kebbell and J. Binns, "Development of a full scale Moth hydrofoil control system test rig," in *The 9th Conference on Computational Methods in Marine Engineering (Marine 2021)*, vol. 1. Edinburgh University Library, 2022.

[9] K. L. Wadlin, C. L. Shuford, Jr., and J. R. McGehee, "A theoretical and experimental investigation of the lift and drag characteristics of hydrofoils at subcritical and supercritical speeds," National Advisory Committee for Aeronautics, NACA Report 1232, 1955, NASA Technical Reports Server document 19930092241.

[10] S.-H. Kim and H. Yamato, "An experimental study of the longitudinal motion control of a fully submerged hydrofoil model in following seas," *Ocean Engineering*, vol. 31, no. 5-6, pp. 523–537, 2004.

[11] J. Bai and Y. Kim, "Control of the vertical motion of a hydrofoil vessel," *Ships and Offshore Structures*, vol. 5, no. 3, pp. 189–198, 2010.

[12] R. J. Johnston, "Hydrofoils," *Naval Engineers Journal*, vol. 97, no. 2, pp. 142–199, 1985.

[13] J. P. Breslin and R. Skalak, "Exploratory study of ventilated flows about yawed surface-piercing struts," National Aeronautics and Space Administration, NASA Memorandum 2-23-59W, Apr. 1959, NASA Technical Reports Server document 19980228299.

[14] S.-H. Kim and H. Yamato, "On the design of a longitudinal motion control system of a fully-submerged hydrofoil craft based on the optimal preview servo system," *Ocean Engineering*, vol. 31, no. 13, pp. 1637–1653, 2004.

[15] B. Beaver and J. Zseleczky, "Full scale measurements on a hydrofoil International Moth," in *SNAME 19th Chesapeake Sailing Yacht Symposium*. SNAME, 2009, paper D021S002R006.

[16] W. J. Pierson and L. Moskowitz, "A proposed spectral form for fully developed wind seas based on the similarity theory of S. A. Kitaigorodskii," *Journal of Geophysical Research*, vol. 69, no. 24, pp. 5181–5190, 1964.

[17] M. J. Griffin, "Vibration discomfort," in *Handbook of Human Vibration*. Academic Press, 1990, pp. 43–123.

[18] R. J. Johnston and W. C. O'Neill, "The development of automatic control systems for hydrofoil craft," in *International Hovercraft, Hydrofoil and Advanced Transit Systems Conference*, Brighton, England, May 1974, pp. 265–273, copy in the library of the International Hydrofoil Society, https://foils.org/development-of-automatic-control-systems-for-hydrofoil-craft/.

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

[20] S. J. Qin and T. A. Badgwell, "A survey of industrial model predictive control technology," *Control Engineering Practice*, vol. 11, no. 7, pp. 733–764, 2003.

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

[22] M. Tomizuka, "'Optimum Linear Preview Control With Application to Vehicle Suspension'—Revisited," *Journal of Dynamic Systems, Measurement, and Control*, vol. 98, no. 3, pp. 309–315, 1976.

[23] T. R. Kuphaldt, *Lessons in Industrial Instrumentation*, version 2.33 ed. Open textbook, Creative Commons Attribution 4.0 license, ibiblio.org/kuphaldt/socratic/sinst, 2024.

[24] D. R. Stark, W. E. Farris, and A. O. Harang, "Hydrofoil ship control and dynamics specifications and criteria, volume III," Boeing Marine Systems, for the David Taylor Naval Ship Research and Development Center, Document D321-51313-1, Mar. 1977, contract N00600-75-C-1107. Copy in the library of the International Hydrofoil Society, https://foils.org/control/.

[25] M. Aguiar Ferreira, C. Navas Rodríguez, G. Jacobi, D. Fiscaletti, A. Greidanus, and J. Westerweel, "On the ventilation of surface-piercing hydrofoils under steady-state conditions," *Journal of Fluid Mechanics*, vol. 1028, p. A25, 2026.

5.10 Exercises

5.1 ⌕ Concept check · 2 min

The PID flies at 20 kn in a 0.2 m sea from astern. Which wave handling gives the lower vertical acceleration at the seat, and why?

(a) Contouring, because following the surface keeps the main foil at one depth under every wave

(b) Contouring, because a small sea from astern is overtaken slowly enough for every wave to be followed

(c) Platforming, because the short waves are met near 5.5 rad/s and are small enough to pass under the hull

(d) Platforming, because the integral action of the PID removes the wave error that contouring adds

5.2 ⌕ Concept check · 1 min

With both flaps fixed at 20 kn, the boat trimmed at a ride height of 0.75 m settles within about a second after a small loss of lift, while the boat trimmed at 0.20 m keeps sinking toward the water. What makes the difference?

(a) At 0.20 m the heave damping is much smaller, so the boat is free to accelerate

(b) At 0.20 m the main foil works close to its stall, so it cannot add lift by angle

(c) At 0.20 m the rear foil carries the weight, and the main foil adds no restoring force

(d) At 0.20 m the main foil is two chords deep, where the surface factor hardly changes with depth

5.3 ⌕ Concept check · 2 min

The pitch loop drives the rear flap from the measured pitch, positive bow up. A positive rear flap lifts the stern. In the convention of Section 1.3, which statement about its action is right?

(a) Direct acting, $a = -1$: more rear flap lowers the pitch, so the output rises when the pitch is above its setpoint

(b) Reverse acting, $a = +1$: more rear flap raises the pitch, so the output rises when the pitch is below its setpoint

(c) Reverse acting, $a = +1$: the height loop on the main flap is reverse acting, and both flaps push the boat upward

(d) Either action works, because the integral term learns the sign of the process gain during the run

5.4 Theory and calculation · 6 min

Take the main foil as carrying the whole weight mg of the default boat ($m = 1600\,\mathrm{kg}$, $g = 9.81\,\mathrm{m/s^2}$) at fixed speed and angle, so that its lift is proportional to the surface factor $f(d/c) = 1 - A_s \mathrm{e}^{-B_s d/c}$ with $A_s = 0.5$, $B_s = 2.5$ and chord $c = 0.35\,\mathrm{m}$. In level flight a rise Δh of the boat reduces the foil depth d by Δh.

(a) Show that the free stiffness is $k = mgA_sB_s\mathrm{e}^{-B_s d/c}/\big(c\,f(d/c)\big)$ and evaluate it at $d = 0.15$, 0.40 and 0.70 m. Give three significant figures. *Checkpoint:* 23.2, 3.31 and 0.379 kN/m.

(b) Compute the ratio of the stiffness at 0.15 m to that at 0.70 m and compare it with the measured stiffnesses of Table 5.3. *Checkpoint:* 61.1 against $24.00/0.39 = 61.5$.

(c) With the measured damping of 27.0 kN s/m and stiffness of 3.43 kN/m at the 0.50 m ride height, find the time constant of the free height. *Checkpoint:* 7.9 s.

5.5 Theory and calculation · 6 min

The sea's peak period is $T_p = c_T\sqrt{H_s}$ with $c_T = 5.0\,\mathrm{s/m^{1/2}}$, a deep-water wave of frequency ω has the wavenumber $k = \omega^2/g$, and a boat at speed V meets it at $\omega_e = |kV - s\omega|$, with $s = +1$ from astern and $s = -1$ from ahead. The boat flies at 20 kn, which is 10.29 m/s.

(a) For $H_s = 0.5\,\mathrm{m}$ find the peak period, the peak frequency and its wavenumber. *Checkpoint:* 3.54 s, 1.78 rad/s and 0.322 rad/m.

(b) Find the encounter frequency of that peak from astern and from ahead, and the encounter period from astern. *Checkpoint:* 1.54 rad/s and 5.09 rad/s; 4.09 s.

(c) Repeat the encounter frequency from astern for $H_s = 0.2\,\mathrm{m}$, and compare all three with the platform observer's 0.6 rad/s and the contour observer's 8 rad/s. *Checkpoint:* 5.47 rad/s.

5.6 Theory and calculation · 6 min

The wand sets the main flap to $\delta_f = \delta_n - r(\varphi - \varphi_n)$, with φ its angle below the horizontal, $\varphi_n = 53^\circ$, the linkage ratio r in degrees of flap per degree of wand, and the pivot 0.30 m above the keel. The boat settles near φ_n, so the pivot flies about $\ell_w \sin\varphi_n$ above the water.

(a) Find the wand's gain $\partial\delta_f/\partial h = r/(\ell_w\cos\varphi)$ in degrees of flap per meter for $r = 0.6$ and $\ell_w = 1.0\,\mathrm{m}$, and compare it with the PID's default $K_c = 20^\circ/\mathrm{m}$. *Checkpoint:* 57.1 °/m, 2.9 times the PID's.

(b) Predict the ride height $\ell_w \sin\varphi_n - 0.30\,\mathrm{m}$ for wand lengths of 0.8, 1.2 and 1.3 m, and compare with the heights measured in case 4: 0.34, 0.65 and 0.72 m. *Checkpoint:* 0.339, 0.658 and 0.738 m.

(c) Find the gain at $r = 2.7$ with the default length. *Checkpoint:* 257 °/m.

5.7 App challenge · 8 min

Load scenario 3, The free stiffness: open loop near the surface and deep. The boat flies at 0.75 m with the main flap fixed by hand, and the flap moves 0.5° up (less lift) at 3 s. Use the default seed and speed 1, run the whole 15 s and export the CSV file.

(a) Read the true ride height, column `h_m`, in the rows at 2.90 s and 8.00 s. *Checkpoint:* 0.750 m and 0.709 m, within 0.003 m.

(b) Press Reset, set the Ride height setpoint to 0.50 m, which is the height the boat starts trimmed at, run the 15 s again and export. Read `h_m` at 4.00, 8.00 and 14.90 s. *Checkpoint:* 0.434, 0.245 and −0.004 m, within 0.005 m.

(c) Explain the two responses with the time constants of Table 5.3.

5.8 App challenge · 9 min

Load scenario 6, Waves: hold the height, or follow the sea. The PID platforms with a cutoff of 0.6 rad/s on a 0.5 m sea from astern at 20 kn. Use the default seed and speed 1.

(a) Run as shipped until the run stops. Read the Flying clean tile and the time on the clock when the run pauses. *Checkpoint:* 11.2 s clean, paused at 11.65 s, within 0.05 s.

(b) Press Reset, set In waves to Contour and run the whole 40 s. Read the Slams, Ventilations and Vertical acceleration rms tiles. *Checkpoint:* 0, 0 and 2.67 m/s^2, within 0.02 m/s^2.

(c) Press the reload button, set the Platform cutoff to 1.5 rad/s, keep Platform and run the whole 40 s. Read the same tiles. *Checkpoint:* 0, 0 and 1.29 m/s^2, within 0.02 m/s^2.

Explain the three results with the encounter frequency of this sea.

5.9 App challenge · 8 min

Load scenario 10, MPC: the water the bow has already measured. The MPC flies case 6's 0.5 m sea from astern with its preview On. Use the default seed and speed 1 and run each case for the whole 40 s. Read the Vertical acceleration rms tile each time, within 0.02 m/s^2.

(a) Run as shipped. *Checkpoint:* 0.74 m/s^2, no slam or ventilation.

(b) Press Reset, set the MPC preview to Off (Advanced level) and run again. *Checkpoint:* 0.98 m/s^2, no slam or ventilation.

(c) Press Reset, set the Height controller to Cascade and run again. *Checkpoint:* 1.30 m/s^2, no slam or ventilation.

(d) Compute the share of the MPC's advantage over the cascade that the preview accounts for, defined as (Off − On)/(Cascade − On) from the three readings. *Checkpoint:* 43 %, within 4 points.

5.10 Application area · 4 min

The declared boat's takeoff speed is 14.48 kn. The sequencer enters HUMP above 60 % of it, returns to DISPLACEMENT below 85 % with the keel down, and lands the boat when the speed setpoint is below 95 %.

(a) Compute the three speeds in knots. *Checkpoint:* 8.69, 12.31 and 13.76 kn.

(b) In case 9's 1.4 m head sea, setting the speed setpoint to 12 kn before Run keeps the hull in the water with no slam, while slowing from flight at 3 s puts the hull down onto a crest at 4.61 s. What does this tell a skipper about when to take the decision?

5.11 Application area · 3 min

The app models ventilation as a latch that sets when the foil is shallower than 0.3 chords with an effective angle above 2°, keeps 30 % of the lift, and clears only below 1.0 chord. Name two things a real ventilating foil depends on that the latch leaves out, and say why the hysteresis matters to a height loop that has just lost lift.

5.12 Review the agent · 5 min

An AI assistant reviews this chapter's results and writes: "On case 6's sea the MPC rides at 0.74 m/s^2 against 1.30 m/s^2 for the cascade, a 43 % cut in seat acceleration. Replace the cascade with the MPC in every following sea up to 1.0 m." Find the flaw and say how you would check the recommendation.

CHAPTER 6

Choosing What Not to Reject in a Dynamic Positioning Simulator

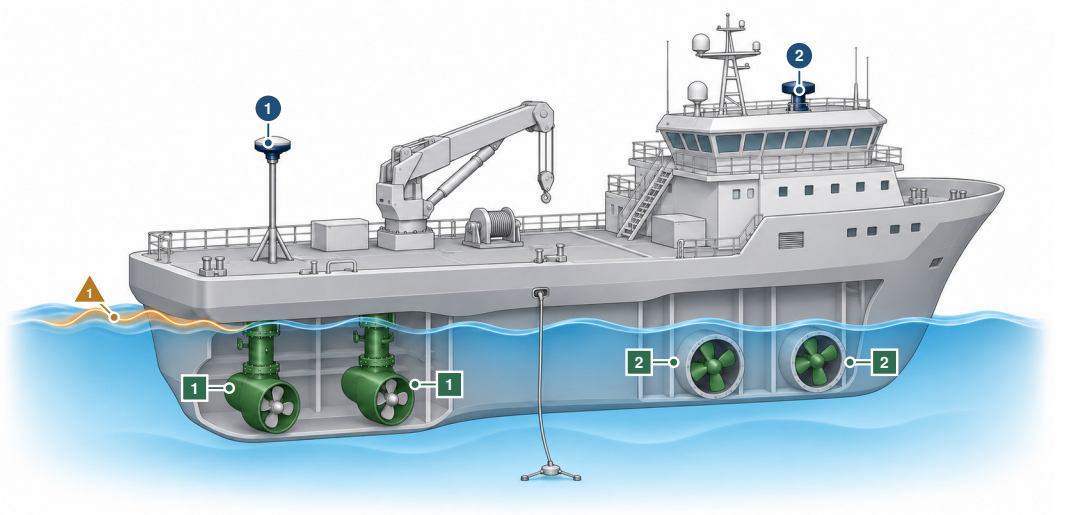

The controller moves 1 two azimuth thrusters aft, which turn to aim their thrust
2 two tunnel thrusters forward, which push sideways

The controller reads 1 position reference 2 heading sensor (gyro)

What disturbs it 1 waves, which rock the hull and push it slowly off station

Notice the line from the hull down to the wellhead the vessel must stay over, and the waves at the waterline: the thrusters should stop the slow drift and leave the rocking with each wave alone. Illustration, not to scale.

IN PLAIN TERMS **Holding a Ship Over One Spot**

The main idea. A work ship at sea often has no anchor down, so its propellers must keep pushing it back over one spot on the seabed. A computer reads where the ship is and tells the propellers how hard to push.

The challenge. The reading also shows the ship rocking with every wave, a motion no propeller can stop and none needs to. A computer that chases the rocking wears out the propellers and holds the ship no better.

Why it matters. Telling the slow drift from the rocking, and turning the bow into the weather, decide how well and how cheaply the ship stays on its spot.

THREE WAYS TO READ THIS CHAPTER

Explore	12 min	the first page, the history of the first deep-water drill ship in the Introduction, the box on the first controllers (Section 6.2), the two facts and the try-it segments of scenarios 3 and 8 (Section 6.6), and the key points
Study	60 min	the whole chapter with its three pause points and the twelve exercises at its end
Practice	20 min	the first page, the measured scenarios (Section 6.6), Technology in Practice (Section 6.7), Limitations (Section 6.8) and the open challenge

The control question The position reference of a vessel on thrusters adds the slow drift that the thrusters must stop to the wave motion that they cannot and need not stop. Which part of that measurement should the controllers be shown, and which decisions outside the three PID loops decide how well and how cheaply the vessel holds station?

Learning objectives

- Derive the inertias and the closed-loop frequency and damping of a position loop on a hull with no restoring force.
- Explain why an observer with bias and wave states, and not a low-pass filter, keeps wave motion out of the thruster commands.
- Compare raw, model-only and full-observer estimates on thruster travel, energy and position, and say which number carries the case.
- Compute the sway and heading capability of the thruster layout, and predict whether thrust, power or heading binds first.
- Evaluate a solved allocation, a heading supervisor and a redundant reference against a fixed split, a fixed heading and a single reference.

Prerequisites Sections 1.6.2 (velocity-form PID), 1.4.3 (position loops and integral action), 1.4.6 (bias states and wave filtering) and 1.4.7 (weighted allocation). Chapter 3 shows estimated-force feedforward on a simpler vehicle.

The app https://apmonitor.com/apps/dynamic-positioning/

See also Volume 5, Chapter 13 for a marine vehicle navigated by a Kalman filter on late position fixes, and Volume 5, Chapter 12 for a sensor whose reading does not say which way the error lies. Volume 4, Chapter 13 and Volume 4, Chapter 12 for allocation across actuators with limits and prices.

6.1 Introduction

A feedback loop is built on the premise that the difference between the measurement and the setpoint is an error, and on most plants the premise is safe. A dynamically positioned vessel is where it fails. The position reference on an offshore construction vessel reports two motions superimposed, the hull pushed off the wellhead by wind and current, and the same hull riding over individual waves. In the 4 m sea of this chapter's results the second reaches 2.39 m at an eleven second period while the first is not permitted to reach the 5 m watch circle at all, so the larger signal on the wire is the one that must be left alone. No thruster stops eight thousand tonnes moving with the sea, and a control system that tries destroys the machinery it depends on while holding station no better for the effort.

The plant underneath is as plain as the measurement is difficult. A vessel on thrusters has no anchor, no mooring and no hydrostatic restoring force in surge, sway or yaw, so all three controlled axes are double integrators carrying a little hydrodynamic damping, and the entire station-keeping capability of the vessel is the control system [1, 2]. The dynamics are trivial to write down and the measurement is the hard part, which is the reverse of the arrangement a first control course usually presents.

> ⧖ HISTORY **A Drill Ship Without Anchors (1961 to 1971)**
>
> Drilling into the deep ocean floor from a dynamically positioned drill ship was first done in March and April of 1961, with radar reflectors and acoustic transducers on taut line buoys as the reference. A single taut line, then in use on drill ships in 600 to 1000 feet of water, was tested and found to have too much error in deep water. The research drill ship accepted in 1968 listened instead to an acoustic beacon dropped on the sea floor, and a computer turned the reading into commands for four tunnel thrusters and two main propellers. It was the first use of tunnel thrusters as the restoring force for dynamic positioning, and the risk of losing the drill string was judged real enough to require redundancy wherever possible [3].

The misconception the simulator corrects is that wave filtering is a filter. A low-pass element on a signal on its way into a controller removes the wave-frequency motion and removes the phase margin with it, because a plant with no restoring force needs derivative action and derivative action needs phase. What industrial dynamic positioning does instead is carry a model of the motion it intends to keep: an observer runs a vessel model, a slowly varying bias state that absorbs the wind and the current, and a lightly damped oscillator tuned to the peak wave period, and divides one innovation among the three so that only the low-frequency estimate reaches the controllers [4, 5]. A block diagram cannot show why that matters, because the cost of getting it wrong is not visible in the controlled variable. It appears in the thruster commands, and this application measures it there. Section 1.4.6 names the idea for every machine in this volume, and Chapter 5 meets it again on a foiling boat that must choose which waves to ride over and which to follow.

Three general classes of problem sit behind the vessel. The first is measurement decomposition, in which part of a measured signal is real, periodic and outside the authority of the actuator, so that the correct response is a structured estimator rather than a filter [5, 6]. The second is disturbance estimation with feedforward, in which a large slow unmeasured load is identified inside the observer, where it can be found quickly without entering the loop, so the feedback can stay slow [7, 8]. The third is control allocation over an actuator set with more independent forces than controlled degrees of freedom, solved across aerospace, automotive and marine applications by the same weighted pseudo-inverse [9]. A reader who has met a level loop on a sloshing surface, or a vehicle with individual wheel actuators, has met one of the three already.

A dynamic positioning system is normally taught from its equations, and the equations are the part that transfers least well: the observer is written in three lines that look like a Kalman filter and hide the one decision the designer actually makes, which is how much of the measurement to believe. This chapter prices each structural decision separately, on an interactive vessel where the estimator, the allocation and the heading policy can each be switched to a naive alternative and the difference read off the same scorecard. Section 6.2 traces the problem from anchor spreads to the Kalman filter. Section 6.3 sets down the eighteen-state vessel, environment and thruster model, with the equations, the constants and the initial conditions in the detail another implementation needs to reproduce the trajectories, and Section 6.4 the observer, the three control loops, the allocation and the heading supervisor. Section 6.5 lists the twelve scenarios, which put the four structural decisions ahead of every gain. Section 6.6 puts a number against each choice on one plant. The sharpest is that the same sea and the same tuning cost 327.5 MN of thruster command travel over the twenty-five minutes from 300 to 1800 s with the raw measurement shown to the loops and 72.0 MN with the full observer, while the vessel holds station no better for the extra effort, 0.55 m off on average against 0.51 m. The most valuable of the four decisions turns out not to be in the control system at all. Section 6.7 names the position references, thrusters and computers of real vessels, and Sections 6.8 to 6.10 state what the model cannot represent and what the measurements say together. The vessel model, the dynamic positioning observer and the pseudo-inverse allocation are all standard results.

6.2 From Anchors to a Kalman Filter

Station keeping began as a mooring problem, solved by naval architects with an anchor spread rather than by control engineers, because the restoring force came from catenary chain and not from feedback. How deep-water drilling first held a ship on thrusters is told in the history box of the Introduction, and how its first controllers kept the sea out of the thrusters in the box below.

> ⧖ HISTORY **A Twenty-Second Lag in Front of the Loop (1967)**
>
> An equipment proposal of 1967 for that drill ship simulated the vessel in three degrees of freedom. The wave motion that leaked into the position data was represented by sine waves of plus or minus 70 feet at a 10 second period, roughly a sea state of 5 or greater. Each loop used rate, position and integral feedback, and the position inputs passed through simple 20 second lags to keep unnecessary demands off the thrusters. Its authors noted that these lags, added to the thruster lags, made a rather slow control loop. After three years at sea the operators asked for a program to which the current and the wind could be fed or sensed, so that the computer could anticipate losing station instead of reacting after moving off station [3].

The modern structure arrived between 1976 and 1983, as an estimation result rather than a control result. The Trondheim group posed dynamic positioning as a Kalman filtering and optimal control problem, with a vessel model, a wave-frequency model and a slowly varying bias estimated together and the controller closed on the low-frequency estimate alone [4, 8]. A British line of work reached the same structure from stochastic optimal control, treating wave-frequency motion as a colored disturbance the controller must not attempt to remove [10]. Both replaced the filter in front of the controller with a model inside an estimator, which made the loop fast enough to be useful and quiet enough to be affordable, and a survey three decades later still finds the same three blocks [2]. Later, a passive nonlinear observer removed the need to linearize about fixed headings [11], estimator tuning was scheduled against the sea state rather than fixed [6], and allocation grew from a fixed split written at commissioning into a solved optimization now standard in ships, aircraft and road vehicles alike [9].

The last piece of the history is not a control result. Dynamic positioning failures are dominated by loss of position reference, by thruster and power failures and by human factors rather than by loop tuning, and the industrial answer has been redundancy, independent reference systems and a formal vessel class rather than a better controller [12]. Scenario 10 is exactly that failure, included because no controller in the collection can solve it.

6.3 Vessel, Environment and Thruster Model

The vessel is a 90 m offshore construction ship of 8000 t displacement holding a wellhead. Two 220 kN azimuth thrusters sit 35 m aft of midships and two 120 kN tunnel thrusters sit forward, with 5200 kW of thruster power available between them. Work on deck stops when the vessel leaves a 5 m watch circle and the riser is disconnected at 9 m. Figure 6.1 defines the two reference frames, the thruster geometry and every symbol the equations use.

> ✦ DID YOU KNOW? **Twenty Thousand Feet Down, Forty Feet Off**
>
> By 1971 the acoustic positioning system of the first deep-sea research drill ship had let it drill and core in water 20,000 feet deep, about six kilometers, and 3888 feet into the ocean floor, with an apparent excursion of only plus or minus 40 feet for four or five days at a time [3].

6.3.1 Classification of Variables

Four loops are declared. Three of them pair a position or heading with a generalized force demand, and the fourth reports the present heading against the cheapest one so that the supervisory decision is visible on the same time axis as the loops it displaces. Table 6.1 names every signal.

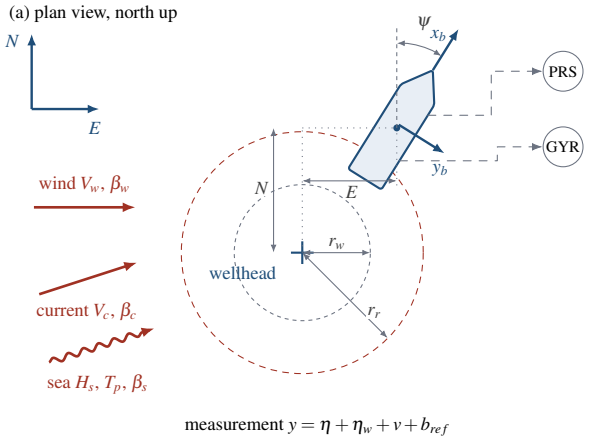

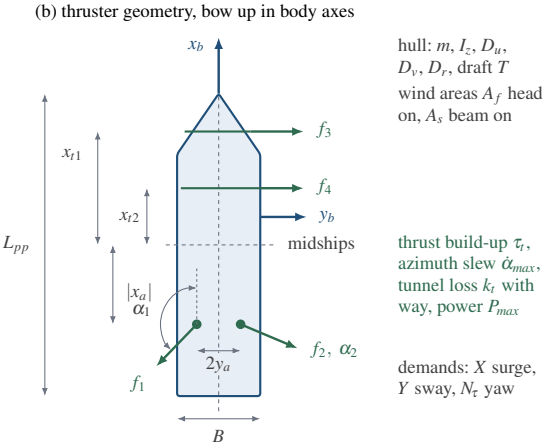

Figure 6.1 Dynamically positioned construction vessel and the geometry the parameters refer to. Panel (a) is the plan view in earth axes: the wellhead at the origin of N and E, the watch circle r_w inside which deck work continues, the disconnect radius r_r, the heading ψ from north, and the wind, current and sea each specified by the direction β they arrive from. The position reference and gyro report the low-frequency motion η, the wave-frequency motion η_w and noise v added together, plus any reference bias b_{ref}. Panel (b) is the same hull in body axes, with two azimuths of thrust f_1, f_2 at angles α_1, α_2, two tunnels of thrust f_3, f_4 on the longest levers on the vessel, and the three demands X, Y and N_τ the allocation resolves onto them.

Table 6.1 Classification of Every Signal. Entries Marked (adv.) Appear Only in the Advanced Control Set.

Symbol	Role	Default	Range
N	CV, north position	0 m	—
E	CV, east position	0 m	—
ψ	CV, heading	0°	—
r	output, radius off wellhead	0 m	—
N_{SP}	setpoint, north	0 m	±30 m
E_{SP}	setpoint, east	0 m	±30 m
ψ_{SP}	setpoint, heading	0°	0 to 359°
X	MV, surge demand	0 kN	±680 kN
Y	MV, sway demand	0 kN	±680 kN
N_τ	MV, yaw demand	0 MN m	±12 MN m
V_w, β_w	DV, wind speed, origin	8, 90°	0 to 30 m/s
V_c, β_c	DV, current speed, origin	0.3, 60°	0 to 2.5 m/s
H_s	DV, significant wave height	1.5 m	0 to 7 m
T_p	DV, peak wave period	9 s	4 to 18 s
β_s	DV, wave origin	95°	0 to 359°
b_{ref}	DV, reference bias (adv.)	0 m	±5 m
—	DV, thruster availability (adv.)	all four	0 or 1 each
N, E, ψ (LF)	states, low-frequency pose	0	—
u, v, r_ψ	states, body velocities	0	—
η_w (6)	states, wave motion and rates	0	—
$f_1..f_4$	states, thruster forces	balanced	$\pm F$ per unit
α_1, α_2	states, azimuth angles	balanced	±180°
K_c^p	tuning, position gain	37	2 to 200 kN/m
τ_I^p	tuning, position reset	600	20 to 3000 s
τ_D^p	tuning, position rate	42	0 to 200 s
τ_f	tuning, derivative filter	9	0 to 60 s
K_c^h	tuning, heading gain	0.34	0.02 to 4 MN m/°
τ_I^h	tuning, heading reset	700	20 to 3000 s
τ_D^h	tuning, heading rate	34	0 to 200 s
K	tuning, observer gain	0.28	0.01 to 1.5/s
K_b	tuning, bias gain multiple	1.00	0 to 4
$\dot{\psi}_{SP}$	tuning, vane slew rate (adv.)	0.25	0.02 to 2 °/s
—	option, estimate	observer	raw, model, observer
—	option, allocation	minimum norm	minimum norm, fixed
—	option, heading	as set	as set, min. thrust
—	option, feedforward (adv.)	on	on, off

6.3.2 Rigid-Body Motion in Three Degrees of Freedom

The vessel is a rigid body in surge, sway and yaw. Body-axis momentum balances with linear damping and the rotating-frame terms give Equation (6.1).

$$\begin{aligned} m_u \frac{\mathrm{d}u}{\mathrm{d}t} &= X_{thr} + X_{wind} + X_{cur} + X_{drift} - D_u u + m_u v r_\psi \\ m_v \frac{\mathrm{d}v}{\mathrm{d}t} &= Y_{thr} + Y_{wind} + Y_{cur} + Y_{drift} - D_v v - m_v u r_\psi \\ I_z \frac{\mathrm{d}r_\psi}{\mathrm{d}t} &= N_{thr} + N_{wind} + N_{cur} + N_{drift} - D_r r_\psi \end{aligned} \tag{6.1}$$

where u and v are the body surge and sway velocities in m/s, r_ψ is the yaw rate in rad/s, D_u, D_v and D_r are the linear damping coefficients, and m_u, m_v and I_z are the inertias including added mass, given by Equation (6.2).

$$\begin{aligned} m_u &= m(1+a_u), \qquad m_v = m(1+a_v) \\ I_z &= m\,(k_g L_{pp})^2 (1+a_r) \end{aligned} \tag{6.2}$$

where m is the displacement, a_u, a_v and a_r are the added-mass fractions and k_g is the yaw radius of gyration as a fraction of L_{pp}. At the constants of Table 6.2, $m_u = 8.64 \times 10^6$ kg, $m_v = 1.48 \times 10^7$ kg and $I_z = 5.47 \times 10^9$ kg m^2, so the same force accelerates the hull 1.71 times faster along the bow than across it. Nothing on the right of (6.1) depends on position except through the environment, so there is no restoring term anywhere and each axis is an integrating plant driven through a second integration.

The pose in earth axes follows from the rotation through the heading, as shown in Equation (6.3).

$$\frac{\mathrm{d}N}{\mathrm{d}t} = u\cos\psi - v\sin\psi, \quad \frac{\mathrm{d}E}{\mathrm{d}t} = u\sin\psi + v\cos\psi, \quad \frac{\mathrm{d}\psi}{\mathrm{d}t} = r_\psi \tag{6.3}$$

The heading is wrapped to $(-\pi, \pi]$ after every step, as are the two azimuth angles.

6.3.3 Wind, Current and Wave Drift

The three environmental loads are algebraic in the disturbance settings and the heading. Wind is a dynamic pressure acting on two projected areas blended with the angle of attack, in the manner of a first-cut coefficient set [13], as shown in Equation (6.4).

$$\begin{aligned} q_a &= \tfrac{1}{2}\rho_a C_a V_w^2, \qquad \gamma_w = \beta_w - \psi \\ X_{wind} &= -q_a A_f \cos\gamma_w |\cos\gamma_w| \\ Y_{wind} &= -q_a A_s \sin\gamma_w |\sin\gamma_w| \\ N_{wind} &= 0.14 L_{pp} \cos\gamma_w \, Y_{wind} \end{aligned} \tag{6.4}$$

where ρ_a is the air density, C_a the drag coefficient, A_f and A_s the head-on and beam-on projected areas and γ_w the wind angle relative to the bow, and the absolute values preserve the sign through the quadrant. The yaw moment acts through a lever that is a fixed fraction of the length and changes sign with the angle of attack, which is what makes the hull weathervane.

Current acts the same way on the underwater profile, driven by the water speed relative to the hull, as shown in Equation (6.5).

$$\begin{aligned} u_r &= -V_c\cos(\beta_c - \psi) - u, \;\; v_r = -V_c\sin(\beta_c - \psi) - v \\ X_{cur} &= \tfrac{1}{2}\rho_w C_w B T \, u_r |u_r| \\ Y_{cur} &= \tfrac{1}{2}\rho_w C_w L_{pp} T \, v_r |v_r| \\ N_{cur} &= 0.11 L_{pp} \, Y_{cur} \end{aligned} \tag{6.5}$$

where ρ_w is the sea water density, C_w the hull drag coefficient, B the beam and T the draft. The relative velocity carries the vessel's own motion, so (6.5) supplies part of the damping as well as the load.

The sea contributes twice, and the two contributions must be kept apart. The steady part is the second-order drift force, which is proportional to the square of the significant wave height and is the part the thrusters genuinely have to carry [14], as shown in Equation (6.6).

$$\begin{aligned} X_{drift} &= -0.35 k_d H_s^2 \cos(\beta_s - \psi) \\ Y_{drift} &= -k_d H_s^2 \sin(\beta_s - \psi) \\ N_{drift} &= 0.06 L_{pp} Y_{drift} \end{aligned} \tag{6.6}$$

where k_d is the drift force per square meter of significant height in N/m^2. At $H_s = 4$ m this is 83.2 kN of beam force, slow enough that the bias estimator treats it as part of the steady environment.

6.3.4 Wave-Frequency Motion and the Measurement

The oscillatory part of the sea moves the hull without displacing it from station. It is represented by three lightly damped second-order oscillators at the peak wave frequency, driven by independent colored noise so that the seaway is irregular and the estimator cannot lock onto a sine, as shown in Equation (6.7).

$$\begin{aligned} \ddot{\eta}_{w,i} + 2\zeta\omega_0\dot{\eta}_{w,i} + \omega_0^2\eta_{w,i} &= 2\zeta\omega_0^2 g_i \, w_i(t) \\ g_N &= G_w H_s \cos(\beta_s - \psi) \\ g_E &= G_w H_s \sin(\beta_s - \psi) \\ g_\psi &= G_\psi H_s \sin(\beta_s - \psi) \end{aligned} \tag{6.7}$$

for $i \in \{N, E, \psi\}$, where $\omega_0 = 2\pi / \max(3\,\mathrm{s}, T_p)$ is the peak wave frequency, ζ is the motion damping ratio, G_w and G_ψ are the first-order motion gains in meters and degrees per meter of significant height, and w_N, w_E and w_ψ are unit-variance colored noise sequences of correlation time 2.2, 2.4 and 2.0 s. The factor 2ζ on the right cancels the resonator's own amplification at ω_0, so G_w means what its name says rather than that value times a quality factor of three. In the 4 m seaway used throughout the results the peak excursion over the observation window of Section 6.6 is 2.39 m (root mean square 0.76 m across the beam) against a watch circle of 5 m.

What the control system reads is the sum, as shown in Equation (6.8).

$$y = \eta + \eta_w + v + b_{ref} \tag{6.8}$$

where $\eta = (N, E, \psi)$ is the low-frequency pose, η_w is the wave-frequency motion of (6.7), v is instrument noise and b_{ref} is a stale differential correction applied to the north channel. Nothing in (6.8) says which term a given meter belongs to, and that is the whole of the control problem. The instrument noise v has reached the measurement since the repair recorded in Section 6.4.6, and every result in this chapter is obtained with it live at 0.35 m on position and 0.12° on the gyro, each a colored sequence with a 1.5 s correlation time. The position reference is sampled on the controller scan rather than the physics step, which is a zero-order hold of 0.5 s, and the controller acts on each sample at the instant it is taken, so no further dead time is applied.

6.3.5 Thrusters

A commanded force reaches each thruster through a first-order lag, and each azimuth swings toward its commanded angle no faster than its slew limit allows. Both appear in Equation (6.9).

$$\begin{aligned} \tau_t \frac{\mathrm{d}f_k}{\mathrm{d}t} &= f_k^{cmd} - f_k \\ \frac{\mathrm{d}\alpha_j}{\mathrm{d}t} &= \mathrm{clamp}\left(\frac{\alpha_j^{cmd} - \alpha_j}{0.5\,\mathrm{s}},\ \pm\dot{\alpha}_{max} \right) \end{aligned} \tag{6.9}$$

for $k = 1 \ldots 4$ and $j = 1, 2$, where the angular error is wrapped to $(-\pi, \pi]$ before the clamp. The forces are clipped to each machine's rating after every step. What the four thrusters produce in the body frame is Equation (6.10).

$$\begin{aligned} \ell &= \mathrm{clamp}(1 - k_t |u|,\ 0.25,\ 1) \\ X_{thr} &= f_1 \cos\alpha_1 + f_2 \cos\alpha_2 \\ Y_{thr} &= f_1 \sin\alpha_1 + f_2 \sin\alpha_2 + (f_3 + f_4)\,\ell \\ N_{thr} &= x_a\,(f_1 \sin\alpha_1 + f_2 \sin\alpha_2) \\ &\quad - y_a\,(f_1 \cos\alpha_1 - f_2 \cos\alpha_2) + (x_{t1} f_3 + x_{t2} f_4)\,\ell \end{aligned} \tag{6.10}$$

where ℓ is the tunnel efficiency, k_t is the tunnel loss per unit of forward speed, x_a is the signed azimuth position relative to midships and y_a is the azimuth offset from the centerline. At half a knot the tunnels retain 85.9 % of their rating and at two knots 43.4 %, so a tunnel thruster is worth appreciably less on a vessel that is moving than on one that is stopped. Shaft power follows the usual three-halves law of Equation (6.11).

$$P = k_P \sum_{k=1}^{4} |f_k|^{3/2} \tag{6.11}$$

where k_P is in kW per $\mathrm{N}^{1.5}$. All four machines at their ratings draw 5182 kW against 5200 kW installed, so the switchboard and the thrust run out at almost the same place. Equation (6.11) also prices the allocation directly, because 200 kN of sway spread over four thrusters costs 801 kW and the same 200 kN through two costs 1132 kW.

6.3.6 Parameters, Initialization and Numerics

Table 6.2 lists every constant. Four are hidden from the parameter panel (the two densities, the power coefficient and the allocation ridge) and the rest are adjustable, a deliberately thin concealment because what the student is asked to infer is structure and not a hidden number. The heading tolerance is read only by the heading badge above the animation, which turns amber while the heading is further than the tolerance from its setpoint.

Table 6.2 Model Parameters. Four Constants Are Hidden From the Parameter Panel and Used by the Model Only. Every Other Row Is Adjustable Over the Range Named in the Application.

Symbol	Meaning	Value	Unit	Visibility
Vessel				
L_{pp}	length between perpendiculars	90	m	adjustable
B	beam	20	m	adjustable
m	displacement	8.0×10^6	kg	adjustable
a_u	added mass, surge	0.08	—	adjustable
a_v	added mass, sway	0.85	—	adjustable
k_g	yaw radius of gyration	0.25	$\times L_{pp}$	adjustable
a_r	added inertia, yaw	0.35	—	adjustable
D_u	surge damping	9.0×10^4	N s/m	adjustable
D_v	sway damping	2.6×10^5	N s/m	adjustable
D_r	yaw damping	1.2×10^8	N m s	adjustable
Environment				
ρ_a	air density	1.225	kg/m^3	hidden
ρ_w	sea water density	1025	kg/m^3	hidden
A_f	wind area, head on	320	m^2	adjustable
A_s	wind area, beam on	900	m^2	adjustable
T	draft	6.0	m	adjustable
C_a	wind drag coefficient	0.85	—	adjustable
C_w	current drag coefficient	1.05	—	adjustable
k_d	wave drift per m^2 of H_s^2	5200	N/m^2	adjustable
G_w	wave motion gain	0.42	m/m	adjustable
G_ψ	wave yaw gain	0.22	°/m	adjustable
ζ	wave motion damping	0.16	—	adjustable
Thrusters				
F_a	azimuth rating, two aft	220	kN	adjustable
F_t	tunnel rating, two forward	120	kN	adjustable
x_a	azimuth position from midships	−35	m	adjustable
y_a	azimuth offset from centerline	6.5	m	adjustable
x_{t1}	forward tunnel position	38	m	adjustable
x_{t2}	second tunnel position	26	m	adjustable
τ_t	thrust build-up time constant	4.0	s	adjustable
$\dot{\alpha}_{max}$	azimuth turning rate	6.0	°/s	adjustable
k_P	power per unit thrust$^{3/2}$	1.79×10^{-5}	kW/N$^{3/2}$	hidden
P_{max}	thruster power available	5200	kW	adjustable
k_t	tunnel loss with way	0.55	s/m	adjustable
τ_α	azimuth angle filter	20	s	adjustable
ε	allocation ridge	0.025	—	hidden
Instrument and limits				
σ_p	position reference noise	0.35	m	adjustable
σ_g	gyro noise	0.12	°	adjustable
r_w	watch circle	5.0	m	adjustable
r_r	disconnect radius	9.0	m	adjustable
—	heading tolerance	3.0	°	adjustable, badge
Numerics				
Δt	integration step (RK4)	0.05	s	—
Δt_c	controller period	0.5	s	—
—	historian rate	2	Hz	—
—	random seed	20260916	—	—

Every scenario opens with the vessel on station and the thrusters carrying the environment: the pose starts at the setpoint, the velocities and wave states at zero, and the thruster forces and azimuth angles at the allocation solution for the steady demand at the initial heading. The control system starts from the same place. The bias estimate opens holding the environmental force at the initial heading in earth axes, and with the feedforward off the three integrals open holding the corresponding demand instead, so the first scan commands the thrust the vessel is already producing. In the first 200 s every automatic scenario except the capability limit and the challenge as loaded stays within 0.95 m of the

wellhead in its low-frequency position. An earlier version of the app started the other way, with the bias estimate and the integrals at zero, and a drive-off carried scenario 8 past the disconnect radius at 91 s. A fixed 0.05 s step of classical fourth-order Runge-Kutta carries the states forward, chosen against the wave period rather than the vessel, which puts two hundred steps in one wave at $T_p = 10\,\text{s}$. The control system executes on a 0.5 s period, the scan rate of an industrial dynamic positioning system, and the historian logs two samples per simulated second. Colored noise and random draws come from one seeded generator, so a run replays identically.

6.3.7 Assumptions

Detail is spent wherever the physics forces a decision on the controller and saved everywhere else. Each simplification below fails somewhere, and the place it fails is named with it.

Three degrees of freedom. Roll, pitch and heave are absent. On a construction vessel it is usually heave rather than horizontal position that stops a crane or a diving operation, and nothing in this application can represent that.

Blended area coefficients for wind and current. Equations (6.4) and (6.5) use two projected areas and a cosine blend with a fixed yaw lever. Measured coefficient sets come from a wind tunnel and are not smooth functions of heading [13], which matters precisely where this application makes its strongest claim, in choosing the cheapest heading.

A single-frequency sea. Equation (6.7) places all of the wave-frequency energy at the peak period, so a resonator tuned there leaves the tails of a real broad-banded spectrum in the position estimate, and the wave direction is applied in earth axes rather than resolved through the hull.

Drift proportional to H_s^2 with a fixed directional shape. The second-order force on a real hull depends on the period as well as the height and can change sign [14], so (6.6) is a magnitude estimate rather than a hull calculation.

Idealized machinery. Equations (6.9) and (6.10) carry a first-order build-up, an azimuth slew rate and the tunnel loss with forward speed. There are no forbidden azimuth zones, no thruster-hull or thruster-thruster interaction and no power management system shedding load before the switchboard trips. When the allocation asks for more than P_{max} the whole solution is scaled back, whereas a real power management system would decide which degree of freedom to surrender. The azimuth angles are filtered ahead of the solve rather than rate-constrained inside it, so the allocation is a two-stage approximation to the constrained problem rather than the constrained problem itself.

One position reference. A vessel of dynamic positioning class 2 or 3 carries three independent references and votes on them [12]. The application has one, which is what makes scenario 10 possible and its remedy impossible to demonstrate.

Repeatability that no control system has. Freezing the integration step and the scan period repeats one trajectory exactly wherever it is run, whereas a real dynamic positioning computer has jitter, dropped telegrams and a scan order that moves.

6.4 Estimator, Controllers and Allocation

The control system is an observer, three PID loops, an allocation and a supervisor above them, arranged as in Figure 6.2. All four execute on the same 0.5 s scan.

6.4.1 The Observer

The observer is a Kalman filter with the gains placed by hand, which is what most industrial dynamic positioning observers are [2]. It joins the bias state of Equation (1.10) and the wave filtering of Section 1.4.6 in one estimator, and Volume 5, Chapter 13 applies the same machinery to an underwater vehicle whose acoustic position fixes arrive late and occasionally wild. It carries a vessel model driven by the thrust actually produced, a bias state for the steady environment and a

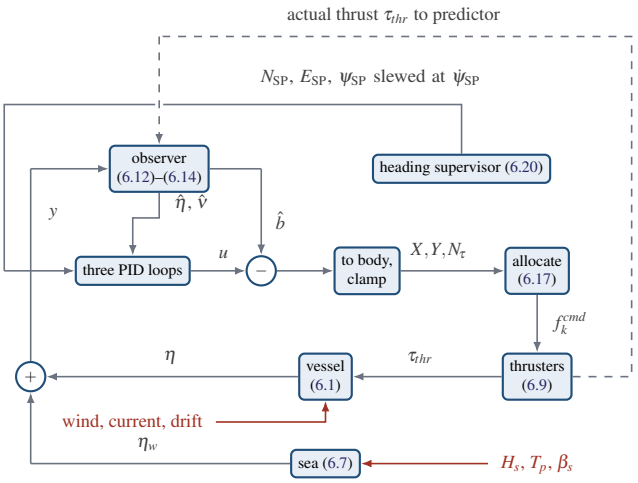

Figure 6.2 Control structure as implemented. One innovation from the position reference y is divided by the observer among the low-frequency pose $\hat{\eta}$, its velocity, the environmental bias $\hat{b}$ and the wave-frequency states, and only $\hat{\eta}$ reaches the three PID loops. The bias estimate is subtracted from the loop output as a feedforward (on every estimator setting, while its own switch is on), the result is rotated into body axes and clamped there, and the allocation resolves the three demands onto six thruster forces. The heading supervisor sits outside every loop.

wave-frequency resonator, is advanced by explicit Euler on the 0.5 s scan, and predicts by (6.12).

$$\begin{aligned} F_b &= (\hat{\tau}_N \cos\hat{\psi} + \hat{\tau}_E \sin\hat{\psi}) - D_u \hat{u}, \qquad \hat{a}_b = F_b/m_u \\ F_s &= (-\hat{\tau}_N \sin\hat{\psi} + \hat{\tau}_E \cos\hat{\psi}) - D_v \hat{w}, \qquad \hat{a}_s = F_s/m_v \\ \dot{\hat{V}}_N &= \hat{a}_b \cos\hat{\psi} - \hat{a}_s \sin\hat{\psi}, \qquad \dot{\hat{V}}_E = \hat{a}_b \sin\hat{\psi} + \hat{a}_s \cos\hat{\psi} \\ \dot{\hat{V}}_\psi &= \left(N_{thr} + \hat{b}_\psi - D_r \hat{V}_\psi\right)/I_z \end{aligned} \tag{6.12}$$

where $\hat{\tau}_N = X_{thr}\cos\hat{\psi} - Y_{thr}\sin\hat{\psi} + \hat{b}_N$ and $\hat{\tau}_E = X_{thr}\sin\hat{\psi} + Y_{thr}\cos\hat{\psi} + \hat{b}_E$ are the thrust and bias in earth axes, $\hat{u}$ and $\hat{w}$ are the estimated body surge and sway velocities obtained by rotating $\hat{V}_N$ and $\hat{V}_E$ into the body, and $\hat{V}$ denotes an estimated earth-frame velocity. The rotation in (6.12) is not cosmetic: the inertia is a body-frame matrix whose two entries differ by a factor of 1.71, so an earth-frame force divided directly by a body-frame mass is exact only with the bow due north. Hull damping belongs in the observer's model for the same reason, because without it every newton of drag has to be absorbed by the bias state, which then hunts.

The wave states are propagated as an undriven resonator at the same peak frequency the plant uses, as shown in Equation (6.13).

$$\ddot{\hat{\eta}}_w = -2\zeta\omega_0\dot{\hat{\eta}}_w - \omega_0^2\hat{\eta}_w \tag{6.13}$$

on each of the three axes. It has no input, because the estimator does not know the seaway, only the frequency at which part of the innovation belongs to it.

The correction step takes one innovation per axis and divides it four ways, as shown in Equa-

tion (6.14).

$$
\begin{aligned}
&e = y - (\hat{\eta} + \hat{\eta}_w), \qquad L_2 = \frac{K^2}{3}, \qquad L_3 = \frac{K^3}{180} K_b \\
&\hat{\eta} \leftarrow \hat{\eta} + K\, e\, \Delta t_c, \qquad \hat{V} \leftarrow \hat{V} + L_2\, e\, \Delta t_c \\
&\dot{\hat{\eta}}_w \leftarrow \dot{\hat{\eta}}_w + 2.2\, K\, \omega_0\, e\, \Delta t_c \\
&\hat{b}_b \leftarrow \hat{b}_b + L_3\, m_u\, e_b\, \Delta t_c, \qquad \hat{b}_s \leftarrow \hat{b}_s + L_3\, m_v\, e_s\, \Delta t_c
\end{aligned}
\tag{6.14}
$$

where K is the single observer gain in 1/s, K_b is a multiplier on the nominal bias gain, and e_b and e_s are the north and east innovations rotated into body axes so that the two inertias scale the correct component. The two body-frame bias increments are rotated back into earth axes before they are added to $\hat{b}_N$ and $\hat{b}_E$, and the yaw bias is scaled by I_z in the same way. Three observer poles placed together at $\omega = K/3$ give $L_1 = K$, $L_2 = K^2/3$ and a third of $K^3/27$. The implementation uses $K^3/180$ instead, placing the bias pole 6.7 times slower, because a bias estimate that moves as fast as the position loop is a second controller working on the same error. The heading innovation is wrapped to $(-\pi, \pi]$ and the bias states are clamped at 1.2 MN and 60 MN m.

The estimator option selects what the controllers are given rather than whether the observer runs. The full setting passes $\hat{\eta}$ from (6.14). The model-only setting holds the wave states at zero, so the sea lands partly in the position estimate and partly in the bias. The raw setting overwrites $\hat{\eta}$ with the measurement y at the end of every scan. The observer runs underneath all three, so the bias estimate subtracted in (6.16) is available on every setting and the feedforward follows its own switch. Switching to the raw reference therefore changes what position the loops are shown and nothing else.

6.4.2 The Three Loops as Implemented

Each loop is the velocity-form PID of Section 1.6.2, written for this application as Equation (6.15). The output before limiting, which Chapter 1 writes v_k, is c_k here because v is the sway velocity, and the output bias u_b is zero on all three loops.

$$
\begin{aligned}
&e_k = a\,(\mathrm{SP}_k - \mathrm{PV}_k), \qquad P_k = K_c e_k \\
&\mathrm{PV}_k^f = \alpha \mathrm{PV}_{k-1}^f + (1-\alpha)\mathrm{PV}_k, \qquad \alpha = \mathrm{e}^{-\Delta t_c/\tau_f} \\
&D_k = -K_c \tau_D\, a\, \frac{\mathrm{PV}_k^f - \mathrm{PV}_{k-1}^f}{\Delta t_c} \\
&S_k = S_{k-1} + (P_k - P_{k-1}) + (D_k - D_{k-1}) \\
&c_k = S_k + I_{k-1}, \qquad u_k = \mathrm{clamp}(c_k, u_{min}, u_{max}) \\
&I_k = I_{k-1} + \left[\frac{K_c}{\tau_I} e_k + \frac{u_k - c_k}{T_t} \right] \Delta t_c
\end{aligned}
\tag{6.15}
$$

where $a = +1$ on all three loops, so the controller is reverse acting in the convention of Section 1.3 and a vessel north of its setpoint receives a southward demand. The derivative acts on the measurement after a first-order filter, so moving a setpoint does not kick the thrusters, and the filter uses the exact discrete update for a first-order lag, stable at any scan period. The integral is advanced after the clamp by back-calculation with the default tracking time $T_t = \tau_I$. There is no rate limit on the demand, because the rate limits in this plant are the thruster build-up τ_t and the azimuth slew $\dot{\alpha}_{max}$ and belong to the machinery. The back-calculation sees only the loop's own clamp, so while the body-frame clamp of (6.16), a thruster rating or the power scaling of Section 6.4.3 binds downstream, the integral keeps accumulating. In manual the heading loop is tracked on the same wrapped-error measurement it uses in automatic, which makes transfer to automatic bumpless on all three loops. When a run opens in automatic with the feedforward off, I starts at the steady demand of Section 6.3.6.

The north and east loops act on $\hat{N}$ and $\hat{E}$ from the observer with $u_{min}, u_{max} = \mp 680\,\mathrm{kN}$, which is the vessel's own sway capability rather than a round number, because a controller allowed to ask for

more than the machinery can produce is a controller whose anti-windup never engages. The heading loop acts on the wrapped heading error with $\mp 12\,\mathrm{MN\,m}$, and because the derivative in (6.15) is taken on the measurement, that wrapped error is supplied as PV_k against a setpoint of zero rather than the other way round, which is what makes τ_D^h effective at all. The three outputs are earth-frame forces, and the environment feedforward is applied there before the rotation, as shown in Equation (6.16).

$$\begin{aligned} X &= \mathrm{clamp}\big[(u_N - \hat{b}_N)\cos\psi + (u_E - \hat{b}_E)\sin\psi,\ \pm X_{lim}\big] \\ Y &= \mathrm{clamp}\big[-(u_N - \hat{b}_N)\sin\psi + (u_E - \hat{b}_E)\cos\psi,\ \pm X_{lim}\big] \\ N_\tau &= \mathrm{clamp}\big[u_\psi - \hat{b}_\psi,\ \pm N_{lim}\big] \end{aligned} \tag{6.16}$$

where $X_{lim} = 680\,\mathrm{kN}$ and $N_{lim} = 12\,\mathrm{MN\,m}$. The clamp is applied after the rotation and not before, because clamping the earth-frame pair separately would let 600 kN north and 600 kN east arrive as 849 kN of sway. Subtracting $\hat{b}$ in (6.16) is what allows τ_I^p to be as long as 600 s. The bias state finds the same force inside the estimator, where it is not in the control loop and can therefore be fast without destabilizing anything.

6.4.3 Thrust Allocation

Four machines produce six independent forces, an x and a y from each azimuth and a y from each tunnel, against three demands, so there is a three-dimensional family of solutions and the application offers two ways to pick one. The minimum-weighted-norm solution is the weighted least-norm allocation of Section 1.4.7, written for this vessel as (6.17) [9].

$$\mathbf{u} = W^{-1}B^T\left(BW^{-1}B^T\right)^{-1}\tau, \qquad \tau = [X,\ Y,\ N_\tau]^T \tag{6.17}$$

where $\mathbf{u} = [f_{x1}, f_{y1}, f_{x2}, f_{y2}, f_3, f_4]^T$, the configuration matrix is Equation (6.18),

$$B = \begin{bmatrix} 1 & 0 & 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 1 & \ell & \ell \\ -y_a & x_a & y_a & x_a & x_{t1}\ell & x_{t2}\ell \end{bmatrix} \tag{6.18}$$

and $W = \mathrm{diag}(F_a^{-2}, F_a^{-2}, F_a^{-2}, F_a^{-2}, F_t^{-2}, F_t^{-2})$ divides each machine by the square of its own rating, which comes close to minimizing power, because the three-halves law of (6.11) makes 200 kN of sway spread over four machines cost 801 kW against 1132 kW through two. The 3×3 system is solved by Gaussian elimination with partial pivoting on every scan, and a failed thruster's entries of W are divided by an availability factor of 10^{-4} and its command is then set to zero, so the same arithmetic spreads its work across whatever remains with no rule written for the case. A constrained version of the same least-squares problem shares the push among the planted feet of the walking robot in Volume 4, Chapter 13, where a foot can push but never pull.

Three details of the azimuth solution are the difference between a working allocation and a limit cycle. An azimuth can produce any force by pointing at it or by pointing away and running astern, and always taking the first turns a demand that changes sign into a 180° slew during which the unit pushes everywhere except where it was asked. The implementation picks whichever representation is nearer the present angle, which is why f_1 and f_2 are signed. The commanded angle is then filtered and the forces re-solved with the angles held, for the reasons and by the arithmetic set out in Section 6.4.6. Finally, if (6.11) exceeds P_{max} the whole vector is scaled by $(P_{max}/P)^{2/3}$ rather than letting one machine absorb the shortfall.

The alternative is a fixed split of the kind written at commissioning. The azimuths take all of the surge and half of the sway, and the tunnels take the remaining sway and the whole yaw moment, as shown in Equation (6.19).

$$f_x = \frac{X}{2}, \quad f_y = \frac{Y}{4}, \quad f_3 = f_4 = \frac{1}{2\ell}\left(\frac{Y}{2} + \frac{N_\tau}{x_{t1}}\right) \tag{6.19}$$

with $f_1 = f_2 = |(f_x, f_y)|$ clipped at F_a and both azimuths pointed along $\operatorname{atan2}(f_y, f_x)$. Equation (6.19) is indistinguishable from (6.17) while all four machines are running. The share belonging to a machine that has stopped is simply not produced.

6.4.4 The Heading Supervisor

The heading is the one setpoint on the panel that the job does not fix, because a wellhead does not care which way the bow points, and both the wind area and the underwater area of (6.4) and (6.5) depend on it. That makes the environmental demand an objective rather than a load, as shown in Equation (6.20).

$$\begin{aligned} \tau_{env}(\psi) &= -\left[\tau_{wind}(\psi) + \tau_{cur}(\psi) + \tau_{drift}(\psi)\right] \\ \psi^{\star} &= \arg\min_{\psi} \| \tau_{env}(\psi) \| \end{aligned} \tag{6.20}$$

where the norm is the magnitude of the surge and sway components. The minimum is scanned at 2° intervals every 30 s rather than solved, because the objective is not convex once the wind and the current disagree and the scan is cheap on a quantity that changes slowly. Figure 6.3 plots the objective for the conditions of scenario 8.

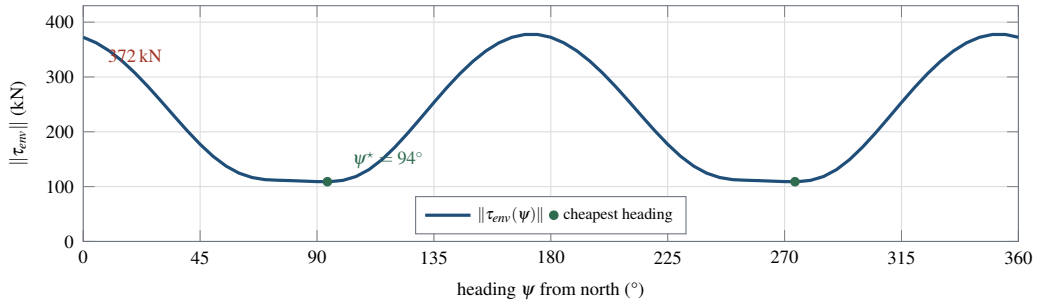

Figure 6.3 Environmental thrust demand against heading for the conditions of scenario 8, evaluated from Equation (6.20) at $V_w = 17\,\mathrm{m/s}$ from 90°, $V_c = 0.8\,\mathrm{m/s}$ from 75° and $H_s = 3.5\,\mathrm{m}$ from 95°. The demand falls from 372 kN beam on to 109 kN at $\psi^{\star} = 94°$, and the power of the exact minimum-norm solution of (6.17) from 2135 kW to 463 kW. The curve repeats at 180° because the projected-area model is symmetric fore and aft.

The supervisor does not command the heading setpoint to $\psi^{\star}$. It walks it there at $\dot{\psi}_{\mathrm{SP}} = 0.25\,^{\circ}/\mathrm{s}$, so the 94° turn of scenario 8 takes 375.5 s. Slewing a setpoint instead of stepping it is the general remedy whenever the transient of a supervisory move is itself the hazard, but on this vessel the slew is not what protects the position. At 2 °/s the setpoint arrives in 46.6 s and the worst excursion during the turn is 6.67 m against 6.46 m at the default rate, only 0.2 m more, because the yaw authority of the heading loop, not the setpoint, sets how fast the bow comes round.

6.4.5 Three Models Kept Apart

The simulation model is Equations (6.1) to (6.11) integrated over eighteen states. The controller's internal model is what appears in (6.12) to (6.14), which is the vessel inertias and damping, the peak wave frequency, and nothing else. It does not know the wind areas, the drag coefficients, the drift constant or the disturbance settings, and it never predicts forward. The heading supervisor is the one block that does hold the environmental model of (6.4), (6.5) and (6.6), which is realistic because a capability calculation on a real vessel is made from a measured coefficient set and a wind sensor, and it is worth noticing that the one place the application uses a plant model is outside every feedback loop.

6.4.6 What the Repaired Measurement Path Changed

Four implementation defects in an earlier version of the application, all repaired in the version this chapter describes, show what each part of the measurement path and the allocation is for, and each was large enough in effect to state.

The position reference noise σ_p and the gyro noise σ_g of Table 6.2 were declared as model parameters and read from the disturbance set, so both resolved to an undefined quantity and no noise reached (6.8) at any setting. Reading them from the parameter set makes both live, and two arguments move with them: the case for the derivative filter τ_f no longer rests on the wave-frequency content of the measurement alone, because instrument noise through $K_c^p \tau_D^p$ by itself drives the surge demand to its clamp in flat calm, and the case for the wave filter is stronger rather than weaker, because the same estimator rejects both motions and the travel ratio rises from 3.3 on a quiet reference to 4.5 on the realistic one. Second, the heading loop was handed the wrapped heading error as its setpoint and a constant zero as its measurement, so the derivative in (6.15) was taken on a signal that never moved and τ_D^h had no effect at any value. Passing the error as the measurement instead makes the loop genuinely three-term, and a 20° heading step in the calm water of scenario 2, made at 300 s, now peaks at 24.1° against 35.7° with the term removed. The manual tracking call kept the old arrangement for a time, which kicked the yaw demand by 1.25 MN m per degree of heading error on transfer to automatic, and now uses the same one. Third, a declared and unread reference lag was deleted, leaving the 0.5 s zero-order hold of the controller scan as the only reference delay the model claims.

The fourth defect was in the allocation and was hidden by the first. Equation (6.17) was solved afresh every scan with the azimuth angles read straight off atan2 of the result, and the direction of a small vector is badly conditioned. The behavior is reproduced here by the current model with the angle filter and the allocation ridge described below both set to zero, a counterfactual run because the ridge is hidden. In the calm water of scenario 2 the port azimuth demand then has a median magnitude of 36 kN, much of it noise through the derivative term, so the demanded direction moves a median of 3.2° per scan and 51.5 % of scans ask for more than the 3° a 6 °/s unit can deliver in half a second. The unit never arrives: over the observation window the port azimuth turns through 7115°, nearly twenty full revolutions, with the rate limiter saturated 66 % of the time, and the heading loop inherits a limit cycle from it, so that a 5 m north step made at 300 s peaks anywhere between 5.95 m and 7.77 m across five seeds, with the heading swinging as far as 10.9°. Setting only the adjustable angle filter to zero in the application gives the same picture, 7048° with the limiter saturated 65 % of the time. A magnitude guard cannot address it, because the magnitude clears any threshold long before the direction stops being noise.

The repair is the one an industrial allocation makes. The commanded azimuth angle is passed through a first-order lag whose time constant is scheduled inversely with the demand magnitude against a reference of 11 kN, so a direction that is mostly noise is filtered hard and one carrying real load is barely filtered at all. The schedule matters: a fixed time constant also delays the redistribution after a thruster trip and costs scenario 9 its point. The forces are then re-solved from the same weighted pseudo-inverse with both angles held, four unknowns against the same three demands, so the tunnels and the azimuth magnitudes make up what the lagging angles cannot. That solve carries a ridge ε of 2.5 % of each axis's free authority, because beam on both azimuths sit near 90°, the surge row collapses, and an exact solve would make a few kilonewtons of surge out of two enormous opposing forces. With the filter at zero the angles are those of the free solution and the forces fall short of it only by the ridge, which beam on in the conditions of scenario 8 delivers 358 kN of the 372 kN of sway asked and leaves the rest to the loops. With the repair as shipped, azimuth travel in calm water falls to 2636° and saturation to 3.2 %, the 5 m step peaks between 5.89 m and 6.70 m across the same five seeds with the heading inside 1.7°, and station keeping improves everywhere. The thruster command travel reported below rises, from 42.8 MN to 72.0 MN in the seaway of scenario 4 over the observation window while the mean offset falls from 0.64 to 0.51 m, because churn once absorbed by

Table 6.3 The Twelve Scenarios and What Each One Isolates

#	Scenario	What it isolates	Kind
1	The environment by hand	a balance is not control	Demonstration
2	Calm water, three loops	derivative action on an undamped plant	Experiment
3	The sea arrives, no filter	chasing a motion that is not an error	Demonstration
4	The same sea, filtered	what the wave model is worth	Experiment
5	Where the waves come back	the observer gain as the estimation trade	Experiment
6	The bias estimate against the integral	disturbance feedforward	Experiment
7	A squall	an unmeasured load found by being pushed	Experiment
8	The heading is free	supervisory choice of a free variable	Experiment
9	A thruster trips	solved allocation against a written one	Experiment
10	The position reference lies	a fault no controller can answer	Demonstration
11	The capability limit	heading capability, not thrust, binds	Demonstration
12	Challenge: a working shift	all four structural decisions at once	Challenge

unphysical spinning now appears where the metric can see it. Three further behaviors of the earlier version have since been repaired in the application and are described in this chapter as they now act: the start-up of Section 6.3.6, the feedforward the raw setting used to withhold (Section 6.4.1) and the heading-loop transfer kick.

6.5 Scenarios

Twelve set-ups are arranged so that the plant is met before the control system and each structural decision is broken alone. Table 6.3 lists them with the mechanism each one isolates. Each carries a briefing and a completion rule in the manner of Section 1.6.4, the challenge is judged on the case as given, so a changed model parameter or disturbance voids the attempt, and the Exercises at the end of the chapter use the scenarios by number.

The reference tuning is $K_c^p = 37\,\mathrm{kN/m}$, $\tau_I^p = 600\,\mathrm{s}$, $\tau_D^p = 42\,\mathrm{s}$, $\tau_f = 9\,\mathrm{s}$, $K_c^h = 0.34\,\mathrm{MN\,m/^\circ}$, $\tau_I^h = 700\,\mathrm{s}$, $\tau_D^h = 34\,\mathrm{s}$, $K = 0.28/\mathrm{s}$ and $K_b = 1.00$. The weak set that opens the manual and challenge scenarios is roughly a quarter of the position gain, a fifth of the derivative time, no derivative filter and an observer gain of 0.9/s.

Scenario 1 is worth running once. The manual helper computes the steady environmental demand of (6.20) at the present heading, which is 164 kN of sway at 24 % of what the thrusters can produce, and sets the three manual demands to it, rounded to 1 kN and 0.1 MN m. Nothing measures the position, so nothing corrects the residual. The heading wanders 9.2° in the first five minutes, the balance computed for the original heading no longer matches the load, and the vessel leaves the watch circle at 267 s, before the weather has changed at all. The wind then freshens at 300 s and veers 30° at 700 s, the vessel passes the disconnect radius at 343 s, and it is 124 m away at twenty minutes. Eight thousand tonnes with no restoring force is a pure double integrator, and every scenario after the first is about what feedback costs rather than what it is for.

6.6 Results

The values that follow come from the application's own scorecard and its exported samples, at the constants of Table 6.2, with each setting made before the run and the shipped seed. Accumulating measures such as thruster command travel and energy are quoted with the horizon over which they accumulated, and averages are quoted with their window, because both continue to move after a transient ends. Peaks and crossing times are taken at the 0.05 s physics step.

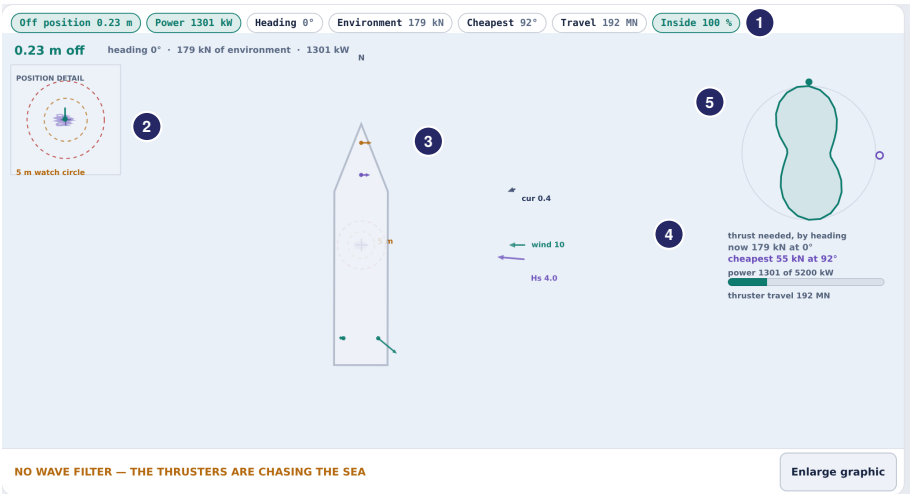

Figure 6.4 Scenario 3 in progress, with the controllers looking straight at the position reference in a 4 m sea at $T_p = 11$ s, cropped to the plant view. (1) The badges give the frame at $t = 900$ s, with the vessel 0.23 m off station on 1301 kW and the command travel at 192 MN, on the way to 327.5 MN over the 300 to 1800 s window used throughout Section 6.6, against 72.0 MN in the identical sea with the wave filter in. (2) The position detail draws the watch circle. (3) The hull carries its thruster arrows, and (4) the current, the wind and the sea arrive from the side. (5) The thrust needed by heading is the objective of the heading supervisor. The banner at the foot names the cause. The trend strips, in which the demands reverse at the wave period, and the scorecard are cut.

6.6.1 The Cost of Chasing the Sea

Scenarios 3 and 4 are the same seaway with the same tuning and the same environment feedforward, differing only in what the controllers are shown. The environment is $H_s = 4$ m at $T_p = 11$ s from 95°, with 10 m/s of wind and 0.4 m/s of current, which is a large sea and a modest steady load on purpose, because a vessel that cannot hold the environment anyway would teach the wrong lesson twice. Table 6.4 reports the three estimator settings, the model-only column being scenario 4 with that setting chosen before the run. Figure 6.5 draws four of its rows as bars, and Figure 6.4 shows scenario 3 in progress.

> **TRY IT IN THE APP** **3 · The Sea Arrives, With No Filter**
>
> Load the case and press Run. Watch the surge and sway demand strips rather than the position: they reverse at the wave period for as long as the sea lasts. Then load scenario 4 and compare the Thruster travel tiles at the same time. Over the chapter's window the travel is 327.5 MN against 72.0 MN, while the vessel sits about half a meter off station in both.

The travel figure is the one to put in a business case, because it is the only one that moves by a factor. Over the observation window the thruster commands travel 327.5 MN with the raw reference and 72.0 MN with the wave model, almost all of the difference in reversals at the wave period on machinery designed for hours of steady running. Energy is not the argument at 9 %, and a reader shown only the fuel bill would conclude the wave filter was optional. The vessel is not better positioned for any of it either. It sits 0.55 m off on average against 0.51 m, because a force that reverses every five seconds does nothing at all to eight thousand tonnes. An earlier version of the application gave 1.34 m against 0.63 m, a gap that came mostly from the feedforward its raw setting withheld and from the start-up of Section 6.3.6, both since repaired.

Table 6.4 The Same 4 m Seaway and the Same Tuning at Three Estimator Settings. Every Figure Is Taken Over the Same Observation Window, the 1500 s From $t = 300$ s to $t = 1800$ s, With Travel and Energy Accumulated Across It, Mean Power as Energy Over the Window, and the Rest Averaged or Maximized Across It. The Environment Feedforward Is On in All Three Columns.

Quantity	Raw	Model only	Observer
Surge demand, standard deviation (kN)	96	72	65
Thruster command travel (MN)	327.5	91.1	72.0
Thruster energy (kW h)	452	461	412
Mean thruster power (kW)	1085	1107	989
Mean off-position, low frequency (m)	0.55	0.51	0.51
Peak off-position (m)	3.02	3.19	3.29
Mean estimate error (m)	0.77	0.55	0.41
Inside the watch circle (%)	100.0	100.0	100.0

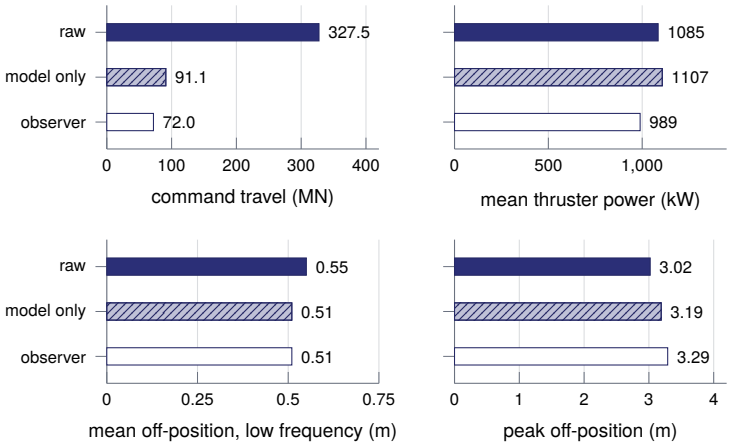

Figure 6.5 Four rows of Table 6.4 drawn as bars for the three estimator settings in the same 4 m sea over the window from 300 to 1800 s: solid for the raw reference, hatched for the model only and white for the full observer. Notice that only the command travel moves by a factor, while the mean power and the mean off-position hardly change and the peak off-position is a little larger with the observer.

> ✦ DID YOU KNOW? **A Quieter Loop With a Worse Worst Moment**
>
> From the simulator: the full observer does not buy a smaller worst moment. Over the window its peak off-position is 3.29 m, against 3.02 m on the raw reference that costs four and a half times the command travel. The observer saves the machinery, not the worst moment (Table 6.4).

The measurement carries two motions the thrusters cannot answer rather than one, and the wave filter rejects both. With σ_p and σ_g set to zero before the run the same comparison gives 82.8 MN against 25.4 MN, a factor of 3.3 against the 4.5 measured with the instruments as specified, so wave-frequency motion alone accounts for the smaller part of the difference. That is a stronger claim than a noise-free reference can support, and a more honest one, because an instrument resolving position to a third of a meter is the realistic case. The silenced generators draw no random numbers, so the sea in those runs is a different realization of the same seaway, with a 2.21 m peak.

The mechanism is the derivative term. With the raw reference the position handed to (6.15) carries the full wave motion, 2.39 m at its peak at $\omega_0 = 0.571\,\mathrm{rad/s}$, a rate amplitude of 1.37 m/s. The derivative filter at $\tau_f = 9$ s attenuates that by $1/\sqrt{1+(\omega_0\tau_f)^2} = 0.19$, and $K_c^p \tau_D^p = 1.55\,\mathrm{MN\,s/m}$ turns the remainder into about 400 kN of commanded amplitude at the peak. The same product acting on the instrument noise is what makes τ_f itself compulsory rather than merely useful: in the calm

water of scenario 2 with the raw reference and $\tau_f = 0$ the surge demand runs from one 680 kN clamp to the other, at a clamp 20 % of the time, with a standard deviation of 504 kN over 300 to 1800 s, against 91 kN at $\tau_f = 9$ s and 2 kN with the noise removed as well.

Removing the derivative term altogether removes the loop with it, which is what scenario 2 exists to show. In its calm water a 5 m north step made at 300 s is 63 % complete at 27.6 s and 95 % at 79.2 s, and the low-frequency position peaks at 6.29 m over the following ten minutes against 5.65 m with the instruments silenced. Stepped at other instants between 200 and 900 s the peak ranges from 6.10 to 6.39 m, so the noise costs between eight and fifteen points of overshoot and little else. At $\tau_D^p = 0$, set before the run, the same step first overshoots to 10.7 m at 53 s, and the swing then grows into a sustained oscillation of about 15 m either side of the setpoint with a period near 102 s, still running an hour later. The derivative term is compulsory on this plant and it is also the term that converts both wave motion and instrument noise into thruster commands, which is precisely why the answer has to be an estimator and not a filter.

🕮 **Pause.** Try Exercise 6.1 before reading on.

The intermediate setting is the useful control. A vessel model and a bias estimate with no wave model give 91.1 MN of travel, capturing 92 % of the reduction from 327.5 MN to 72.0 MN and holding station as well as the full observer to within a centimeter. Almost all of the benefit is in having any model at all, and the remaining 19.1 MN is what the wave states are worth. That is a weaker case for the wave states than the same comparison on a noise-free reference, where the model-only setting captures only 60 %, and it is the honest one.

6.6.2 The Observer Gain

Scenario 5 walks the observer gain across its range in the same heavy sea, raising it to 1.2/s at 300 s, dropping it to 0.05/s at 900 s and restoring the tuned 0.28/s at 1500 s. Held at the tuned gain over 300 to 1800 s (scenario 4), the estimate sits 0.41 m from the true low-frequency position on average, the vessel holds 0.81 m from the wellhead and the commands travel 2.88 MN per minute. At $K = 1.2/\text{s}$ the travel is 11.93 MN per minute in scenario 5 over 300 to 900 s and 12.18 MN held over the whole window, nearly all the way back to the 13.10 MN of no filter at all, because an estimator that believes the measurement completely is not an estimator and the measurement contains the sea and the instrument noise together. In the ten minutes at $K = 0.05/\text{s}$ the estimate is smooth and 1.78 m from the truth on average, 2.30 m at worst, and the vessel wanders to 4.14 m while the trace on the screen reports that everything is fine. Held at 0.05 from the start of scenario 4, where the estimator opens on the true load and nothing changes it, the estimate averages only 0.21 m from the truth over 300 to 1800 s. The drift in scenario 5 is the low gain failing to correct the velocity error the high-gain spell leaves behind, 0.19 m/s in the last ten seconds before the switch against 0.03 m/s over the first 300 s at the tuned gain.

The price of filtering is stated cleanly by scenario 10, which steps the position reference 3 m at $t = 400$ s. Measured as the estimate's offset from the true position, the estimate covers 63 % of that step in 5.05 s at the tuned gain and in 17.35 s with the gain set to 0.05/s before the run. Anything the estimator puts into the wave states and the vessel model it does not put into the position, so an estimator tuned for a heavy sea is slower to notice a real excursion. That is the trade in every filtering problem, and it is why industrial observers schedule the gain against the sea state rather than fixing it [6].

6.6.3 Bias Feedforward Against Integral Action

Scenario 6 removes the feedforward at 200 s and restores it at 900 s. Until then the bias estimate holds the environmental force it opened with, and the controller subtracts it before the loops see any error at all. When the load changes the estimate has to find the change, and after the squall of scenario 7 it covers 63 % of it in 237 s. With the feedforward off the vessel is pushed 4.38 m downwind by 387 s and is still 1.78 m out when the feedforward returns at 900 s, because the only mechanism left that

can find a steady force is an integral with a ten-minute time constant, long deliberately, because a fast integral on a plant with this little damping rings. The return is not clean. The integral still holds most of the force it found, so the force is counted twice: the vessel crosses inside half a meter at 919 s, carries on to 3.28 m off the other way at 1073 s, and is still 1.41 m out at 1500 s while the same slow integral unwinds.

Scenario 7 measures what the arrangement is worth when something happens. A squall takes the wind from 12 to 24 m/s and the sea to 4 m, which raises the environmental demand at zero heading from 164 to 417 kN and the mean thruster power from 15 % of installed over 100 to 300 s to 51 % over 500 to 1500 s. The vessel is pushed 6.7 m off station before the estimator works out what has happened, and is back inside a meter 6.9 min after the squall arrives. With the feedforward off before the run the integrals carry the original load and the same squall takes the vessel 7.7 m off, enough with the wave motion on top to pass the 9 m disconnect radius at 435 s, and it is still 3.4 m off at 900 s and not back within a meter until 1195 s, a quarter of an hour after the squall. The bias converges to within 2 % of the true environment about ten and a half minutes after the squall, which places the feedforward exactly where it belongs. It is fast enough to matter and not fast enough to be the answer on its own.

6.6.4 The Heading Is Worth More Than the Gains

Scenario 8 holds the vessel beam on to 17 m/s of wind, 0.8 m/s of current and 3.5 m of sea, an environment of 372 kN that the exact minimum-norm solution carries on 2135 kW. The scripted selection of the minimum-thrust heading at 400 s slews the setpoint through 94° in 375.5 s at 0.25 °/s, and the heading has held within 5° of the cheapest one for two minutes by 981 s. Afterward the environment asks for 109 kN and the mean thruster power over 1400 to 2400 s is 738 kW against 2297 kW over 100 to 400 s before the turn, a saving of 68 %. Nothing about the control system changed. The vessel presents about a third of its wind area and a fifth of its underwater area to the weather, which is the whole of the mechanism and is visible in Figure 6.3.

> **TRY IT IN THE APP 8 · The Heading Is Free**
>
> Load the case and press Run. At 400 s the minimum-thrust heading is selected: watch the strip Heading against the cheapest one as the setpoint walks through 94° in 375.5 s. Compare the Power tile before the turn and after 1400 s. The chapter's means are 2297 kW and 738 kW, with the control system unchanged.

This is the result the application is organized around. A single supervisory choice over a variable the job leaves free is worth a factor of three in thruster power, and no position loop can make it, because it has no way to know that the setpoint it is not being given is free. Weather optimal positioning of this kind has been designed as a feedback law and demonstrated experimentally [15], and the version here is the simpler supervisory form.

6.6.5 A Solved Allocation Against a Written One

Scenario 9 trips the port azimuth at 400 s in a brisk sea, and Table 6.5 compares the two allocations, the fixed split chosen before the run, over the ten minutes from 300 to 900 s that bracket the trip.

Table 6.5 Port Azimuth Failure at $t = 400$ s. Both Columns Use the Same Tuning and the Same Environment, and Every Statistic Is Taken Over the 600 s From $t = 300$ s to $t = 900$ s. The Median Starboard Azimuth Thrust Over 200 to 400 s, Before the Trip, Is 65 kN.

Quantity	Minimum norm	Fixed split
Peak off-position (m)	1.87	5.40
Mean off-position (m)	0.70	2.86
Median starboard azimuth after the trip (kN)	128	105

The pseudo-inverse solves the same problem with one fewer column and hands the work to whatever is left, and nobody wrote a rule for it, because a failed machine is a weight of zero in (6.17) and the arithmetic redistributes on its own. The fixed split of (6.19) does have a rule, the rule is still followed with one azimuth gone, and the share belonging to the missing machine is simply not produced, so the remaining azimuth never rises past a median of 105 kN against the 128 kN the pseudo-inverse asks of it. The vessel is four times as far off, sustained, in an environment that happens to be almost all sway. Placing the weather on the bow makes the same failure considerably more expensive, because the surge axis has only the two azimuths.

6.6.6 A Fault No Controller Can Answer

Scenario 10 applies 3 m of reference bias to the north channel at 400 s and removes it at 1200 s. The vessel drives 3 m the other way, exactly the bias (3.18 m south on average over 1000 to 1200 s), after a 5.9 m overshoot at 429 s on the way, to correct an error that does not exist, and every trace available on the bridge agrees that it is on position. Filtering does not help: a steady offset is precisely what a bias estimator and an integral are built to find, and they find it perfectly. The vessel is more than a meter off the wellhead from 410 s to 1340 s, 15.5 min spent off station for a reason nobody aboard could have seen. The remedy is not in the control system. It is three independent position references and a voting scheme that rejects the one that disagrees, which is why a vessel of dynamic positioning class 2 carries them, and loss of position reference is among the dominant contributors to real incidents [12].

🕮 **Pause.** Try Exercise 6.2 before reading on.

6.6.7 What Actually Runs Out

Scenario 11 holds the vessel beam on while the weather builds through 25 m/s at 300 s, 29 m/s at 700 s and a current of 1.6 m/s at 1100 s. Three things could bind, the sway thrust, the thruster power, or the heading, and the answer is settled before the first scheduled event. The heading leaves a 3° band at 17 s, the vessel leaves the watch circle at 50 s and passes the disconnect radius at 75 s, while the power sits between 54 and 68 % throughout and never approaches its limit.

The arithmetic follows from (6.18) alone and is the result an operator can carry off this application and use. With both tunnels at their 120 kN rating and the azimuths free, the vessel can make $2F_a + 2F_t = 680\,\text{kN}$ of sway if it does not care where the bow points. Requiring zero net yaw moment couples the two, because $x_a \Sigma f_{y,azi} + x_{t1} f_3 + x_{t2} f_4 = 0$ forces the azimuths to produce $(38 f_3 + 26 f_4)/35 = 219\,\text{kN}$ of sway purely to cancel the tunnels, leaving 459 kN of useful beam thrust. Requiring the vessel also to cancel the 3.1 MN m of yaw moment the environment itself applies takes the azimuth share down to 130 kN and the total to 370 kN, against a beam demand of 541 kN. The vessel therefore cannot hold sway and heading at once, the heading loses first, and once the heading has gone the projected areas in (6.4) and (6.5) change and the rest follows, at 21 m off station at two and a half minutes, 320 m at eighteen and 904 m at thirty.

🕮 **Pause.** Try Exercise 6.3 before reading on.

Heading capability rather than installed thrust or power is therefore what binds, which is why a capability plot is drawn per heading. Once the tunnels are pinned there is no recovery: with the minimum-thrust heading chosen before the run, the setpoint walks toward the weather but the bow stalls near 28°, the disconnect radius is reached at 68 s against 75 s, and the vessel is 816 m off at thirty minutes. The operational answer is the turn of scenario 8, made before the weather arrives, while the vessel still has yaw authority to turn with.

6.6.8 The Graded Challenge

The challenge is a forty-minute shift opening with the weak tuning and four things set the wrong way: a model-only estimate with no wave states, a fixed allocation, no environment feedforward and a heading nobody chose. The sea builds at 300 s, a squall arrives at 700 s, the port azimuth trips

at 1400 s and the squall passes at 1900 s. The target, judged at 2400 s, is to hold within 3 m of the wellhead throughout the last third, to stay inside the watch circle for at least 97 % of the shift, never to run over the power limit, and to finish under 130 MN of thruster travel. As loaded the shift fails badly, 23.0 m off at worst in the last third and 20.4 % inside the circle, while meeting the travel and power clauses on 73 MN, which is why travel alone is not the target. No one of the four structural changes passes on its own, and the four together on the weak gains still fail at 10.2 m and 46 % inside.

With all four decisions right the shift finishes on 102.5 MN of travel, 100.0 % inside the watch circle, a peak thruster power of 60 % of installed and 2.64 m off station at worst over the last third. Against the target these are margins of 0.36 m, 3.0 points and 28 MN, and the application gives the same run whether the suggested settings are revealed before the run or carried through Reset. Putting each one back the wrong way, after revealing the suggested settings and before the run, separates their contributions. Leaving the heading fixed leaves the vessel 9.9 m off on average over the last third and outside the watch circle for 40 % of the shift. Leaving the wave model out gives 140.7 MN of travel against a target of 130 MN, and 3.22 m at worst, and the raw reference gives 459.2 MN. The minimum-norm allocation is what makes the trip survivable, because the fixed split finishes 53.7 % inside the circle and 4.71 m off at worst, with the power peaking at 99.7 % of installed. The power clause is a guard rather than a test, because the allocation scales the thrust back before the switchboard is overloaded, so no line-up measured here spends any time over it. The feedforward is the one of the four that could be omitted, still a pass at 100.0 % inside but only just, 2.98 m at worst against the 3 m limit on the shipped seed. None of the four is a gain, and the tuning that is left is nearly what it was.

6.7 Technology in Practice

The vessel of Section 6.3 is read by one position reference and a gyro, and moved by two azimuths and two tunnel thrusters. Figure 6.6 draws its loops with real devices in their place.

What Measures It

The reference of Equation (6.8), with its 0.35 m of colored noise, its 0.5 s hold and its optional bias, stands for one real position reference. The first deep-water drill ship carried two acoustic systems in parallel, and the redundancy was used whenever one of them was being repaired [3]. A vessel of class 2 or 3 now carries at least three references built on different principles, and three gyrocompasses where its control system depends on the heading, beside sensors for its motions and the wind [16], where the model has one reference and one gyro.

> **TECHNOLOGY** **Sensor: Hydroacoustic Position Reference**
>
> A battery beacon on the sea floor sends short sound pulses, and hydrophones at the four corners of a square in the hull hear each pulse at slightly different times. The differences are changed to north and east offsets from the beacon, and two small gyros correct them for roll and pitch. On the first deep-water drill ship, thruster cavitation drowned the pulses at high speed, beacons faded after 40 to 60 hours when their floats took on water, and one wiring fault left the computer holding station on the last acceptable signal it had received [3]. On a trend such a fault is a reading that freezes or jumps while the vessel moves, the reference bias of scenario 10 in another form.

What Moves It

The model drives each thruster through a 4 s build-up, turns the azimuths at 6 °/s and takes thrust from the tunnels as the vessel gathers way, so a lost or weakened unit is a weight in the allocation of Section 6.4.3. The first deep-water drill ship chose fixed propellers in tunnels for reliability and for protection from the drilling work and from storms. Fixed thrusters need more connected power than steerable ones, since they cannot point their thrust exactly where it is wanted, but its designers found the power actually used not appreciably different [3].

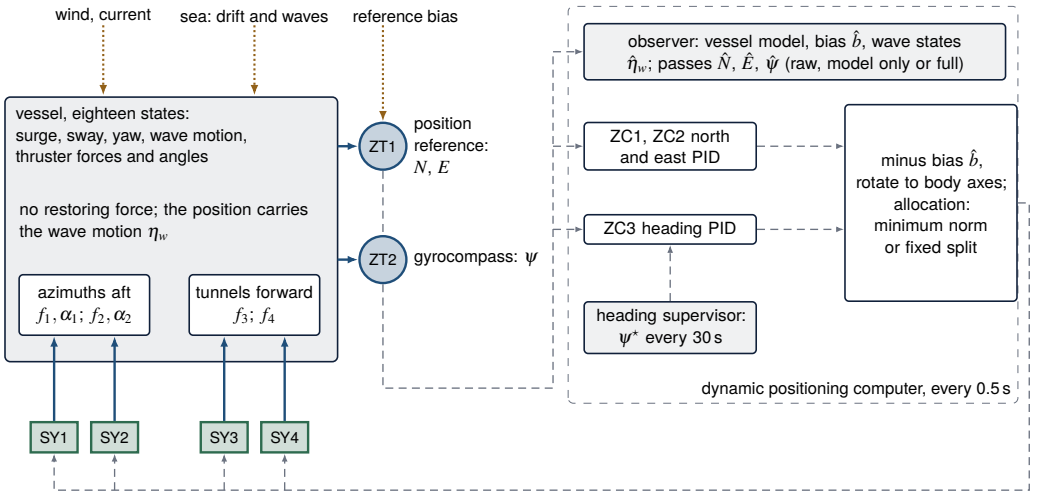

SY1, SY2: azimuth thruster drives, 220 kN each, thrust lag 4 s, turning at 6 °/s. SY3, SY4: tunnel thruster drives, 120 kN each, thrust lag 4 s, less thrust with way on. All four share 5200 kW.

Figure 6.6 The loops of the chapter with real devices in place of the schematic's symbols. ZT1 is the position reference (on a real vessel several, built on different principles, in the model one with 0.35 m of noise) and ZT2 the gyrocompass. One dynamic positioning computer (dashed frame) runs the observer, the north, east and heading PID loops (ZC1 to ZC3), the heading supervisor and the allocation every 0.5 s. SY1 and SY2 drive the two azimuths aft, SY3 and SY4 the two tunnels forward. Notice that the loops see only the observer's estimate, never the reference itself unless the raw setting is chosen, and that the reference bias enters at the instrument, where no loop can tell it from a real move.

TECHNOLOGY Final Element: Tunnel Thruster

A tunnel thruster is a fixed propeller in a tube across the hull that pushes the bow or the stern sideways. On the first deep-water drill ship each was a four-bladed propeller of 59 inches turned through a right-angle gear by an electric motor, for about 18,000 pounds of thrust. The long bow tunnels cut that thrust until the blades were reshaped, cavitation set in at full speed, and the bow units, which ran longer and harder than the stern ones, failed several times in service [3]. Thrust is lost as well to interference between thrusters and the hull [16]. On a trend a lost thruster shows as a step in the others' commands, as the starboard azimuth of scenario 9 rises from a median of 65 kN to 128 kN (Table 6.5).

Where the Controller Runs

Supervisory and lower-level control loops run with cycle times from milliseconds to minutes [17]. The chapter's observer, three PID loops and allocation share one 0.5 s scan, and the heading supervisor rescans every 30 s on top of them. In 1968 that work was done by a computer fed with digital angles from the acoustic receivers [3]. In class 2 it is done by at least two independent computers that transfer control automatically after a detected failure, each on its own uninterruptible supply, and their software also runs the consequence analysis that checks the thrust left after the worst failure [16].

What Is Modernizing It

The guideline of 1994 was replaced for vessels built from June 2017, and one flag administration recommends the new operational requirements for existing vessels as well [18]. The computers of a redundant system now talk over networks, and a regulator has warned that a dual bus network can be silenced by one faulty module where a dual star network need not share a medium at all [19]. Test platforms join a six-degree-of-freedom simulation, with the frequency-dependent hydrodynamics the

chapter's three-degree model leaves out, to a scale model tested at the same Froude number, and use the pair to find a vessel's operational time window [20].

> **IN PRACTICE** **Equipment Class and the Worst Failure**
>
> An international guideline of 1994 sorts dynamic positioning systems into three equipment classes by their worst failure. In class 1 a single fault may lose position. In class 2 no single fault in an active component, such as a generator, a thruster or a switchboard, may lose it, and class 3 adds the loss of everything in one compartment to fire or flooding. Classes 2 and 3 carry at least three position references, not all of one type, and three gyrocompasses. A consequence analysis warns when the thrusters left after the worst failure could no longer hold the vessel, the work stops before the weather reaches that point, and position keeping after single failures is tested every year [16].

6.8 Limitations

The application teaches an estimator structure and does not size a vessel. Blended area coefficients, a single-frequency sea and no thruster interaction put a real hull's capability footprint, its consequence analysis and its dynamic positioning class out of reach, and the capability arithmetic of Section 6.6 is exact for (6.18) and approximate for any real vessel. Three modeling choices in Section 6.3.7 cost more than the rest: the absent heave, which is what most often stops offshore construction work, the smooth coefficient sets, and the single position reference. Three implementation choices cost almost as much. The observer is hand tuned with one gain and a fixed pole pattern rather than carrying a covariance, so the adaptation industrial systems perform against the sea state [6] is left to a slider. Neither the passive nonlinear observer that removes the linearization about a fixed heading [11] nor the constrained allocation [9] is implemented, so the power limit scales the whole solution back rather than choosing a degree of freedom to surrender. The azimuth angles are settled by a filter outside the solve rather than a rate constraint inside it, and every automatic run opens with the estimator already converged on the true load (Section 6.3.6), a state a commissioned system reaches only by letting its bias estimate converge before it takes control.

> **OPEN CHALLENGE** **One Fault in a Redundant Network**
>
> Redundancy protects a vessel only if no single fault can reach both halves. A safety alert of 2013 reports a class 2 diving support vessel whose positioning system failed while divers were working below, so that the vessel drifted off position. The probable cause was a single fault that blocked the system's internal data communications: one module with access to both buses of a dual bus network can block both. The regulator asked owners and suppliers who claim class 2 or better to check their communication architecture [19]. Progress would be designs and tests that show, before a vessel works, that no single fault can silence every path between the references, the computers and the thrusters.

6.9 Discussion

Dynamic positioning is what makes deep-water work possible at all, and the control system on a mid-size construction vessel is the arrangement this chapter describes, a hand-tuned Kalman filter, three PID loops and a pseudo-inverse allocation, because that is what most of them are [2]. The measurements reported above say where the engineering value in that arrangement sits, and it is not where a tuning exercise would put it. Three PID loops account for none of the factor of 4.5 in thruster travel, none of the 68 % saving in thruster power and none of the difference between surviving a thruster failure and drifting off station. The estimator, the allocation and one supervisory decision account for all of it.

Of the three general questions named in the introduction, the third is the rarest and the most valuable, and it is the one a control course is least likely to raise. A variable the process leaves free is an optimization variable, and the gap between its present value and its best one is frequently larger than anything feedback can recover.

The model of Section 6.3 also supports methods beyond the loops implemented in it. Equations (6.1) to (6.10) are already the prediction model a nonlinear model predictive controller would carry, and the constraints a PID loop can meet only by clipping are what a constrained formulation represents naturally, including in the allocation, where a constrained solve handles the saturation, the azimuth rate and the forbidden zones that the pseudo-inverse and the filter of Section 6.4.6 only approximate [9].

6.10 Conclusion

An offshore vessel holding station on thrusters serves in this chapter as an interactive laboratory for the problem of deciding what part of a measurement is an error. The measured comparisons are collected in Section 6.6 and are not repeated here. What they say together is that the three PID loops in the middle of a dynamic positioning system account for almost none of the difference between a good one and a poor one. The estimator that decides what the measurement means, the allocation that turns three demands into six machines, and the one supervisory choice over a variable the job leaves free account for all of it, and the largest of the three is the one no loop can make. Three defects in the measurement path and one in the allocation were found and repaired along the way, three more in the controller since, and the wave-filter argument is stronger on the repaired engine than it was on the quiet one. Further work would add heave, a broad-banded sea, a voting scheme over three position references, and a power management system that surrenders a degree of freedom rather than scaling the whole allocation. The vessel is an accessible browser platform for a lesson that arrives before any tuning does, which is that part of a measurement can be real and still not be an error.

KEY POINTS

- A vessel on thrusters has no anchor and no restoring force, so its position reference shows both the slow drift the thrusters must stop and the wave motion they cannot and need not stop.
- Shown the raw reference, the thruster commands travel 327.5 MN over the window against 72.0 MN with the full observer, and the vessel holds station no better.
- Choosing the heading with the least environmental load saves 68 % of the thruster power in scenario 8, a choice no position loop can make.
- After a thruster trip the solved allocation keeps the vessel 0.70 m off on average against 2.86 m for a fixed split.

Where else this applies. Chapter 5 chooses which waves a foiling boat rides over and which it follows, and Volume 5, Chapter 13 filters the late position fixes of an underwater vehicle.

References

[1] T. I. Fossen, *Handbook of Marine Craft Hydrodynamics and Motion Control*. Chichester, UK: John Wiley & Sons, 2011.

[2] A. J. Sørensen, "A survey of dynamic positioning control systems," *Annual Reviews in Control*, vol. 35, no. 1, pp. 123–136, 2011.

[3] Deep Sea Drilling Project, "Dynamic positioning," Deep Sea Drilling Project, Tech. Rep., 1971.

[4] J. G. Balchen, N. A. Jenssen, E. Mathisen, and S. Sælid, "A dynamic positioning system based on Kalman filtering and optimal control," *Modeling, Identification and Control*, vol. 1, no. 3, pp. 135–163, 1980.

[5] T. I. Fossen and T. Pérez, "Kalman filtering for positioning and heading control of ships and offshore rigs," *IEEE Control Systems*, vol. 29, no. 6, pp. 32–46, 2009.

[6] T. D. Nguyen, A. J. Sørensen, and S. T. Quek, "Design of hybrid controller for dynamic positioning from calm to extreme sea conditions," *Automatica*, vol. 43, no. 5, pp. 768–785, 2007.

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

[8] S. Saelid, N. Jenssen, and J. Balchen, "Design and analysis of a dynamic positioning system based on Kalman filtering and optimal control," *IEEE Transactions on Automatic Control*, vol. 28, no. 3, pp. 331–339, 1983.

[9] T. A. Johansen and T. I. Fossen, "Control allocation—a survey," *Automatica*, vol. 49, no. 5, pp. 1087–1103, 2013.

[10] M. J. Grimble, R. J. Patton, and D. A. Wise, "The design of dynamic ship positioning control systems using stochastic optimal control theory," *Optimal Control Applications and Methods*, vol. 1, no. 2, pp. 167–202, 1980.

[11] T. I. Fossen and J. P. Strand, "Passive nonlinear observer design for ships using Lyapunov methods: full-scale experiments with a supply vessel," *Automatica*, vol. 35, no. 1, pp. 3–16, 1999.

[12] H. Chen, T. Moan, and H. Verhoeven, "Safety of dynamic positioning operations on mobile offshore drilling units," *Reliability Engineering & System Safety*, vol. 93, no. 7, pp. 1072–1090, 2008.

[13] W. Blendermann, "Parameter identification of wind loads on ships," *Journal of Wind Engineering and Industrial Aerodynamics*, vol. 51, no. 3, pp. 339–351, 1994.

[14] J. A. Pinkster, "Mean and low frequency wave drifting forces on floating structures," *Ocean Engineering*, vol. 6, no. 6, pp. 593–615, 1979.

[15] T. I. Fossen and J. P. Strand, "Nonlinear passive weather optimal positioning control (WOPC) system for ships and rigs: experimental results," *Automatica*, vol. 37, no. 5, pp. 701–715, 2001.

[16] International Maritime Organization, "Guidelines for vessels with dynamic positioning systems," MSC/Circ.645, Maritime Safety Committee, London, 1994, adopted 6 June 1994.

[17] 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.

[18] Commissariat aux affaires maritimes, "Ships fitted with dynamic positioning systems," Circular CAM 03/2020, Luxembourg, 2020. [Online]. Available: https://maritime.public.lu/dam-assets/circulaires/003-Dynamic-Positioning-Systems.pdf

[19] Health and Safety Executive, "Warning to offshore industry on blocking of data communications in dynamic positioning systems," Safety alert OSD 1-2013, Offshore Safety Division, 2013. [Online]. Available: https://www.hse.gov.uk/safetybulletins/dynamic-positioning-systems.htm

[20] C. Hu, Q. Shi, X. Li, and X. Guo, "Development of a hybrid simulation and experiment test platform for dynamic positioning vessels," *Ships and Offshore Structures*, vol. 19, no. 10, pp. 1548–1562, 2024.

6.11 Exercises

6.1 ○ Concept check · 2 min

A designer proposes to keep the sea out of the thruster commands by placing a first-order low-pass filter, with a time constant of about one wave period, on the position reference ahead of the three PID loops. Why does an industrial dynamic positioning system use an observer with a wave model instead?

(a) A low-pass filter cannot attenuate a motion as slow as an eleven-second wave
(b) The filter's phase lag removes the damping that the derivative term must supply
(c) A low-pass filter also removes the steady environmental load the loops must carry
(d) The observer executes on a faster clock than any filter placed in the loop

6.2 ○ Concept check · 1 min

In scenario 10 the differential correction goes stale at 400 s and the north reading moves 3 m away from the truth. Which change would have kept the vessel on the wellhead?

(a) A higher observer gain, so the estimate follows the reference sooner
(b) A lower observer gain, so the estimate ignores the sudden jump
(c) Independent position references with voting that rejects the one that disagrees
(d) A shorter position integral time, so the offset is removed sooner

6.3 ○ Concept check · 2 min

In scenario 11 the weather builds on the beam with the heading as set. The beam demand is 541 kN and the four thrusters can make 680 kN of sway between them. Which limit binds first?

(a) The thruster power, because the wind load grows with the square of the wind speed
(b) The sway thrust, because the beam demand exceeds the installed sway at the start
(c) The position integral, which winds up against the demand clamp of the north loop
(d) The heading, because the pinned tunnels cannot hold the sway and the yaw at once

6.4 Theory and calculation · 8 min

The vessel of Table 6.2 has $m = 8.0 \times 10^6$ kg, $a_u = 0.08$, $a_v = 0.85$, $k_g = 0.25$, $L_{pp} = 90$ m, $a_r = 0.35$ and $D_u = 9.0 \times 10^4$ N s/m. At $\psi = 0$ the north loop acts in surge with $K_c^p = 37$ kN/m and $\tau_D^p = 42$ s.

(a) Find m_u, m_v and I_z from Equation (6.2), and the ratio m_v/m_u. *Checkpoint:* 8.64×10^6 kg, 1.48×10^7 kg, 5.47×10^9 kg m^2 and 1.71.

(b) Treat the north loop as a proportional-derivative law on the true position acting on m_u with damping D_u, and ignore the derivative filter, the integral, the observer and the thruster lag. Find the natural frequency, its period and the damping ratio. *Checkpoint:* 0.0654 rad/s, 96 s and 1.45.

(c) Find the damping ratio with $\tau_D^p = 0$, and compare the period with the sustained oscillation of about 102 s that scenario 2 shows at $\tau_D^p = 0$. *Checkpoint:* 0.080. The oscillation is a little slower than 96 s.

6.5 Theory and calculation · 6 min

In scenario 3 the raw position reference carries wave motion of 2.39 m peak at $T_p = 11$ s. The position loop has $K_c^p = 37$ kN/m, $\tau_D^p = 42$ s and a derivative filter $\tau_f = 9$ s.

(a) Find the peak wave frequency ω_0 and the peak rate of the wave motion. *Checkpoint:* 0.571 rad/s and 1.37 m/s.

(b) Find the gain of the derivative filter at ω_0 and the amplitude of the commanded force that the derivative term makes from the wave motion. *Checkpoint:* 0.19 and about 405 kN.

(c) Repeat (b) with $\tau_f = 0$ and compare the result with the 680 kN demand clamp. *Checkpoint:* About 2.12 MN, about three times the clamp.

6.6 Theory and calculation · 7 min

The two azimuths are rated 220 kN each at $x_a = -35$ m and the two tunnels 120 kN each at $x_{t1} = 38$ m and $x_{t2} = 26$ m. The vessel is stopped, so $\ell = 1$. The azimuths can point their thrust anywhere.

(a) Find the most sway the four machines can make if the yaw moment does not matter, and the most they can make with zero net yaw moment while both tunnels push at their rating. *Checkpoint:* 680 kN, and 459 kN with the azimuths making 219 kN.

(b) The environment also applies a yaw moment of 3.1 MN m in the same sense as the azimuths' moment, so the thrusters together must make 3.1 MN m in the sense of the tunnels' moment to cancel it. Find the sway left, and compare it with the beam demand of 541 kN in scenario 11. *Checkpoint:* About 371 kN, short of the demand by about 170 kN.

(c) With $k_P = 1.79 \times 10^{-5}$ kW per N$^{1.5}$, find the power of all four machines at their ratings and say whether power can be the binding limit. *Checkpoint:* 5182 kW against 5200 kW installed.

6.7 App challenge · 8 min

Load scenario 3, *The sea arrives, with no filter*, as shipped (Position estimate on Raw), default seed. Run to 1800 s at speed 20 and export the CSV. Read `metric_travel_kN` and `metric_energy_kWh` in the rows where `Time (min)` is 5 and 30. Then load scenario 4, *The same sea, filtered*, as shipped (Wave filter) and repeat.

(a) Find the thruster command travel over 300 to 1800 s for each run, in MN, within 1 MN. *Checkpoint:* 327.5 MN on Raw and 72.0 MN with the wave filter.

(b) Find the energy over the same window for each run, within 2 kW h, and the ratio of Raw to Wave filter for travel and for energy. *Checkpoint:* 452 kW h and 412 kW h; ratios 4.5 and 1.10.

6.8 App challenge · 8 min

Load scenario 8, *The heading is free*, as shipped, default seed. Run to 2400 s at speed 20 and read the Cheapest heading tile. Export the CSV.

(a) Average `power_kW` over the rows with `Time (min)` from 1.66 to 6.67 (100 to 400 s, before the turn), within 10 kW. *Checkpoint:* About 2296 kW.

(b) Average `power_kW` over the rows from 23.33 to 40 (1400 to 2400 s), within 10 kW, and find the saving as a percentage of the power before the turn. *Checkpoint:* About 738 kW, a saving of 68 %.

(c) Explain the saving from Figure 6.3 and say which setting of the control loops changed during the run. *Checkpoint:* Cheapest heading 94°. No loop setting changed.

6.9 App challenge · 8 min

Load scenario 9, *A thruster trips*, as shipped (Thrust allocation on Minimum norm), default seed. Run to 900 s at speed 20 and export the CSV. Press Reset, set Thrust allocation to Fixed split before Run, run to 900 s again and export a second file. For each file use the rows with `Time (min)` from 5 to 15 (300 to 900 s).

(a) Find the mean and the largest value of `radius_m`, within 0.05 m. *Checkpoint:* Minimum norm 0.70 m and 1.87 m; fixed split 2.86 m and 5.40 m.

(b) Find the median of `f2_kN`, the starboard azimuth, over the rows after the trip (`Time (min)` above 6.67), within 2 kN, and explain the difference between the two runs. *Checkpoint:* 128 kN against 105 kN.

6.10 Application area · 3 min

A vessel manager reviews the comparison of scenarios 3 and 4 and notes that the wave filter saves only about 9 % of the thruster energy. The manager asks whether the wave filter is worth keeping. Which number makes the case, and what does it stand for on the vessel?

6.11 Application area · 4 min

Every automatic run in the application starts with the bias estimate already holding the true environmental load. On a real vessel the DP system is started with the vessel already on station and then switched to automatic. What should the operator wait for before relying on automatic control, and what does this chapter's estimator say about how long that takes?

6.12 Review the agent · 5 min

An AI assistant, shown scenario 4 at the tuned observer gain, writes: "The thrusters are still busy in this sea. Set the Observer gain to 0.05 and leave it there in every weather. The estimate on the North strip becomes almost flat, which shows the vessel is now holding station better, and the thruster travel falls further." Find the flaw in this advice and say how you would check it in the application.

CHAPTER 7

CubeSat Attitude Simulator Whose Wheels Keep Every Torque They Absorb

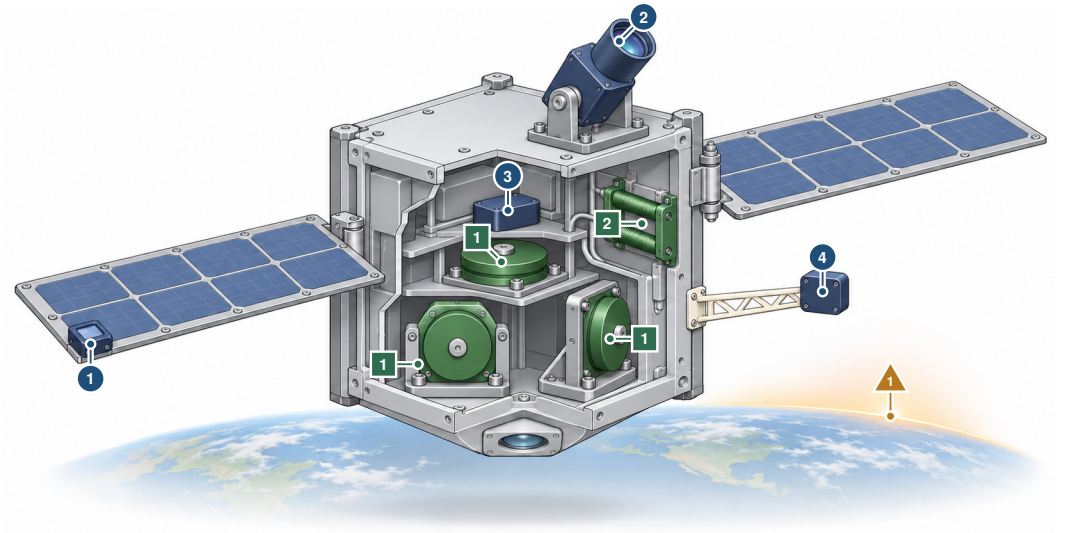

The controller moves 1 three reaction wheels, one on each axis 2 magnetic torque coils

The controller reads 1 sun sensor on a solar array 2 star tracker, looking away from the Earth 3 gyro 4 magnetometer on a short boom

What disturbs it 1 the Sun, which blinds the star tracker when it comes near its line of sight

Notice the camera window facing the Earth and the star tracker looking the other way, while the wheels inside can turn the satellite but keep every push they absorb. Illustration, not to scale.

IN PLAIN TERMS Pointing a Small Satellite With Spinning Wheels

The main idea. A small satellite turns its camera by spinning little wheels inside it the other way, as a person on a free office chair turns by swinging an arm. Magnetic coils push on the Earth's field to slow the wheels again.

The challenge. The wheels soak up every push of the thin air until they spin their fastest, and the coils push only across the Earth's field. The star camera that shows where the satellite points goes blind when the Sun shines into it.

Why it matters. Sharing the work between wheels, coils and the best guess of where the satellite points decides whether the camera holds its spot on the ground.

THREE WAYS TO READ THIS CHAPTER

⊘	**Explore**	12 min	the first page, the fact in the Introduction, the two history segments of Section 7.2, the try-it segments of scenarios 3 and 6 and the fact on the challenge (Section 7.6), and the key points
🕮	**Study**	60 min	the whole chapter with its three pause points and the twelve exercises at its end
⛰	**Practice**	20 min	the first page, the measured cases (Section 7.6), Technology in Practice (Section 7.7), Limitations (Section 7.8) and the open challenge

The control question Reaction wheels turn a small satellite but keep every torque they absorb, and magnetic coils can empty them only across the Earth's field. How must pointing, momentum unloading and attitude estimation share the work so that the camera still holds a ground site through a pass with its star tracker half blind?

Learning objectives

- Compute how fast a constant disturbance torque fills a reaction wheel, and explain why no pointing gain changes it.
- Predict the orbit-averaged unloading rates of the cross-product law from the field's direction, and explain why a strong B-dot gain detumbles slowly.
- Compare an eigenaxis slew with three single-axis loops by the angle turned and the peak body rate.
- Explain why an attitude filter carries the gyro bias as a state, and measure what leaving it out costs during a blind window.
- Judge what a predictive pointing law adds on a handover: the limits it keeps rather than its preview.

Prerequisites Sections 1.4.2 (eigenaxis and stored momentum), 1.4.6 (bias state), 1.4.7 (directional actuators) and 1.6.4 (scenarios and the scorecard). Section 1.6.3 is needed only for the predictive law.

The app https://apmonitor.com/apps/satellite-attitude/

See also Volume 4: Volume 4, Chapter 13 for actuators that push only inside a cone and Volume 4, Chapter 12 for allocation against a consumable budget. Volume 5: Volume 5, Chapter 13 for drift between position fixes and Volume 5, Chapter 12 for a store that is never measured.

7.1 Introduction

An Earth-observation satellite the size of a shoebox now carries the same kinds of actuators and sensors as a large one. The CubeSat standard fixed the outside of such a spacecraft and its interface to a shared deployer, so that many of them could ride as secondary payloads and share a development cost [1], and a 3U CubeSat designed in 2019 was to point better than 0.25° with reaction wheels and a star tracker, its attitude filter running on a microcontroller [2]. Those actuators behave unlike a valve or a heater. A reaction wheel turns the body by spinning the other way, so it moves angular momentum between itself and the spacecraft and removes none of it, and a magnetorquer coil feels a torque $\mathbf{m} \times \mathbf{B}$ that is always perpendicular to the local field. The physics is textbook material, with sensors and actuators, attitude filtering and attitude control each given a chapter of a standard text [3], but its consequences for a controller are easy to miss.

> ✦ DID YOU KNOW? **Despun by the Field in Two Years**
>
> Eddy currents that the Earth's magnetic field drove in the spinning outer shell of Vanguard I slowed it from 2.7 revolutions a second at launch to less than 0.2 revolutions a second two years later [4]. The coils of this chapter use the same field on purpose, to unload the wheels.

Intuition carried over from process loops misleads on this plant, and the simulator puts a number on each surprise. A wheel controller that holds the camera within a few hundredths of a degree looks able to hold it indefinitely, yet at nadir a drag torque of 2 µN m, fixed in the body about pitch, fills the pitch wheel by 37.85 % of its capacity every orbit, and after 245.0 min that wheel can absorb nothing more and the attitude is lost. A stronger magnetic law looks faster, yet B-dot at 8 mN m s needs 85.5 min to bring a 10.1 °/s tumble under 0.5 °/s where 0.5 mN m s needs 26.8, because the strong law leaves the spacecraft turning about the field line, where no coil has any authority. An attitude held on its setpoint looks like an attitude achieved, yet while the Sun blinds the star tracker for 11.6 min a filter that leaves out the gyro bias lets the camera drift 2.35° from nadir, with the pointing loop holding its own estimate on nadir throughout. The double integrator $1/(Js^2)$ from wheel torque to angle contains none of this, because the storage, the direction of the field and the estimate all lie outside it.

None of the three is peculiar to spacecraft. An actuator that stores what it applies needs a second actuator to unload it, and the unloading becomes a slow loop sized against the secular disturbances, which for magnetic laws has been analyzed by averaging their effect over an orbit [5]. Magnetic actuation adds a control direction that the orbit keeps turning, which makes pointing with coils alone a time-varying problem, although it has been shown to have global solutions by attitude and rate feedback [6], and detumbling with coils alone has a convergence proof and a gain criterion that allows for saturated coils [7]. A walking robot meets the same restriction in feet that push only inside a friction cone (Volume 4, Chapter 13). The estimation half integrates a biased rate sensor between absolute fixes that come and go, the arrangement in which the gyro bias became a state of attitude filters [8], and the one in which an underwater vehicle drifts between acoustic fixes (Volume 5, Chapter 13). A surge drum under averaging level control shares the first half of this shape, storing every flow imbalance until a slower loop returns the inventory.

The literature studies these pieces apart, the unloading law by its orbit average [5], the magnetic laws by their convergence [6, 7] and the filter by its error representation [9]. This chapter puts them on one browser simulator of a 6U CubeSat whose flight software sees only its instruments. On it the wheels can be watched filling, the momentum beyond the coils' reach can be read at every instant, and one orbit can be flown with and without the gyro bias in the filter, so that pointing, stored momentum and imaged pass can be apportioned among the layers of an attitude control system. Quaternion feedback, B-dot, the cross-product unloading law and the multiplicative filter are standard laws taken from the literature. Two predictions check the engine, the orbit-averaged unloading rates, under which a pitch excess falls to a tenth in 36.2 min against a predicted 35.1 and an excess in the orbit plane in 64.4 to 70.3 min against 67.9 to 68.1, and the gravity-gradient libration of the free spacecraft, 84.65 min against 84.63. The spacecraft has ten states with its orbit, Sun, shadow and field, instruments that are the controller's only view of the truth, two filters fed the same readings in parallel and a mode supervisor, written down completely enough to be rebuilt. Ten cases and a challenge isolate storage, field geometry, slewing, estimation, an honestly infeasible wheel failure and a handover slewed by an optional model predictive controller, whose measured comparison with the shipped laws credits it with the limits it keeps rather than with speed.

Section 7.2 traces how spacecraft came to be pointed. Section 7.3 sets out the spacecraft, its orbit and its instruments, and Section 7.4 what the inertia and the field permit any controller. Section 7.5 describes the flight software, with the predictive law in Section 7.5.5. Section 7.6 measures the ten cases and the challenge, and Section 7.6.11 sets the predictive law beside the shipped laws. Section 7.7 names the devices of real spacecraft, Section 7.8 states the limitations and Section 7.9 draws the lessons together. The background comes from Chapter 1: the eigenaxis and the momentum stored in wheels from Section 1.4.2, the bias state from Section 1.4.6, actuators that act in only some directions from Section 1.4.7, and the simulator framework from Section 1.6.

7.2 From Spinning Satellites to Three-Axis CubeSats

A spinning body keeps its axis fixed in space, but its environment wears the spin down and moves the axis, as the history box below tells for the first satellites that met the Earth's field. A report of 1961 from the Applied Physics Laboratory set out how the attitude of an Earth satellite could be stabilized by magnetic and by gravity torques, and how the attitude acts on the solar power and the thermal balance [10]. By 1978 attitude determination, prediction and control had grown into a reference of 858 pages prepared by the Computer Sciences Corporation, whose foreword, from the head of the attitude section at NASA's Goddard Space Flight Center, explained that comprehensive reference material had been difficult to obtain [11].

HISTORY The Field as Disturbance and as Tool (1962 to 1969)

In the weather satellite Tiros I, a large magnetic dipole inside the spacecraft turned against the Earth's field, precessed the spin axis and changed the direction of the cameras in ways nobody had predicted. Designers soon used the same torques on purpose. Transit 1B carried a permanent magnet to hold its axis along the local field, and later Tiros satellites aimed their cameras by changing their magnetic moment on command from the ground [4]. On the spin-stabilized Telstar satellite of 1962, eddy-current and magnetic torques slowed the spin with a characteristic time estimated at about 330 days and drifted the spin axis, while a damper cut the precession angle by a factor of e within 30 min at most [12].

Active magnetic control used the same field. An elementary magnetic attitude control system was published in 1976 [13]. Inertial pointing with coils alone was shown in 2002 to have a global solution by static attitude and rate feedback and by dynamic attitude feedback [6], and a rigorous proof that coils alone bring a rigid spacecraft to rest from arbitrary tumbling in a time-varying field came in 2012, together with a simple criterion for a gain that detumbles in quasi-minimum time while the coils saturate [7].

Reaction wheels gave three-axis control without propellant and created the problem of unloading them. Analytical formulas for the orbit-averaged behavior of magnetic laws that unload a reaction wheel system against secular environmental torques were derived in 1980 for an inertially fixed spacecraft whose dominant disturbance was the gravity gradient, and they agreed within 10 % with detailed simulations of the Solar Maximum Mission spacecraft [5]. Section 7.4 applies the same averaging to the Earth-pointing spacecraft of the simulator.

Large slews came to be flown on quaternion errors. The stability of large-angle reorientations under quaternion feedback, with a strapdown inertial reference supplying the attitude change and reaction jets under pulse-width pulse-frequency modulation supplying nearly proportional torques, was analyzed with Liapunov's theorem in 1985 [14], and a quaternion feedback regulator for rotations about the eigenaxis followed in 1989 [15]. Attitude knowledge took a parallel path. A review of 1982 gathered the Kalman filters for spacecraft that carry three-axis gyros and attitude sensors, in which the gyro data replace a dynamic model, the attitude state is augmented with the gyro biases, and the gyro noise enters as state noise rather than as observation noise [8]. Because the quaternion must keep a unit norm, the filter keeps a quaternion for the global attitude and a three-component representation for its error, the form called the multiplicative extended Kalman filter [9]. One absolute reference for such a filter is the autonomous star tracker, an electronic camera whose processor identifies stars against an internal catalog and computes the attitude quaternion by itself [16].

The CubeSat brought these pieces to small spacecraft. A standard deployer developed with Cal Poly students was designed to release CubeSats as secondary payloads on a wide range of launchers, provided each conformed to common physical requirements and shared the deployer interface [1]. The 3U design of 2019 combined reaction wheels, a star tracker, cold-gas propulsion and an unscented Kalman filter on a microcontroller [2]. Wheels still fail, and a spacecraft that loses them becomes a different spacecraft, as the second history box of this section tells. In the simulator all of these layers fly together, and each is judged against a true attitude that only the scoring can read.

Table 7.1 Signals of the Spacecraft by Role, With Their Reference Values

Symbol	Role	Reference	Range
$\hat{\phi}$, $\hat{\theta}$, $\hat{\psi}$	CVs of pointing, estimated 3-2-1 angles from the orbit frame	0	offsets
e_{cam}	scored CV, true camera from the true line of sight	$\leq 0.1°$	—
$\lvert\Delta\mathbf{h}\rvert$	CV of unloading, excess momentum	0	on 40 %, off 10 %
$\lvert\omega_m\rvert$	CV of detumbling, gyro rate	0.5 °/s SP	0.1–3 °/s
τ_x, τ_y, τ_z	MVs, wheel torques	0	±3 mN m
m_x, m_y, m_z	MVs, coil dipoles	0	$\pm 0.5\,\mathrm{A\,m^2}$
τ_a	DV, drag torque at nadir	2 µN m	0–10
—	DV events, wheel failure, tracker outage, kick, handover	off	events
β_B	DV, magnetometer bias, unmeasured	0	±5 µT
q, ω, $\mathbf{h}$	states	case start	$\lvert h_i\rvert \leq h_{max}$
$\mathbf{b}$	hidden state, gyro bias	7, −9, 5 °/h	walks
ω_m, q_{st}	measured, gyro, star tracker	10 Hz, 1 Hz	blind spells
$\mathbf{B}_m$, $\hat{\mathbf{s}}_m$	measured, magnetometer, sun sensor	10 Hz	none in shadow
$\hat{q}$, $\hat{\mathbf{b}}$	filter states	warm or cold	—

HISTORY Two Wheels and the Pressure of Sunlight (2013 to 2014)

The Kepler space telescope pointed on four reaction wheels until the second of them failed in May 2013, which ended its original mission. Over the next year its teams formulated and tested a new one. Two wheels held the telescope's boresight, and the pressure of sunlight on the solar panels was balanced about the roll axis. Like a pencil balanced on a finger, the balance is unstable, so thrusters fired on a schedule of six hours to keep the roll within bounds, and the momentum that built up on the two wheels was dumped every two days. Roll at the edges of the field grew to two orders of magnitude more than in the original mission, and a rebuilt data pipeline restored much of the photometric precision [17].

7.3 The Spacecraft, Its Orbit and Its Instruments

A rigid 6U of 10 kg carries three reaction wheels on its principal axes, three magnetorquers, a camera with a 2° field of view and four instruments around a circular sun-synchronous orbit at 500 km, as shown in Figure 7.1. Its body axes are x for roll, along the track at nadir, y for pitch, against the orbit normal, and z for yaw, the camera's boresight. Ten states are integrated, the attitude quaternion q of the body against inertial space, the body rate ω in body axes and the momenta $\mathbf{h}$ of the three wheels. Because the engine counts time in minutes, the code multiplies each rate below by 60 s/min. Each signal's role is listed in Table 7.1.

7.3.1 Orbit, Sun and Field

The orbit is circular with radius $R = R_E + 500\,\mathrm{km}$, inclined $i = 97.40°$, and its node turns with the Sun, as shown in Equation (7.1).

$$\begin{aligned} \mathbf{r} &= R\,(\cos u\,\hat{\mathbf{P}} + \sin u\,\hat{\mathbf{Q}}), \qquad u = nt, \qquad n = \sqrt{\mu/R^3} \\ \hat{\mathbf{P}} &= (\cos\Omega,\ \sin\Omega,\ 0), \qquad \Omega = \Omega_0 + \dot{\Omega}t \\ \hat{\mathbf{Q}} &= (-\cos i \sin\Omega,\ \cos i \cos\Omega,\ \sin i) \end{aligned} \tag{7.1}$$

where t is the mission time, $\mu = 3.986\,004\,418 \times 10^{14}\,\mathrm{m^3/s^2}$, $R_E = 6378.137\,\mathrm{km}$, $\dot{\Omega}$ is one turn in 365.2422 days and $\Omega_0 = 157.5°$ puts the ascending node at 22:30 local time with the Sun along the inertial X axis, the March equinox. Then $n = 1.106\,783 \times 10^{-3}\,\mathrm{rad/s}$, the period is 94.62 min, the Sun stands 22.3° from the orbit plane, and a cylindrical shadow of the Earth's radius, without penumbra, covers 34.8 min of every orbit. The orbit frame has $\hat{\mathbf{z}}_O = -\hat{\mathbf{r}}$, $\hat{\mathbf{y}}_O = -\mathbf{r} \times \mathbf{v}/|\mathbf{r} \times \mathbf{v}|$ and $\hat{\mathbf{x}}_O = \hat{\mathbf{y}}_O \times \hat{\mathbf{z}}_O$, and it turns at n about the orbit normal plus the node's slow turn. The Earth turns at $\omega_E = 7.292\,115\,9 \times 10^{-5}\,\mathrm{rad/s}$ with Greenwich on the Sun's line at $t = 0$, and its field is the dipole

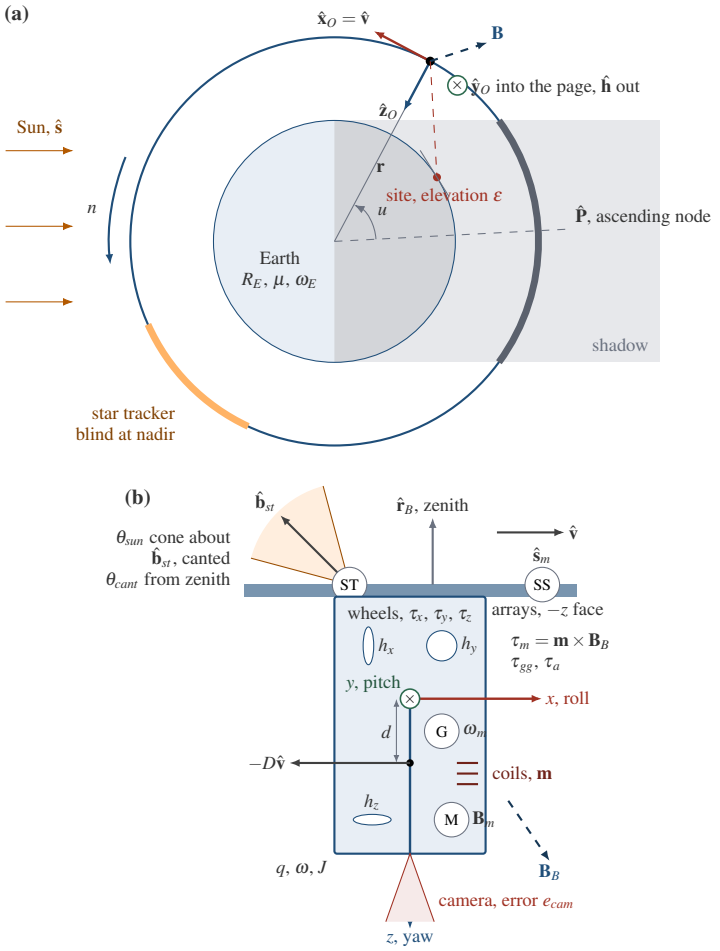

Figure 7.1 The orbit and the spacecraft as modeled, not to scale. (a) The orbit seen from its normal $\hat{\mathbf{h}}$, with the Earth half in night, its shadow away from the Sun $\hat{\mathbf{s}}$, the node $\hat{\mathbf{P}}$, the position $\mathbf{r}$ at argument of latitude u, the orbit frame ($\hat{\mathbf{x}}_O$ along the velocity, $\hat{\mathbf{y}}_O = -\hat{\mathbf{h}}$, $\hat{\mathbf{z}}_O$ to nadir), the field $\mathbf{B}$, a site at elevation ε, and the arcs in shadow (dark) and where the Sun blinds the star tracker at nadir (amber, lighter and labeled). (b) The 6U in the orbit frame, with the wheels' momenta h_i driven by τ_i, the coils $\mathbf{m}$ in the field $\mathbf{B}_B$, the zenith $\hat{\mathbf{r}}_B$, the drag $-D\hat{\mathbf{v}}$ at a center of pressure d from the center of mass along z, the sun sensor (SS) on the arrays, the star tracker (ST) canted θ_{cant} toward the wake inside its Sun cone θ_{sun}, the gyro (G), the magnetometer (M) and the camera with its scored error e_{cam}.

of the thirteenth International Geomagnetic Reference Field for 2020, as shown in Equation (7.2).

$$\mathbf{B} = \left(\frac{a}{r}\right)^3 \left[3(\mathbf{M}\cdot\hat{\mathbf{r}})\hat{\mathbf{r}} - \mathbf{M}\right], \qquad \mathbf{M} = R_Z(\omega_E t)\,(g_1^1,\ h_1^1,\ g_1^0) \tag{7.2}$$

where $g_1^0 = -29\,404.8\,\text{nT}$, $g_1^1 = -1450.9\,\text{nT}$, $h_1^1 = 4652.5\,\text{nT}$ and $a = 6371.2\,\text{km}$ [18], and R_Z turns the Earth-fixed dipole about the spin axis. The dipole is 9.41° off that axis, and along this orbit $|\mathbf{B}|$ runs from 23.69 to 47.37 μT. Two sites fixed to the Earth are the camera's targets, Site A at geocentric 44.36° S, 116.82° W and Site B at 53.15° S, 120.35° W for a handover. A site is in a pass while it sees the satellite above the imaging mask of 5°, and only passes that rise above 20° are tasked,

because lower ones are imaged from more than 60° off nadir. Site A's pass on the fourth orbit lasts 9.10 min and reaches 77.8°.

7.3.2 Body and Wheels

The attitude and the body rate obey Equation (7.3).

$$\begin{aligned} \dot{q} &= \tfrac{1}{2}\, q \otimes \begin{bmatrix} 0 \\ \omega \end{bmatrix}, \qquad \mathbf{v}_I = R(q)\, \mathbf{v}_B \\ J\dot{\omega} &= -\omega \times (J\omega + \mathbf{h}) - \dot{\mathbf{h}} + \tau_{gg} + \tau_a + \mathbf{m} \times \mathbf{B}_B \end{aligned} \tag{7.3}$$

where q is scalar first, $R(q)$ takes body vectors into inertial ones, $J = \mathrm{diag}(0.10, 0.12, 0.05)\,\mathrm{kg\,m^2}$ assigns three moments so that the camera axis has the smallest and pitch the largest, and $\mathbf{B}_B = R(q)^{\mathsf{T}}\mathbf{B}$. The wheels' torque $-\dot{\mathbf{h}}$ is internal, so the inertial total $R(q)(J\omega + \mathbf{h})$ changes only through the external torques. A working wheel delivers its command clipped to τ_{max}, and a torque that would speed it up fades over its last 10 % of speed, as shown in Equation (7.4).

$$\begin{aligned} \dot{h}_i &= \mathrm{sat}(\tau_i,\ \tau_{max}\lambda_i), \qquad \lambda_i = \min\Big(1,\ \frac{h_{max} - |h_i|}{0.1\, h_{max}}\Big) \\ \dot{h}_i &= -c_v\Omega_i - c_c \tanh\Omega_i, \qquad \Omega_i = h_i / I_w \end{aligned} \tag{7.4}$$

where the second line holds for a failed wheel, λ_i applies only when $\tau_i h_i > 0$ and is 1 otherwise, $\tau_{max} = 3\,\mathrm{mN\,m}$ and $h_{max} = 30\,\mathrm{mN\,m\,s}$, so $|h_i|$ never passes h_{max}. At capacity a wheel turns at 625 rad/s, 5968 rpm, on a rotor inertia I_w of $4.8 \times 10^{-5}\,\mathrm{kg\,m^2}$. A working wheel's drive cancels its own bearing friction, so only a failed wheel feels $c_v = 8 \times 10^{-8}\,\mathrm{N\,m\,s/rad}$ and $c_c = 1\,\mu\mathrm{N\,m}$, coasting down and handing its momentum to the body.

7.3.3 Torques From Outside the Spacecraft

Gravity, drag and the coils act on the body, as shown in Equation (7.5).

$$\begin{aligned} \tau_{gg} &= 3n^2\, \hat{\mathbf{r}}_B \times J\hat{\mathbf{r}}_B, \qquad \tau_m = \mathbf{m} \times \mathbf{B}_B \\ \tau_a &= \tau_a\,(\hat{v}_y,\ -\hat{v}_x,\ 0) \end{aligned} \tag{7.5}$$

where $\hat{\mathbf{r}}_B$ and $\hat{\mathbf{v}}$ are the zenith and the velocity direction in body axes. The drag torque is the moment of a drag force $-D\hat{\mathbf{v}}$ applied at a center of pressure displaced d along the camera axis, with $\tau_a = dD = 2\,\mu\mathrm{N\,m}$, so at nadir it is fixed in the body about $-y$, the orbit normal. With a principal axis at nadir the gravity-gradient torque τ_{gg} vanishes. The position, velocity, field, shadow and orbit frame are evaluated at the start of each 0.1 s step and held through it, while the field turns 0.013° on average.

7.3.4 What the Instruments Report

The instruments alone carry the truth into the flight software. Their noise is drawn inside them once per control step, and the gyro's bias walk, part of the hidden truth, is drawn after each integration step. The gyro and the star tracker report as shown in Equation (7.6).

$$\begin{aligned} \omega_{m,k} &= \frac{\varphi(q_{k-1}^{*} \otimes q_k)}{\Delta t} + \mathbf{b}_k + \frac{\sigma_v}{\sqrt{\Delta t}}\, \nu_k \\ \mathbf{b}_{k+1} &= \mathbf{b}_k + \sigma_u \sqrt{\Delta t}\, \nu_k' \\ q_{st} &= q \otimes \delta q\big(\sigma_{st}(\nu_1 \hat{\mathbf{e}}_1 + \nu_2 \hat{\mathbf{e}}_2) + 6\sigma_{st}\nu_3 \hat{\mathbf{b}}_{st}\big) \end{aligned} \tag{7.6}$$

where φ is the rotation vector, so the gyro reports the rotation the body made over its 0.1 s interval Δt, $\sigma_v = 0.05°/\sqrt{\mathrm{h}}$ is its angle random walk, σ_u its bias walk of 0.5 °/h per $\sqrt{\mathrm{h}}$, $\mathbf{b}$ starts at (7, −9,

5) °/h, the ν are standard normal deviates from the case's seeded generator, and $\hat{\mathbf{e}}_1$, $\hat{\mathbf{e}}_2$ span the plane across the tracker's boresight $\hat{\mathbf{b}}_{st}$. The tracker looks out of the zenith face canted $\theta_{cant} = 45^\circ$ toward the wake and reports once a second with $\sigma_{st} = 10$ arcsec across its boresight and six times that about it. It reports nothing while the Sun is within $\theta_{sun} = 30^\circ$ of its boresight outside the shadow, while its boresight is less than $\theta_{limb} = 20^\circ$ above the Earth's limb, the Earth's disc being 68.02° in radius at this height, while the body turns faster than 1 °/s, or during an outage. At nadir the Sun blinds it for 10.85 min of every orbit and the limb never does. The magnetometer reads $\mathbf{B}_B$ every step with 0.1 μT of noise per axis and a bias $(\beta_B, 0, -\beta_B/2)$ when the disturbance sets one, the coils' own field being assumed blanked, and the coarse sun sensor gives the Sun's direction to 1.5° and nothing in the shadow.

7.3.5 The Pass and Its Score

At every step of a tasked pass the score accumulates, from the truth alone, the pass time T_p, the time T_{ok} during which the true camera axis lies within 0.1° of the true line of sight to the site, and the worst such angle. The share imaged is T_{ok}/T_p, and the largest wheel momentum before the pass's first instant is kept apart from the rest of the run.

7.3.6 Epochs, Starting States and Integration

Every case flies the same orbit over the same Earth, starting at its own epoch, the mission time at its $t = 0$. The body starts in the orbit frame or at an offset from it, turning with it, plus any tumble the case adds, with the wheels at zero unless the case loads them. A warm filter starts with an attitude error of (4, −3, 6) μrad and a bias error of (0.04, −0.05, 0.03) °/h under standard deviations of 5 μrad and 0.1 °/h, and a cold one waits for the first tracker reading. A classical Runge-Kutta scheme of fourth order advances the ten states on a fixed 0.1 s step, which is also the period of the flight software, the tracker being read on every tenth step, and every 5 s of flight adds a row to the historian. Each step ends by projecting the quaternion onto unit norm and walking the gyro bias. All random numbers come from one generator seeded at 20261002, as Section 1.6.1 describes for every app. Without the projection Runge-Kutta keeps $|q|$ within 1.9×10^{-10} of one over an orbit of a 10 °/s tumble, and with every external torque off the inertial momentum holds to 6.1×10^{-11} of the 14.2 mN m s in play, 16.0 times closer at half the step. Halving the physics step, with the flight software still at 0.1 s, moves the creep after an orbit by 0.020 % and the worst pointing of case 7 under perfect states by less than 10^{-5}°. All 65 of the application's own checks pass.

7.3.7 Assumptions

A rigid spacecraft with declared inertia. Its principal axes are the body axes and the laws use the declared J, so tripling every moment (yaw's 2.5 times) slows case 5 from 102.3 s to 136.4 s. Flexible arrays, slosh and misaligned wheels would add modes and cross-coupling the model lacks.

Ideal wheels. Inside their torque limit they are torque sources with a linear back-EMF fade, and their motor lag, stiction near zero speed and jitter are absent.

A Keplerian orbit and a dipole field. The orbit ignores the zonal harmonics apart from the node's turn, drag on the orbit and eccentricity, which suits a few orbits, and the Sun is fixed for the five hours of the longest case. The field is the dipole alone, without the terms of degree 2 to 13 that the full reference model carries [18].

One constant drag torque. The drag torque is one declared constant at nadir, turning only with the velocity's direction in the body, while real drag follows the air density. Solar pressure and a residual dipole are not modeled.

Generic instruments. The noise levels stand for no product. The tracker has no latency, stray light or acquisition time, and the orbit and the time are known exactly to the flight software. The filters assume the declared tracker noise, and at 200 arcsec case 7 images 78.0 % of its pass.

Teaching thresholds. The 0.1° requirement, the 80 % wheel limit, the masks, the dead band and the supervisor's thresholds were chosen for teaching and are not a mission specification.

Table 7.2 Model Constants by Class: V for a Show Model Slider, H for a Hidden Parameter and C for a Constant of the Source

Symbol	Meaning	Value	Class
Spacecraft, wheels and coils			
J_x, J_y, J_z	principal moments	0.10, 0.12, 0.05 kg m^2	V
h_{max}, τ_{max}	wheel capacity, torque	30 mN m s, 3 mN m	V
I_w	rotor inertia	4.8×10^{-5} kg m^2	H
c_v, c_c	failed-wheel friction	8×10^{-8} N m s, 1 μN m	V
m_{max}	dipole per coil	0.5 A m^2	V
—	back-EMF fade	last 10 % of speed	C
Instruments			
σ_v	gyro angle random walk	$0.05°/\sqrt{\text{h}}$	V
σ_u	gyro bias walk	0.5 °/h per $\sqrt{\text{h}}$	V
$\mathbf{b}_0$	gyro bias at switch-on	7, −9, 5 °/h	V
σ_{st}	tracker noise across	10 arcsec	V
—	about the boresight	6 σ_{st}	H
θ_{sun}, θ_{limb}, θ_{cant}	tracker exclusions, cant	30, 20, 45°	V
—	tracker rate limit, rate	1 °/s, 1 Hz	H, C
σ_B, σ_{css}	magnetometer, sun sensor	0.1 μT, 1.5°	V
Orbit, field and mission			
$R - R_E$, i, Ω_0	altitude, inclination, node	500 km, 97.40°, 157.5°	C
g_1^0, g_1^1, h_1^1, a	IGRF-13 dipole, 2020	see Eq. (7.2)	C
τ_a	drag torque at nadir	2 μN m	DV
ε_{min}, ε_{task}	imaging mask, tasking	5, 20°	V
—	camera field, requirement	2, 0.1°	V
Δt	step, flight software	0.1 s	C

7.4 What the Inertia and the Field Permit

Whatever controller flies this spacecraft, four properties of Equations (7.2) to (7.5) bound what it can do, and each is checked here on the engine. Their constants are collected in Table 7.2.

7.4.1 A Budget That Fills Every Orbit

The drag torque at nadir is fixed in the body, so one wheel integrates all of it. In one orbit of period T it gains $\tau_a T = 2\,\mu\text{N m} \times 5677\,\text{s} = 11.35\,\text{mN m s}$, 37.85 % of h_{max}, whatever the pointing gains. Only an external torque can return it, and the coils can. A dipole of 0.5 A m^2 in a field of 23.7 to 47.4 μT produces up to 11.8 to 23.7 μN m, six to twelve times the drag, but only across the field. The unloading law of Section 7.5 asks for $|\Delta\mathbf{h}_\perp|/\tau_d$, up to 13.3 μN m at the 40 % excess that switches it on.

🕮 **Pause.** Try Exercise 7.1 before reading on.

7.4.2 The Gravity Gradient Holds the Free Spacecraft

The ordering $J_y > J_x > J_z$ makes nadir pointing stable under the gravity gradient alone. Linearized about the orbit frame, pitch separates from roll and yaw, as shown in Equation (7.7).

$$\begin{aligned} \omega_\theta^2 &= 3n^2\,\frac{J_x - J_z}{J_y} \\ 0 &= s^4 + n^2(1 + 3k_R + k_R k_Y)\,s^2 + 4n^4 k_R k_Y \\ k_R &= (J_y - J_z)/J_x, \qquad k_Y = (J_y - J_x)/J_z \end{aligned} \tag{7.7}$$

where s is the Laplace variable of the coupled roll and yaw motion. With $k_R = 0.7$ and $k_Y = 0.4$ every root is imaginary, and the spacecraft librates at $1.118n$ in pitch and at $0.610n$ and $1.734n$ in roll and yaw, with periods of 84.63, 155.0 and 54.6 min. The engine agrees. Left in manual with no drag and a 2° pitch offset, a start that no case offers, the spacecraft librates in pitch with a period of 84.65 min for four orbits while roll and yaw stay within 0.05°, and from a 2° roll offset roll and yaw

exchange through the faster coupled mode, with zero crossings every 54.63 min, yaw reaching 2.79°. The restoring torque in pitch is $3n^2(J_x - J_z) = 3.2\,\text{nN m}$ per degree, about 1/600 of the drag, so the ordering matters to a spacecraft left alone and very little to one under control.

7.4.3 The Field Against the Stored Momentum

Two frames meet in magnetic unloading. With the body held on nadir the wheels' momentum is fixed in inertial space, apart from what external torques add. In case 4 with unloading off and the drag torque at zero, the excess keeps its inertial direction within 0.03° over two orbits (0.01 to 0.05° on five other seeds), while the roll wheel's 15 mN m s moves to the yaw wheel a quarter of an orbit later as the body turns under it. The field lies close to the orbit plane, on average 1.6 to 9.1° out of it in the first four orbits, radial over the magnetic poles and along the track over the magnetic equator. It therefore makes one turn an orbit in the orbit frame and two in inertial space, and so two against the stored momentum. With $\Delta\mathbf{H} = R(q)\,\Delta\mathbf{h}$ the excess in inertial axes and $\hat{\mathbf{b}}$ the field's direction, the law's torque $-k\Delta\mathbf{h}_\perp$ and its orbit average are given by Equation (7.8).

$$\frac{\mathrm{d}\Delta\mathbf{H}}{\mathrm{d}t} = -k\,(I - \hat{\mathbf{b}}\hat{\mathbf{b}}^\mathsf{T})\,\Delta\mathbf{H} \;\approx\; -k\,\langle I - \hat{\mathbf{b}}\hat{\mathbf{b}}^\mathsf{T}\rangle\,\Delta\mathbf{H} \tag{7.8}$$

where $k = 1/\tau_d$ and the angle brackets average over the orbit, the form of analysis applied to an inertially fixed spacecraft in [5]. Over the first three orbits $\langle\hat{\mathbf{b}}\hat{\mathbf{b}}^\mathsf{T}\rangle$ in the axes $(\hat{\mathbf{P}}, \hat{\mathbf{Q}}, \hat{\mathbf{h}})$ is diag(0.493, 0.491, 0.016), with no off-diagonal term above 0.004. An excess along the orbit normal therefore decays with a time constant of 15.24 min, close to $\tau_d = 15\,\text{min}$, and an excess in the plane with 29.5 to 29.6 min, twice as long. The engine follows, in runs started from a state that no case offers. With 30 % on one wheel, the coils on with the dead band open and no drag, the pitch excess falls to a tenth in 36.2 min against a predicted 35.1, the roll excess, along $\hat{\mathbf{Q}}$ at the start, in 64.4 min against 67.9, and the yaw excess, along $-\hat{\mathbf{P}}$, in 70.3 min against 68.1, as shown in Figure 7.2(a). The excess in the plane falls in steps, quickly while the field lies across it and hardly at all while the field passes along it, four times an orbit. The drag's excess builds up along the orbit normal, the direction the law reaches most easily.

7.4.4 B-dot and the Rate Along the Field

For a spin fast against the field's turning, $\dot{\mathbf{B}}_B \approx -\boldsymbol{\omega} \times \mathbf{B}_B$, and the B-dot law of Section 7.5 gives the torque of Equation (7.9).

$$\mathbf{m} = -\frac{k_B}{|\mathbf{B}_B|^2}\,\dot{\mathbf{B}}_B \quad\Rightarrow\quad \tau_m = \mathbf{m} \times \mathbf{B}_B \approx -k_B\,\omega_\perp \tag{7.9}$$

where $\omega_\perp$ is the part of the rate across the field. The law removes that part and cannot touch the rate along the field. A strong gain removes it quickly and leaves the body turning about the field line, which then has to wait for the field to turn. Figure 7.2(b) shows the price, with each gain set before the run. Case 2's tumble falls under 0.5 °/s fastest at 0.4 and 0.5 mN m s, in 26.655 and 26.81 min, against 48.285 and 46.19 min at 0.20 and 0.25 mN m s, the settings either side of the 0.22 mN m s that the application's guide evaluates from the criterion of [7], 69.67 min at 2 and 85.1 to 86.1 min from 6 to 20. Above 0.5 mN m s the coils spend longer at their limit, 8.2 min at 0.5, 25.0 at 2 and 69.8 at 8, so a larger gain buys no more torque, only more alignment.

🕮 **Pause.** Try Exercise 7.2 before reading on.

7.5 The Flight Software as Written

The flight software works through four stages on each 0.1 s step, the instruments, the filters, the supervisor and the laws. Apart from a teaching switch, Perfect, that hands the laws the true attitude and rates, no stage of it sees the truth. The velocity-form PID of Section 1.6.2, which most apps share, has no role here. No law integrates an error, so nothing can wind up against a wheel's limit and no anti-windup is needed. Table 7.3 holds the settings, their defaults and their ranges.

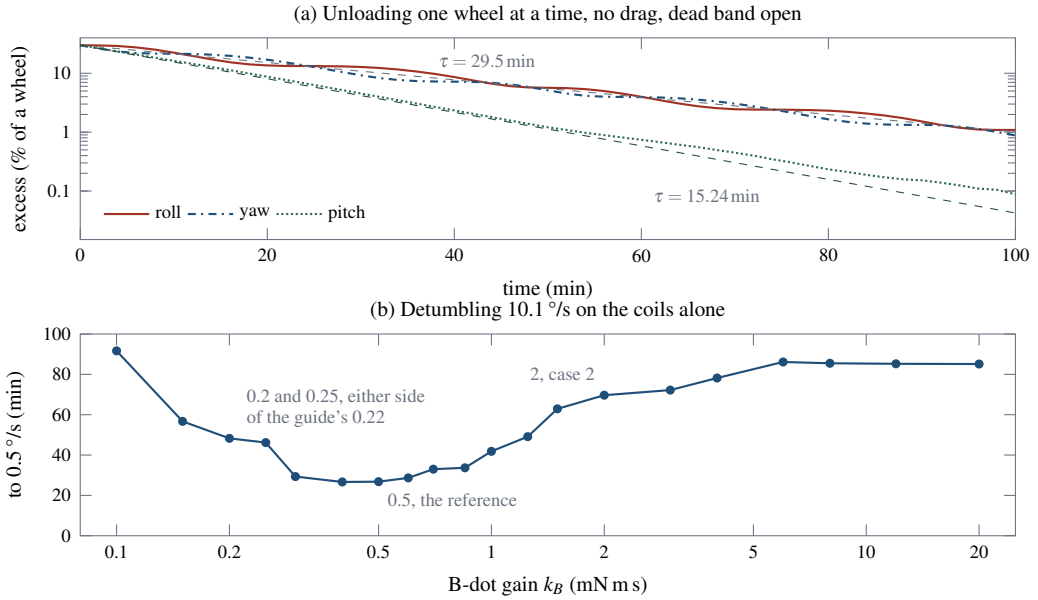

Figure 7.2 What the field allows, measured on the engine. (a) Excess momentum after 30 % is placed on one wheel, with the coils on, no dead band and no drag (roll solid, yaw dash-dotted, pitch dotted), against the orbit-averaged decays of Equation (7.8) (thin dashed). (b) Time for B-dot to bring case 2's 10.1 °/s tumble under 0.5 °/s, against the gain set before the run.

Table 7.3 Flight Software Settings, With Their Defaults and Ranges

Symbol	Setting	Default	Range
—	mode	Auto	one mode
—	attitude estimate	EKF + bias	perfect, no bias
K_p, K_d	attitude and rate gains	0.0625 s^{-2}, 0.45 s^{-1}	to 0.5, 2
ω_{lim}	slew rate limit	1 °/s	0.1–6
—	pointing law	eigenaxis	three axes, MPC
—	MPC preview of the site's path	on	off
—	roll, pitch, yaw offsets	0	±180, ±89, ±180
—	momentum unloading	on	off
τ_d	unloading time constant, $1/k$	15 min	2–120
—	dead band, on and off	40 %, 10 %	1–90, 0–80
$h_{y,SP}$	pitch wheel momentum target	0	±60 %
k_B	B-dot gain	0.5 mN m s	0.05–20
—	detumbled below, for 60 s	0.5 °/s	0.1–3
—	lead before a pass	2 min	0.5–6
$\hat{\sigma}_v$, $\hat{\sigma}_u$	filter's gyro noise, bias walk	as the gyro	to 1, 10
$\hat{\sigma}_{b0}$	filter's bias allowance at start	20 °/h	0.1–100

7.5.1 Two Filters on One Gyro

Two multiplicative filters work in parallel on the same readings, one that estimates the gyro bias as a state, as Section 1.4.6 describes, and one that does not, and the setting selects which one the laws use.

The filter with the bias propagates and updates as shown in Equation (7.10).

$$
\begin{aligned}
&\hat{q} \leftarrow \hat{q} \otimes \delta q\big((\omega_m - \hat{\mathbf{b}})\Delta t\big), \qquad P \leftarrow \Phi P \Phi^{\mathsf{T}} + Q \\
&\Phi = \begin{bmatrix} I - [\hat{\omega}\times]\Delta t + \frac{1}{2}[\hat{\omega}\times]^2\Delta t^2 & -I\Delta t \\ 0 & I \end{bmatrix} \\
&Q = \begin{bmatrix} (\sigma_v^2\Delta t + \sigma_u^2\Delta t^3/3)\, I & -\frac{1}{2}\sigma_u^2\Delta t^2\, I \\ -\frac{1}{2}\sigma_u^2\Delta t^2\, I & \sigma_u^2\Delta t\, I \end{bmatrix} \\
&\mathbf{z} = 2\,\mathrm{sgn}(w)\,\mathrm{vec}(\hat{q}^* \otimes q_{st}) \\
&K = PH^{\mathsf{T}}(HPH^{\mathsf{T}} + R_{st})^{-1}, \qquad \begin{bmatrix} \delta\theta \\ \delta\mathbf{b} \end{bmatrix} = K\mathbf{z} \\
&\hat{q} \leftarrow \hat{q} \otimes \delta q(\delta\theta), \qquad \hat{\mathbf{b}} \leftarrow \hat{\mathbf{b}} + \delta\mathbf{b}
\end{aligned}
\tag{7.10}
$$

where $\hat{\omega} = \omega_m - \hat{\mathbf{b}}$, w is the scalar part of $\hat{q}^* \otimes q_{st}$, $H = [I\ 0]$, the covariance is updated in the Joseph form $P \leftarrow (I - KH)P(I - KH)^{\mathsf{T}} + KR_{st}K^{\mathsf{T}}$, and $R_{st} = \sigma_{st}^2(I - \hat{\mathbf{b}}_{st}\hat{\mathbf{b}}_{st}^{\mathsf{T}}) + (6\sigma_{st})^2\hat{\mathbf{b}}_{st}\hat{\mathbf{b}}_{st}^{\mathsf{T}}$ matches the declared tracker. The filter assumes the gyro's own σ_v and σ_u. A cold filter starts at the first tracker reading with $\hat{q} = q_{st}$, $\hat{\mathbf{b}} = 0$, an attitude variance of σ_{st}^2 and a bias allowance of 20 °/h. The other filter drops the bias rows and treats the gyro as unbiased, so whenever the tracker is blind it integrates the bias into the attitude. The laws use $\hat{q}$ and the rate $\omega_m - \hat{\mathbf{b}}$. A bias state serves the same purpose in the vessel observer of Chapter 6, where it absorbs the wind and the current.

7.5.2 Supervisor and Modes

In Auto the supervisor moves through four modes. Detumble runs B-dot with the wheels idle until the gyro has read under 0.5 °/s for 60 s. Safe turns the arrays onto the measured Sun until the filter is running with its own standard deviation under 0.05° and the Sun is within 10° of the arrays for 30 s. Nadir flies the orbit frame turned by the three offsets, and target tracking points the camera at the tasked site from 2 min before a tasked pass rises above the mask until 30 s after it sets. The target reference puts z on the line of sight with x as near the velocity as the line of sight allows, and its rate and acceleration come from central differences of that frame with a step of 1 s. Without a running filter the pointing modes fall back to safe.

7.5.3 Pointing About the Eigenaxis

The pointing law closes the error of the estimate against the reference through a rate command and an inner rate loop, as shown in Equation (7.11).

$$
\begin{aligned}
&q_e = q_r^* \otimes \hat{q}, \qquad \phi = \varphi(q_e) \\
&\omega_c = -\mathrm{sat}_{\|\cdot\|}\Big(\frac{K_p}{K_d}\phi,\ \omega_{lim}\Big) \\
&\mathbf{u} = J\big[K_d(\omega_r + \omega_c - \hat{\omega}) + \alpha_r\big] + \hat{\omega} \times (J\hat{\omega} + \mathbf{h}) \\
&\tau_i = \mathrm{sat}\big((\mathbf{m} \times \mathbf{B}_m - \mathbf{u})_i,\ \tau_{max}\big)
\end{aligned}
\tag{7.11}
$$

where q_r is the reference, ϕ the rotation vector of the error taken the short way, $\mathrm{sat}_{\|\cdot\|}$ limits the length of the vector and not its parts, as Section 1.4.2 explains, and ω_r and α_r are the reference's rate and acceleration in body axes. Unsaturated, $\mathbf{u}$ is the quaternion feedback $-JK_p\phi - JK_d(\hat{\omega} - \omega_r) + J\alpha_r$ with the gyroscopic term cancelled, and the defaults give $\omega_n = 0.25\,\mathrm{rad/s}$ with $\zeta = 0.9$. Limiting the length keeps a slew on the eigenaxis. The three-axis alternative limits each 3-2-1 angle of q_e in its own loop, as three single-axis controllers would. The wheels deliver $\mathbf{u}$ and also take up the torque the coils are believed to apply, computed from the measured field. A positive error about any axis is driven back by a negative torque about it, as the application's direction checks confirm. In the

convention of Section 1.3 the pointing loop is therefore reverse acting on the body torque $\mathbf{u}$ and direct acting on the wheel torques, whose positive sense, as the application labels it, spins the body the other way. Safe mode uses the same gains on the sun sensor and the raw gyro, with ω_c turning $-z$ onto the measured Sun and $\omega_c = 0$ in the shadow, a rate hold that drifts with the bias.

7.5.4 The Coils

The coils serve three laws, as shown in Equation (7.12).

$$\begin{aligned}
\dot{\mathbf{B}}_{f,k} &= \gamma \dot{\mathbf{B}}_{f,k-1} + (1-\gamma)\,\frac{\mathbf{B}_{m,k} - \mathbf{B}_{m,k-1}}{\Delta t} \\
\mathbf{m}_B &= -\frac{k_B}{|\mathbf{B}_m|^2}\,\dot{\mathbf{B}}_f \\
\mathbf{m}_U &= \frac{k}{|\mathbf{B}_m|^2}\,(\Delta\mathbf{h} \times \mathbf{B}_m) \\
\mathbf{m}_F &= u_j\,\frac{\mathbf{B}_m \times \hat{\mathbf{e}}_j}{|\mathbf{B}_m \times \hat{\mathbf{e}}_j|^2}
\end{aligned} \tag{7.12}$$

where $\gamma = \mathrm{e}^{-\Delta t/\tau_f}$ with $\tau_f = 1\,\mathrm{s}$ filters the field's rate for B-dot ($\mathbf{m}_B$), $\Delta\mathbf{h} = \mathbf{h} - (0, h_{y,SP}, 0)$ from the wheels' tachometers is the excess the unloading law ($\mathbf{m}_U$) acts on, with its torque $-k\Delta\mathbf{h}_\perp$, and $\mathbf{m}_F$ is the least dipole that gives the torque u_j about the axis $\hat{\mathbf{e}}_j$ of a failed wheel, the directional actuator of Section 1.4.7. Unloading switches on above an excess of 40 % of one wheel and off below 10 %, and it is suspended while the coils stand in for a wheel. Every dipole is scaled along its own direction so that no coil passes $0.5\,\mathrm{A\,m^2}$. The order of the cross product matters. Written as $\mathbf{B} \times \Delta\mathbf{h}$, as the application's first design had it, the law with $\tau_m = \mathbf{m} \times \mathbf{B}$ gives $+k\Delta\mathbf{h}_\perp$ and fills the wheels. The code uses the order that empties them, which the application's checks confirm by showing that the torque opposes the excess and never has a part along $\mathbf{B}$.

7.5.5 A Predictive Pointing Law

The pointing law's third value, MPC, replaces $\mathbf{u}$ of Equation (7.11) in nadir and target tracking by the first move of a plan re-solved at every decision, the receding-horizon form of Section 1.6.3, whose stability and optimality are reviewed in [19], applied to spacecraft attitude in explicit form in [20] and surveyed for aerospace systems in [21]. The wheels, or two wheels and $\mathbf{m}_F$ after a failure, deliver its $\mathbf{u}$ through the last line of Equation (7.11), and detumble, safe mode, unloading and manual mode are unchanged. Its model is Equation (7.3) at the declared inertia $J_0 = \mathrm{diag}(0.10, 0.12, 0.05)\,\mathrm{kg\,m^2}$, frozen when the application loads, written about the supervisor's reference at each decision as shown in Equation (7.13).

$$\begin{aligned}
\dot{\phi} &= \delta\omega, \qquad \delta\omega = \hat{\omega} - \omega_r, \qquad \dot{\mathbf{h}} = \tau_m - \mathbf{u} \\
\delta\dot{\omega} &= J_0^{-1}(\mathbf{u} - \mathbf{g}) + \hat{\mathbf{d}} - \alpha_r, \qquad \mathbf{g} = \hat{\omega} \times (J_0\hat{\omega} + \mathbf{h})
\end{aligned} \tag{7.13}$$

where ϕ, ω_r and α_r are those of Equation (7.11), $\mathbf{h}$ comes from the tachometers, $\tau_m = \mathbf{m} \times \mathbf{B}_m$ and the gyroscopic torque $\mathbf{g}$ are held over the horizon, and $\hat{\mathbf{d}}$ is an estimated disturbance acceleration. The term $\phi \times \delta\omega$ is dropped, which is exact along a fixed eigenaxis. With wheel j failed, u_j comes from the coils, the other wheels take $\dot{h}_i = -u_i - a_i u_j$ with $a_i = B_{m,i}B_{m,j}/|\mathbf{B}_m \times \hat{\mathbf{e}}_j|^2$, and h_j leaves the model. It comes from the physics and not from a fit, since stepping the torque could reveal only $1/J_0$, 0.573, 0.477 and $1.146\,^\circ/\mathrm{s^2}$ per mN m about roll, pitch and yaw. Two checks confirm it. With no bound or band active the first move equals the discrete LQR's on the same weights to 10^{-6}, and a fresh controller handed only what each decision received reproduces 599 of 599 moves bit for bit on a spacecraft with $J_x = 0.2\,\mathrm{kg\,m^2}$, drag and magnetometer bias.

Estimator. On each axis a two-state Kalman filter with a constant gain, solved once from the Riccati equation, runs on $[\omega_i, d_i]$ at the 0.1 s step, driven by the torque applied and measuring $\hat{\omega}_i$, with

process variances $(2\times10^{-6}\,\mathrm{rad/s})^2$ and $(2\times10^{-8}\,\mathrm{rad/s^2})^2$ a step and the gyro's $\sigma_v^2/\Delta t$. Integrating disturbances of this kind make a predictive controller offset-free [22], and after 10 min of case 3 the pitch estimate stands at −1.986 µN m against the 2 µN m of drag. A jump of more than 1 °/s in one step while every wheel is under $0.9h_{max}$ is read as a kick, which resets the rate and leaves $\hat{\mathbf{d}}$ alone. The filter runs under every law, so a switch starts converged. Because the gyro averages over its last step, the plan's starting rate is advanced half a step with the acceleration the applied torque and $\hat{\mathbf{d}}$ give, 0.17 °/s on yaw at 3 mN m.

The plan. Every 0.5 s on an integer step clock, and at once whenever the relative rate leaves the plan's own prediction by more than 0.02 °/s, the controller solves Equation (7.14) over $N = 120$ steps of $T_p = 0.5$ s, 60 s in all, with the input held in ten move blocks (lengths 1, 1, 2, 3, 5, 8, 12, 18, 30 and 40), and applies the first.

$$\begin{aligned}\min\ &\sum_{k=1}^{N}\left(Q_\phi|\phi_k|^2+Q_\omega|\delta\omega_k|^2\right)+\|\phi_N\|_{P_\phi}^2+\|\delta\omega_N\|_{P_\omega}^2\\&+S\sum_b|\Delta\mathbf{u}_b|^2+R\sum_{k=0}^{N-1}|\mathbf{u}_k-\mathbf{u}_0^h|^2\\&+\sum_\ell\left(\rho_\ell s_\ell^2+\sigma_\ell s_\ell\right)\\\text{s.t. }&|\tau_{m,i}-u_{k,i}|\le\tau_{max},\quad |h_{k,i}|\le0.9\,h_{max}+s_{h,i}\\&|\delta\omega_{k,i}|\le\omega_{lim}\max(|\hat{e}_i|,0.1)+s_{k,i}\end{aligned}\tag{7.14}$$

in degrees, °/s, mN m and mN m s, where $Q_\phi = 3$, $Q_\omega = 0.1$ and $S = R = 0.01$, $\|\mathbf{x}\|_P^2 = \mathbf{x}^\mathsf{T}P\mathbf{x}$, P_ϕ and P_ω are the diagonal of the discrete LQR's Riccati solution for each axis, $\Delta\mathbf{u}_b$ is the change at each block boundary from the move in force, $\mathbf{u}_0^h = \mathbf{g} - J_0\hat{\mathbf{d}}$ is the torque that holds the reference, and $\hat{\mathbf{e}}$ is the present error's eigenaxis. The torque bound is hard, as the allocation clamps it, and with a wheel failed the coils' authority about its axis on the measured field bounds u_j. The rate box, whose corner lies at ω_{lim} along $\hat{\mathbf{e}}$ so that no direction exceeds $1.01\omega_{lim}$, is soft with a slack at every judged step and $\rho = \sigma = 10^6$. The wheels' ceiling at the start of the motor's fade is soft with one slack, $\rho = 10^5$ and $\sigma = 10^4$, and is not judged on a wheel where the momentum needed to hold the reference alone comes within 1 % of capacity of it, because only the coils can remove that momentum. Bands are judged at block ends, where a held torque puts their extremes, and no move cap is needed, because the wheels have no torque-rate limit. While every wheel works the problem separates exactly into three 3-state, 10-variable quadratic programs given the coupled problem's ridge of 0.028, whose moves equal the coupled solution's bit for bit on 1,579 recorded decisions, and a failed wheel couples the axes through the coils, so the coupled program is solved. A plan is accepted only with an optimal status, finite moves, predictions, cost and slacks, every planned torque inside its box and the right shapes. Otherwise the eigenaxis law, computed at every step, takes the torque at once, the historian is marked, and control returns to the plan, blended over 1 s, with the next accepted solution. The preview, which may be switched off, is the tasked site's look-frame acceleration over the horizon from the orbit the flight software knows.

Tuning and cost. The weights were chosen by a rule fixed before any grid: the most of case 10's pass within 0.1°, with the true relative rate under $1.02\omega_{lim}$ on cases 5, 7 and 10, no fallback, case 7 whole and case 5 settled no later than the eigenaxis law. Its first winner, $Q_\omega = 0.3$, imaged 5.0 % of the pass with J_x at its slider's low end of 0.02 kg m^2, because a torque planned for J_0 acts up to six times harder than planned for the 0.5 s it is held. The re-plan on a broken prediction was then added and the rule amended, openly, to require 90 % imaged and under 1 s above $1.05\omega_{lim}$ with each moment at 0.02 kg m^2 and with all at 0.3 kg m^2. Of 36 sets 33 meet it, and the winner above images 94.26 % against the first winner's 94.24 %. A decision costs about 0.6 ms in Chrome's own frame loop on a loaded host, up to 1.3 ms at low speed. The speeds run from 0.25 to 60 simulated minutes per second, so the default of 1 min/s, 60 times real time, asks for 120 decisions a wall second, which

Table 7.4 The Ten Cases and the Graded Challenge

#	Change from the reference line	Teaching point
1	manual, 1 mN m on the roll wheel from 0.5 min	a wheel moves momentum
2	detumble from 5.8, −5.8, 5.8 °/s, $k_B = 2$ mN m s	a strong gain is slow
3	nadir, unloading off, three orbits	the wheels fill
4	nadir, wheels at 50, −60, 0 %	across a turning field
5	slew at 0.5 min to roll 70, pitch 60, yaw 90°	eigenaxis or three axes
6	nadir, no-bias filter, filters cold, 14 min before the Sun's window	coasting on the gyro
7	Site A's pass, 8 min after the start	a half-blind tracker
8	as 7, roll wheel fails at 12.5 min	honest infeasibility
9	safe mode from 30, −40, 20° offsets, through the shadow	power on two sensors
10	as 7, handover to Site B at 12 min, MPC	a slew planned in its limits
C	three orbits then the pass, unloading off, no-bias filter	graded challenge

took 7.5 to 15 % of a core in two measurements, past the tenth of a core allowed to every app at 1×, where real time costs 0.15 %. A frame guard stops stepping after 25 ms, so a request for 20 or 60 min/s runs at 11 to 13 min/s and 32 to 37 frames a second, against 15.2 and 16.2 min/s at 15.3 and 5.5 frames without it. With a wheel failed the coupled program costs 2.0 to 2.2 ms, 27 % of a core at 1 min/s, and runs at 3.4 min/s at the top speeds.

7.5.6 Hand Control, and What Each Model Knows

Under hand control the operator sets the three wheel torques and the three dipoles directly, and in Auto the manual settings follow the commands after every control step, rounded to the 0.01 mN m and 0.01 A m^2 that the sliders display, so a switch to manual starts from the last commands. Three models are kept apart. Equations (7.1) to (7.6) simulate the spacecraft. The flight software carries a lighter one, the rigid body with the declared J, an exact orbit, Earth rotation and site geometry, the filters' gyro model and the field it measures, with nothing of the drag or of the gravity gradient. Its predictive law reduces that model to the linear error dynamics of Equation (7.13), which of the run's parameters reads only the actuator limits that the allocation also reads. Equations (7.7) to (7.9) serve the analysis, to interpret the runs, and nothing is identified from data.

7.6 Ten Cases and a Five-Hour Challenge, Measured

Each case takes the defaults of Table 7.3 as its starting line, Auto with the bias filter, the eigenaxis law, unloading on and 2 µN m of drag, and alters only the items its own row of Table 7.4 lists. All of the results in this section come from the engine alone, stepped as the browser steps it, at the default seed. Each shipped scenario gives identical logs whichever way it is reached, and Reset and reload behave as Section 1.6.4 describes. Apart from case 1's kick and switch to Auto, each variant below makes its change before the run, which gives the same run as the change carried through a restart. Where the instruments' noise matters, the text gives the range over five other seeds. The results are collected in Table 7.5.

7.6.1 A Wheel Moves Momentum

Case 1 starts at nadir in manual with the wheels still, and at 0.5 min a script puts 1 mN m on the roll wheel and takes it away 10 s later. The pulse leaves the wheel with 10.0 mN m s, 33.3 % of its capacity, and the body rolling the other way at 5.73 °/s, still 5.74 °/s at 10 min. The total angular momentum starts at 0.13 mN m s, the spacecraft's share of the orbit's own rotation, $J_y n$, reaches 0.19 by the pulse as the drag pushes on a body nothing holds, and is 0.22 after it and 0.19 at 10 min. In a restarted run a deployment kick at 1 min is external and lifts the total to 3.91 mN m s. Switched to Auto at 3 min, the pointing law has the camera within 0.1° of nadir at 5.995 min, and the wheels then hold 3.90 mN m s between them at 12 min, the kick's momentum stored and not removed. Nothing in manual depends on the instruments' noise, so every seed gives these values.

Table 7.5 Each Case's Measured Result, Default Seed

Case and configuration	Result
1, as shipped	pulse 10.0 s, wheel 10.0 mN m s, body −5.73 °/s, total 0.13, 0.19, 0.22 mN m s
1, kick at 1, Auto at 3 min	total 3.91 mN m s, within 0.1° at 5.995 min, wheels hold 3.90 mN m s
2, $k_B = 2$	1 °/s at 34.5 min, 0.5 °/s at 69.7 min, safe at 70.8, nadir at 80.6 min
2, $k_B = 0.5$ and 8	0.5 °/s at 26.8 and 85.5 min
3	37.84 % and 75.69 % after one and two orbits, 98 % at 245.0 min, 1° off at 250.8 min
4	78.1 % to 10 % by 58.6 min, along **B** up to 84 % of it, on again at 139.1 min
4, unloading off	97.75 % at 60 min
5, eigenaxis	90.62° flown for 90.45°, 102.3 s, 0.94 °/s, 0.85 mN m
5, three axes	105.28° flown, 86.4 s, 1.69 °/s, axis wanders 25.3°
5, eigenaxis at 6 °/s	5.10 mN m asked, 1.7 s at the limit, 30.9 s
6, no bias	blind 11.55 min, worst 2.35°, 0.046° a minute after
6, bias	blind 10.84 min, worst 0.054°, bias within 1 °/h at 0.60 min
7, bias, no bias, perfect	100 %, 51.15 %, 100 % within 0.1°, worst 0.026, 0.905, 0.015°
8	81.97 %, field within 3.46° of roll at 16.82 min, worst 5.82°
8, four other settings	81.56 to 84.75 %
9	within 5° at 1.02 min, 8.65° at sunrise, 6.28° with no drag
10, MPC, eigenaxis	on Site B in 39.6 and 52.1 s, 94.26 % and 92.44 % imaged

7.6.2 A Tumble Against the Field

Case 2 releases the spacecraft at 10.1 °/s with B-dot at 2 mN m s. The rate halves in 5.7 min and is 2.51 °/s at 12 min, with the coils on their limit for 25.0 min, but the last degree per second is slow, 1 °/s at 34.5 min and 0.5 °/s only at 69.7 min. Between 30 and 60 min the cosine of the angle between the rate and the field averages 0.984, so the spacecraft is turning almost exactly about the field line, where Equation (7.9) gives it no torque. The supervisor hands over to safe mode at 70.8 min, between 70.9 and 71.2 min on five other seeds, and to nadir at 80.6 min, once the filter, started cold at its first tracker reading, is under 0.05° with the arrays on the Sun.

7.6.3 The Wheels Fill

Case 3 holds nadir for three orbits with unloading off. The camera stays within 0.018° of nadir on average, 0.076° at worst, while the pitch wheel gains 37.84 % after one orbit and 75.69 % after two, within 0.01 % of $\tau_a T / h_{max}$ and its double. At 245.0 min, 2.59 orbits in, the wheel passes 98 % of its capacity and the fading motor can absorb no more. The drag then turns the body, the camera is 1° off nadir at 250.8 min, and the body then turns freely, 136.5° off at the end of the third orbit and anywhere from 17.9 to 143.5° on five other seeds. No pointing gain changes this, because the momentum has to leave the spacecraft.

> **TRY IT IN THE APP** **3 · Nadir Hold: The Wheels Creep**
>
> Load the case, choose 20 min/s and press Run. Open the Momentum view and watch the pitch wheel's bar grow while the camera stays on nadir. The chapter reports 37.84 % of capacity after one orbit and 75.69 % after two. Past 98 %, at 245.0 min, the wheel can absorb no more and the camera leaves nadir.

7.6.4 Unloading Across a Turning Field

Case 4 starts with the roll wheel at 50 % and the pitch wheel at −60 %, an excess of 78.1 % of one wheel with 32.5 % of a wheel along the field. The coils sit at their limit for the first 7.4 min, and the law brings the excess under 10 % in 58.6 min. While the law runs, the share of the excess that lies along the field swings from a peak to almost nothing in 8 to 19 min, its peaks come 22 to 26 min apart, and it reaches 84 %. The part along the field is 27.8 % of a wheel at 22 min, as shown in Figure 7.3, and the excess falls fastest when the field has turned across it. After the first descent the drag refills the pitch wheel by about 38 % an orbit and the dead band cycles, on again at 139.1 min and off at 172.9 min, with the excess never above the 40.0 % that switches the law on. With unloading off the same start reaches 97.75 % at 60 min. The seed moves none of these times by more than 0.1 min.

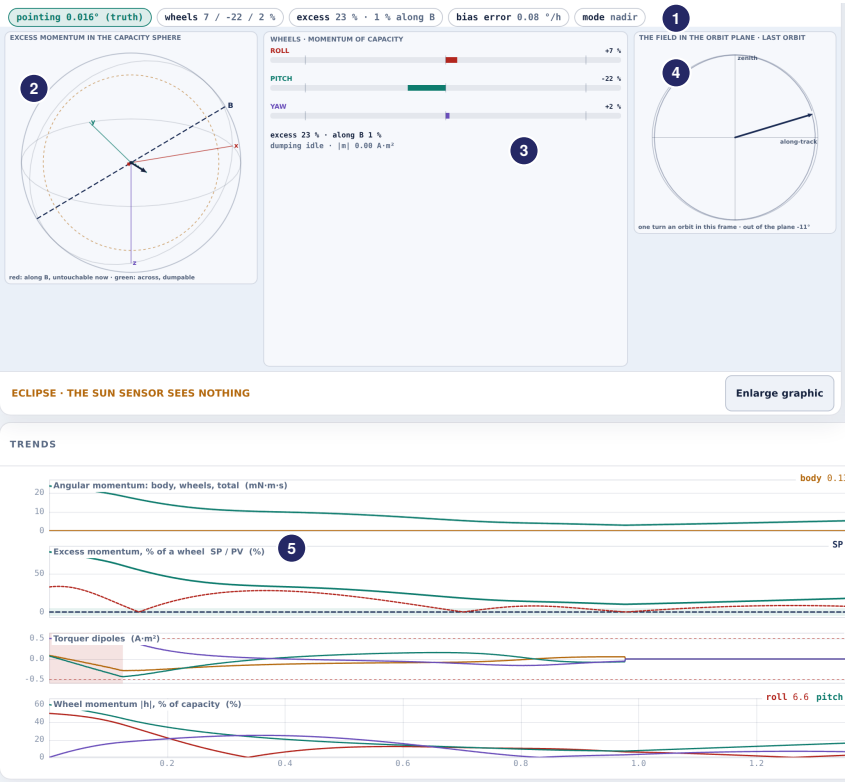

Figure 7.3 The spacecraft in case 4, Momentum dumping, at $t = 94.62$ min, one orbit (the clock reads 1.58 h), with no control changed, in the Momentum view with the 1.58 h trend window, cropped to the plant view and the trends: (1) the badges, (2) the capacity sphere, (3) the wheels' bars, (4) the field panel and (5) the excess strip. The badges read the camera 0.016° off its reference (truth), the wheels at 7, −22 and 2 %, an excess of 23 % with 1 % along **B**, a bias error of 0.08 °/h and nadir mode, and the status line reads Eclipse, the sun sensor sees nothing. The capacity sphere holds the small excess, almost all across the dashed field line, with the law idle in its dead band, and the field panel traces one turn of the field in the orbit plane over the last orbit, the field now −11° out of the plane. The strips, cut short of the orbit's end, show the excess falling from 78.1 % to 10 % by 58.6 min, its part along **B** (dashed) peaking at 27.8 % of a wheel at 22 min, the dipoles at their limit (shaded) for 7.4 min and idle after the law switches off, and the roll wheel's momentum passing to the yaw wheel, 0.01 % and 25.0 % near 21 min, as the body turns under the stored momentum.

7.6.5 One Axis or Three

Case 5 commands roll 70°, pitch 60° and yaw 90° from the orbit frame at 0.5 min, which compose a single rotation of 90.45° about one axis. The eigenaxis law turns the body through 90.62°, 90.48° with perfect states, at up to 0.94 °/s with its error axis within 0.3° of where it started and a peak wheel torque of 0.85 mN m, and the camera settles within 0.1° of its reference 102.3 s after the command. The three-axis law settles sooner, at 86.4 s, but turns the body through 105.28°, 16 % more, at up to 1.69 °/s and with an axis that wanders 25.3°, because three loops each held to 1 °/s add up to nearly $\sqrt{3}$ times the limit in a direction that is not the eigenaxis. At a limit of 6 °/s the eigenaxis slew asks one wheel for 5.10 mN m against its 3, spends 1.7 s at the torque limit and settles in 30.9 s with its axis bent by only 1.0°. On five other seeds the settling times move by at most 0.5 s.

🕮 **Pause.** Try Exercise 7.3 before reading on.

7.6.6 Coasting Through the Blind Window

Case 6 holds nadir with both filters starting cold, 14 min before the Sun enters the tracker's cone. The bias filter, started at the first tracker reading 1 s in, has the bias within 1 °/h at 0.60 min and within 0.3 °/h at 1.37 min, and through a window of 10.84 min it holds the camera within 0.054°, its bias estimate 0.18 °/h off as the window opens and 0.35 °/h as it closes. The filter that leaves the bias out integrates the gyro as it reads, about 11.4 °/h across the camera axis, and the pointing loop holds that drifting estimate on nadir. The camera reaches 2.35° off, more than its whole field, and the window lasts 11.55 min because the drifting attitude keeps the Sun in the cone. A minute after the tracker returns the camera is back within 0.046°. Perfect states hold 0.015°. The noise moves these results more than any other case's. On five other seeds the worst angle in the window is 2.32 to 2.46° without the bias state and 0.036 to 0.102° with it, and the bias estimate comes within 0.3 °/h anywhere from 0.65 to 4.19 min after the start.

> **TRY IT IN THE APP** **6 · The Gyro Bias and the Blind Window**
>
> Load the case and press Run. The shipped filter leaves the gyro bias out: watch the pointing badge while the Sun blinds the star tracker, and the camera drifts to 2.35° off nadir. Then press Reset, set the Attitude estimate to EKF + bias on the MODE card before the window opens at about 14 min, and run again. The worst angle in the window falls to 0.054°.

7.6.7 A Pass With the Tracker Half Blind

Case 7 enters target tracking 2 min before Site A rises and reaches the moving reference 80.7 s later. Through the 9.10 min pass the body turns at up to 0.86 °/s, the camera swings to 67.5° off nadir at the ends, and the wheels use 4.9 % of their capacity. The tracker is blind for 46.9 % of the pass, 34.9 % by the limb while the camera looks far along the track and 12.0 % by the Sun near the top. With the bias filter the camera holds the site within 0.1° for the whole pass, 0.026° at worst, and without the bias state for 51.15 %, 0.905° at worst. On five other seeds the bias filter again holds the whole pass, 0.023 to 0.043° at worst, and the filter without it holds 50.40 to 51.63 %.

7.6.8 A Wheel Lost Mid-Pass

Case 8 fails the roll wheel 4.5 min into the same pass, holding −0.1 % of its capacity, so its coast hands the body almost nothing. The supervisor gives roll to the coils, whose authority about roll scales with the sine of the field's angle from the roll axis. The camera passes 0.1° at 15.46 min, the field comes within 3.46° of the roll axis at 16.82 min, and the camera is 5.82° off at worst, with 81.97 % of the pass imaged. A stiffer loop ($K_p = 0.25$, $K_d = 0.9$), a softer one (0.015, 0.22), the three-axis law and perfect states all land between 81.56 and 84.75 %, and on five other seeds all five settings stay between 81.5 and 85.0 %, which is the honest infeasible case. Continuous three-axis pointing needs a torque about every axis at every instant, and with a wheel gone that depends on the field. Back at nadir two wheels and the coils hold the camera 0.013, 0.016 and 0.015° off at 30, 60 and 100 min, but with the coils busy on roll nothing unloads the others, and the pitch wheel holds 33.8 % of its capacity at 110 min, an orbit after the failure.

7.6.9 Safe Mode in the Shadow

Case 9 starts in safe mode with the arrays 63.66° off the Sun. They are within 5° of it 1.02 min later and, in sunlight, within 0.58°. The shadow lasts 34.77 min, the Sun rising at 71.2 min, and holding the gyro's reading at zero holds the true rate at minus the bias, while the drag pushes through the rate loop, so the arrays are 8.65° off the Sun at sunrise and back within 5° in 6 s. With no drag they are 6.28° off, and on five other seeds the two angles are 8.4 to 9.1° and 6.0 to 6.6°. Unloading never switches on in this case, because the excess stays under 9.7 %.

Table 7.6 The Challenge in Five Configurations at the Default Seed

Configuration	Imaged (%)	Worst (°)	Fullest (%)	Excess (%)
as shipped	24.12	27.50	100.0	100.0
unloading only	49.81	1.065	40.0	33.0
bias filter only	20.64	70.25	100.0	100.0
suggested, both	100.00	0.022	40.0	33.0
both, perfect states	100.00	0.015	40.0	33.0

Imaged within 0.1°, fullest wheel before the pass and excess at the end. Seeds 1 to 5: suggested 100 % (worst 0.021 to 0.041°), shipped 24.70 to 27.76 %, unloading only 50.06 to 53.03 %.

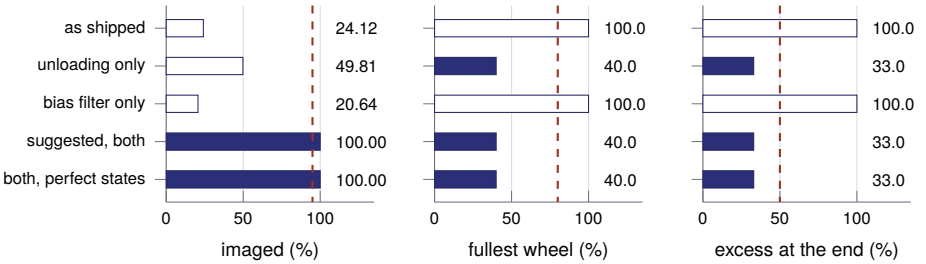

Figure 7.4 Three columns of Table 7.6 drawn as bars for the five configurations of the challenge, with its targets as dashed lines: at least 95 % of the pass imaged, no wheel past 80 % before the pass and an excess under 50 % at the end. Solid bars meet the target of their panel and white bars miss it. Notice that unloading alone passes the two momentum targets and still images only half the pass, and that the bias filter alone images less than the shipped configuration.

7.6.10 Three Orbits, Then the Pass

The challenge holds nadir for three orbits and then images Site A's pass, with unloading off and the no-bias filter as shipped. Its target asks for the true camera within 0.1° of the site for at least 95 % of the pass, no wheel past 80 % in the three orbits before it, and an excess under 50 % of a wheel when the run ends 10 min after the pass, at 302.95 min. It counts only on the case as given, so a changed model parameter or disturbance, or an extra disturbance fired, voids the attempt. At 20 simulated minutes per second the run takes about 15 s. Table 7.6 compares the configurations, and Figure 7.4 draws three of its columns as bars. As shipped the pitch wheel saturates at 245.0 min, 38.9 min before the pass, and the camera holds the site for 24.12 % of it. Each change alone fails. Unloading keeps the fullest wheel at 40.0 %, but the no-bias filter drifts through the blind half of the pass and images 49.81 %, and the bias state alone images 20.64 % because the attitude was lost before the site rose. With both the camera holds the site for the whole pass, 0.022° at worst, with the excess at 33.0 % at the end, which passes the three conditions by margins of 5, 40 and 17 percentage points, and on five other seeds the suggested configuration images 100 % each time while the shipped one images 24.70 to 27.76 %.

> **+ DID YOU KNOW? A Better Filter, a Worse Worst Moment**
>
> From the simulator: in the challenge, the gyro bias state alone leaves the camera 70.25° from the site at worst, against 27.50° with the filter as shipped. The pitch wheel still fills before the pass, so the better estimate arrives with the attitude already lost (Table 7.6).

7.6.11 A Handover Planned Inside the Limits

Case 10 flies case 7's pass and hands the camera over to Site B at 12 min, 4 min into the pass, a turn of 38.6° onto a second moving site, with the MPC pinned. Table 7.7 sets it beside the shipped laws, the eigenaxis law retuned by the MPC's own rule, once, across all of its knobs, that law fed what the

Table 7.7 The Predictive Law Against the Pointing Laws, Default Seed

Case and measure	Eigenaxis	Three axes	Retuned	Retuned, MPC signals	MPC	K_p 4, K_d 3.4
10, handover: on Site B (s), imaged (%)	52.1, 92.4	51.9, 92.5	41.9, 93.9	41.7, 94.0	**39.6, 94.3**	39.9, 94.2
10: past τ_{max} (s), torque travel (mN m)	0, 84	0, 87	0.4, 230	0.3, 236	**0**, 173	1.9, 699
10: peak relative rate (°/s)	0.999	1.335	1.002	1.002	1.015	—
10 at $J_y = 0.02$, all at 0.3: imaged (%)	80.1, 61.1	81.2, 60.9	93.6, 92.9	—	**94.2, 93.4**	94.2, 93.8
10, tracker out 11 to 17 min: on Site B (s), imaged (%)	52.2, 92.4	51.8, 92.5	42.0, 93.9	41.8, 93.9	**39.7, 94.2**	—
7: on Site A (s)	80.6	79.4	71.0	70.8	**68.7**	—
7, tracker out 10 to 16 min: worst (°), all 100 %	0.040	0.025	0.027	0.026	0.025	—
5: settled (s), peak relative rate (°/s)	102.2, 1.00	86.3, 1.73	93.1, 1.00	92.9, 1.00	**90.8**, 1.01	—
8, roll wheel fails: imaged (%), worst (°)	82.0, 5.82	**84.7, 3.63**	82.1, 5.64	82.1, 5.64	81.5, 5.74	—
3, first orbit: mean error (°), travel (mN m)	0.0188, 369	0.0188, 369	0.0077, 1008	**0.0073**, 1038	0.0076, 687	0.0070, 3199
3, three orbits: 0.1° lost (min), above 98 % (min)	250.3, 35.0	250.3, 36.6	250.3, 35.1	250.3, 38.9	250.2, 38.8	—
Challenge as shipped: imaged (%)	24.1	13.3	26.8	17.1	**28.1**	—
Challenge suggested	met	met	met	met	met	—

Retuned: the eigenaxis law with $K_p = 0.5$ (its slider's end), $K_d = 1.2$, $\omega_{lim} = 1\,^\circ/\mathrm{s}$, chosen by the MPC's own rule. MPC signals: that law with $\hat{\mathbf{d}}$ as a feedforward torque and the half-step rate correction, the strongest classical form of what the plan knows. K_p 4, K_d 3.4: the eigenaxis law past its sliders. The MPC solves 2,400 times in 20 min with no fallback, and with its preview off every row is the same to the printed digit. Travel is the summed change of the wheel commands.

plan knows, and the law past its sliders, all on the same seed, instruments and estimate. The MPC holds the camera within 0.1° of Site B 39.6 s after the handover, 1.0 s more than 38.6° takes at exactly 1 °/s, and images 94.26 % of the 11.47 min of the two passes, against 52.1 s and 92.44 % under the eigenaxis law, as shown in Figure 7.5. It spends the pitch wheel's full 3 mN m for the first second, reaches the rate limit 1.5 s after the handover, rides it at up to 1.015 °/s and brakes as late as its model allows, while the law's rate command eases out along its exponential and leaves the camera 2.90° off at 40 s.

Most of the 12.5 s is tuning, not prediction. Retuned once, the eigenaxis law settles in 41.9 s and images 93.92 %, and fed the plan's signals it settles in 41.7 s. The 2.3 s left after the retune is bandwidth that the sliders do not offer, and $K_p = 4$ with $K_d = 3.4$ closes it, 39.9 s and 94.22 %, but those gains ask the wheels for more than 3 mN m during 1.9 s of the run and move them 4.0 times as far as the plan, 699 against 173 mN m of torque travel, and 4.7 times as far over an orbit at nadir. The retuned law asks past 3 mN m for 0.4 s, and the plan never does, because the bound is a constraint of its problem. The preview is worth nothing measurable, because the reference's angular acceleration through the pass stays under 0.007 83 °/s^2 and the law's feedforward of its present value already carries it. With the preview off, every row of Table 7.7 is unchanged to the printed digit.

What the plan measurably adds is the limits, and steadiness when the spacecraft differs from its declaration. With the pitch moment at its slider's low end of 0.02 kg m^2, or every moment at 0.3 kg m^2, it images 94.22 and 93.42 % of case 10, against 93.65 and 92.89 % for the retuned law and 80.07 and 61.12 % for the eigenaxis law, which since the correction of Section 7.6.13 keeps the declared inertia just as the plan does. The stiff law images 94.24 and 93.79 % there but at $J_y = 0.02\,\mathrm{kg\,m^2}$ chatters, asking past 3 mN m for 1096 s of the 20 min and running above $1.05\omega_{lim}$ for 130.6 s.

Elsewhere the plan wins the slews and loses the failed wheel. It reaches Site A in case 7 11.9 s sooner than the law. At nadir its estimate of the drag gives the retuned law's mean error, 0.0076° against 0.0077°, with 687 against 1008 mN m of torque travel in the first orbit. Case 8 is its loss, 81.5 % imaged against 82.0 % for the law and 84.7 % for three axes. After the roll wheel fails the plan holds the coils' present authority over its horizon and cannot see the field turn toward the roll axis, so it neither saves nor spends roll authority ahead of the field's pass. The challenge as shipped saturates the pitch wheel before the pass under every law, and the plan images 28.1 % of it against 24.1 %.

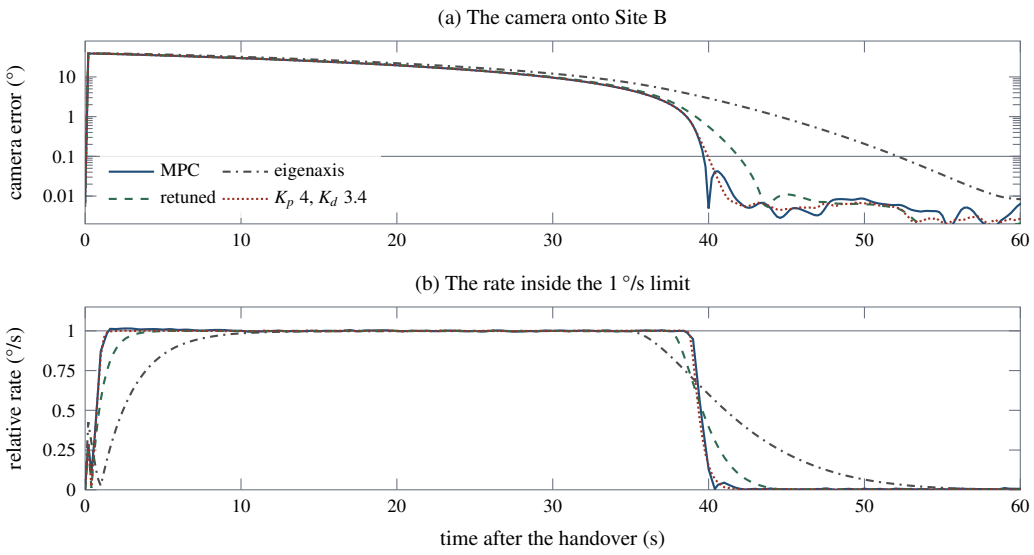

Figure 7.5 Case 10's handover to Site B under four pointing laws, on the engine with the default seed: the MPC solid, the shipped eigenaxis law dash-dotted, the law retuned by the MPC's rule ($K_p = 0.5$, $K_d = 1.2$) dashed and the stiff law ($K_p = 4$, $K_d = 3.4$) dotted. (a) The true camera error from its reference, with the 0.1° requirement as a thin gray line. (b) The true relative rate, with the 1 °/s limit as a thin gray line. The MPC and the stiff law both ride the limit and stop within 40 s, while the shipped law eases in and out and the retuned law lies between them.

7.6.12 Worst Cases Under Every Pointing Law

Every law flew every failure, outage, kick and extreme setting the application offers, as listed in Table 7.8. No run fell back, no wheel delivered more than τ_{max} or stored more than h_{max}, and dumping, safe mode, the failed axis's coils and the wheel clamps act under the MPC as under the laws. The plan is better in the torque-limited slew at 6 °/s, settled in 20.3 s against 30.8 s without asking past τ_{max}, and with a 0.5 mN m torque limit, 47.0 s against 54.5 s. It is worse at two check levels. The tracker's 1 °/s rate limit is also the slew limit the plan rides, so after a kick at nadir the tracker is blind by rate 3.2 s longer than under the law, and in the slew from wheels at 85 % 11.0 s longer, with 62.6 s over $1.02\omega_{lim}$ against 30.0 s. Every law fills a wheel in that slew. When the roll wheel fails at the handover, the plan holds the pitch command on its bound for 0.8 s of the first second and spends 3.1 s between 1.020 and 1.027 °/s on the return to nadir 7.8 min later, where the law spends none.

7.6.13 What the Review Corrected, and at What Price

A technical review of the predictive law's first build found two behaviors in which the plan was less safe than the eigenaxis law, and each was corrected and measured before and after. With one slack for the rate band over the whole horizon, a plan that started outside the box, as on the entry to nadir from detumble with a 108.8° error, paid for the excess once and rode 1.127 °/s for 14.0 s, where the law peaked at 1.001 °/s. A slack at every judged step gives 1.011 °/s and no time over $1.02\omega_{lim}$, at the price of 0.1 ms a decision, 28.1 % of the challenge imaged as shipped in place of 30.7 %, and 39.7 s in place of 39.6 s for case 10 at $J_z = 0.02\,\mathrm{kg\,m^2}$. A 20 mN m s kick, an impulse, was read by the disturbance filter as torques of 147 and 110 µN m about roll and yaw, against 2 µN m of drag, and the plan countered them with the relative rate 6 to 8 % over its limit for 39.1 s after a kick at nadir, against 5.6 s for the law. The kick test of Section 7.5.5 cuts that to 4.7 s and the tracker's rate blindness after a kick at nadir from 39.3 to 34.0 s, for 92.1 % of the pass imaged after a kick inside it

Table 7.8 Worst Cases Under Every Pointing Law, Default Seed

Case and measure	Eigenaxis	Three axes	MPC
Pitch or yaw wheel fails in the pass: %, worst °	100, 0.026	100, 0.025	100, 0.025
Roll wheel fails at the handover: s, %	52.1, 92.4	52.2, 92.4	39.6, 94.3
Tracker out the whole pass: worst °	0.040	0.039	0.038
Tracker out, no-bias filter, handover: %	0	0	0
Kick of 20 mN m s in the pass: %, s over, s blind	89.3, 5.5, 34.5	90.7, 46.0, 31.2	92.1, 4.7, 33.6
Kick at nadir: worst °, s over, s blind	0.036, 5.6, 30.8	0.018, 46.5, 47.3	0.023, 4.7, 34.0
Case 5 at 6 °/s: settled s, s past τ_{max}, peak °/s	30.8, 1.7, 5.94	34.2, 2.1, 7.66	20.3, 0, 6.00
Case 5, wheels at 85 %: settled s, s at limit, s over, s blind	111.0, 54.5, 30.0, 29.3	210.0, 181.8, 188.8, 187.3	97.0, 51.7, 62.6, 40.3
Challenge suggested, drag 10 µN m or bias 5 µT	met	met	met
Case 10, $J_x = 0.2$: s, %	52.2, 92.4	52.1, 92.4	39.7, 94.2
Case 10, $J_z = 0.02$: s, %	52.1, 92.4	52.2, 92.4	39.7, 94.2
Case 10, $\tau_{max} = 0.5$ mN m: s, %	54.5, 92.1	54.4, 92.1	47.0, 93.3
Case 2 into nadir: peak °/s, s over, s blind	1.001, 0, 1064.9	1.732, 38.1, 1094.7	1.011, 0, 1064.9
Case 9, safe mode through the shadow	safe	safe	safe

s over: true relative rate above $1.02\omega_{lim}$ in the pointing modes. s blind: tracker blind by its 1 °/s rate limit. s: time onto Site B. No fallback in any row.

in place of 92.5 %. A first form, on every innovation, slowed the slew from wheels at 85 % from 97.0 to 108.8 s, because a wheel in its fade gives real innovations up to 7.3 °/s, and was replaced.

The soft wheel ceiling stays priced with one slack, as built. Priced on every step, the $0.9h_{max}$ ceiling made the plan give up nadir at 224.8 min of case 3, 25.5 min before the law, while the wheel reached 98 % at 248.5 min anyway, so the ceiling binds only the plan's own moves.

One defect that predated the controller was fixed in the same round. The shipped laws took the run's inertia and the filters the run's tracker noise, while the plan used the declared values. All now use the declared values (Section 7.3.7), no shipped or suggested run moved, and case 10 at $J_y = 0.02\,\text{kg m}^2$ or with every moment at $0.3\,\text{kg m}^2$ falls under the eigenaxis law from 92.2 and 92.5 % imaged to 80.1 and 61.1 %.

7.7 Technology in Practice

The spacecraft of Section 7.3 is read by a star tracker, a gyro, a magnetometer and a sun sensor, and moved by three reaction wheels and three coils. Figure 7.6 draws its loops with real devices in their place.

What Measures It

The tracker of Equation (7.6), with 10 arcsec of noise across its boresight and six times that about it, blind within 30° of the Sun and above 1 °/s, stands for one real star tracker. Real trackers too know the roll about the boresight less well than the boresight, and they lose their solution at a body rate that depends on the design. Missions that depend on accurate attitude carry a Kalman filter propagated on the inertial sensors to keep the attitude when the tracker loses track [23], the role of the chapter's bias filter. The magnetometer is placed well apart from the wheels, whose motors and rotors are not magnetically clean, because the coils' commands are computed from the field it measures [23]. The model assumes that the coils' own field is blanked from its readings.

> **TECHNOLOGY** **Sensor: Star Tracker**
>
> A star tracker compares a digital image of the sky with a star catalog on board and reports the full three-axis attitude, up to several times a second. The body's rate, stray light and glare can corrupt or invalidate a solution. Most modern trackers know their boresight to within a few arcseconds and the roll about it less well. A low-cost tracker may hold a solution only up to 0.3 °/s of body rate, a costlier one up to 3.0 °/s [23]. On a trend a lost solution is a gap in the tracker's readings, which a filter propagated on the gyro must bridge, as in cases 6 and 7.

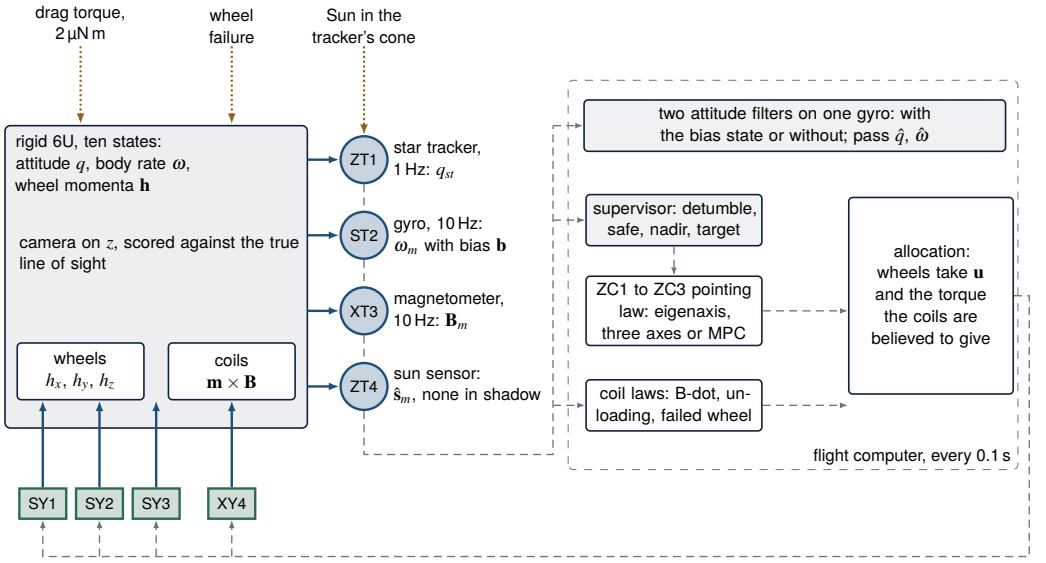

SY1 to SY3: reaction wheel drives, 3 mN m and 30 mN m s each, torque fading over the last 10 % of speed; tachometers return $\mathbf{h}$ to the coil laws. XY4: three magnetorquer coils, 0.5 A m^2 each, torque only across the field.

Figure 7.6 The loops of the chapter with real devices in place of the schematic's symbols. ZT1 is the star tracker, ST2 the gyro, XT3 the magnetometer and ZT4 the sun sensor. One flight computer (dashed frame) runs the two filters, the supervisor, the pointing law on three axes (ZC1 to ZC3) and the coil laws every 0.1 s. SY1 to SY3 drive the three wheels and XY4 the three coils. Notice that the laws see only the filter's estimate and the measured field, and that only external torques such as the coils' and the drag's change the total momentum, which the wheels can store but not remove.

What Moves It

The wheels of Equation (7.4) deliver up to 3 mN m, fade over their last 10 % of speed and have no stiction near zero speed, where real wheels lose performance. Three wheels on orthogonal axes are the least that gives full three-axis control, and for a small spacecraft in Earth orbit the magnetic torquers are the main means of keeping the wheels out of saturation. A torquer is either an air coil, whose residual field dies quickly once it is switched off, or a rod with a ferrous core, which gives more dipole for the same power but can keep a residual field that upsets a magnetometer. Three torquers on orthogonal axes usually need no spare, because their windings give internal redundancy [23]. The model's coils are ideal dipoles of 0.5 A m^2 with no residual field.

> **TECHNOLOGY** **Final Element: Reaction Wheel**
>
> A reaction wheel gives torque and stores momentum along its spin axis, and the spacecraft turns the other way as the wheel speeds up. At its top speed a wheel is saturated and can neither give torque nor store more momentum in that direction, so the wheels are desaturated from time to time by an external torque from thrusters or magnetic torquers. Most wheels perform worse near zero speed, so some missions bias the wheels at half speed for finer control and less jitter, at the price of half the capacity. Static and dynamic imbalance in a wheel disturbs the spacecraft it steers [23]. On a trend saturation is a wheel momentum flat at its limit while the attitude error grows, as in case 3.

Where the Controller Runs

The chapter's flight software runs its filters, supervisor and laws on one 0.1 s step and reads the tracker once a second, where real trackers report up to several times a second. On a small spacecraft

that work is often done by an integrated attitude unit, which joins the wheels, magnetometer, torquers, sun sensors and star tracker with a processor and software for Sun tracking, inertial pointing and Earth target tracking [23], the safe, nadir and target modes of Section 7.5. The avionics of a small spacecraft are usually single-string, so one failed element can stop the whole system, and designers now add radiation-hardened or radiation-tolerant designs for long missions [23].

What Is Modernizing It

Star trackers are becoming faster, more accurate, smaller and lighter, and cheaper designs now give full three-axis attitude to missions that could not have afforded one. A magnetically levitated reaction wheel in development replaces ball bearings with active magnetic bearings to remove mechanical wear and stiction [23]. The pointing law itself can be a predictive controller that keeps the wheels' limits as constraints, as in Section 7.5.5.

> **IN PRACTICE** **Sizing the Wheels Against the Air**
>
> NASA's design criteria of 1971 for aerodynamic torques ask that the drag torque, together with every other disturbance, be shown not to degrade the attitude control. Below about 600 km aerodynamic forces usually predominate. The torque is to be computed at enough points of the orbit and times of the year to give its profile, in a model atmosphere whose density follows the sunspot cycle. The distance from the center of mass to the center of pressure is estimated conservatively large, and the result sizes the actuators' peak torque, momentum storage and momentum transfer [24].

7.8 Limitations

The spacecraft is rigid with its inertia declared to the flight software, the wheels are ideal torque sources with a linear fade, and the coils are ideal dipoles. The orbit is Keplerian apart from the node's turn, the field is the 2020 dipole without its higher harmonics, the shadow has no penumbra, and the drag torque is one constant, so the model teaches the structure of attitude control rather than the budget of a particular mission. The orbit itself, flown with drag and the Earth's oblateness, is the subject of Chapter 8.

The instruments are generic and the star tracker has no latency, stray light or acquisition time. The filter fuses the tracker alone, not the magnetometer or the sun sensor, so a long outage leaves it on the gyro, and the orbit, the time and the site are known exactly. The laws are the textbook ones, without the momentum-bias or optimal designs a mission may prefer apart from one predictive law, and every threshold of the supervisor and the scorecard is a teaching value.

The predictive law's model holds the gyroscopic torque, the coils' torque and, after a wheel failure, the coils' authority at their present values for its 60 s horizon, which is why it loses case 8. It drops the coupling of the error with the reference's rate, under 0.4 % of the torque authority at the pass's 0.86 °/s, sees the gravity gradient and the drag only through $\hat{\mathbf{d}}$, stands a soft ceiling at $0.9h_{max}$ in for the motor's fade, anticipates no switch of mode or of tasking, and leaves out the tracker's 1 °/s rate limit, which it reaches sooner than the law.

OPEN CHALLENGE Pointing on Fewer Wheels

Reaction wheel failure is a common mode of failure in many missions. A large spacecraft may carry three spare wheels, while a small one with no room for spares often mounts four wheels skewed, so that after any single failure a reduced torque remains about every axis. Magnetic torquers alone cannot usually stabilize three axes, and the coarse three-axis pointing shown with them in research takes several orbits and an orbit inclined enough for the field to vary [23]. In case 8 two wheels and the coils image at most 84.75 % of the pass in every tuning tried. Progress would be laws that plan the coils' authority as the field turns.

7.9 Conclusion

On this simulator a 6U CubeSat points a camera with wheels that store every torque they absorb, coils that act only across a turning field and a filter that coasts on a biased gyro, and the camera is scored against a truth the flight software never sees. A drag torque of 2 µN m fills the pitch wheel by 37.85 % an orbit and costs the attitude at 245.0 min. The coils empty a 78.1 % excess in 58.6 min while up to 84 % of it lies out of their reach, and over an orbit they remove an excess along the orbit normal twice as fast as one in the plane. B-dot is fastest near 0.5 mN m s and three times slower at 8. The eigenaxis law flies a 90.45° slew through 90.62° where three separate axes fly 105.28°. Without the bias state the camera drifts 2.35° in the Sun's blind window, and a failed wheel caps a pass at 84.75 % in every tuning tried. The challenge rises from 24.12 % of the pass imaged to 100 % only with both unloading and the bias state. The model omits flexibility, the field's higher harmonics and solar pressure, and a fuller model would add a momentum-biased design, a star tracker with latency and stray light, a filter that fuses the field and the Sun, and unloading scheduled around the passes. The spacecraft is a browser laboratory for actuators that keep what they apply, and its optional predictive law shows that a plan's measured gain on a handover, 39.6 against 52.1 s, comes from staying inside the wheels' limits at full bandwidth rather than from looking ahead.

KEY POINTS

- A wheel turns the spacecraft by spinning the other way, so it stores every torque it absorbs, and only the coils, pushing across the Earth's field, can empty it.
- A drag torque of 2 µN m fills the pitch wheel by 37.85 % an orbit and costs the attitude at 245.0 min.
- B-dot is fastest near 0.5 mN m s and three times slower at 8, because a strong law leaves the body turning about the field line.
- The challenge rises from 24.12 % of the pass imaged to 100 % only with both unloading and the bias state.

Where else this applies. Volume 4, Chapter 13 meets the same limit in feet that push only inside a friction cone, and Volume 5, Chapter 13 drifts between position fixes as this filter drifts between star tracker readings.

References

[1] J. Puig-Suari, C. Turner, and W. Ahlgren, "Development of the standard CubeSat deployer and a CubeSat class PicoSatellite," in *2001 IEEE Aerospace Conference Proceedings*, vol. 1, 2001, pp. 347–353.

[2] I. Ofodile, J. Kütt, J. Kivastik, M. Kaspar Nigol, A. Parelo, E. Ilbis, H. Ehrpais, and A. Slavinskis, "ESTCube-2 attitude determination and control: Step towards interplanetary CubeSats," in *2019 IEEE Aerospace Conference*, 2019, pp. 1–12.

[3] F. L. Markley and J. L. Crassidis, *Fundamentals of Spacecraft Attitude Determination and Control*, ser. Space Technology Library. New York, NY: Springer, 2014.

[4] National Aeronautics and Space Administration, "Spacecraft magnetic torques," National Aeronautics and Space Administration, Washington, DC, Space Vehicle Design Criteria (Guidance and Control) NASA SP-8018, Mar. 1969. [Online]. Available: https://ntrs.nasa.gov/citations/19690020961

[5] P. J. Camillo and F. L. Markley, "Orbit-averaged behavior of magnetic control laws for momentum unloading," *Journal of Guidance and Control*, vol. 3, no. 6, pp. 563–568, 1980.

[6] A. Astolfi and M. Lovera, "Global spacecraft attitude control using magnetic actuators," in *Proceedings of the 2002 American Control Conference*, vol. 2, 2002, pp. 1331–1335.

[7] G. Avanzini and F. Giulietti, "Magnetic detumbling of a rigid spacecraft," *Journal of Guidance, Control, and Dynamics*, vol. 35, no. 4, pp. 1326–1334, 2012.

[8] E. J. Lefferts, F. L. Markley, and M. D. Shuster, "Kalman filtering for spacecraft attitude estimation," *Journal of Guidance, Control, and Dynamics*, vol. 5, no. 5, pp. 417–429, 1982.

[9] F. L. Markley, "Attitude error representations for Kalman filtering," *Journal of Guidance, Control, and Dynamics*, vol. 26, no. 2, pp. 311–317, 2003.

[10] R. E. Fischell, "Magnetic and gravity attitude stabilization of Earth satellites," Applied Physics Laboratory, The Johns Hopkins University, Report AD0271383, Defense Technical Information Center, May 1961.

[11] J. R. Wertz, Ed., *Spacecraft Attitude Determination and Control*, ser. Astrophysics and Space Science Library. Dordrecht: D. Reidel, 1978.

[12] E. Y. Yu, "Spin decay, spin-precession damping, and spin-axis drift of the Telstar satellite," *Bell System Technical Journal*, vol. 42, no. 5, pp. 2169–2193, 1963.

[13] A. C. Stickler and K. T. Alfriend, "Elementary magnetic attitude control system," *Journal of Spacecraft and Rockets*, vol. 13, no. 5, pp. 282–287, 1976.

[14] B. Wie and P. M. Barba, "Quaternion feedback for spacecraft large angle maneuvers," *Journal of Guidance, Control, and Dynamics*, vol. 8, no. 3, pp. 360–365, 1985.

[15] B. Wie, H. Weiss, and A. Arapostathis, "Quaternion feedback regulator for spacecraft eigenaxis rotations," *Journal of Guidance, Control, and Dynamics*, vol. 12, no. 3, pp. 375–380, 1989.

[16] C. C. Liebe, "Accuracy performance of star trackers - a tutorial," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 38, no. 2, pp. 587–599, 2002.

[17] J. E. Van Cleve, S. B. Howell, J. C. Smith, B. D. Clarke, S. E. Thompson, S. T. Bryson, M. N. Lund, R. Handberg, and W. J. Chaplin, "That's how we roll: The NASA K2 mission science products and their performance metrics," *Publications of the Astronomical Society of the Pacific*, vol. 128, no. 965, p. 075002, 2016.

[18] P. Alken, E. Thébault, C. D. Beggan *et al.*, "International Geomagnetic Reference Field: the thirteenth generation," *Earth, Planets and Space*, vol. 73, p. 49, 2021.

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

[20] Ø. Hegrenæs, J. T. Gravdahl, and P. Tøndel, "Spacecraft attitude control using explicit model predictive control," *Automatica*, vol. 41, no. 12, pp. 2107–2114, 2005.

[21] U. Eren, A. Prach, B. B. Koçer, S. V. Raković, E. Kayacan, and B. Açıkmeşe, "Model predictive control in aerospace systems: Current state and opportunities," *Journal of Guidance, Control, and Dynamics*, vol. 40, no. 7, pp. 1541–1566, 2017.

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

[23] S. V. Weston, C. D. Burkhard, J. M. Stupl, R. L. Ticknor, B. D. Yost, R. A. Austin, P. Galchenko, L. K. Newman, and L. Santos Soto, "State-of-the-art small spacecraft technology," National Aeronautics and Space Administration, Ames Research Center, Moffett Field, CA, Technical Publication NASA/TP-20250000142, Jan. 2025. [Online]. Available: https://ntrs.nasa.gov/citations/20250000142

[24] National Aeronautics and Space Administration, "Spacecraft aerodynamic torques," National Aeronautics and Space Administration, Washington, DC, Space Vehicle Design Criteria (Guidance and Control) NASA SP-8058, Jan. 1971. [Online]. Available: https://ntrs.nasa.gov/citations/19710016459

7.10 Exercises

7.1 ⌕ Concept check · 2 min

In scenario 3 the camera holds nadir to 0.018° on average for more than two orbits, and then, at 245.0 min, the attitude is lost although no setting has changed. What is the cause?

(a) The pointing gains are too low to resist the drag torque once the pitch wheel spins fast

(b) The pitch wheel has stored the drag's momentum orbit after orbit and, near its capacity, can absorb no more torque

(c) The gravity-gradient torque grows with the pitch wheel's speed until it overpowers the pointing law

(d) The bias estimate of the filter drifts over three orbits until the estimated attitude is no longer trusted

7.2 Concept check · 2 min

In scenario 2 the B-dot law detumbles the spacecraft from 10.1 °/s. With a gain of 0.5 mN m s the rate falls under 0.5 °/s in 26.8 min, and with 8 mN m s in 85.5 min. Why is the stronger law slower?

(a) The larger dipoles saturate the magnetometer, so the law loses its measurement of the field's rate

(b) The larger gain makes the filter on the field's rate unstable, so the dipoles chatter

(c) The strong law quickly removes the rate across the field and leaves the body turning about the field line, where the coils make no torque

(d) The larger dipoles add a torque along the field that spins the body up again

7.3 Concept check · 2 min

In scenario 5 both pointing laws use the same gains and the same 1 °/s slew rate limit to make a change of attitude of 90.45°. The eigenaxis law turns the body through 90.62° at up to 0.94 °/s. The three-axis law turns it through 105.28° at up to 1.69 °/s. Which explanation is right?

(a) The three-axis law applies higher gains on the axis with the largest error, so it overshoots

(b) The three-axis law leaves out the gyroscopic term, so the body is pushed off its path

(c) The three Euler angles cannot describe a rotation this large, so the three-axis law mistakes the error

(d) Each loop limits only its own component, so the rate vector can reach nearly $\sqrt{3}$ times the limit along a direction off the eigenaxis

7.4 Theory and calculation · 5 min

The spacecraft holds nadir with unloading off. A drag torque of 2 µN m, fixed in the body about pitch, acts the whole time. The orbit's mean motion is $n = 1.106\,783 \times 10^{-3}\,\mathrm{rad/s}$, a wheel holds at most $h_{max} = 30\,\mathrm{mN\,m\,s}$, and each coil's dipole is limited to $0.5\,\mathrm{A\,m^2}$.

(a) Find the orbital period and the momentum the pitch wheel gains in one orbit, as a share of its capacity. *Checkpoint:* 5677 s (94.62 min) and 37.85 %.

(b) Find the time from an empty wheel to 98 % of capacity, in orbits and in minutes. *Checkpoint:* 2.59 orbits, 245.0 min.

(c) The field along the orbit runs from 23.69 to 47.37 µT. Find the largest torque one coil at its limit can make across it, and the torque the unloading law asks with $\tau_d = 15\,\mathrm{min}$ at the 40 % excess that switches it on. *Checkpoint:* 11.8 to 23.7 µN m, and 13.3 µN m.

7.5 Theory and calculation · 6 min

Over the first three orbits the orbit average of $\hat{\mathbf{b}}\hat{\mathbf{b}}^{\mathsf{T}}$, the field's direction times itself, is diag(0.493, 0.491, 0.016) in the axes $(\hat{\mathbf{P}}, \hat{\mathbf{Q}}, \hat{\mathbf{h}})$ of Equation (7.8). The unloading law runs with $\tau_d = 15\,\mathrm{min}$ and no dead band.

(a) Find the time constant with which an excess along the orbit normal decays, and the time constants of an excess along $\hat{\mathbf{P}}$ and along $\hat{\mathbf{Q}}$. *Checkpoint:* 15.24 min, and 29.59 and 29.47 min.

(b) Find the time each excess takes to fall to a tenth, and compare with the engine's 36.2, 70.3 and 64.4 min for the pitch, yaw and roll excess, which start along $\hat{\mathbf{h}}$, $-\hat{\mathbf{P}}$ and $\hat{\mathbf{Q}}$. *Checkpoint:* 35.1, 68.1 and 67.9 min.

(c) Explain why the drag's excess at nadir is the easiest one to unload, and why an excess in the orbit plane falls in steps. *Checkpoint:* The field lies close to the orbit plane.

7.6 Theory and calculation · 6 min

The pointing law of Equation (7.11) runs with $K_p = 0.0625/\mathrm{s}^2$, $K_d = 0.45/\mathrm{s}$ and $\omega_{lim} = 1\,^\circ/\mathrm{s}$.

(a) Find the natural frequency and the damping ratio of the unsaturated loop. *Checkpoint:* 0.25 rad/s and 0.9.

(b) Find the error angle below which the rate command leaves its limit, for the shipped gains and for the gains $K_p = 0.5$, $K_d = 1.2$ that the MPC's tuning rule chose for the eigenaxis law. *Checkpoint:* 7.2° and 2.4°.

(c) Case 10 hands the camera over to a site 38.6° away. Find the shortest time any law held to the rate limit could take, and the largest body rate three single-axis loops held to the same limit could reach. *Checkpoint:* 38.6 s and 1.73 °/s.

7.7 App challenge · 8 min

Load scenario 3, Nadir hold: the wheels creep, and change no setting. Run at 20 simulated minutes per second to 285 min, three orbits, with the default seed, and export the CSV file. The file's time column is in hours.

(a) Read the pitch wheel's momentum, column `hy_abs_pct`, in the row nearest one orbit, 94.62 min. *Checkpoint:* 37.84 %, within 0.05 %.

(b) Find the first row in which `hy_abs_pct` reaches 98 %, and the first row in which the camera's true error, column `point_err_deg`, exceeds 1°. Give both times in minutes. *Checkpoint:* 245.0 min and 250.8 min, each within 0.1 min.

(c) Press the reload button, set Momentum dumping On on the MOM card before Run, and run again to 285 min. Read the Fullest wheel tile. *Checkpoint:* 40.0 % (pitch).

Compare (a) and (b) with the drag's budget $\tau_a T / h_{max}$.

7.8 App challenge · 8 min

Load scenario 6, The gyro bias and the blind window, as shipped, with the Attitude estimate at EKF, no bias. Run to 27 min at 1 simulated minute per second with the default seed and export the CSV file, whose time column is in minutes. The column `tracker_blind` is 1 while the star tracker reports nothing, and `point_err_deg` is the camera's true error.

(a) Find the row in which the tracker goes blind, the row in which it reports again, and the largest `point_err_deg` between them. *Checkpoint:* 14.08 to 25.58 min, and 2.34°, within 0.05°.

(b) Press Reset, set the Attitude estimate to EKF + bias before Run, and repeat. *Checkpoint:* 14.08 to 24.92 min, and 0.054°, within 0.005°.

(c) The true gyro bias is 7, −9 and 5 °/h about roll, pitch and yaw. Estimate the drift of the no-bias filter's attitude across the camera axis over the window of (a), and explain why that window is longer. *Checkpoint:* About 2.2°.

7.9 App challenge · 9 min

Load the challenge, Challenge: three orbits, then the pass. Use the default seed and let each run end by itself at 302.95 min, at 20 simulated minutes per second. After each run read the Pass imaged, Fullest wheel and Excess momentum now tiles.

(a) Run the challenge as shipped, with unloading off and the Attitude estimate at EKF, no bias. *Checkpoint:* 24.1 %, 100.0 % and 100.0 %.

(b) Press Reset, set Momentum dumping On and run. Then press the reload button, set the Attitude estimate to EKF + bias and run. *Checkpoint:* 49.8 %, 40.0 and 33.0 %, then 20.6 %, 100.0 and 100.0 %.

(c) Press Reset, set Momentum dumping On as well and run. *Checkpoint:* 100.0 %, 40.0 and 33.0 %.

Explain why each change alone fails the target.

7.10 Application area · 4 min

The drag torque at nadir follows the air density, which changes by several times over a solar cycle. The app's drag slider reaches 10 µN m. Use $T = 5677\,\mathrm{s}$, $h_{max} = 30\,\mathrm{mN\,m\,s}$ and $\tau_d = 15\,\mathrm{min}$.

(a) With unloading off, find the momentum the pitch wheel gains per orbit and the time from an empty wheel to 98 %. *Checkpoint:* 189.2 % of capacity, 49.0 min.

(b) With unloading on all the time, the excess along the orbit normal settles where the averaged law's torque equals the drag. Find that excess and say what it means for the dead band of 40 and 10 %. Name two changes that would add margin. *Checkpoint:* About 30.5 % of a wheel.

7.11 Application area · 3 min

In case 8 the roll wheel fails during the pass, and the coils take over roll. Their roll torque is $\mathbf{m} \cdot (\mathbf{B} \times \hat{\mathbf{x}})$, and the field comes within 3.46° of the roll axis at 16.82 min.

(a) With a dipole of $0.5\,\mathrm{A\,m^2}$ along $\mathbf{B} \times \hat{\mathbf{x}}$ and a field of 23.69 to 47.37 µT, find the largest roll torque at that instant, and compare it with a wheel's 3 mN m. *Checkpoint:* 0.71 to 1.43 µN m, about 2100 to 4200 times less.

(b) Name two ways an operator could keep imaging with two wheels. *Checkpoint:* Plan passes and attitudes around the field.

7.12 Review the agent · 5 min

An AI assistant reviews scenario 10 and writes: "The MPC puts the camera on Site B in 39.6 s against 52.1 s for the eigenaxis law because its 60 s preview of the site's path lets it brake at the right moment. Prediction is what the classical law lacks. Switch every pass, and the challenge as shipped, to the MPC with its preview on, and the same gain will follow." Find the flaw and say how you would check it on the app.

CHAPTER 8

A Low Earth Orbit Simulator Flown on Mean Elements and Collision Probability

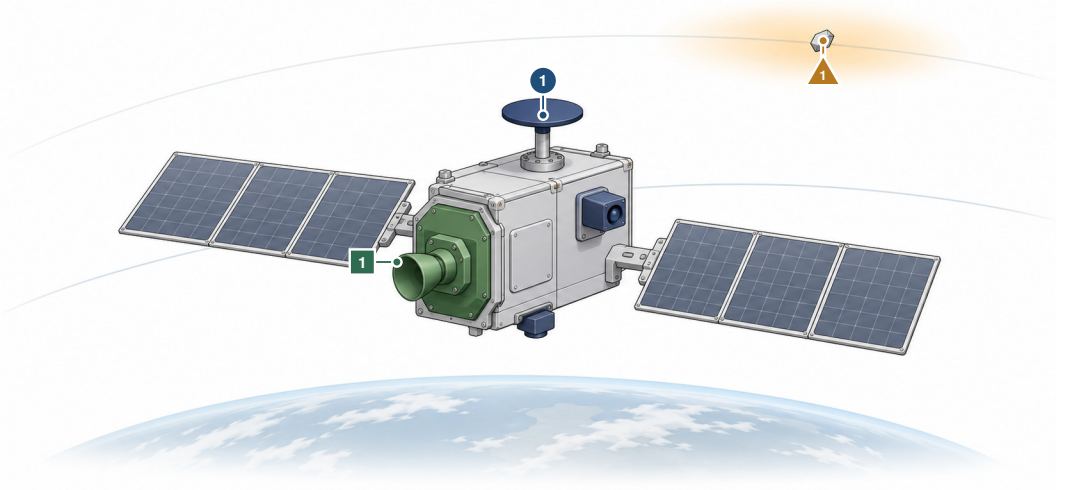

The controller moves 1 electric (Hall) thruster that fires every burn

The controller reads 1 satellite navigation antenna for the position fixes

What disturbs it 1 debris fragment, its position known only within the amber haze

Notice the fragment far from the satellite: whether its pass is safe depends on how uncertain its position is, not on the distance alone. Illustration, not to scale.

IN PLAIN TERMS Keeping a Satellite Up and Out of Harm's Way

The main idea. Thin air at the top of the sky slowly drags a low satellite down, so a small engine gives it a gentle push now and then. Its height is judged by the average over one lap, since the bulge of the Earth makes it swing on every lap.

The challenge. Bits of old rockets cross its path, and nobody knows exactly where they are. A piece that looks far away can still be a danger, so each push is planned hours ahead on the odds of a hit.

Why it matters. The fuel is small and every crash makes more debris, so each push must count.

THREE WAYS TO READ THIS CHAPTER

⊘	**Explore**	12 min	the first page, the two history segments of Section 8.2, the try-it segment of scenario 1 (Section 8.4), the fact on the miss distance (Section 8.5), the try-it segment of scenario 3 and the fact on the campaign (Section 8.7), 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 8.7), Technology in Practice (Section 8.8), Limitations (Section 8.9) and the open challenge

The control question A satellite in low Earth orbit is flown on a mean altitude that no instrument reports, a drag that has to be forecast and the collision probability of a pass hours ahead. When must a burn be decided, how large must it be, and on which prediction should it be judged?

Learning objectives

- Explain why orbit maintenance works on the one-orbit mean altitude and why one prograde burn raises the far side of the orbit.
- Compute the change of the mean, the burn time and the xenon of a reboost, and the drift that an along-track burn buys over a lead time.
- Evaluate the collision probability of a pass from its miss and its covariance, and show why a large miss distance can still be unsafe.
- Compare deadband and predictive maintenance, and nominal and uncertainty-aware avoidance, on the measured cases and the campaign.

Prerequisites Sections 1.4.6 (Kalman filter) and 1.4.10 (bound or commit), with Sections 1.6.1 and 1.6.4 for the runs. Chapter 7 helps but is not required.

The app https://apmonitor.com/apps/leo-orbit-control/

See also Volume 5, Chapter 13 for a vehicle steered on a filter's estimate of a position nothing aboard measures. Volume 9, Chapter 7 for a decision taken days ahead on an uncertain forecast under a chance constraint, and Volume 9, Chapter 5 for releases made one travel time ahead of the demand.

8.1 Introduction

A feedback loop compares a measurement with a setpoint. A satellite in low Earth orbit is flown instead on three quantities that no instrument reports. The first is the mean altitude, an average over the last revolution, because the oblateness of the Earth adds periodic terms to every orbital element and the motion of a close satellite is written in mean elements with those terms added [1]. The second is the drag that lowers the mean, set by an upper-atmosphere density that the first satellites were used to measure from the contraction of their own orbits [2] and that density models still have to supply [3]. The third is the collision probability of a debris pass announced hours ahead, computed from two uncertain orbits in the plane normal to their relative velocity [4, 5].

The simulated satellite is a 120 kg Earth-observation spacecraft in a sun-synchronous orbit at a mean altitude of 500 km and an inclination of 97.4°, pushed by a 20 mN Hall thruster that fires for a minute to change its velocity by 1 cm/s. Three properties separate it from a loop with a delay. With nothing but gravity acting, its altitude swings between 498.42 and 501.58 km twice an orbit while the mean holds within 1 m. A burn changes the altitude only on the far side of the orbit, half a revolution later, and moves the satellite along its track by about three times the velocity change for every second that follows [6]. The criterion for a safe pass is a probability, and a predicted miss distance does not measure it.

The simulator addresses two beliefs. One is that the altitude a satellite reports is the altitude it controls, and the other is that a pass predicted to miss by a wide margin is safe. The second has a history in operations, since Space Shuttle avoidance was commanded whenever a conjunction fell inside a box around the vehicle, while the International Space Station was designed to maneuver on a probability of collision [4]. Data of lower quality can even make a pass look safer, the effect called probability dilution [7]. The class of problem is general. A decision must satisfy a constraint with a stated probability, on a model whose uncertainty is carried forward to the moment the constraint applies, which is chance-constrained predictive control [8]. It bounds the unknown and commits to a prediction, two of the four ways of Section 1.4.10. A reservoir released against an inflow forecast, as in Volume 9, Chapter 7, and a power system committed against a load forecast have the same shape. A transfer function with a dead time represents none of it, because the controlled variables here are a running average and a probability.

The operational tools behind these decisions run on real catalogs and tracking data held by agencies and operators [9], and none of them is a plant on which a student can make the decisions and watch their consequences. This chapter casts orbit maintenance and collision avoidance as control problems whose controlled variables are predictions, in a browser simulator of a low Earth orbit satellite whose orbit filter sees only navigation fixes and whose catalog knows only its own tracking. The oblateness and drag models, the extended Kalman filter, the short-encounter collision probability and burn sizing on the linear relative motion are standard. Around them the simulator keeps a truth model apart from the operator's side, designs synthetic debris against the true trajectory with declared covariances that tracking shrinks, carries in the satellite's covariance the drag forecast and the execution errors of burns still to fly, and gives two planners that screen every burn against every object. Every setting the runs need is given, so that they can be reproduced. Eight cases and a four-day campaign meet forecasting problems and feedback problems one at a time. The debris is synthetic and the atmosphere illustrative, so nothing in the app is a live orbit forecast.

Section 8.2 reviews how drag, mean orbits and conjunction assessment came to be handled as they are. Section 8.3 gives the truth model of the satellite, the atmosphere and the synthetic debris, and Section 8.4 derives four consequences of that model for a planner: the oscillation, the decay, the response to a burn and the geometry of a pass. Section 8.5 gives the orbit filter, the screening and the collision probability, and Section 8.6 the maintenance and avoidance planners and the lower layer that flies their burns. Section 8.7 reports the eight cases and the campaign, Section 8.8 names the devices of real satellites, and Section 8.9 states the limits of the model.

8.2 Drag, Mean Orbits and Conjunction Assessment

Orbit decay turned the first satellites into instruments of the atmosphere they flew through. The contraction of the early orbits gave the density of the upper atmosphere at their heights and its variation from 1957 to 1961 [2]. The theory of a close satellite in the oblate field of the Earth was worked out in the same years, with the perturbations of the six elements expressed as functions of mean orbital elements and time, periodic terms of the first order and secular terms to the second [1]. The secular turning of the node is what makes a sun-synchronous orbit possible, and the split between periodic and secular motion is why maintenance is flown on mean elements. Methods of orbit determination and prediction, including the force models and the filters used here, are collected in [10].

> ⧖ HISTORY **Hall Thrusters Go to Orbit (1971 to 2019)**
>
> The first pair of Hall thrusters flew on a Soviet spacecraft in 1971, and more than 300 others flew between 1971 and 2018. The first built outside Russia flew in 2003, and the first built in the United States in 2006. Launches grew sharply in 2019, when two large constellations began to carry a Hall thruster on every satellite. The reasons given are a simple design, a reliability proven in flight, good efficiency, a high specific impulse and a high thrust for the power drawn. Every Hall thruster flown had run on xenon, the propellant of this chapter's satellite, until one of those constellations turned to krypton, which is far cheaper in the quantities that thousands of satellites need [11].

Density stayed the weak link. Thermospheric mass density, its measurement and its models are reviewed in [3]. In February 2022 a minor to moderate geomagnetic storm raised the density between 200 and 400 km by 50 to 125 %, and 38 of 49 newly launched satellites were lost, while pre-launch analyses with an empirical density model had underestimated the effect [12]. Station keeping has also moved on board. One study of a remote-sensing spacecraft at low altitude used a low-power Hall thruster for continuous in-track control, an extended Kalman filter with satellite navigation among its measurements, and a Lyapunov-based orbit law [13], an arrangement close to the one simulated here except that the app plans discrete burns.

The debris problem was framed in 1978 as a growth process, in which each collision produces fragments that raise the probability of the next, so that a belt of debris could form unless launch constraints and operational procedures delayed it [14]. The first collision between satellites came on 10 February 2009, when the active Iridium 33 met the retired Cosmos 2251, and by 1 December 2009 the U.S. tracking system had catalogued 1632 fragments from it [15]. Avoidance moved from distances to probabilities over the same decades. The Space Shuttle maneuvered when a conjunction fell inside a box extending 5 km along track and 2 km radially on either side of its predicted position, and the space station was designed to maneuver on a probability computed in the conjunction plane [4]. The theory of short encounters behind that probability is collected in [5], and an operational service for the collision risks of a fleet of satellites processes tracking data and updates ephemerides, covariances, conjunction geometries and risk estimates [9].

The weak points of that chain are the ones the app exposes. Uncertainty in the atmospheric density is usually left out when a covariance is carried to the time of closest approach, and it can change the probability of collision significantly, as the effect of a geomagnetic storm illustrates [16]. A covariance that is too large dilutes the probability, so that worse data appear to lower the risk, a symptom analyzed as false confidence [7]. The burn follows from a linear relation between the impulse and the relative motion in the encounter plane, which turns its design into the search for the least velocity change that meets a miss distance or a probability [17]. The linear relative motion is the terminal-guidance model of 1960 [6].

> ⧖ HISTORY **A Catalog Rebuilt for Probabilities (1986 to 2009)**
>
> The prediction of close approaches between objects in orbit began in support of the Space Shuttle's return to flight after the Challenger accident of 1986, drawing on a military network that tracked objects in orbit for missile defense. Its catalog of element sets covers tracked objects larger than about 10 cm in low orbit, but the analytic theory behind it gives no usable covariance. To compute a probability of collision for the space station, a more accurate catalog was built on special perturbations theory, with the covariances that the computation needs. Screening against that catalog has since grown, especially after a collision of two satellites in 2009, to cover every operational satellite in orbit [18].

8.3 Satellite, Atmosphere and Debris

The truth model carries the satellite in Earth-centered inertial coordinates with seven states, its position $\mathbf{r}$, its velocity $\mathbf{v}$ and the xenon m_p left in the tank. The orbit, the satellite's radial, along-track and cross-track axes $(\hat{\mathbf{R}}, \hat{\mathbf{T}}, \hat{\mathbf{N}})$, the forces, the navigation fix, the crossing geometry of a pass and its encounter plane are shown in Figure 8.1. The engine's clock runs in hours, so every rate below is multiplied by 3600 s/h in the code.

Table 8.1 assigns each signal its role. The controlled quantities are not states. The mean altitude is an average of the state over the last orbit, and the collision probability and the miss distance are predictions made for each pass.

8.3.1 Equations of Motion

Gravity is the central term plus the J_2 term of the oblateness, drag acts against the velocity relative to an atmosphere that turns with the Earth, and the thruster pushes along a direction fixed in the satellite axes, as shown in Equation (8.1).

$$
\begin{aligned}
\ddot{\mathbf{r}} &= -\frac{\mu}{r^3}\mathbf{r} - \frac{3 J_2 \mu R_E^2}{2 r^5}\begin{bmatrix} x(1-5z^2/r^2) \\ y(1-5z^2/r^2) \\ z(3-5z^2/r^2) \end{bmatrix} + \mathbf{a}_D + \mathbf{a}_T \\
\mathbf{a}_D &= -\tfrac{1}{2}\rho B |\mathbf{v}_r| \mathbf{v}_r, \qquad \mathbf{v}_r = \mathbf{v} - \boldsymbol{\omega}_E \times \mathbf{r} \\
\mathbf{a}_T &= \frac{\eta g_b F \ell}{m_{dry} + m_p}\left(\hat{u}_R \hat{\mathbf{R}} + \hat{u}_T \hat{\mathbf{T}} + \hat{u}_N \hat{\mathbf{N}}\right)
\end{aligned}
\tag{8.1}
$$

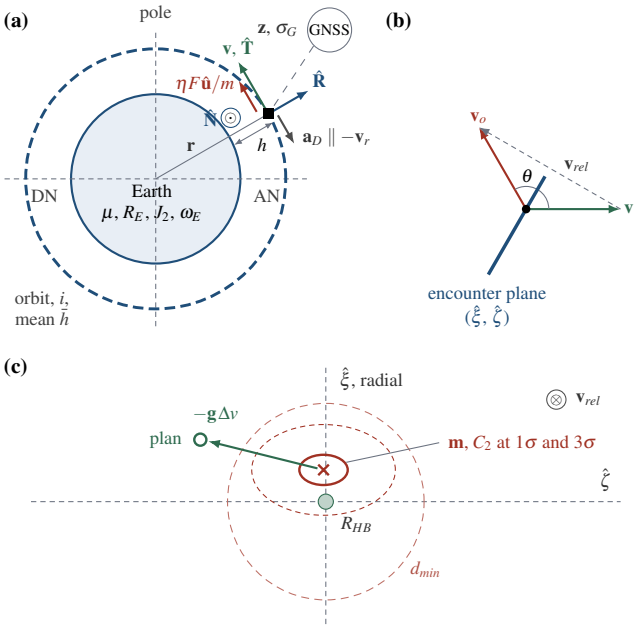

Figure 8.1 The simulated satellite, drawn out of scale. (a) The orbit plane with $\hat{\mathbf{N}}$ out of the page and the height exaggerated as in the app's orbit view, showing the equator through the ascending and descending nodes, the position $\mathbf{r}$ and altitude h, the satellite axes, the thrust along $\hat{\mathbf{u}}$, the drag against the wind relative to the turning atmosphere and the navigation fix $\mathbf{z}$. (b) A pass in the local horizontal plane, $\hat{\mathbf{R}}$ out of the page. The object crosses at the angle θ with the same speed, and the encounter plane is normal to the relative velocity. (c) The encounter plane at the time of closest approach (TCA), centered on the satellite and its combined hard-body circle, with the predicted miss $\mathbf{m}$ inside its combined covariance C_2, the miss-distance floor and the position a planned burn gives.

Table 8.1 Variables of the Satellite Plant by Role

Symbol	Role	Default	Range
$\bar{h}$	CV, true mean altitude	500.000 km	$\bar{h}_{SP} \pm \Delta h$
$\hat{\bar{h}}$	filter mean, planned on	—	—
P_c, d	CVs per pass	—	$\le P_{c,max}$, $\ge d_{min}$
$\bar{h}_{SP}$	band center	500 km	480–520 km
Δh	band half-width	250 m	0.05–2 km
u	MV in manual, thrust	0 %	±100 % of F
burns	MV in automatic	—	$\Delta v \le \Delta v_{max}$
s	DV, solar density scale	1	0–5
f_{st}	DV event, storm	1	peak 3
η	DV, thruster output	100 %	20–100 %
objects	DV, synthetic debris	per case	Table 8.6
$\mathbf{r}$, $\mathbf{v}$	states, inertial	mean circular	—
m_p	state, xenon	2.0 kg	0.2–10 kg
$\hat{\mathbf{x}}$	filter states	$\hat{k} = \hat{\eta} = 1$	clamped

where μ is the gravitational parameter, R_E the equatorial radius, ω_E the rotation of the Earth about its axis, $B = C_D A/(m_{dry} + m_p)$ the drag factor, F the rated thrust, $\ell \in [0,1]$ the share of the step the thruster fires, η its true output, g_b the magnitude factor of the burn and $\hat{\mathbf{u}}$ its direction with the pointing error included. The axes follow from the state, $\hat{\mathbf{R}} = \mathbf{r}/r$, $\hat{\mathbf{N}} = \mathbf{r} \times \mathbf{v}/|\mathbf{r} \times \mathbf{v}|$ and $\hat{\mathbf{T}} = \hat{\mathbf{N}} \times \hat{\mathbf{R}}$.

The thrust stops when the tank is empty. Integration is fixed-step fourth-order Runge-Kutta at 10 s. Halving the step moves the position after one day by 0.43 m, and an orbit under gravity alone keeps its energy to a relative 4.4×10^{-10} over a day.

8.3.2 Atmosphere and Storm

The density is a piecewise-exponential table after CIRA-72 for moderate solar activity, evaluated at the height above the ellipsoid with a geocentric latitude. It is scaled by the true solar-activity factor s and a storm profile, as shown in Equation (8.2).

$$\begin{aligned} \rho &= s\, f_{st}(t)\, \rho_j\, \mathrm{e}^{-(h_e - h_j)/H_j}, \qquad h_j \le h_e < h_{j+1} \\ h_e &= r - R_E\left(1 - f\, z^2/r^2\right) \end{aligned} \tag{8.2}$$

where f is the flattening and (h_j, ρ_j, H_j) the base height, base density and scale height of band j in Table 8.3. The storm multiplier rises linearly from 1 to its peak $p = 3$ over 2 h, holds for 12 h and returns over 6 h. The operator never sees s or f_{st}. Their only trace is in the orbit.

8.3.3 Thruster, Propellant and Battery

Planned burns fire the Hall thruster at its rated thrust, and in manual mode the operator's slider sets a share of it. The xenon flow follows the command, not the thrust delivered, and the battery pays for the power the thruster draws beyond the surplus of the solar arrays, as shown in Equation (8.3).

$$\dot{m}_p = -\frac{F\,\ell}{I_{sp}\, g_0}, \qquad \dot{E} = \begin{cases} -(P_{thr} - P_{sur})\ell & \text{firing} \\ P_{sur} & \text{idle} \end{cases} \tag{8.3}$$

with E clamped to $[0, E_{bat}]$. At 1400 s the flow is 1.46 mg/s, so each 1 cm/s of the 120 kg satellite costs 87.4 mg of xenon. A full battery allows 1.0 h of continuous firing, and the arrays sustain a duty cycle of 43 %. A thruster at $\eta = 0.55$ spends the same xenon and delivers 55 % of the change.

8.3.4 Mean Elements and the Initial State

The mean altitude and the mean eccentricity vector are running averages of the osculating radius and eccentricity vector over one period T_0 of the initial mean orbit, sampled every step, as shown in Equation (8.4).

$$\bar{h} = \frac{1}{L}\left(\sum_{j=0}^{N-1} |\mathbf{r}_{k-j}| + \phi\, |\mathbf{r}_{k-N}|\right) - R_E, \qquad L = \frac{T_0}{\Delta t} = N + \phi \tag{8.4}$$

where $\Delta t = 10$ s and the oldest sample carries the fractional weight ϕ. The mean eccentricity $\bar{e}$ is the magnitude of the averaged vector, and the mean perigee and apogee are $\bar{h} \mp (\bar{h} + R_E)\bar{e}$. The average removes the oscillation that J_2 puts into every element and lags half an orbit behind a burn. The initial state sits at the ascending node and is found by Newton's method on the same average, so that after one orbit of gravity alone the mean radius is $R_E + \bar{h}_0$ and the mean eccentricity vector is zero. The history before $t = 0$ is filled by coasting that state backward for one orbit. Starting instead on an osculating circle of radius $R_E + 500$ km would leave the mean altitude 4.87 km low and the mean eccentricity at 4.8×10^{-4}.

8.3.5 Synthetic Debris and Tracking

No real object or catalog is represented. Each object is created at its alert against the satellite's true trajectory, propagated from that moment with the truth model and the burns already committed. At the design TCA the object's velocity is the satellite's turned about $\hat{\mathbf{R}}$ by the crossing angle θ, so the relative velocity is horizontal, and its catalog position sits at the declared miss vector in the encounter plane. The true position differs from the catalog position by an offset δ drawn from the final catalog covariance, or from a tighter spread when the case declares one, in the object's own radial, along-track

Table 8.2 Satellite and Environment Parameters, All Shown and Adjustable in the Parameter Panel

Symbol	Meaning	Value	Range
m_{dry}, $m_{p,0}$	dry mass, xenon	118, 2.0 kg	50–500, 0.2–10
A, C_D	drag area, coefficient	2.4 m^2, 2.2	0.5–10, 1.5–3
F, I_{sp}	thrust, specific impulse	20 mN, 1400 s	5–100, 300–3000
P_{thr}, P_{sur}	power, array surplus	350, 150 W	100–1500, 20–1000
E_{bat}	battery for thrusting	200 Wh	20–2000
i	inclination	97.4°	0–180
$\bar{h}_0$	initial mean altitude	500 km	400–700
R_{HB}	combined hard-body radius	15 m	2–50
σ_G	fix noise per axis	5 m	0.5–50
σ_θ, σ_F	pointing, magnitude error	1°, 2 %	0–10, 0–20
t_{slew}	slew before a burn	5 min	0–30
σ_{drag}	drag forecast error	10 %	0–50
m_{alloc}	xenon allocated	30 g	1–1000

Table 8.3 Fixed Constants and the Density Bands, in the Source Only

$\mu = 3.986\,004\,418 \times 10^{14}\,m^3/s^2$	$R_E = 6378.137\,km$
$J_2 = 1.082\,626\,68 \times 10^{-3}$	$f = 1/298.257\,223\,563$
$\omega_E = 7.292\,115\,9 \times 10^{-5}\,rad/s$	$g_0 = 9.806\,65\,m/s^2$

h_j (km)	ρ_j (kg/m^3)	H_j (km)	h_j (km)	ρ_j (kg/m^3)	H_j (km)
300	2.418×10^{-11}	53.628	600	1.454×10^{-13}	71.835
350	9.518×10^{-12}	53.298	700	3.614×10^{-14}	88.667
400	3.725×10^{-12}	58.515	800	1.170×10^{-14}	124.64
450	1.585×10^{-12}	60.828	900	5.245×10^{-15}	181.05
500	6.967×10^{-13}	63.822	1000	3.019×10^{-15}	268.00

and cross-track axes. Both then move on two-body orbits from their states at TCA. The catalog gives the operator a mean offset $\hat{\delta}$ and a diagonal standard deviation σ per axis. Tracking updates, evenly spaced from the alert to a last update before TCA (one hour by default), shrink σ geometrically from σ_0 to σ_F through Kalman updates on a measurement of the true offset, as shown in Equation (8.5).

$$\begin{aligned} \sigma_k &= \sigma_0 \left(\sigma_F/\sigma_0\right)^{k/N_u}, \qquad r_k^2 = \left(\sigma_k^{-2} - \sigma_{k-1}^{-2}\right)^{-1} \\ \hat{\delta}_k &= \sigma_k^2 \left(\frac{\hat{\delta}_{k-1}}{\sigma_{k-1}^2} + \frac{\delta + r_k \varepsilon_k}{r_k^2} \right) \end{aligned} \tag{8.5}$$

where N_u counts the tracking updates and ε_k is a standard normal deviate. The operator sees $\hat{\delta}$ and σ and never δ. The declarations of every object are listed in Table 8.6.

8.3.6 Numbers

The parameters are listed in Table 8.2 and the constants and the density bands in Table 8.3. The dry mass and the xenon give $B = 0.044\,m^2/kg$, and 97.4° is the sun-synchronous inclination at 500 km. Every case starts from the mean circular state at $\bar{h}_0$ with the node at zero right ascension, a full battery, the tank at $m_{p,0}$ and the random streams seeded from one fixed seed, so every run replays bit for bit (Section 1.6.1).

8.3.7 Assumptions

Gravity stops at J_2. Higher harmonics, the Sun, the Moon and radiation pressure drive eccentricity and inclination over weeks and are absent, so multi-week eccentricity control and ground-track keeping

are outside the model.

The atmosphere is a table times two scales, illustrative rather than a forecast. Real density at these heights follows solar and geomagnetic activity [3], and the storm of February 2022 raised it by 50 to 125 % between 200 and 400 km [12], less than the tripling case 4 imposes. The height uses a geocentric latitude, within 25 m of the geodetic height, which moves the density by 0.04 %.

The satellite is a point mass with a constant drag area. Attitude, the slew itself and thermal limits are reduced to a fixed slew time and a pointing error drawn once per burn, so a burn never fails to start. The attitude of a small Earth-observation satellite is the subject of Chapter 7.

The debris is synthetic and moves on two-body orbits, with its own drag and J_2 folded into declared covariances. The design places every pass where the case wants it, so the geometry of a pass is chosen, not found by a catalog-wide screening.

Navigation uses white position fixes with no outages or clock errors, while an operational orbit determination system also processes velocity or range measurements with a richer force model [10].

8.4 The Orbit the Planners See

Four consequences of Equation (8.1) decide what a planner can do. The oscillation comes first. In case 1, with no air and no thrust, the altitude $|\mathbf{r}| - R_E$ swings between 498.42 and 501.58 km, the height above the ellipsoid between 501.6 and 519.5 km, and the osculating semimajor axis between 486.0 and 504.9 km above R_E, twice an orbit. The mean altitude stays within 0.98 m of 500 km for a day and the mean eccentricity at most 6.2×10^{-7}. The one-orbit average of the semimajor axis sits 4.56 km below the mean radius, so the app defines the mean on the radius, the quantity its altitude band refers to. The node turns at the rate of Equation (8.6).

$$\dot{\Omega} = -\tfrac{3}{2} n J_2 \left(\frac{R_E}{a} \right)^2 \cos i \tag{8.6}$$

A straight-line fit to the node over three days gives 0.9875° a day. Equation (8.6) gives 0.9854° with a at the mean radius and 0.9877° at the mean semimajor axis, within 0.02 % of the simulation, and the Sun appears to move 0.9856° a day.

> **TRY IT IN THE APP** **1 · Unforced Coast**
>
> Load the case, press Run and stay on the Orbit view. Watch the Altitude now trend swing up and down twice every orbit, between 498.42 and 501.58 km, while the Mean altitude trend stays flat. The chapter reports the mean within 0.98 m of 500 km for a day. Nothing pushes or drags this satellite, and the swing comes from the bulge of the Earth.

The second is the decay. Averaging the drag of Equation (8.1) over a near-circular orbit gives the fall of the mean semimajor axis in Equation (8.7), which is also the forecast the maintenance planner uses.

$$\frac{\mathrm{d}\bar{a}}{\mathrm{d}t} = -k\,\hat{\rho}_s\,\langle \rho \rangle \sqrt{\mu \bar{a}}\, B \left(1 - \frac{\omega_E \bar{a} \cos i}{v} \right)^2 \tag{8.7}$$

where $v = \sqrt{\mu/\bar{a}}$ and $\langle \rho \rangle$ is the table density averaged over 24 points around the orbit at the height above the ellipsoid. In the planner $k = \hat{k}$ and $\hat{\rho}_s$ is the operator's density scale, and with both at 1 the law gives the decay in the default atmosphere. At the tabulated density for exactly 500 km, $\rho\sqrt{\mu a} B$ gives 138.7 m a day. The orbit flies 1.6 to 19.5 km higher above the ellipsoid, where the air is thinner, and the average gives 120.4 m a day, while case 2 loses 115.4, 121.2 and 120.3 m on its first three days. The same law gives 270.9 and 630.3 m a day at 450 and 400 km.

The third is the burn. An along-track change Δv changes the mean semimajor axis and, for a burn made τ before a time of interest, the radial and along-track position at that time as the linear relative

motion [6] gives in Equation (8.8).

$$\Delta a = \frac{2a^2 v \Delta v}{\mu}, \qquad \begin{aligned} x &= \tfrac{2\Delta v}{n}\,(1-\cos n\tau) \\ y &= \tfrac{\Delta v}{n}\,(4\sin n\tau - 3n\tau) \end{aligned} \tag{8.8}$$

where $n = 1.1068 \times 10^{-3}\,\mathrm{rad/s}$ is the mean motion and the period is 94.62 min. At 500 km one m/s raises the mean 1.807 km. The point where a prograde burn fires stays where it was, and the far side rises by $2\Delta a$, so a single reboost makes the orbit eccentric. The secular term $-3\Delta v\,\tau$ is the along-track drift that avoidance buys, 10.8 m per hour for each mm/s.

Pause. Try Exercise 8.3 before reading on.

The fourth is the geometry of a pass. Two objects of equal speed crossing at the angle θ have a relative velocity that makes the angle $90^\circ - \theta/2$ with each along-track axis, so an along-track error enters the encounter plane multiplied by the factor of Equation (8.9).

$$\left|\hat{\mathbf{T}} \times \hat{\mathbf{v}}_{rel}\right| = \cos\frac{\theta}{2} \tag{8.9}$$

At 160° the factor is 0.174, so a nearly head-on pass hides most of an along-track error in the timing of the pass. At 100° it is 0.643 and the along-track error lands in the plane. The same projection applies to the along-track part of an avoidance burn, while its radial part enters whole.

8.5 Navigation, Screening and Collision Probability

8.5.1 Orbit Filter

The operator knows the orbit through an extended Kalman filter with eight states, $\hat{\mathbf{x}} = [\hat{\mathbf{r}}, \hat{\mathbf{v}}, \hat{k}, \hat{\eta}]$, where $\hat{k}$ scales the drag the operator's model predicts and $\hat{\eta}$ the thrust the lower layer commanded. The prediction integrates the operator's model with the same step, with J_2, drag at $\hat{k}\hat{\rho}_s$ times the table and the commanded thrust times $\hat{\eta}$. The covariance follows Equation (8.10).

$$\begin{aligned} &\Phi = I + A\Delta t + \tfrac{1}{2}(A\Delta t)^2, \qquad P \leftarrow \Phi P \Phi^{\top} + Q \\ &K = PH^{\top}(HPH^{\top} + \sigma_G^2 I)^{-1}, \qquad \hat{\mathbf{x}} \leftarrow \hat{\mathbf{x}} + K(\mathbf{z} - H\hat{\mathbf{x}}) \end{aligned} \tag{8.10}$$

where A holds the two-body gravity gradient $\mu r^{-3}(3\hat{\mathbf{R}}\hat{\mathbf{R}}^{\top} - I)$ and the sensitivities of the acceleration to $\hat{k}$ and $\hat{\eta}$, $H = [I_3\ 0]$ and $\mathbf{z} = \mathbf{r} + \varepsilon$ is a fix every 10 s with 5 m of white noise per axis. The covariance update uses the Joseph form, and $\hat{k}$ and $\hat{\eta}$ are clamped to $[0.2, 6]$ and $[0.2, 1.3]$ after each update. The process noise adds $q_a^2 \Delta t$ to each velocity variance, with $q_a = 2 \times 10^{-9}\,\mathrm{m/s^{1.5}}$, and while a burn fires another $60\Delta t\,(0.05 F\ell/\hat{m})^2$. It adds $q_k^2 \Delta t$ to the drag scale with $q_k = 1 \times 10^{-3}\,\mathrm{s^{-1/2}}$ and $q_\eta^2 \Delta t$ to the efficiency with $q_\eta = 1 \times 10^{-6}\,\mathrm{s^{-1/2}}$. The initial estimate is the truth plus random errors with standard deviations of 20 m and 0.02 m/s per axis, and the covariance starts diagonal at 20, 0.02, 0.3 and 0.1. The filter's own mean $\hat{\bar{h}}$ is the average of Equation (8.4) applied to $\hat{\mathbf{r}}$, and it is what the maintenance planner works on. Over hours 24 to 72 of case 2 the filter holds the position to 0.53 m and the velocity to 0.58 mm/s RMS while the fixes scatter 8.66 m RMS in three dimensions, and $\hat{k}$ stays between 0.92 and 1.07 against a true 1. The filter is the variance-chosen form of the fusion filters of Section 1.4.6, with the drag scale and the thrust efficiency as bias states. Volume 5, Chapter 13 steers an underwater vehicle on the same kind of filter, with fixes that arrive late.

8.5.2 Screening and the Time of Closest Approach

Every 30 min while a pass is pending, at every alert and tracking update and 15 min after every burn, the planner propagates the filter's estimate with the operator's model and the burns already committed. This operational trajectory runs to just past the last pass due within 30 h. For each catalogued object the closest approach is found on the 10 s grid within 2 min of the catalog TCA and refined on

straight-line relative motion. The encounter frame takes $\hat{\mathbf{e}}_z = \mathbf{v}_{rel}/|\mathbf{v}_{rel}|$, the satellite's radial direction projected into the plane as $\hat{\xi}$, and $\hat{\zeta} = \hat{\mathbf{e}}_z \times \hat{\xi}$, and the miss vector $\mathbf{m}$ is the object's position relative to the satellite resolved on those two axes. The last assessment 10 min before TCA is the one each pass is scored on.

8.5.3 The Satellite's Own Covariance at TCA

The filter covariance is rotated into the satellite's rotating axes and carried to TCA by the linear state transition $\Phi_{CW}(\tau)$ of the relative motion [6], and two terms are added, as shown in Equation (8.11).

$$C_{own} = \left[\Phi_{CW} P_H \Phi_{CW}^{\top}\right]_{rr} + \sigma_\kappa^2 \hat{a}_D^2 \, \mathbf{g}_a \mathbf{g}_a^{\top} + \sum_b G_b S_b G_b^{\top} \tag{8.11}$$

where $\mathbf{g}_a$ is the response at TCA to a unit along-track acceleration held for τ, $\hat{a}_D$ the modeled drag, and $\sigma_\kappa = (P_{kk}/\hat{k}^2 + \sigma_{drag}^2)^{1/2}$ the relative drag uncertainty of the filter and the forecast together. The sum, used only by the uncertainty-aware planner, carries the execution error of each burn still to fly through the impulse responses G_b, with S_b holding $(\sigma_F \Delta v)^2$ along the thrust, the filter's efficiency variance added, and $(\sigma_\theta \Delta v)^2$ across it. The drag term dominates early. At the alert of case 5, 9.7 h before TCA, it gives 407 m of along-track uncertainty at TCA, and at 4.2 h it gives 59 m. The combined covariance adds the catalog's diagonal covariance in the object's axes and is projected onto the plane, $C_2 = [\hat{\xi}\ \hat{\zeta}]^{\top}(C_{own} + C_{obj})[\hat{\xi}\ \hat{\zeta}]$.

8.5.4 Collision Probability in the Encounter Plane

The probability follows the short-encounter model [4, 5]. At 10.8 to 15.0 km/s the 30 m hard-body diameter is crossed in 2 to 3 ms. The relative motion during the encounter is therefore a straight line at the relative velocity found at TCA, the position errors of both objects are Gaussian and frozen while it lasts, velocity errors are not carried, and the two errors are independent so that their covariances add. Both objects are spheres whose radii add to R_{HB}, and the probability is the mass of a two-dimensional Gaussian inside a disc, as shown in Equation (8.12).

$$P_c = \iint_{|\mathbf{x}| \le R_{HB}} \frac{\exp\!\left(-\tfrac{1}{2}(\mathbf{x}-\mathbf{m})^{\top} C_2^{-1} (\mathbf{x}-\mathbf{m})\right)}{2\pi\sqrt{|C_2|}} \, \mathrm{d}\mathbf{x} \tag{8.12}$$

The integral is a polar quadrature with 12 Gauss-Legendre points in radius on one interval and 24 midpoint angles, refined to three intervals and 48 angles when the smallest standard deviation of C_2 is below R_{HB}. For a zero miss and an isotropic σ it reproduces $1 - \exp(-R_{HB}^2/2\sigma^2)$ to 10^{-9} for σ from 5 to 1000 m, and for an offset, correlated case it agrees with a 1500×1500 grid within 10^{-4}, two of the application's 43 passing checks. The model stops being valid for slow, co-orbital encounters that last minutes, where the relative motion curves.

> **+ DID YOU KNOW? Ten Kilometers to Be Sure**
>
> In a NASA comparison of conjunctions, the passes with a collision probability above 1 in 10,000 had predicted misses from under 100 m to over 10 km. A rule on the miss alone would need a 10 km limit to catch 95 % of them, and it would also flag more than half of the harmless passes [18].

8.6 Implemented Planners and Lower Layer

No feedback loop closes on the altitude, and this application does not instantiate the framework's velocity-form PID controller of Section 1.6.2. An altitude loop would fire twice an orbit against the J_2 oscillation, and a miss distance is not a quantity a loop can regulate. Every burn is instead decided by one of two planners and executed by a lower layer, all on the operator's side.

8.6.1 Timing

The filter and the lower layer run every 10 s, one physics step (Section 1.6.1), and the order within a step is the controller, the integration, then the post-step updates of the battery, the true mean elements and the true closest approaches. The conjunction planner runs at the screening times above, and the maintenance planner at its decision interval T_{plan}, 6 h by default. One sample per simulated minute goes to the historian, and the app runs from 0.05 to 24 simulated hours per second, one by default.

8.6.2 Maintenance Planner

At each decision the planner reads the filter's mean $\hat{\bar{h}}$ against the band $\bar{h}_{SP} \pm \Delta h$. Above the band it lowers the orbit to $\bar{h}_{SP}$. The deadband mode reboosts to $\bar{h}_{SP} + 0.8\Delta h$ only when $\hat{\bar{h}}$ is already below the band. The predictive mode integrates Equation (8.7) over the horizon H_f in 15 min steps, with $k = \hat{k}$ when the drag scale is estimated. It reboosts to the same target when the forecast crosses $\bar{h}_{SP} - 0.9\Delta h$ within the lead of Equation (8.13).

$$t_{lead} = 2T + \frac{\Delta v}{a_{max}} + t_{slew} + T_{plan}, \qquad \Delta v = \frac{v\,\Delta a}{2a} \tag{8.13}$$

where T is the period of the estimated mean orbit and $a_{max} = \hat{\eta} F / \hat{m}$ when burns are sized with the estimated efficiency. A reboost starts one slew time and one minute after the decision, and with two burns selected each half fires half an orbit apart. The planner then holds for 1.1 orbits after the last burn so the mean settles, and it stays silent while an avoidance or its return is pending.

8.6.3 Avoidance Planner

The planner acts on one pass at a time, the first in the catalog that needs action. The nominal mode acts when the predicted miss is under the floor d_{min} and sizes the burn on the miss alone. The uncertainty-aware mode acts when $P_c > P_{c,max}$ or the miss is under the floor, plans to the target $P^* = P_{c,max}/k_P$, and keeps the floor with $1.5\sigma_{cat}$ in hand, where σ_{cat} is the largest standard deviation of the catalog covariance in the plane. Candidate burns, prograde or retrograde and radial only when allowed, are centered every sixteenth of an orbit from 20 min before TCA back to the end of the slew or to $t_{lead,max}$ before TCA, whichever is later. A candidate is kept only if it starts after the slew and ends at least 20 min before TCA. For each one the linear response of Equation (8.8), averaged over the finite burn by Simpson's rule and resolved on $(\hat{\xi}, \hat{\zeta})$, gives the in-plane displacement $\mathbf{g}$ per unit Δv. With a small disc the probability is $P_c \approx K \exp(-d^2/2)$, where $K = R_{HB}^2/(2\sqrt{|C_2|})$ and d is the Mahalanobis distance of the miss. Both constraints are then quadratic in Δv, and the least velocity change follows in closed form, as shown in Equation (8.14).

$$\Delta v = \frac{b + \sqrt{b^2 - ac}}{a}, \qquad \begin{aligned} a &= \mathbf{g}^\top W \mathbf{g}, \quad b = \mathbf{g}^\top W \mathbf{m} \\ c &= \mathbf{m}^\top W \mathbf{m} - D^2 \end{aligned} \tag{8.14}$$

with $W = I$ and $D = d_{min} + 1.5\sigma_{cat}$ ($D = d_{min}$ in the nominal mode) for the floor, and $W = C_2^{-1}$ and $D^2 = 2\ln(K/P^*)$ for the probability, the larger of the two setting the burn. The cheapest candidate under Δv_{max} is chosen, and with re-screening on it must also leave every other screened object clear on the same linear model. The uncertainty-aware mode then refines its size by bisection on the exact P_c with the burn's own execution error included. Both modes verify the plan by propagating the full model with the finite burn, and a plan that falls short is scaled once, by 1.25 or in the nominal mode by the ratio of floor to miss, at most doubling it. The resulting cost against lead time is shown in Figure 8.2.

When nothing meets the criterion, the planner searches the longest full-thrust burns the clock still allows, up to 2 min before TCA, in both directions, and flies the best of them only if it at least halves the probability (or raises the miss by half in the nominal mode). The pass is then flagged infeasible with its residual probability, whether or not a burn flies.

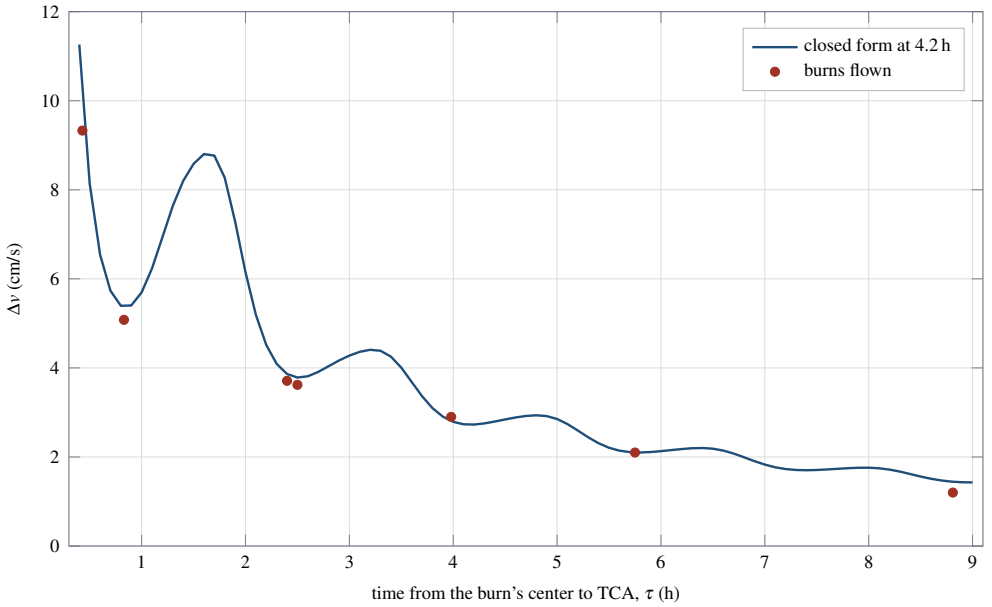

Figure 8.2 Least along-track velocity change that moves Fragment 1 of case 5 to the planner's floor of 272.7 m, against the time from the burn's center to TCA. The curve is Equation (8.14) on the case-5 prediction at 4.2 h, retrograde and averaged over the finite burn. The points are the burns flown with the earliest burn set from 0.5 to 12 h, each planned on the tracking available when it was committed.

8.6.4 Return, Re-Screening and Verification

A return burn, equal and opposite and a whole number of orbits after the avoidance, at least 15 min after TCA, restores the orbit. It is omitted when the avoidance was a prograde burn from below $\bar{h}_{SP}$ that leaves the mean under $\bar{h}_{SP} + 0.8\Delta h$, which is then a free reboost. With re-screening on, the plan and up to three placements of its return are checked against every screened object. Every return or reboost is also screened again on the full model when it is committed, and it is moved one orbit later, at most three times, if it would create a pass over the planning target or under the floor. With verification on, the uncertainty-aware mode re-plans 15 min after an avoidance on the orbit the filter now sees and flies a trim if the pass still needs one. The nominal mode fires and forgets. In every mode the return is resized, until it is committed, from what the filter believes the avoidance and any trim delivered.

8.6.5 Lower Layer and Manual Operation

A burn is committed one slew time before it starts, and its errors are drawn then, $g_b = 1 + \sigma_F \varepsilon$ and a pointing error of σ_θ on each of the two axes across the thrust. It fires at the rated thrust, the last step at the share that completes the impulse, and it pauses when the battery falls to 5 % of E_{bat} and resumes above 25 %. No burn fires once the xenon allocated to the run is spent. The lower layer counts commanded on-time and cannot see the thrust delivered. In manual mode the operator's along-track slider sets ℓ and the sign of the thrust, and two buttons schedule ±5 cm/s burns after the slew. The planners still assess every pass in manual mode, but they plan nothing, and automatic burns wait. A burn whose window passes before the return to automatic expires unflown, and both planners then plan again at once from the orbit as it is. Switching a planner off drops the burns it has planned but not committed.

Table 8.4 Planner and Filter Settings. Controls Are Adjustable, the Rest Fixed in the Source

Setting	Default	Range
maintenance planner	predictive	off, deadband
reboost as	two burns	one or two
horizon H_f, decision T_{plan}	24 h, 6 h	2–72, 1/6–24 h
avoidance planner	aware	off, nominal
$P_{c,max}$	10^{-4}	10^{-6}–10^{-2}
d_{min}, k_P	200 m, 3	0–2000, 1–10
Δv_{max}, $t_{lead,max}$	30 cm/s, 6 h	1–100, 0.5–24
re-screen, verify	on, on	switches
radial burns	off	advanced switch
size with $\hat{\eta}$, estimate $\hat{k}$	on, on	switches
$\hat{\rho}_s$, q_k	1, 0.06 $\mathrm{h}^{-1/2}$	0–5, 0–0.6
screening interval	30 min	fixed
final assessment	TCA − 10 min	fixed
last end of a planned burn	TCA − 20 min	fixed
floor margin, spacing	$1.5\sigma_{cat}$, $T/16$	fixed
battery pause, resume	5, 25 %	fixed
q_a ($\mathrm{m/s}^{1.5}$), q_η ($\mathrm{s}^{-1/2}$)	2×10^{-9}, 10^{-6}	fixed

Table 8.5 Scenario Sequence and the Changes From the Defaults

#	Scenario and change	Teaching point
1	coast, $s = 0$, $\hat{\rho}_s = 0$	mean against altitude
2	decay, both planners off	120 m a day
3	$\bar{h}_0 = 499.80$ km, deadband every 10 min, one burn	the far side rises
4	$\bar{h}_0 = 499.95$ km, deadband, storm at 6 h	a storm between decisions
5	Fragment 1, late fragment	lead time buys the pass
6	rocket body, nominal planner	distance is not risk
7	Fragment 2, $\eta = 55$ %, verification off	only the orbit shows
8	Fragment 3, Satellite B, re-screening off	the return makes a pass
9	96 h campaign, 8 g allocated	graded challenge

8.6.6 Truth, Operator and Analysis Models

The truth model is Equations (8.1)–(8.5) with the true density scale, the storm, the true thruster output and the true debris. The operator's model is the same force model with $\hat{k}\hat{\rho}_s$ for the density, the commanded thrust times $\hat{\eta}$ along the nominal direction, and the catalog, plus the linear relative motion for sizing burns and carrying covariances and Equation (8.7) for the maintenance forecast. The analysis model is Equations (8.6), (8.8) and (8.9), which this chapter uses to interpret the runs. The settings of the planners and the filter are listed in Table 8.4.

8.7 Cases and Measured Outcomes

Table 8.5 lists the eight cases and the campaign, and Table 8.6 their objects. Each case starts from the defaults of Table 8.4 except for the changes in Table 8.5. The numbers below come from headless engine runs at the default seed, in hours of simulated time, and a configuration gives one run whether it is set before the run starts or carried through Reset (Section 1.6.4). A pass counts as resolved when its final assessment gives $P_c \leq 10^{-4}$ and a predicted miss of at least 200 m. These thresholds, like the 250 m band, are the app's design choices, and Table 8.7 collects the outcomes.

Table 8.6 Synthetic Objects. Miss Vector (ξ, ζ) in Meters, Standard Deviations in the Object's Radial, Along-Track and Cross-Track Axes

Case, object	Alert, TCA (h)	θ (°)	Miss	σ_0 (m)	σ_F (m)	N_u
5, Fragment 1	0.5, 10.2	160	60, −40	40/300/40	15/80/15	4
5, late fragment	20, 20.42	150	−15, 8	80/400/80	60/250/60	1[a]
6, rocket body	0.5, 14.1	100	120, 700	300/3000/300	150/1500/150[b]	4
7, Fragment 2	0.5, 10.2	160	60, −40	40/300/40	15/80/15	4
8, Fragment 3	0.5, 10.2	160	60, −40	40/300/40	15/80/15	3[c]
8, Satellite B	4.86, 15.39[d]	90	3, 2	6/30/6	4/20/4[e]	3
9, Fragment 1	0.5, 10.0	160	60, −40	40/300/40	15/80/15	4
9, rocket body	38, 52.0	100	120, 700	300/3000/300	150/1500/150[b]	4
9, Fragment 2	62, 72.0	150	−50, 30	40/300/40	15/80/15	4

[a] last update 0.25 h before TCA. [b] true offset drawn from 20/100/20. [c] last update 6.2 h before TCA. [d] catalogued 0.5 h after Fragment 3's avoidance and placed three orbits after its planned return. [e] true offset drawn from 2/5/2.

Table 8.7 Results of the Cases in Headless Runs at Seed 20261102

Case	Measured outcome
1	$\lvert\mathbf{r}\rvert - R_E$ 498.42–501.58 km, $\bar{h}$ within 0.98 m for 24 h
2	$\bar{h}$ 499.885, 499.763, 499.643 km at 24, 48, 72 h, out at 50.79 h
3, one burn	24.91 cm/s at 10.94 h, perigee 499.770, apogee 500.575 km
3, two burns	2×12.45 cm/s, perigee 500.176, apogee 500.179 km
4, deadband	lowest 499.677 km, 6.08 h outside, 2.52 g
4, predictive	reboost 12.10 h, lowest 499.848 km, none outside, 1.68 g
5, Fragment 1	2.10 cm/s at 4.43 h, final 245 m, true 278 m
5, late, best effort	18.11 cm/s, P_c $1.6 \times 10^{-2} \to 1.1 \times 10^{-3}$, infeasible
5, nominal	1.46 cm/s, final 171 m (true 204), unresolved
6, nominal	no burn, final 1622 m, $P_c = 1.8 \times 10^{-4}$, true 807 m
6, aware	3.66 cm/s at 8.12 h, final 3209 m, $P_c = 2.5 \times 10^{-6}$
7, no verification	final 138 m (true 171), unresolved, 4.20 cm/s
7, verification	trim 1.18 cm/s, final 232 m (true 265), 7.96 cm/s
8, unscreened	B at 3.1 m, $P_c = 0.23$, then 0.62 cm/s for B, 4.35 cm/s
8, re-screened	return one orbit later, B at 216 m, 3.73 cm/s

8.7.1 The Mean Orbit and the Decay

The coast of case 1 and the decay of case 2 show the plant before any planner acts. The first shows the 1.6 km oscillation of the altitude against a mean that holds to 0.98 m, and the second a mean that falls to 499.885, 499.763 and 499.643 km at 24, 48 and 72 h and leaves its band at 50.79 h. The decay is nearly steady because 360 m is a small part of the 64 km scale height at 500 km. A controller chasing the altitude of case 1 would fire twice an orbit to cancel an oscillation that costs nothing.

8.7.2 A Reboost and a Storm

Case 3 starts 50 m above the bottom of the band with a deadband deciding every ten minutes. The true mean crosses 499.75 km at 10.84 h, and the decision at the start of that 10 s step already finds the filter's mean 2 cm below the edge, so 24.91 cm/s fires from 10.94 h for 24.9 min, 2.18 g of xenon. At 16 h the mean stands at 500.173 km, but the perigee stays at 499.770 km, where the burn fired, and the apogee has risen to 500.575 km, 325 m above the band. Split into two burns of 12.45 cm/s at 10.94 and 11.72 h, the reboost leaves the perigee at 500.176 km and the apogee at 500.179 km, an eccentricity of 2.6×10^{-7}. The mean spends 0.36 h below the band either way, because the burn starts a slew time and a minute after the decision and the mean lags half an orbit behind it. The filter's mean is so close to the edge at that decision that on five of six other seeds the reboost waits for the 11.00 h decision and the mean spends 0.53 to 0.54 h below the band.

> **TRY IT IN THE APP 3 · Scheduled Reboost**
>
> Load the case and run to 16 h. On the Mean-orbit perigee and apogee trend the single reboost leaves the perigee at 499.770 km, where the burn fired, and lifts the apogee to 500.575 km. Press Reset, set Reboost as to Two burns before pressing Run, and run again: the perigee and apogee end at 500.176 and 500.179 km.

Case 4 triples the density from 6 h, and the filter's drag scale passes 2.5 at 8.72 h from the navigation fixes alone. With maintenance off the decay steepens to 361 m a day between 10 and 20 h. The deadband decides every 6 h. At its 18 h decision the mean reads 499.761 km, 11 m above the edge, so it waits. The mean crosses at 18.76 h, and the reboost, two burns of 14.42 cm/s, comes only at the 24 h decision. The mean bottoms at 499.677 km, 73 m below the band, and spends 6.08 h outside. The predictive planner forecasts a crossing at 17.25 h at its 12 h decision, reboosts at 12.10 h with two burns of 9.63 cm/s, keeps the mean above 499.848 km and uses 1.68 g against 2.52 g. The deadband's failure comes from waiting, not from its target.

8.7.3 Lead Time and a Late Alert

Case 5 warns at 0.5 h of Fragment 1, predicted 70.7 m away at 10.20 h, nearly head-on at 15.0 km/s, with $P_c = 1.0 \times 10^{-2}$. The uncertainty-aware planner, allowed to burn no earlier than 6 h before TCA, flies 2.10 cm/s retrograde from 4.43 h for 126 s, and a matching return after the pass. The final prediction is 245 m with a probability below 10^{-12}, and the truth revealed after the pass is 278 m. The state of this case at 4.2 h, before the burn is committed, is shown in Figure 8.3. Allowed 12 h, the planner flies 1.20 cm/s centered 8.81 h ahead. Held to 3 h it needs 3.62 cm/s, held anywhere from 2 to 0.9 h it needs 5.08 cm/s centered 0.83 h ahead, about half an orbit, and held to 0.5 h it needs 9.33 cm/s, as shown in Figure 8.2. The closed-form curve falls roughly as $1/\tau$ with a ripple of one orbital period, its minima about one orbit apart and the first near half an orbit, where the radial displacement of Equation (8.8) is largest. The points depart from the curve because each burn was planned on the tracking of its own moment.

Pause. Try Exercise 8.1 before reading on.

At 20 h a late fragment is reported 25 min before a pass predicted at 17.1 m. With the 5 min slew and no planned burn allowed in the last 20 min, nothing meets the criterion. The best effort is 18 min of full thrust, 18.11 cm/s prograde, ending 2 min before TCA, and it cuts the final probability from the 1.6×10^{-2} of a pass with no burn to 1.1×10^{-3}, eleven times the threshold. The pass is flagged infeasible under both planners. The nominal planner handles Fragment 1 with 1.46 cm/s sized to the 200 m floor on the prediction of its moment, and ends at a predicted 171 m, inside the floor, once later tracking moves the fragment.

8.7.4 A Miss Distance That Hides Its Risk

Case 6 predicts a rocket body 725 m away at 14.1 h with a 3 km along-track standard deviation that tracking brings to 1.5 km. The crossing angle is 100°, so Equation (8.9) puts 0.643 of that error into the plane. At the final assessment the combined ellipse is 971 by 150 m, the predicted miss 1622 m, eight times the floor, and $P_c = 1.8 \times 10^{-4}$, over the threshold. The nominal planner flies nothing. The rocket body passes 807 m away, 839 m from where it was predicted. The uncertainty-aware planner burns 3.66 cm/s prograde at 8.12 h, and the pass ends at a predicted 3209 m with $P_c = 2.5 \times 10^{-6}$ and a true miss of 2386 m. The miss stays within 2° of the long axis of the ellipse in both runs, and the burn doubles its Mahalanobis distance from 1.71 to 3.38.

Pause. Try Exercise 8.2 before reading on.

The same case shows dilution in numbers. At the alert the ellipse is 2010 by 301 m and $P_c = 1.6 \times 10^{-4}$, while scaling that covariance would give at most 4.7×10^{-4}, with its standard deviations at 0.38 of their size. As tracking shrinks the ellipse to 1155 by 178 m the probability rises to 2.8×10^{-4}, and it falls to 1.8×10^{-4} only when the last update moves the predicted miss out to 1622 m. At that

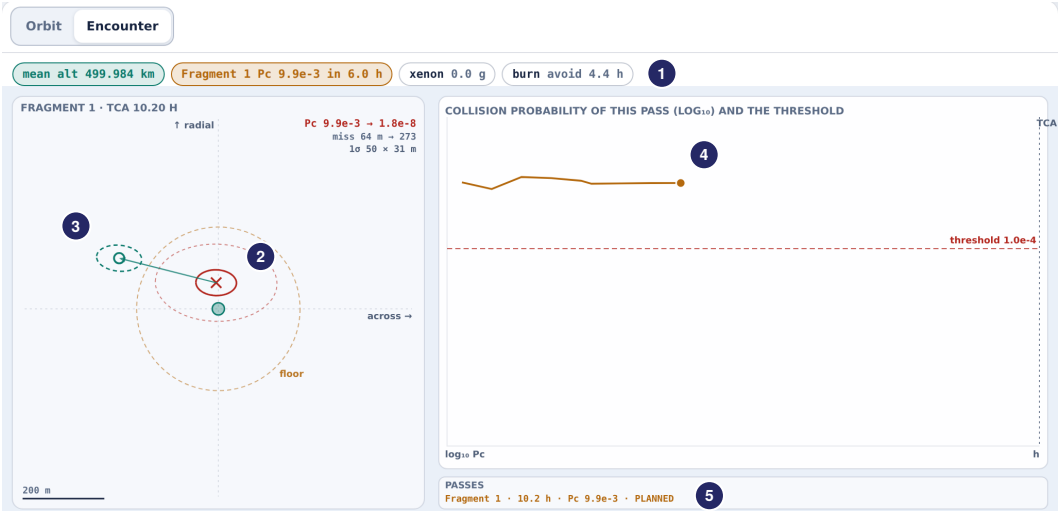

Figure 8.3 The application in case 5, a well-observed conjunction, at $t = 4.20$ h with the shipped settings, 6.0 h before Fragment 1 reaches its closest approach at 10.20 h, cropped to the plant view. (1) The badges give the next burn as an avoidance at 4.4 h. (2) The Encounter view shows the plane of the pass, centered on the satellite and its 15 m hard-body circle with radial up, with the fragment 64 m away inside its combined 1σ ellipse of 50 by 31 m and its 3σ ellipse, and the dashed 200 m floor. (3) An arrow leads to the ring where the planned burn will put it, 273 m away inside its own dashed ellipse. The readout gives P_c from 9.9×10^{-3} now to 1.8×10^{-8} with the plan. (4) Beside the plane, the probability of this pass since the alert at 0.5 h stays between 6.5×10^{-3} and 1.5×10^{-2}, two decades above the 10^{-4} threshold, and (5) the list of passes marks Fragment 1 as planned.

point the probability is 92 % of the largest any scaling of the covariance gives, reached at 1.21 times its standard deviations. A covariance larger than the worst case makes a pass look safer than better data will, the behavior analyzed in [7].

8.7.5 A Thruster That Underdelivers

Case 7 repeats Fragment 1 with the thruster delivering 55 % of its command and verification off. The 2.10 cm/s burn delivers 0.548 of it. The filter's efficiency reads 0.82, 0.68 and 0.63 at 8, 14 and 22 min after the burn ends. Fired and forgotten, the pass ends at a predicted 138 m, inside the floor, with a true miss of 171 m. With verification on, the planner re-screens 15 min after the burn on the orbit the filter now sees and flies a trim at 5.02 h sized to deliver 1.18 cm/s with the estimated efficiency. The pass ends at 232 m (true 265 m), and the return is sized from what the filter believes both burns delivered, 7.96 cm/s commanded in all against 4.20 cm/s. The xenon flow of Equation (8.3) follows the command, so the weak thruster spends the full propellant for every burn.

8.7.6 A Return Burn That Creates a Pass

Case 8 avoids Fragment 3 with 1.86 cm/s at 4.33 h and plans an equal return at 10.64 h. Half an hour after the avoidance the catalog lists Satellite B, an active spacecraft with a shared ephemeris, predicted 670 m away at 15.39 h. The world placed it three orbits after the planned return, so it is invisible on the operational trajectory, which holds only committed burns. Committed unscreened at 10.56 h, the return brings B to a predicted 8.4 m, and to 3.1 m with $P_c = 0.23$ once it has flown. With the planner switched off at that moment the true miss would be 14.8 m, inside the hard-body radius, and the run logs a collision. With the planner on and 4.8 h to go, a third burn of 0.62 cm/s clears B

Table 8.8 Campaign Configurations Over 96 h. Out Is the Time the True Mean Spent Outside the Band

Configuration	Unres.	Worst P_c	Out (h)	Xenon (g)	Burns
shipped, deadband, nominal	3	1.7×10^{-3}	5.53	2.83	7
predictive, nominal[b]	3	2.1×10^{-3}	0	4.16	10
deadband, aware	0	4.9×10^{-5}	5.51	4.59	9
deadband, planner off	3	2.4×10^{-2}	5.59	2.35	1
suggested[a]	0	7.7×10^{-5}	0	6.04	12
predictive one burn, aware	0	4.9×10^{-5}	0	5.96	10

[a] predictive, two burns, uncertainty-aware, 69.1 cm/s in all

[b] two-burn reboosts

to 241 m, 4.35 cm/s in all. With re-screening on, the return is screened before it is committed and moved one orbit later, and B passes 216 m away with no extra burn, 3.73 cm/s in all.

8.7.7 The Four-Day Campaign

The campaign decays the orbit from 500.05 km, triples the drag from 14 h, sends Fragment 1 at 10 h, the rocket body at 52 h and, after the thruster falls to 55 % at 58 h, Fragment 2 at 72 h, with 8 g of xenon allocated. Its target is every pass resolved, no time with the true mean outside the band, the mean in the band at the end and the xenon within the allocation. The six configurations measured appear in Table 8.8, and Figure 8.4 draws three of its columns as bars. In its delivered form, a one-burn deadband with nominal avoidance leaves all three passes unresolved, Fragment 2 at a predicted 76 m, with a worst final P_c of 1.7×10^{-3}, 17 times the threshold. It spends 5.53 h outside the band and 2.83 g of xenon. In the configurations measured, predictive maintenance is what keeps the band and the uncertainty-aware planner is what resolves the passes, so the two failures have separate remedies. With both, the worst final probability is 7.7×10^{-5}, 23 % under the threshold, for the rocket body at 1885 m, and Fragment 2 needs two trims after the thruster fault. The smallest final miss is 212 m, 12 m over the floor, the lowest mean is 499.800 km, 50 m inside the band, and the four days use 6.04 g of the 8 g. A one-burn predictive reboost also meets the target, because the target counts the mean and not the apogee. Over six other seeds, nominal avoidance leaves one to three passes unresolved and the deadband 5.47 to 9.00 h outside the band, while the uncertainty-aware planner resolves every pass and the suggested configuration meets the target each time with 3.22 to 5.23 g of xenon.

> ✦ DID YOU KNOW? **Safety Costs Xenon**
>
> From the simulator: in the four-day campaign the shipped configuration spends 2.83 g of xenon in 7 burns and leaves all three passes unresolved. The suggested configuration, which meets the target, spends 6.04 g in 12 burns, more than twice as much (Table 8.8).

8.8 Technology in Practice

The satellite of Section 8.3 is read by navigation fixes and a debris catalog, and it is moved by one Hall thruster. Figure 8.5 draws its loop with real devices in the place of the schematic's symbols.

What Measures It

The fixes of Equation (8.10), 5 m of white noise on each axis every 10 s, stand for a satellite navigation receiver, now the main means of orbit determination in low orbit [11]. Nothing aboard measures the mean altitude, the drag or the thrust delivered, and the filter infers all three from the fixes. An operator's own prediction of its orbit is often the best one available, because the operator holds the receiver's fixes, a good model of the drag area and the planned maneuvers. The handbook of practice

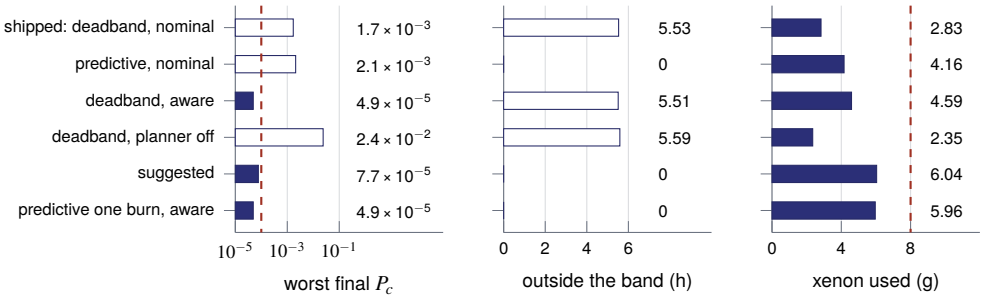

Figure 8.4 Three columns of Table 8.8 drawn as bars for the six configurations of the four-day campaign, with its limits as dashed lines: the 10^{-4} threshold on the worst final collision probability (logarithmic axis) and the 8 g of xenon allocated. Solid bars meet the target of their panel, which for the time outside the band is none, and white bars miss it. Notice that the passes are resolved only by the uncertainty-aware planner and the band is kept only by predictive maintenance, so only the configurations with both meet the target, and they spend the most xenon.

asks that every point of that prediction carry a realistic covariance [18], the role of the satellite's own covariance at TCA in Equation (8.11). The catalog of the debris, with its declared standard deviations, stands for the tracking network that the second history segment describes.

> **TECHNOLOGY** **Sensor: Satellite Navigation Receiver**
>
> For a satellite in low orbit, a receiver of the Global Positioning System (GPS) is now the main means of orbit determination, in place of tracking from the ground. It times the signals of at least four navigation satellites: three fix the position and the fourth corrects the receiver's clock. Its accuracy is limited by the passage of the signals through the upper atmosphere and by the precision of the civil code. Jamming and spoofing of these signals have been detected in orbit, and a satellite should be able to tolerate their loss [11]. On a trend a lost receiver is a gap in the fixes, across which the filter runs on its model alone.

What Moves It

The thruster of Table 8.2, 20 mN at 350 W with a specific impulse of 1400 s, gives little thrust for its power, so its burns are long. A reboost of 24.91 cm/s fires for 24.9 min, more than a quarter of an orbit, which is why the avoidance planner averages each burn's response over its length. The magnitude and pointing errors drawn for each burn in the lower layer stand for the performance drift and the walk of the thrust vector of a real thruster [11], and the output η of case 7 for a thruster that has lost part of its thrust. The model has no attitude, so the slight roll torque that a Hall thruster puts on the spacecraft over a long firing, which reaction wheels may have to absorb [11], is left to Chapter 7.

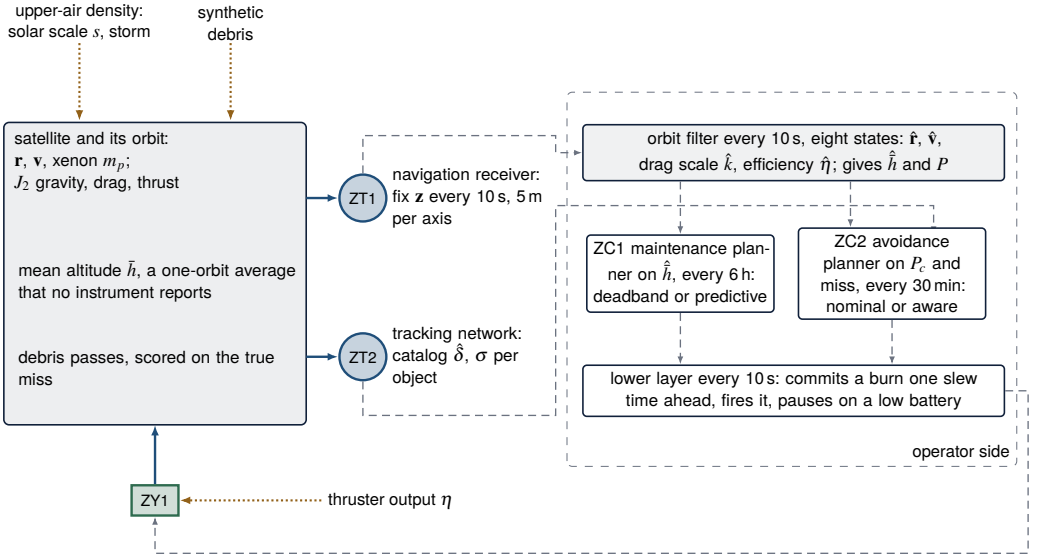

ZY1: Hall thruster drive, 20 mN at 1400 s on xenon, 350 W; the lower layer counts commanded on-time and cannot see the thrust delivered; each burn carries a 2 % magnitude and a 1° pointing error.

Figure 8.5 The loop of the chapter with real devices in place of the schematic's symbols. ZT1 is the satellite navigation receiver and ZT2 the tracking network behind the debris catalog. On the operator's side (dashed frame) the orbit filter feeds the maintenance planner (ZC1) and the avoidance planner (ZC2), which also reads the catalog, and the lower layer flies their burns on the Hall thruster drive ZY1. Notice that neither controlled quantity is measured: the mean altitude is an average of the filter's estimate and the collision probability a prediction for each pass.

TECHNOLOGY Final Element: Hall Thruster

A Hall thruster sets a strong axial electric field and a radial magnetic field at the exit of its discharge chamber. Electrons held in a current around the chamber ionize the incoming propellant, the ions are accelerated by the electric field, and a cathode neutralizes the plume. Electric thrusters give less than about 75 mN of thrust for each kilowatt, so they may fire for hundreds of hours where a chemical thruster needs seconds or minutes, which suits station keeping. Erosion of the plasma-wetted surfaces changes the performance over the thruster's life, and the thrust vector can walk while the thruster warms up [11]. A weak thruster shows only in the orbit, as in case 7.

Where the Controller Runs

The chapter's filter and lower layer run every 10 s, the avoidance planner every 30 min, and the maintenance planner every 6 h. In operations an operator's flight dynamics system produces the predicted orbit with its planned maneuvers and sends it at least daily to a screening authority, which checks it against the catalog [11]. The structure of this chapter therefore lives in decision layers above any loop, not in a function block of a controller (Section 8.6).

What Is Modernizing It

Collision avoidance is being tried on board. A NASA swarm demonstration launched in 2023 carried software that computes the probability of collision and proposes maneuvers to reduce the risk, and a ground hub is planned to screen such proposed trajectories in near real time and return the results to every operator involved [11]. Hall thrusters fed with krypton, less efficient than xenon but much cheaper, are expected to become more common, above all for large constellations [11]. The

planner itself is a chance-constrained decision on a carried-forward uncertainty, the form of stochastic predictive control [8], and the uncertainty-aware planner of Section 8.6 is a small instance of it.

> **IN PRACTICE** **A Threshold, a Margin and a Deadline**
>
> NASA's best practices for conjunction assessment record a broad convergence on mitigating a pass when its probability of collision exceeds 1 in 10,000, and also when the predicted miss is smaller than the hard-body radius. A mitigation should lower the probability by at least 1.5 orders of magnitude and create no new conjunction above the threshold, which is checked by screening the planned trajectory again. The decision falls due at a maneuver commitment point before the closest approach. An earlier burn can be smaller, while waiting may show that no burn is needed [18].

8.9 Limitations

The debris is synthetic, placed to produce each pass, with covariances declared rather than derived from tracking, and the atmosphere is a table times two scales. No number in the app is an orbit forecast or a statement about any real object, and its thresholds are design choices rather than validated practice.

The collision probability is the short-encounter model with independent Gaussian errors, and the satellite's covariance is carried with the linear relative motion, which ignores the effect of J_2 on that motion. Burns are sized on the same linear model and verified on the full one, which corrects their size but not the covariance.

The satellite has no attitude dynamics, navigation is position only, and the loop scorecard of the framework, computed for the mean altitude, means little for a plant flown by planners. This chapter therefore reports passes, times, altitudes and xenon instead. A burn already committed when its planner is switched off still flies, because its slew has begun, and the challenge score reads the mean altitude only, so it cannot tell a round orbit from an eccentric one.

> **OPEN CHALLENGE** **Covariances Worth Trusting**
>
> A probability of collision is only as good as the covariances that feed it. NASA's handbook found the realism of the covariances in operators' own predicted orbits much more uneven than in the catalog, with few published studies and no practical guides for testing it. Drag adds an error that is hard to bound, because the space weather indices that drive density models are difficult to predict [18]. In case 6 the probability rose as tracking shrank the ellipse. Progress would be covariances whose realism is tested routinely against later tracking, with the density forecast error carried inside them.

8.10 Conclusion

The simulator of this chapter flies a low Earth orbit satellite on a running mean of its altitude and the collision probability of each pass, both kept apart from the truth that produces them. With gravity alone the altitude swings 1.6 km twice an orbit while the mean holds to 0.98 m, and drag lowers that mean about 120 m a day. One reboost burn lifts the apogee 325 m above the band, and a deadband deciding every six hours spends 6.08 h below it through a storm that a predictive planner never lets the mean reach. A 71 m pass costs 1.20 cm/s with 8.8 h of lead, 2.10 cm/s with 5.8 h and 5.08 cm/s with less than an orbit, while 18.11 cm/s with 25 min of warning leaves a probability eleven times the threshold. A rocket body predicted 1.6 km away carries $P_c = 1.8 \times 10^{-4}$ because its ellipse is 971 by 150 m. A 55 % thruster leaves a pass at 138 m unless the planner checks the orbit and trims, and an unscreened return brings an active satellite to 3.1 m. The model stops at J_2 with a tabulated atmosphere, synthetic debris and a linear own covariance. A fuller model would add a density forecast with its errors, catalog-wide screening, attitude and slew dynamics and velocity measurements in the filter. The application is a browser plant on which a satellite is flown, as in operations, on what its operators predict rather than on what they measure.

KEY POINTS

- A satellite in low orbit is flown on the average of its height over one orbit and on the odds of each collision, and no instrument reports either.
- With gravity alone the altitude swings 1.6 km twice an orbit while the mean holds to 0.98 m, and drag lowers that mean about 120 m a day.
- A deadband deciding every six hours spends 6.08 h below the band through a storm that a predictive planner never lets the mean reach.
- A rocket body predicted 1.6 km away carries $P_c = 1.8 \times 10^{-4}$, over the threshold, because its ellipse is 971 by 150 m.

Where else this applies. Volume 9, Chapter 7 commits its releases days ahead on an uncertain forecast under a chance constraint, and Volume 5, Chapter 13 steers on a filter's estimate of a position nothing aboard measures.

References

[1] Y. Kozai, "The motion of a close earth satellite," *The Astronomical Journal*, vol. 64, p. 367, 1959.

[2] D. G. King-Hele and D. M. C. Walker, "Density of the upper atmosphere determined from satellite orbits, and its variation during 1957–61," *Nature*, vol. 191, no. 4784, pp. 114–118, 1961.

[3] J. T. Emmert, "Thermospheric mass density: A review," *Advances in Space Research*, vol. 56, no. 5, pp. 773–824, 2015.

[4] M. R. Akella and K. T. Alfriend, "Probability of collision between space objects," *Journal of Guidance, Control, and Dynamics*, vol. 23, no. 5, pp. 769–772, 2000.

[5] F. K. Chan, *Spacecraft Collision Probability*. American Institute of Aeronautics and Astronautics, Inc., 2008.

[6] W. H. Clohessy and R. S. Wiltshire, "Terminal guidance system for satellite rendezvous," *Journal of the Aerospace Sciences*, vol. 27, no. 9, pp. 653–658, 1960.

[7] M. S. Balch, R. Martin, and S. Ferson, "Satellite conjunction analysis and the false confidence theorem," *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, vol. 475, no. 2227, p. 20180565, 2019.

[8] M. Farina, L. Giulioni, and R. Scattolini, "Stochastic linear Model Predictive Control with chance constraints – a review," *Journal of Process Control*, vol. 44, pp. 53–67, 2016.

[9] T. Flohrer, H. Krag, and H. Klinkrad, "ESA's process for the identification and assessment of high-risk conjunction events," *Advances in Space Research*, vol. 44, no. 3, pp. 355–363, 2009.

[10] O. Montenbruck and E. Gill, *Satellite Orbits*. Springer Berlin Heidelberg, 2000.

[11] S. V. Weston, C. D. Burkhard, J. M. Stupl, R. L. Ticknor, B. D. Yost, R. A. Austin, P. Galchenko, L. K. Newman, and L. Santos Soto, "State-of-the-art small spacecraft technology," National Aeronautics and Space Administration, Ames Research Center, Moffett Field, CA, Technical Publication NASA/TP-20250000142, Jan. 2025. [Online]. Available: https://ntrs.nasa.gov/citations/20250000142

[12] T.-W. Fang, A. Kubaryk, D. Goldstein, Z. Li, T. Fuller-Rowell, G. Millward, H. J. Singer, R. Steenburgh, S. Westerman, and E. Babcock, "Space weather environment during the SpaceX Starlink satellite loss in February 2022," *Space Weather*, vol. 20, no. 11, p. e2022SW003193, 2022.

[13] A. Garulli, A. Giannitrapani, M. Leomanni, and F. Scortecci, "Autonomous low-earth-orbit station-keeping with electric propulsion," *Journal of Guidance, Control, and Dynamics*, vol. 34, no. 6, pp. 1683–1693, 2011.

[14] D. J. Kessler and B. G. Cour-Palais, "Collision frequency of artificial satellites: The creation of a debris belt," *Journal of Geophysical Research: Space Physics*, vol. 83, no. A6, pp. 2637–2646, 1978.

[15] T. Wang, "Analysis of debris from the collision of the Cosmos 2251 and the Iridium 33 satellites," *Science & Global Security*, vol. 18, no. 2, pp. 87–118, 2010.

[16] C. D. Bussy-Virat, A. J. Ridley, and J. W. Getchius, "Effects of uncertainties in the atmospheric density on the probability of collision between space objects," *Space Weather*, vol. 16, no. 5, pp. 519–537, 2018.

[17] C. Bombardelli and J. Hernando-Ayuso, "Optimal impulsive collision avoidance in low earth orbit," *Journal of Guidance, Control, and Dynamics*, vol. 38, no. 2, pp. 217–225, 2015.

[18] National Aeronautics and Space Administration, "NASA spacecraft conjunction assessment and collision avoidance best practices handbook," National Aeronautics and Space Administration, Washington, DC, Special Publication NASA/SP-20205011318, Dec. 2020. [Online]. Available: https://ntrs.nasa.gov/citations/20205011318

8.11 Exercises

8.1 Concept check · 2 min

In case 5 the uncertainty-aware planner clears Fragment 1 with 1.20 cm/s when its burn may be centered 8.8 h before the closest approach, and needs 5.08 cm/s when it is held to under one orbit. What chiefly explains the difference?

(a) An along-track burn changes the period, so the shift it buys along the track grows with every second before the pass

(b) The catalog covariance is smaller early in the alert, so an early burn has a smaller miss to reach

(c) The thruster delivers more of its command early in the run, while the battery is still full

(d) Drag acts on the lowered orbit for longer and multiplies the effect of an early retrograde burn

8.2 Concept check · 2 min

At the final assessment of case 6 the rocket body is predicted to miss by 1622 m, eight times the 200 m floor. The combined 1σ ellipse in the encounter plane is 971 by 150 m. Which statement is right?

(a) The pass is safe, because the predicted miss is larger than the whole 1σ ellipse

(b) The pass needs a burn, because the miss lies near the long axis, only 1.71 standard deviations out

(c) The pass is safe, because more tracking would shrink the ellipse and could only lower the probability

(d) The pass needs a burn, because the hard-body radius is close to the short axis of the ellipse

8.3 Concept check · 1 min

In case 3 a single prograde reboost of 24.91 cm/s fires while the mean altitude is at the bottom of its band. Once the mean has settled, where is the orbit higher than before?

(a) Everywhere by the same amount, because the mean has risen by about 450 m

(b) At the point where the burn fired, which rises by about twice the change in the mean

(c) Half an orbit from the burn, which rises by about twice the change in the mean

(d) Nowhere until a second burn fires, because one burn changes only the eccentricity

8.4 Theory and calculation · 6 min

The satellite of this chapter has a mass of 120 kg, a Hall thruster of 20 mN at a specific impulse of 1400 s and a mean orbit of radius $a = R_E + 500\,\mathrm{km} = 6878.137\,\mathrm{km}$, with $\mu = 3.986004418 \times 10^{14}\,\mathrm{m^3/s^2}$ and $g_0 = 9.80665\,\mathrm{m/s^2}$.

(a) Find the change in the mean radius per m/s of along-track velocity change, and the rise of the mean from the 24.91 cm/s reboost of case 3. *Checkpoint:* 1.807 km per m/s and 450 m.

(b) Find the time this burn takes at the rated thrust and the xenon it uses. *Checkpoint:* 1495 s (24.9 min) and 2.18 g.

(c) An impulse along the track makes the perigee-to-apogee difference $2\Delta a$. A burn spread evenly over an arc φ of the orbit, centered on the same point, makes it smaller by the factor $\sin(\varphi/2)/(\varphi/2)$. With a period of 94.62 min, estimate the difference this burn leaves and compare it with the mean perigee and apogee that case 3 reports. *Checkpoint:* about 801 m, against 805 m measured.

8.5 Theory and calculation · 7 min

At the final assessment of case 6 the combined 1σ ellipse of the rocket body is 971 by 150 m, the predicted miss lies at a Mahalanobis distance $d = 1.71$, and the combined hard-body radius is $R_{HB} = 15\,\mathrm{m}$. Use the small-disc form $P_c \approx K \exp(-d^2/2)$ with $K = R_{HB}^2/(2\sigma_1\sigma_2)$.

(a) Compute the collision probability. *Checkpoint:* 1.79×10^{-4}.

(b) Scale both standard deviations by a factor s, which divides d by s and K by s^2. Find the s that makes the probability largest and that largest value, and the share of it that the actual covariance gives. *Checkpoint:* $s = 1.21$, 1.94×10^{-4} and 92 %.

(c) The uncertainty-aware planner plans to $P^* = P_{c,max}/k_P$ with $P_{c,max} = 10^{-4}$ and $k_P = 3$. Find the Mahalanobis distance that meets P^* on this covariance. *Checkpoint:* 2.51.

8.6 Theory and calculation · 6 min

An along-track velocity change of 1 cm/s is made at 500 km, where the mean motion is $n = 1.1068 \times 10^{-3}$ rad/s. Use Equation (8.8) for the radial displacement x and the along-track displacement y a time τ after the burn.

(a) Find x and the size of y half an orbit after the burn ($n\tau = \pi$) and one orbit after it ($n\tau = 2\pi$). *Checkpoint:* 36.1 m and 85.2 m; 0 m and 170.3 m.

(b) The pass crosses at 160°, so the radial displacement enters the encounter plane whole and the along-track one multiplied by $\cos(\theta/2)$ (Equation (8.9)). Find the in-plane displacement at both times, and say which lead the burn should prefer. *Checkpoint:* 39.0 m at half an orbit, 29.6 m at one orbit.

8.7 App challenge · 8 min

Load scenario 4, Density surge, and change nothing: the maintenance planner is a deadband deciding every 6 h, and a storm triples the density from 6 h. Use the default seed and the default speed, one simulated hour per second.

(a) Run to 48 h. Read the Time outside the band and the Propellant used tiles, and find in the exported file the lowest mean altitude, column `h_mean_km`, and the first row at which the filter's drag scale, column `drag_scale_est`, exceeds 2.5. *Checkpoint:* 6.08 h, 2.5 g; 499.677 km at 24.3 h; 8.73 h, within 0.02 h.

(b) Press Reset, set Maintenance planner to Predictive before pressing Run, and run to 48 h again. Read the same tiles and the lowest mean altitude. *Checkpoint:* 0.00 h, 1.7 g; 499.848 km.

Accept altitudes within 0.002 km. Explain which part of the deadband's rule lets the mean leave the band.

8.8 App challenge · 8 min

Load scenario 5, Well-observed conjunction. Use the default seed and the default speed. For each run read the Δv commanded tile at any time between 9.5 and 10.6 h, after the avoidance burn and before the return burn is committed, and in the exported file the predicted miss, column `miss_pred_m`, in the row at 10.10 h, after the final assessment.

(a) Run as shipped, with Earliest burn before TCA at 6 h, and pause at about 10.6 h. *Checkpoint:* 2.1 cm/s and 245 m.

(b) Press Reset, set Earliest burn before TCA to 12 h before pressing Run, and run to 10.6 h. *Checkpoint:* 1.2 cm/s and 204 m.

(c) Press Reset, set Earliest burn before TCA to 1.5 h, and run to 10.6 h again. *Checkpoint:* 5.1 cm/s and 231 m.

Accept a miss within 1 m. Column `dv_cmd_mps` gives the burn in m/s to more digits.

8.9 App challenge · 8 min

Load the Challenge, four days in the band. Set the speed to 12×, so that each 96 h run takes 8 s, use the default seed, and read the Unresolved passes, Worst P_c at TCA, Time outside the band and Propellant used tiles when the run stops at 96 h. Make every change before pressing Run.

(a) Run the challenge as shipped: one-burn deadband maintenance and nominal avoidance. *Checkpoint:* 3 of 3, 1.7×10^{-3}, 5.53 h, 2.8 g.

(b) Press Reset, set Avoidance planner to Uncertainty-aware and run again. *Checkpoint:* 0 of 3, 4.9×10^{-5}, 5.51 h, 4.6 g.

(c) Press Reset, also set Maintenance planner to Predictive and Reboost as to Two burns, and run again. *Checkpoint:* 0 of 3, 7.7×10^{-5}, 0.00 h, 6.0 g.

Say which change cures which failure, and whether the target is met.

8.10 Application area · 4 min

A mission planner sizes the xenon of this satellite for the default atmosphere, in which drag lowers the mean altitude by an orbit average of 120.4 m a day at 500 km. Use $a = 6878.137$ km, $v = 7612.6$ m/s and 87.4 mg of xenon per cm/s.

(a) Find the velocity change per year that holds the mean altitude. *Checkpoint:* about 24.3 m/s a year.

(b) Find the xenon per year and the years that the 2.0 kg tank lasts, and name two effects of the Limitations section that make this an upper bound on the lifetime. *Checkpoint:* about 213 g a year and 9.4 years.

8.11 Application area · 3 min

The challenge target counts passes, time with the true mean altitude outside the band, the mean at the end and xenon. Predictive maintenance with one-burn reboosts and the uncertainty-aware planner meets it with 5.96 g in 10 burns, against 6.04 g in 12 burns with two-burn reboosts.

(a) Using case 3, say what the one-burn plan does to the orbit that the score cannot see. *Checkpoint:* the apogee rises 325 m above the band.

(b) An Earth-observation customer needs the altitude inside the band everywhere around the orbit. Name the quantities you would add to the target and where the app reports them. *Checkpoint:* the mean perigee and apogee.

8.12 Review the agent · 5 min

An AI assistant is asked to automate station keeping and replies as follows. “The mean altitude lags the orbit by half a revolution, so a planner on it is always late. The navigation fix gives the altitude every 10 s with only 5 m of noise. Close a PI loop from the measured altitude to the along-track thrust with a setpoint of 500 km, and let the thruster rest while the altitude is within the ±250 m band. The loop will hold the altitude with far less delay than any planner.” Find the flaw in the proposal and say how it could be checked.

CHAPTER 9

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

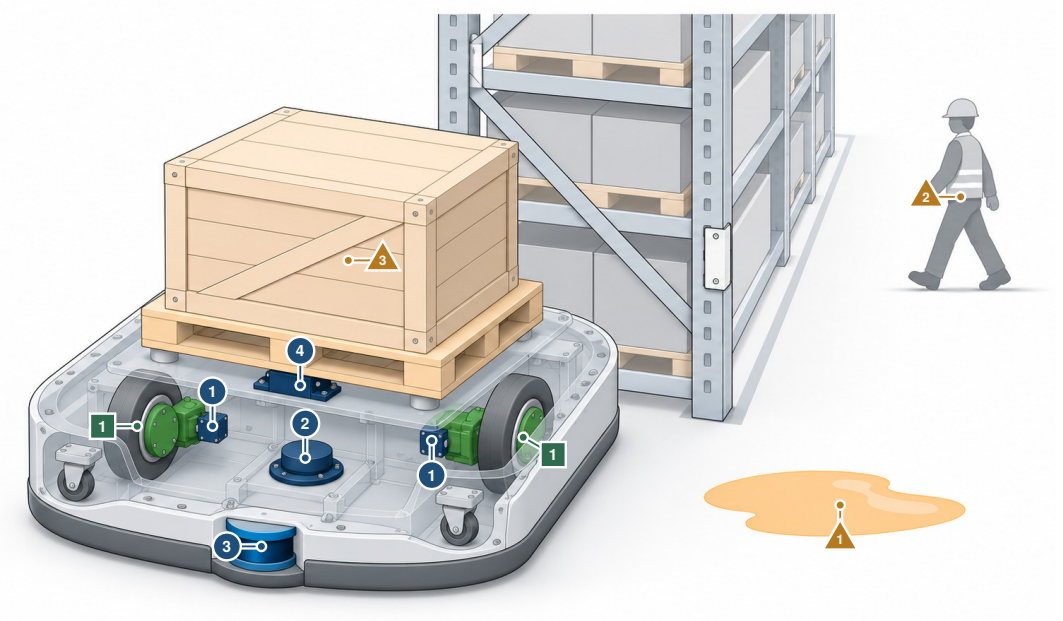

The controller moves 1 left and right wheel motors, one torque each

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

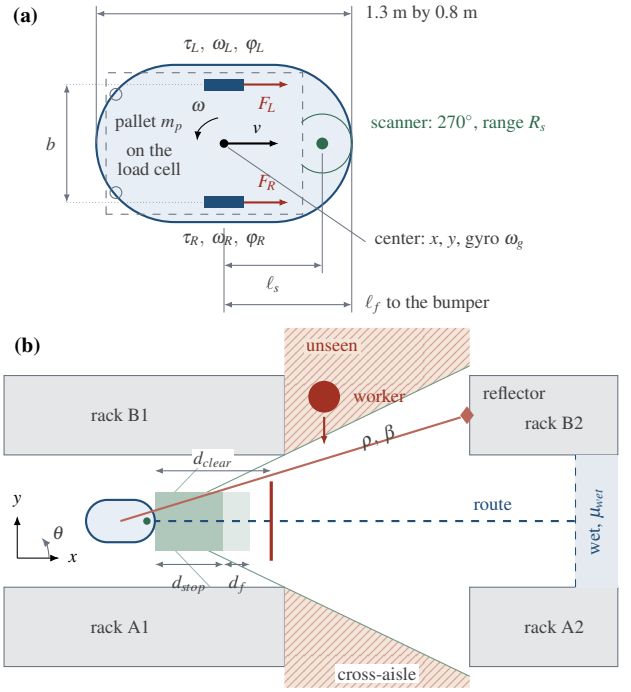

Figure 9.1 The robot and a blind cross-aisle. (a) The robot heading right, with its pallet m_p on a load cell, drive wheels of radius r on track b with torques τ_L, τ_R, speeds ω_L, ω_R, encoder angles φ_L, φ_R and traction forces F_L, F_R, the speed v and yaw rate ω at the center (x, y) where the gyro 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

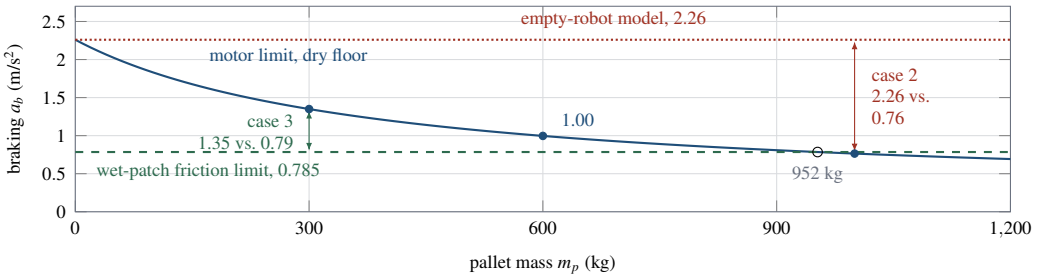

Figure 9.2 Braking capability against pallet mass from Equation (9.4). The motor limit (solid) falls with the load, the wet patch caps braking at 0.785 m/s^2 (dashed), and the empty-robot model believes 2.26 m/s^2 whatever the load (dotted). Filled dots mark three pallet masses and the open circle the 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

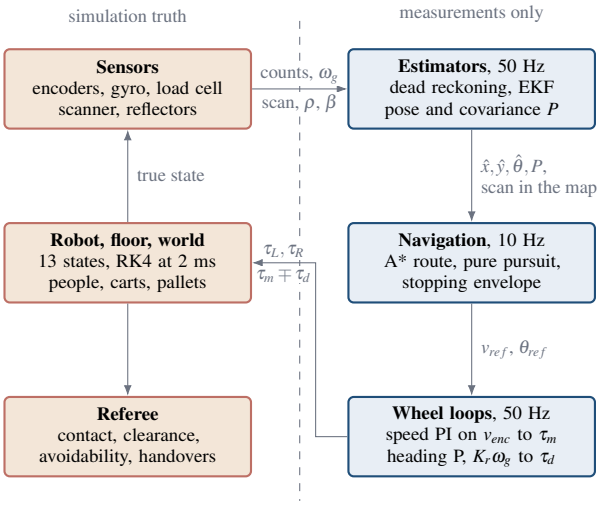

Figure 9.3 The layers of the controller. Only the sensor layer reads the simulation truth, and the referee scores the run from the truth without feeding anything back. The encoders, gyro, load cell and wheel loops run at 50 Hz, and the scanner, reflectors, navigation and supervisor at 10 Hz.

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).

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

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

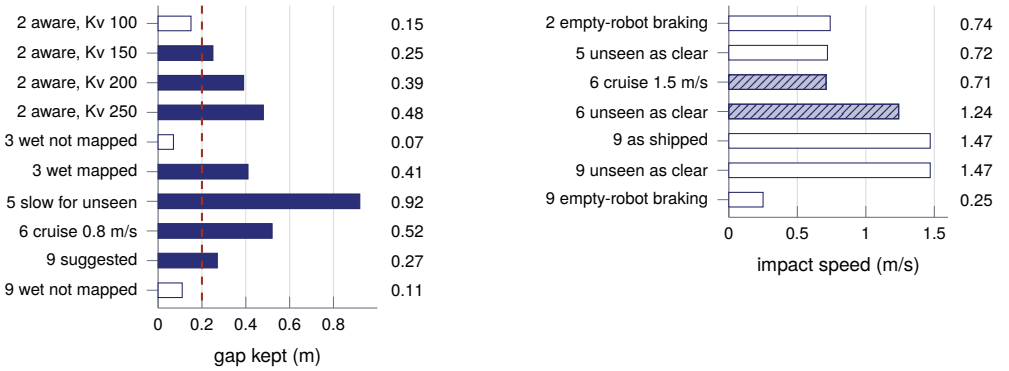

Figure 9.4 Two outcomes of Table 9.6 drawn as bars, labeled by case. Left, the gap each run kept to the cart, worker or pallet ahead, with the 0.2 m clearance as a dashed line: solid bars keep it and white bars violate it. Right, the impact speed of each run that collided, with the two collisions the 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.

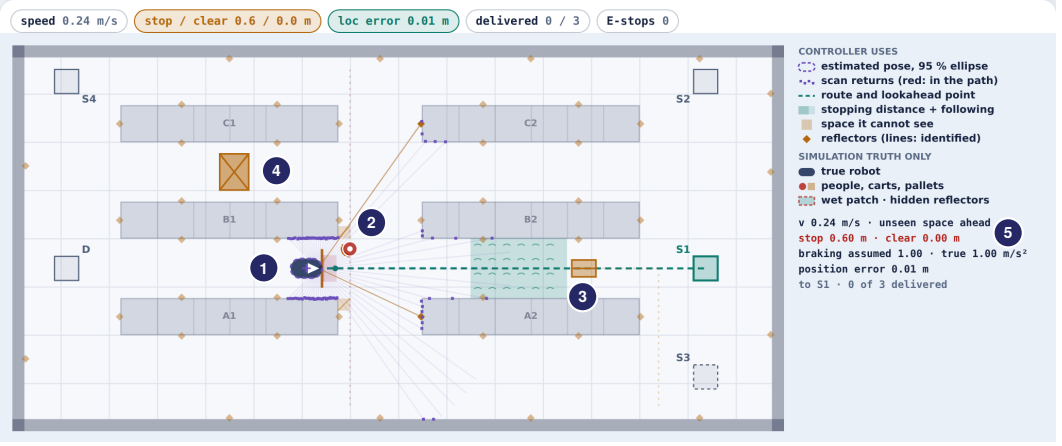

Figure 9.5 The graded shift with the suggested configuration at $t = 11.00$ s, shortly before the shipped configuration's collision at 11.09 s, cropped to the plant view. The 600 kg robot (1, solid, with its estimate dashed) approaches the first cross-aisle at 0.24 m/s. Shaded amber cells mark unseen 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.

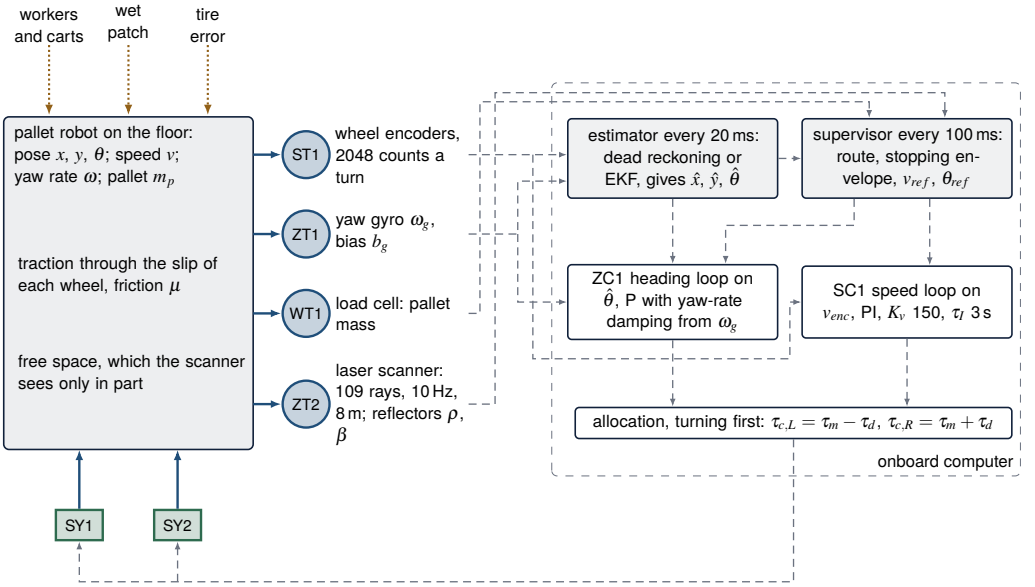

SY1, SY2: left and right wheel drives, an electric motor and gearbox on each driven wheel, 40 N m each through a 0.05 s torque lag; casters carry the rest of the weight. The emergency stop brakes both flat out.

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.

CHAPTER 10

Robot Conveyor Picking Simulator for Late Camera Reports and Moving-Target Interception

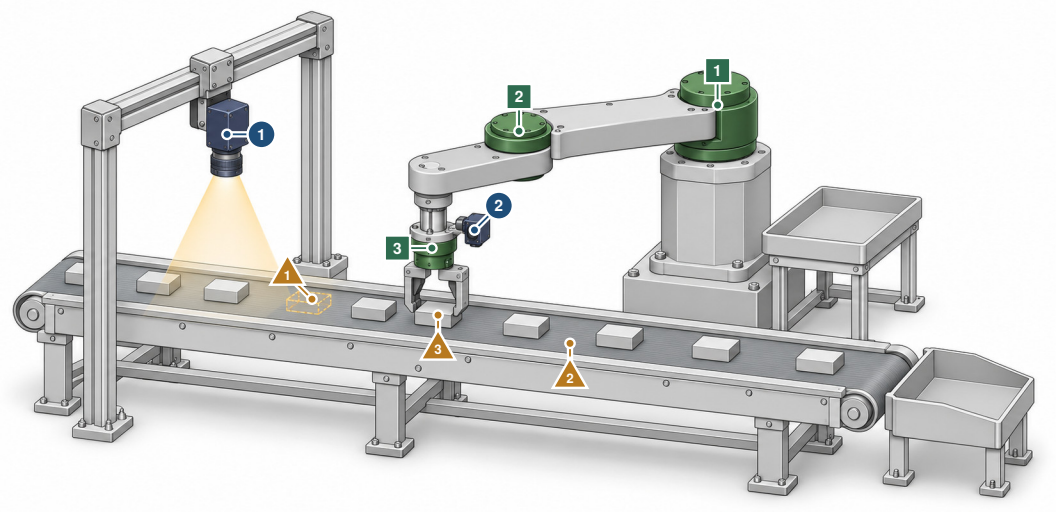

The controller moves 1 shoulder motor, one torque 2 elbow motor, one torque 3 gripper, which closes and lets go

The controller reads 1 camera upstream, which reports each box late 2 switch that tells a box is held

What disturbs it 1 where the camera last saw a box 2 belt speed 3 mass of the box

Notice that the camera watches the belt upstream of the arm, so every report describes a box that has already moved on, and the gripper (a vacuum cup in the chapter, drawn here with jaws) must close where the box will be. Illustration, not to scale.

IN PLAIN TERMS Catching Boxes on a Moving Belt

The main idea. A robot arm picks small boxes off a moving belt and sets them in a tray. A camera upstream finds each box, but the box has moved on by the time the arm gets there.

The challenge. The arm has to grab each box where it will be, and its hand has to move with the belt as it closes. It must also carry the box gently, or its grip lets go.

Why it matters. On a fast belt, an arm that aims at old reports or waits standing still puts no box in the tray.

THREE WAYS TO READ THIS CHAPTER

Explore	11 min	the first page, the history segment and the fact on early robots (Section 10.2), the try-it segments of scenarios 4 and 2 (Section 10.4), the fact on lost frames (Section 10.6), and the key points
Study	45 min	the whole chapter with its three pause points and the twelve exercises at its end
Practice	15 min	the first page, the measured scenarios (Section 10.6), Technology in Practice (Section 10.7), Limitations (Section 10.8) and the open challenge

The control question A camera reports each product on a moving belt after it has moved on, and a grasp takes only if the gripper is in the right place, nearly still relative to the product and able to hold it, all at one instant. How does a picking arm meet a product where it will be and at the speed it is going, with limited torque and grip?

Learning objectives

- Compute the distance a report-aimed gripper closes behind its product from the belt speed, the latency and the time to the grasp.
- Explain why a gripper that waits at rest fails above the relative speed tolerance, and how arriving at belt velocity removes that failure.
- Bound the carry acceleration by the grip force and the product mass, and explain why the usable limit sits below the bound.
- Predict how a payload left out of the controller's arm model breaks the torque budget and the tracking.
- Configure the cell for a mixed production run and justify each change by the grasp budget it restores.

Prerequisites Sections 1.4.2 (arm dynamics and the payload term), 1.4.4 (quintic profiles and feedforward), 1.4.6 (the age of a report), 1.6.1 and 1.6.4 (running and reading a scenario).

The app https://apmonitor.com/apps/robot-conveyor-picking/

See also Volume 5, Chapter 13 applies each late acoustic fix at the instant it describes, Volume 5, Chapter 7 follows material that has moved on past a slow analyzer, and Volume 5, Chapter 8 closes a robot loop through a delayed link.

10.1 Introduction

Many machines act on an object they measured a moment earlier. A visual servo closes a robot loop through a camera, a hand-eye system that grasps a moving part has to predict where the part will be, and a state estimator that receives a reading late has to incorporate it although it describes an earlier state [1–3]. In each case the measurement says where the object was, the action takes effect where the object will be, and the gap between the two is set by the latency of the sensor, the time the actuator needs to get there and the time it needs to act.

Conveyor picking makes that gap concrete. A camera cannot watch the part of the belt where the arm works, so it looks upstream and reports each product after capture, transfer and recognition. The product rides on while the arm moves to it and while the gripper closes, and the grasp takes only if the gripper is in the right place, nearly still relative to the product and able to hold it, all at one instant. The simulator targets two beliefs. One holds that camera latency is the whole of the timing error. With no latency at all, a gripper sent to the camera's last report on a 0.15 m/s belt still closes 63.0 mm behind the product it was sent for, because that report is 0.42 s old when the grasp completes. The other holds that an accurate position is enough. An arm that waits at rest where a product will be, 0.96 mm from it on average, fails all 60 of its grasps on a 0.30 m/s belt, because the product slides under the gripper at the belt speed.

The class of problem is general. Choosing the point on a predicted path at which to meet a moving object is an optimization of its own [4], prediction and replanning for interception have been developed and tested for objects moving in industrial settings [5], and a Kalman filter can incorporate a measurement that arrives late [3], as the underwater navigation filter of Volume 5, Chapter 13 does with its acoustic fixes. A flying shear that cuts a moving strip, a labeler matched to a product stream and a radar track carried between sweeps have the same shape, and so does a continuous line whose analyzer describes material that has moved on, as in Volume 5, Chapter 7. In each, a measurement is moved to the moment of action by a model of what happened in between, and the actuator must arrive at the object's speed as well as its position.

The behavior matters industrially. Robots pack products in the confectionery industry [6], and on a line that never stops, the number of products an arm catches depends on the order in which it takes them [7] and on how several arms along one belt share the stream [8]. Model-based servos also make

the payload part of the control problem, because the torque a move needs depends on what the arm carries [9].

This chapter treats a conveyor picking simulator in which the camera latency, the belt speed, the product mass, the grip and the arrival pattern can be changed and their consequences read as products placed, missed, failed and dropped. Conveyor tracking, quintic trajectories, computed-torque feedforward and earliest-deadline scheduling are standard methods. The simulator gives a model complete enough to rebuild, in which the controller reads only its own sensors, the products and the camera draw from separate random streams so that two controller settings meet the same products, and the planner checks torque with the controller's own arm model, so that a payload the model omits is invisible to the check. Each grasp condition reduces to a bound that holds before any tuning. An aim that ignores the closing interval misses by at least the belt speed times the closing time, a gripper that waits at rest meets the product at the belt speed, and the grip caps the carry acceleration at the holding force divided by the mass. Eight cases and a graded run expose these bounds one at a time.

Section 10.2 traces how robots came to pick from moving belts. Section 10.3 gives the arm, belt, camera and grasp model, and Section 10.4 turns the grasp into four budgets. Section 10.5 describes the controller layer by layer, Section 10.6 reports the measured scenarios and the production run, Section 10.7 names the devices of real picking cells, and Sections 10.8 and 10.9 state the limitations and the conclusions.

10.2 Picking Products From a Moving Belt

The arm simulated here is a SCARA, a selective compliance assembly robot arm, whose origin the history segment below tells. Both of its joints turn about vertical axes, so the arm sweeps across a plane while gravity loads neither joint. The Delta parallel robot was designed by a research team at the École Polytechnique Fédérale de Lausanne, and a complete model of it, from the kinematics to the inverse dynamics, was written to need few arithmetic and trigonometric operations [10]. Robots for packaging in the confectionery industry, where chocolates are handled, have been reviewed in their own right [6].

> ⧖ HISTORY **An Arm Like a Folding Screen (1978 to 1981)**
>
> The selective compliance assembly robot arm was invented at the University of Yamanashi for assembly work. Its arm gives easily in the horizontal plane and stays stiff in the vertical direction, so a peg slightly out of line can slide into its hole, and its links fold about vertical axes like the panels of a Japanese folding screen. A research group of five companies, formed in April 1978, made the first prototype possible. It was built that October with a servo motor and an optical encoder at each joint, a second prototype followed in 1980, and in 1981 some of the partners began to market their own versions [11].

A moving product first has to be found. A hand-eye system that tracked a moving object with cameras, predicted its motion and closed a gripper on it was reported in 1993 [2], and a tutorial on visual servo control organized the ways a camera can close a robot's loop, with the camera fixed in the workcell or carried on the arm [1]. Interception then became a planning problem in its own right. The meeting point on a predicted path can be selected optimally [4], and prediction, planning and execution were combined into interception strategies tested experimentally for objects moving in industrial settings [5]. A measurement that reaches an estimator late can be incorporated into a discrete-time Kalman filter although it describes an earlier state [3], which is the estimation side of the same timing problem.

A stream of products added the question of order. When items on a belt arrive faster than one robot can take them, the order of the picks changes how many are caught, and online scheduling algorithms were proposed to improve pick-and-place performance on a moving conveyor [7]. Several robots along one conveyor must then coordinate which items each one picks [8].

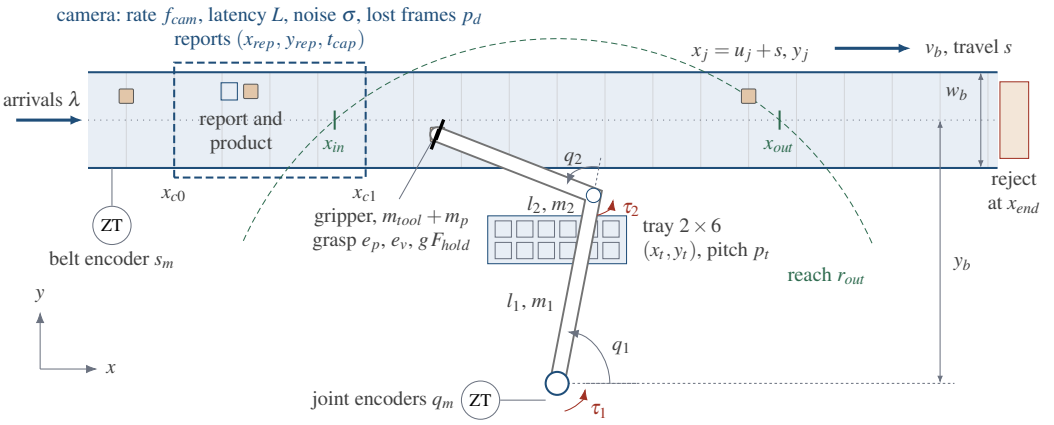

Figure 10.1 The picking cell seen from above, drawn to scale. The shoulder is the origin, the upper arm of length l_1 turns through q_1 from the belt direction and the forearm of length l_2 through q_2 from the upper arm, driven by the joint torques τ_1 and τ_2. Products arrive at rate λ and ride the belt at $x_j = u_j + s$. The camera reports those inside $[x_{c0}, x_{c1}]$ after a latency L, so each report (hollow square) lies behind its product, and the planner meets the products inside the reach r_{out}, between x_{in} and x_{out} on the centerline. A product carried beyond $l_1 + l_2$ is missed, and one that reaches x_{end} goes to the reject bin. The part-present switch sits in the gripper.

> **DID YOU KNOW? Fewer Than Forty Thousand Robots**
>
> At the end of 1981 there were fewer than 40,000 industrial robots in the world, and by 1992 there were 560,000. The first industrial use of the new horizontal arm put eight of them to work assembling audio circuit boards [11].

The joint servos came from model-based control of fast arms. Resolved acceleration control computes the joint torques that produce a commanded acceleration from the manipulator dynamics and corrects the remaining error by feedback [12]. Real-time experiments on a direct-drive arm sampled every 2 ms later compared computed-torque control, which uses the dynamics inside the feedback loop to linearize and decouple the joints, with feedforward compensation, which adds the model torque in the feedforward path, and showed that the off-diagonal terms of the manipulator inertia matrix belong in the torque computation [13]. Both forms depend on the load, whose mass and center of mass can be estimated from general motion of the arm, while its moments of inertia are harder to obtain [9]. A vacuum cup holds its load by astrictive force, one of the forms of prehension that gripper design classifies [14], and that force has to exceed the inertial load of the handling motion. Minimum-time motion along a specified path with known torque limits was solved in 1985 [15], and trajectories generated online from whatever state the arm is in let a controller react at once to a new measurement [16]. The application of this chapter assembles these pieces at teaching scale.

10.3 The Picking Cell

The cell is drawn from above in Figure 10.1. The shoulder of the arm is the origin, the belt runs along $+x$ with its centerline at $y_b = 0.55\,\text{m}$, the camera watches the belt between $x_{c0} = -0.80\,\text{m}$ and $x_{c1} = -0.40\,\text{m}$, upstream of the arm, and a tray of twelve cavities lies between the shoulder and the belt. The six integrated states are the joint angles q_1 and q_2, their rates ω_1 and ω_2, the belt travel s and the belt speed v_b. Everything discrete, which covers the products, the camera frames and their delivery, the gripper, the slip test and the tray, is updated after each integration step and never inside the derivative function.

Table 10.1 Classification of the Signals in the Picking Cell

Symbol	Role	Default	Range
q_1, q_2	CV, joint angles	ready pose	$\lvert q_2\rvert \le 150^\circ$
ω_1, ω_2	states, joint rates	0	—
s, v_b	states, belt travel and speed	0, v_{set}	—
τ_1, τ_2	MV, joint torques	0	± 40, ± 17 N m
—	MV, gripper command	open	open, close
v_{set}	DV, belt speed setting	0.15 m/s	0.05–0.50
λ	DV, arrival rate	24 min^{-1}	0–120
L	DV, camera latency	0.15 s	0–0.5
p_d	DV, lost frames	5 %	0–90
σ	DV, camera noise (1σ)	2 mm	0–10
m_p	DV, product mass	0.10 kg	0.02–2.0
g	DV, grip strength	100 %	10–100
—	DV, extra product, burst of 12	—	events
—	DV, 3 s blackout, 0.6 rad/s bump	—	events
q_m	measured, joint encoders	17 bit	—
s_m	measured, belt encoder	0.1 mm	—
x_{rep}, y_{rep}, t_{cap}	measured, camera report	15 Hz	—
—	measured, part present	off	on, off
—	outputs, product counts	0	—

The signals are classified in Table 10.1. The two joint angles are the controlled variables of the servo loops, whose references come from the planner, while the outcome of the task is scored as products placed, missed, failed and dropped.

10.3.1 Arm Dynamics

Both joints turn about vertical axes, so gravity does no work on the links and the rigid-body dynamics take the form of Equation (10.1), the arm balance of Equation (1.2) in Chapter 1.

$$M(q)\,\ddot{q} + C(q,\dot{q})\,\dot{q} + F(\dot{q}) = \tau, \qquad |\tau_i| \le \tau_{i,\max} \tag{10.1}$$

where $q = [q_1, q_2]^T$ holds the shoulder angle, measured from the belt direction, and the elbow angle, measured from the upper arm, both counterclockwise seen from above, τ holds the two joint torques and F is the joint friction. The upper arm and the forearm are uniform rods of lengths l_1 and l_2 and masses m_1 and m_2, and a point mass $m_{tip} = m_{tool} + m_p$ at the tip stands for the gripper and its vertical quill plus any product of mass m_p being held. With $l_{ci} = l_i/2$ and $I_i = m_i l_i^2/12$, the inertia coefficients are given by Equation (10.2).

$$\begin{aligned} a_1 &= m_1 l_{c1}^2 + I_1 + (m_2 + m_{tip})\, l_1^2 + J_{1m} \\ a_2 &= m_2 l_{c2}^2 + I_2 + m_{tip}\, l_2^2 \\ a_3 &= (m_2 l_{c2} + m_{tip}\, l_2)\, l_1 \end{aligned} \tag{10.2}$$

where J_{1m} and J_{2m} are the rotor inertias of the two motors reflected through their gears. The inertia matrix, the velocity terms and the friction follow as Equation (10.3).

$$\begin{aligned} M &= \begin{bmatrix} a_1 + a_2 + 2a_3\cos q_2 & a_2 + a_3\cos q_2 \\ a_2 + a_3\cos q_2 & a_2 + J_{2m} \end{bmatrix} \\ C\dot{q} &= a_3 \sin q_2 \begin{bmatrix} -(2\dot{q}_1\dot{q}_2 + \dot{q}_2^2) \\ \dot{q}_1^2 \end{bmatrix} \\ F_i &= b_i\,\dot{q}_i + f_{ci}\tanh(\dot{q}_i/0.05) \end{aligned} \tag{10.3}$$

where b_i and f_{ci} are viscous and Coulomb friction coefficients and the tanh smooths the Coulomb term over 0.05 rad/s. Every row carries units of N m. The elbow motor's gyroscopic coupling is left out. The tip position and velocity are given by Equation (10.4).

$$p_{tip} = \begin{bmatrix} l_1 \cos q_1 + l_2 \cos(q_1 + q_2) \\ l_1 \sin q_1 + l_2 \sin(q_1 + q_2) \end{bmatrix}, \qquad v_{tip} = J(q)\,\dot{q} \tag{10.4}$$

where J is the tip Jacobian, and $a_{tip} = J\ddot{q} + \dot{J}\dot{q}$. A held product adds exactly $m_p J^T J$ to M, which one of the application's verification checks confirms to within 10^{-12}. At the ready pose, $q_2 = 62.6°$, the empty arm has $M_{11} = 0.9996\,\mathrm{kg\,m^2}$ and $M_{22} = 0.1579\,\mathrm{kg\,m^2}$, and a 1.5 kg case raises them to 1.6167 and $0.3417\,\mathrm{kg\,m^2}$. The elbow stops dead at $|q_2| = q_{2,\max} = 150°$, where its rate into the stop is set to zero and an end-stop event is marked. The planner works inside $r_{out} = l_1 + l_2 - 0.03\,\mathrm{m} = 0.72\,\mathrm{m}$, short of the straight arm so that the elbow never has to straighten, and at the elbow limit the inner radius of the reach is 0.20 m.

10.3.2 Belt, Products and Camera

The belt drive follows its speed setting through a first-order response, and each product rides at a fixed belt coordinate, Equation (10.5).

$$\dot{s} = v_b, \qquad \tau_b \dot{v}_b = v_{set} - v_b, \qquad x_j = u_j + s \tag{10.5}$$

where $\tau_b = 0.5\,\mathrm{s}$ and u_j is the belt coordinate of product j. Products enter at $x = -1.45\,\mathrm{m}$ as a Poisson stream of rate λ per minute, never closer than 0.07 m along the belt, with a lateral position y_j uniform within $y_b \pm 0.07\,\mathrm{m}$ and the mass m_p in force when they enter. At reset the belt is preloaded upstream of the camera, the first product at $x = -0.95\,\mathrm{m}$. A product is missed when it passes $x > 0$ beyond the full reach $l_1 + l_2 = 0.75\,\mathrm{m}$, and it leaves the belt at $x_{end} = 0.95\,\mathrm{m}$.

Every $1/f_{cam}$ seconds the camera takes a frame at time t_k. The whole frame is lost with probability p_d, or while a blackout lasts, and otherwise every product inside $x_{c0} \le x_j(t_k) \le x_{c1}$ is reported as shown in Equation (10.6).

$$(x_{rep}, y_{rep}, t_{cap}) = \big(x_j(t_k) + n_x,\; y_j + n_y,\; t_k\big) \tag{10.6}$$

where n_x and n_y are independent normal deviates of standard deviation σ. The report is posted at the first integration step that ends at or after $t_k + L$, the controller reads it at its next scan, and it carries no identity and no velocity. The belt encoder counts s_m in steps of $\Delta s = 0.1\,\mathrm{mm}$, and the joint encoders count q_m in steps of $2\pi/131072 = 4.79 \times 10^{-5}\,\mathrm{rad}$, about 0.036 mm at the full reach. The arrivals and the camera draw from two separate random streams, both seeded from the scenario seed at reset, so a reset replays the same products, noise and lost frames, and a change of controller setting meets exactly the same products. Reset and reload of a scenario behave as Section 1.6.1 describes.

10.3.3 Grasp, Carry and Placement

The gripper closes in $t_{close} = 0.10\,\mathrm{s}$, a time that also stands for the quill's dip to the belt, and opens in $t_{rel} = 0.08\,\mathrm{s}$. When it finishes closing, at time t_g, the grasp takes only if the three conditions of Equation (10.7) hold for the nearest product on the belt.

$$\begin{aligned} \|p_{tip} - p_j\| &\le e_p \\ \|v_{tip} - [v_b, 0]^T\| &\le e_v \\ m_p \|a_{tip}\| &\le g\,F_{hold} \end{aligned} \tag{10.7}$$

where $e_p = 8\,\mathrm{mm}$, $e_v = 0.15\,\mathrm{m/s}$, $F_{hold} = 10\,\mathrm{N}$ is the holding force at full grip and g is the grip strength. A failed grasp is charged to the first condition that fails, in the order position, speed and grip.

A successful catch is inelastic and conserves generalized momentum, as shown in Equation (10.8).

$$M(q; m_p)\,\dot{q}^{+} = M(q; 0)\,\dot{q}^{-} + m_p J^{T}(q) \begin{bmatrix} v_b \\ 0 \end{bmatrix} \tag{10.8}$$

where $M(q; m)$ is the inertia matrix with mass m held. A gripper already moving at belt velocity therefore catches the product without any change in its joint rates, which a verification check confirms to within 10^{-12}. While a product is held it slips out, and counts as a drop, whenever $m_p \|a_{tip}\| > g F_{hold}$. On release the product is placed if the tip is within 12 mm of the nearest tray cavity and that cavity is empty, it is put back on the belt if the tip is over the belt, and otherwise it is dropped. The tray holds 2×6 cavities at a 45 mm pitch, and the cell program replaces it at once by an empty one after every twelfth release over it, placed or not, so a missed cavity cannot put the program's count out of step with the tray.

10.3.4 Parameters, Numerics and Scoring

The parameters are listed in Table 10.2. The Show Model dialog offers a slider for every parameter marked adjustable, the control panel holds the tuning and the options, and the parameters marked hidden set the geometry of the cell. The six states advance by fourth-order Runge-Kutta integration on a fixed 1 ms step, the controller runs every 2 ms, and the historian logs at 50 Hz, in the step order of Section 1.6.1, with the end stop and the discrete events handled after the integration. This chapter writes Δt for the integration step and h for the control period, the reverse of the symbols in that section. Under a fixed torque history, halving the step to 0.5 ms changes the joint angles after 1 s by 5.2×10^{-8} rad, and the next halving shrinks that difference by a factor of 14.1, close to the 16 expected of a fourth-order method. Case 6 run for 60 s at 0.5 ms reproduces the placed, missed, failed and dropped counts and all 11 torque-limit hits, with the energy within 2.3 J of 1397.4 J and the tracking error at 6.190 against 6.202 mm rms.

The scorecard counts products placed, missed, failed and dropped, together with the picks per minute, the torque-limit hits, the shaft work $\sum_i \int |\tau_i \omega_i|\, dt$ with braking counted as spent, and the rms distance of the tip from the tip position of its reference in automatic mode. A torque-limit hit is counted once per episode, and a new episode begins only after 50 ms clear of the limit. These scores use the true state, and the controller never reads them.

10.3.5 Assumptions

A planar arm. Both joints turn about vertical axes and the quill's dip to the belt is folded into t_{close} and t_{rel}. A real cell also plans that vertical axis, and its settling adds to the cycle.

Rigid links and a point payload. Real arms flex at tens of m/s^2, and a carton has rotational inertia of its own, which is the part of a load that is hardest to estimate [9].

Ideal drives up to a limit. Each drive delivers the commanded torque up to $\tau_{i,\max}$, and the speed ratings enter only as planner limits. A real motor loses torque with speed, and its current loop and gearbox add lag and backlash.

One clock and exact time stamps. The camera and the encoder share a clock and every frame is stamped exactly. On a real line, clock offset and trigger jitter between the two add a position error proportional to the belt speed.

An ideal recognizer. The camera reports positions only, with Gaussian noise, and finds every product in its view unless the whole frame is lost. Misrecognition, overlapping products and occlusion by the arm are absent.

A threshold grasp. Equation (10.7) either holds or not, and a slip loses the product at once. Real vacuum cups creep and skid first, which turns the three thresholds into a probability.

An exact controller model. The controller's arm model matches the plant apart from the payload switch. An identified model would leave some tracking error that only feedback removes.

Table 10.2 Model, Controller and Numerical Constants. Class A Is Adjustable in the Interface, H Is a Hidden Parameter and C Is Fixed in the Code

Symbol	Meaning	Value	Class
Arm and drives			
l_1, l_2	link lengths	0.40, 0.35 m	A
m_1, m_2	link masses	4.0, 2.0 kg	A
m_{tool}	gripper and quill at the tip	0.5 kg	A
$q_{2,\max}$	elbow travel	150°	A
$\tau_{1,\max}$, $\tau_{2,\max}$	torque limits	40, 17 N m	A
$\omega_{1,\max}$, $\omega_{2,\max}$	speed ratings (planner)	7.0, 12.0 rad/s	A
J_{1m}, J_{2m}	reflected rotor inertias	0.05, 0.015 kg m^2	A
b_1, b_2	viscous friction	0.6, 0.3 N m s/rad	A
f_{c1}, f_{c2}	Coulomb friction	0.6, 0.3 N m	A
—	joint encoder count	$2\pi/131072$ rad	H
Line and camera			
y_b, w_b	belt centerline, width	0.55, 0.20 m	H
τ_b	belt drive time constant	0.5 s	A
Δs	belt encoder count	0.1 mm	A
—	entry point, reject bin	−1.45, 0.95 m	H
—	minimum spacing on the belt	0.07 m	A
f_{cam}	camera frame rate	15 Hz	A
x_{c0}, x_{c1}	camera field of view	−0.80, –0.40 m	H
Gripper and tray			
t_{close}, t_{rel}	closing and release times	0.10, 0.08 s	A
e_p, e_v	grasp tolerances	8 mm, 0.15 m/s	A
F_{hold}	holding force at full grip	10 N	A
x_t, y_t, p_t	tray center and pitch	0, 0.30, 0.045 m	H
—	placement tolerance	12 mm	C
Controller			
K_{p1}, K_{p2}	position gains	400, 100 N m/rad	A
K_{d1}, K_{d2}	rate gains	25, 6 N m s/rad	A
$\alpha_{1,\max}$, $\alpha_{2,\max}$	acceleration limits	40, 80 rad/s^2	A
a_{carry}	carry acceleration limit	30 m/s^2	A
v_a	assumed belt speed	0.15 m/s	A
—	torque share for planning	90 %	C
τ_f	joint rate filter	1 ms	C
—	replan and schedule period	20 ms	C
Numerics			
Δt, h	integration step, control period	1, 2 ms	C
—	historian rate	50 Hz	C

10.4 Four Budgets for One Grasp

Equation (10.7) is what makes conveyor picking different from trajectory tracking. The three conditions must hold at one instant, each is met by a different part of the controller, and each yields a bound that decides feasibility before any gain is tuned. A fourth budget, the drive torque, is where the payload enters.

10.4.1 Position and the Age of a Report

A report describes the product at its capture time. An estimator that aims at the report as if it were current sends the gripper to where the product was, and when the grasp completes the product has moved on by the distance of Equation (10.9), the error vL of Equation (1.11) with the age carried on to the grasp.

$$d = v_b\,(t_g - t_{cap}) = v_b\,(L + T_d) \tag{10.9}$$

where T_d is the time from the delivery of the report to the completed grasp. T_d contains at least the closing time, so even an instantaneous camera leaves $d \geq v_b t_{close}$. With $e_p = 8\,\text{mm}$ and $t_{close} = 0.10\,\text{s}$, an aim that does not predict the closing interval fails above $e_p / t_{close} = 0.08\,\text{m/s}$ whatever the latency. In case 4 on a 0.15 m/s belt, $T_d = 0.42\,\text{s}$ on average, and over latencies from 0 to 0.5 s the distance from the gripper to the product it was sent for follows $d = 64.4 + 148.2L$ in mm with no residual above 3.2 mm, a slope 1.2 % below the belt speed, as shown in Figure 10.2. The two points that sit 3 mm above the line, at 0.15 and 0.20 s, are runs in which one grasp of 42 lands on a following product, which shifts the later attempts and raises their mean T_d to 0.45 s. The remedy is the time stamp itself. Moved on by the belt travel since t_{cap} and predicted to t_g, the estimate places 24 of 25 products at every latency in that range, with a mean position error at the grasp between 0.81 and 1.06 mm.

> **TRY IT IN THE APP** **4 · Camera Delay**
>
> Load the case and press Run. The controller aims at the camera's last report, which arrives 300 ms late. Watch where the gripper closes and where its product is by then: the gripper closes on empty belt behind it, and none of 42 grasps in the first minute take. Press Reset, set Camera latency to 0 and run again. The gripper still closes behind its product.

10.4.2 Relative Speed at the Grasp

A gripper that arrives early and waits at rest meets the product at a relative speed equal to the belt speed, so a rest intercept is feasible only while $v_b < e_v$. Case 2 at 30 arrivals a minute shows a sharp edge, with the belt speed set before the run so that the belt starts at that speed. Waiting at rest places every product that reaches the arm, with no failed grasp, up to 0.14 m/s, fails 12 grasps at 0.15 m/s, where the relative speed at the grasp spans 0.1476 to 0.1509 m/s about the tolerance, and places none from 0.16 m/s upward. Across the sweep the mean relative speed is 0.98 to 1.002 times the belt speed. Arriving at belt velocity instead places 23 to 30 products at every speed from 0.05 to 0.50 m/s, with a mean relative speed of 0.0024 to 0.0032 m/s at the grasp.

> **TRY IT IN THE APP** **2 · Higher Belt Speed**
>
> Run the case as loaded. The arm arrives early and waits at rest, and every grasp fails on speed as the product slides under the gripper at the belt's 0.30 m/s. Press Reset, switch Meet the chocolate to Move with it and run again. The gripper now follows each product while it closes, and 30 are placed in the first minute with no failure.

10.4.3 Holding Force During the Carry

Once a product is held, the only force that can pull it off the cup is its own acceleration, so the carry must keep $\|a_{tip}\| \leq g F_{hold} / m_p$. For 0.10 kg at 15 % grip the bound is $15\,\text{m/s}^2$, and for the 0.6 kg boxes of the challenge at 60 % it is $10\,\text{m/s}^2$. The carry limit a_{carry} bounds the planned acceleration, while the real one also carries the tracking error, so the usable limit sits below the bound. In case 7 the largest force any carry needs is 1.445 N at a $14\,\text{m/s}^2$ limit and 1.230 N at $12\,\text{m/s}^2$, 3.2 and 2.5 % above the planned value, and a $15\,\text{m/s}^2$ limit drops 20 of 24 products.

10.4.4 Torque and the Payload

The payload also enters the torque budget, because it adds $m_p J^T J$ to M. A 1.5 kg case raises M_{11} at the ready pose by 62 % and M_{22} by 116 %. The planner checks every trajectory against 90 % of each drive with the controller's own arm model, so a model that leaves the case out approves carries that the drives cannot deliver.

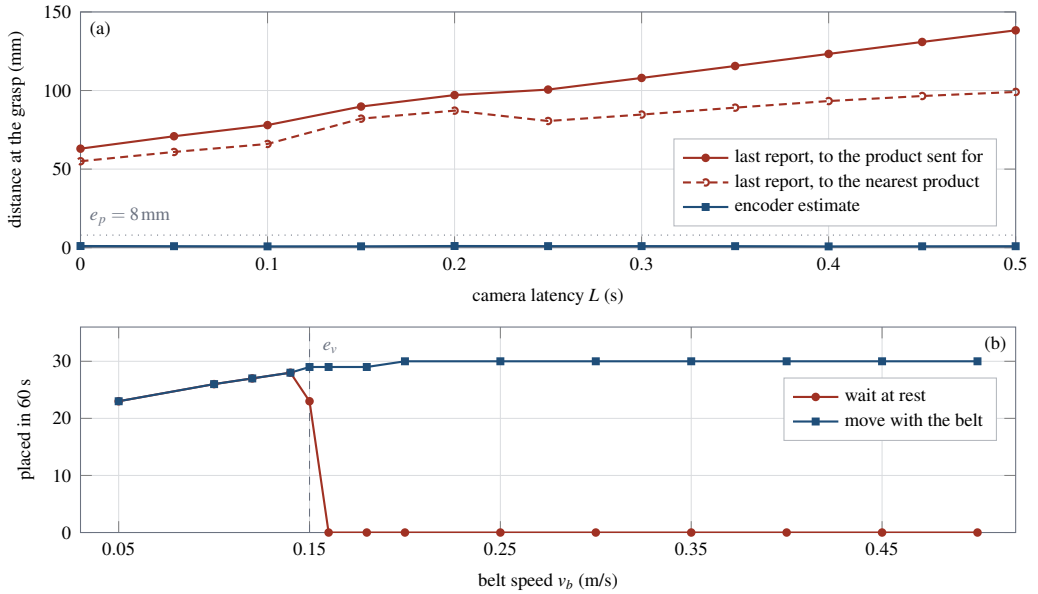

Figure 10.2 Grasp outcome against camera latency and belt speed, measured on the application. (a) Case 4 at 0.15 m/s. Aimed at the last report, the gripper closes behind the product it was sent for (solid, filled circles) by $64.4 + 148.2L$ in mm. The application's own grasp error, the distance to the nearest product (dashed, open circles), grows more slowly because a following product is sometimes nearer. Moved on by the encoder (squares), the error stays between 0.81 and 1.06 mm and 24 of 25 products are placed at every latency. (b) Case 2 at 30 arrivals a minute, the belt speed set before each run. Waiting at rest (circles) places nothing once the belt runs faster than the relative speed tolerance e_v (dashed line), while arriving at belt velocity (squares) places 23 to 30 at every speed.

10.5 Implemented Controller

The controller reads only its own sensors, which are the joint encoders, the belt encoder, the delivered camera reports and the part-present switch in the gripper. A verification check wraps the simulated world so that any other read by the control law fails, and it passes over four scenarios. The controller is built in four layers, as shown in Figure 10.3. None of them is a PID loop. The application does not use the velocity-form PID of Section 1.6.2, and with no integral action anywhere there is nothing for anti-windup to act on.

10.5.1 Sensors

The joint rates are a filtered backward difference of the encoder angles, $\hat{\omega}_k = \mathrm{e}^{-h/\tau_f}\,\hat{\omega}_{k-1} + (1 - \mathrm{e}^{-h/\tau_f})(q_{m,k} - q_{m,k-1})/h$, with $h = 2$ ms and $\tau_f = 1$ ms. The belt speed $\hat{v}_b$ is the slope of the belt encoder count over the last 50 ms, passed through a first-order filter of 0.1 s. Every belt encoder count is also pushed, with its time, into a history of 1024 samples at the control period, which covers 2.05 s.

10.5.2 Estimator and Tracks

Three estimates of a product's position are offered, Equation (10.10).

$$\hat{x}(t) = \begin{cases} x_{rep} & \text{last report} \\ x_{rep} + v_a\,(t - t_{cap}) & \text{assumed speed} \\ x_{rep} + s_m(t) - s_m(t_{cap}) & \text{encoder} \end{cases} \tag{10.10}$$

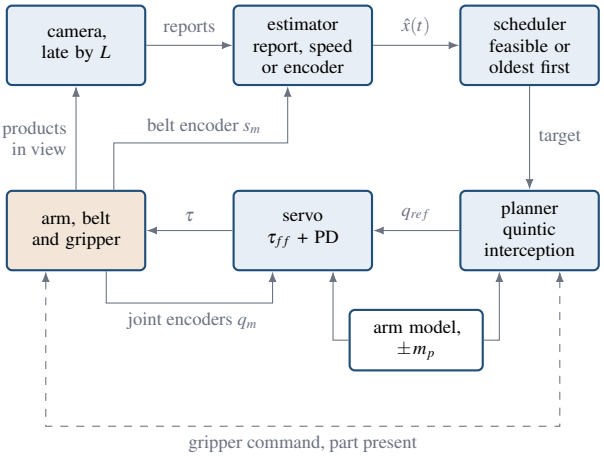

Figure 10.3 Layers of the picking controller. The estimator turns late camera reports and the belt encoder count into positions at any time, the scheduler chooses the next product, the planner builds a minimum-time quintic that meets it, and the servo follows that reference with computed-torque feedforward and a PD correction. One arm model, with or without the payload m_p, serves both the feedforward and the planner's torque check. The task sequence in the planner commands the gripper and reads the part-present switch.

where $s_m(t_{cap})$ is read from the encoder history by linear interpolation and v_a is a typed-in belt speed. A stored history serves the same purpose in Chapter 9, which forms each innovation of its pose filter against the pose stored for the moment of the measurement. The last two hold each product as a track in belt coordinates, $u = x_{rep} - s_m(t_{cap})$ or $u = x_{rep} - v_a t_{cap}$, where a product does not move, and average successive reports as $u \leftarrow u + k_n(u_{rep} - u)$ with $k_n = \max(1/n, 0.25)$ after the nth report. The average is exact for the first four reports and exponential after that, and its stationary variance under independent noise is $0.25/(2-0.25) = 1/7$ of the camera's. The 2 mm camera noise therefore becomes 0.76 mm per axis and a mean distance of 0.95 mm, against the 0.99 mm measured in case 1. Predicting to a later time t_a, the encoder track adds $\hat{v}_b(t_a - t)$. A report joins the track whose position at t_{cap} lies within 30 mm along the belt and 25 mm across it, or it starts a new track. The last-report estimator matches a report to a track when the report lies between 20 mm upstream and 80 mm downstream of that track's previous report. A change of estimator or of v_a restarts every track and abandons an approach in progress.

10.5.3 Scheduler

While the arm is free, the scheduler runs every 20 ms over the tracks not yet picked and tried fewer than twice. Feasible first plans an interception for each candidate, drops any it cannot meet inside the reach before the grasp finishes, and takes the one with the earliest deadline, the time left before the product passes $x_{out} - \hat{v}(t_{close} + 0.03\,\mathrm{s})$, where $\hat{v}$ is the belt speed the estimator uses ($\hat{v}_b$ or v_a). Oldest first takes the lowest track number whether it can be caught or not. With the last-report estimator the scheduler takes the most downstream report inside the reach.

10.5.4 Planner

Every move is one quintic polynomial per joint, from the present reference state of angle, rate and acceleration to a goal state, over one duration T shared by both joints. The reference is therefore continuous in acceleration across replans, and the feedforward torque never steps [16, 17]. A rest-to-rest move of Δq peaks at the rate $1.875\,\Delta q/T$ and the acceleration $(10/\sqrt{3})\,\Delta q/T^2 = 5.774\,\Delta q/T^2$ of Section 1.4.4, which gives the starting duration of Equation (10.11), the quintic duration of

Equation (1.5) for each joint with two more bounds.

$$T_0 = \max_i \left\{ 0.12,\ \sqrt{\frac{5.774\,|\Delta q_i|}{\alpha_{i,\max}}},\ \frac{1.875\,|\Delta q_i|}{\omega_{i,\max}},\ \frac{|\Delta \dot{q}_i|}{\alpha_{i,\max}} \right\} \tag{10.11}$$

where T_0 is in seconds. T is then rescaled, at most ten times, by the largest of the ratios

$$\sqrt{\frac{\max|\ddot{q}_i|}{\alpha_{i,\max}}},\quad \frac{\max|\dot{q}_i|}{\omega_{i,\max}},\quad \sqrt{\frac{\max|\hat{\tau}_i|}{0.9\,\tau_{i,\max}}},\quad \sqrt{\frac{\max\|a_{tip}\|}{a_{carry}}}$$

the last only while a product is held, each peak taken from 21 samples of the polynomial, until the tightest ratio lies between 0.97 and 1, and never beyond $8T_0$. The torque $\hat{\tau}$ comes from the controller's arm model, which is how a payload left out of that model escapes the check.

The goal is where the product will be when the grasp takes, which depends on T, which depends on the goal. Up to eight fixed-point iterations settle Equation (10.12), because the product moves slowly compared with the arm [4]. A table tennis ball moves fast compared with the arm, and Chapter 11 predicts its crossing point and swings within the time left.

$$T = T_{\min}\big(q_{ref} \to q(\hat{p}(t+T+\Delta))\big) \tag{10.12}$$

where $\hat{p}$ is the predicted product position and Δ is the delay from arrival to grasp. Waiting at rest, the arm arrives 0.03 s early and $\Delta = 0.03 + t_{close}$, so the product reaches the stationary gripper as the grasp takes. Moving with the belt, $\Delta = 0$, the goal carries the joint rates $\dot{q} = J^{-1}[\hat{v}, 0]^T$, with $\hat{v}$ the estimator's belt speed as in the scheduler, and the accelerations that keep the tip on a straight line, and while the gripper closes the reference follows the product along the belt at $\hat{v}$, stopping at the edge of the reach. If the product has not yet entered the reach, the arrival is delayed until it does. An arrival more than 0.25 s beyond the minimum time is made in two moves, first at minimum time to a staging point under the meeting point, or 0.12 s of belt travel behind it when moving with the belt, then a final approach launched when the time left falls to its own minimum plus 40 ms. In case 1, 13 of 19 approaches were made this way. Every 20 ms the goal is refined toward the latest estimate at the same arrival time, except in the last 60 ms, and a correction that would break a limit is absorbed by arriving later. A final approach more than twice over a limit is abandoned and the arm holds, which happens only at speed ratings too low to follow the belt.

10.5.5 Servo

Each joint follows its reference with the law of Equation (10.13), executed every $h = 2$ ms.

$$\tau_i = \tau_{ff,i} + K_{p,i}\,(q_{ref,i} - q_{m,i}) + K_{d,i}\,\big(\dot{q}_{ref,i} - \delta\,\ddot{q}_{ref,i} - \hat{\omega}_i\big) \tag{10.13}$$

where $\tau_{ff} = \hat{M}\ddot{q}^* + \hat{C}\dot{q}^* + \hat{F}$ is the inverse dynamics of the controller's arm model, the reference feedforward of Section 1.4.3 written for the arm, friction included, along the reference q^* half a control period ahead, because the torque computed now is held for the whole period, and along the reference itself while the gripper closes. The shift $\delta = h/2 + \tau_f = 2$ ms moves the reference rate back to the instant that the filtered backward difference $\hat{\omega}_i$ describes. The arm model either knows only the arm or also counts the product, at its recipe mass, while the part-present switch is on. The command is clipped at $\pm\tau_{i,\max}$, the only output limit, with no rate limit on the torque. The defaults are $K_{p1} = 400$ N m/rad, $K_{d1} = 25$ N m s/rad, $K_{p2} = 100$ N m/rad and $K_{d2} = 6$ N m s/rad. Against the diagonal inertias at the ready pose they give a shoulder loop of 20.0 rad/s at a damping ratio of 0.63 and an elbow loop of 25.2 rad/s at 0.75, and a 1.5 kg case lowers these to 15.7 rad/s at 0.49 and 17.1 rad/s at 0.51. In the convention of Section 1.3 each loop is reverse acting ($a = +1$), because an angle short of its reference receives a positive torque and a positive torque accelerates the joint forward, as the application's sign checks confirm. The application gives the loops no action label. There is no integral term, because with no gravity on the links there is no steady torque for one to find, and the feedforward already supplies the inertia and friction torques the motion needs.

Table 10.3 Scenario Sequence and the Changes Each Makes to the Base Line

#	Change	Teaching point
1	0.10 m/s, 20 min^{-1}	the cell working, report against product
2	0.30 m/s, 30 min^{-1}	speed at the grasp
3	0.20 to 0.30 m/s at 20 s, assumed 0.20 m/s, move with it	a typed-in speed against the encoder
4	$L = 0.30$ s, last report, move with it	the age of a report
5	0.25 m/s, 30 min^{-1}, 50 % lost, 3 s blackout at 27.9 s, move with it	one look is enough, none is not
6	1.5 kg, $F_{hold} = 60$ N, move with it	the payload in the model
7	$g = 15$ %, move with it	carry acceleration against grip
8	0.25 m/s, 20 min^{-1}, 12 more at 10 s, oldest first, move with it	scheduling and the ceiling
9	100 boxes of 0.6 kg at 36 min^{-1}, 0.22 to 0.30 m/s at 60 s, $L = 0.2$ s, 10 % lost, $g = 60$ %, assumed 0.22 m/s	graded production run

10.5.6 Manual Mode and Transfer

In manual mode the operator sets τ_1 within ±40 N m and τ_2 within ±17 N m, and the arrow keys add ±3 N m and ±1 N m while held. The scheduler, planner and gripper stay idle, the estimator keeps tracking, and the reference follows the measured arm. A return to automatic mode plans a stop from the reference of the last manual scan, so the initial servo error is only the arm's travel over one control period, 2.7 mrad for a shoulder turning at 1.35 rad/s. In automatic mode the manual torques follow the clipped command after every scan, so a switch to manual starts from the last automatic torque. Manual mode exposes the inertial coupling of Equation (10.3). At the ready pose, 1 N m at the shoulder accelerates the shoulder at 1.57 rad/s^2 and the elbow at −2.39 rad/s^2, while 1 N m at the elbow accelerates the elbow at 9.95 rad/s^2 and the shoulder at −2.39 rad/s^2.

10.5.7 Plant, Controller and Analysis Models

Equations (10.1) to (10.8), with the true payload, friction and limits, make up the simulation model. The controller's model is the same arm dynamics, exact apart from the payload switch and used by both the feedforward and the torque check, together with one of the three position estimates of Equation (10.10). The analysis model is the set of budgets of Section 10.4 and the quintic peak coefficients, which this chapter uses to read the measurements and which appear nowhere in the source. No model is identified from data.

10.6 Scenarios and Measured Results

Table 10.3 lists the eight cases and the challenge. Every case starts from the same line, a 0.15 m/s belt, 24 arrivals a minute, 150 ms of latency, 5 % of frames lost, 2 mm of noise, 0.10 kg products and full grip, with the encoder estimator, waiting at rest, the arm-only model, feasible-first scheduling and a carry limit of 30 m/s^2. The values in Table 10.4 were measured by running the application's own engine at its default seed over the first 60 s unless another span is named, with each change made before the run starts. Position errors are means over the grasps concerned, the failed ones where a case fails, and estimate errors are means over the logged samples in which a target is assigned. Over six other seeds the counts move by a few products, and every comparison below keeps its direction except where the text says otherwise. Figure 10.4 draws the products placed as bars.

10.6.1 The Cell Working, and the Age of a Report

On the slow belt of case 1 the cell places 18 of the 24 products that arrive in the first minute, and the other six are still on their way. No product is missed and no grasp fails. While the arm goes for a product, the camera's last report is 20.2 mm behind it on average and the encoder estimate is 0.99 mm from it. Switched to the last report, the gripper closes 46.8 mm from the nearest product on average

Table 10.4 Measured Results

Case and condition	Result
1, encoder, at rest	18 of 24 placed, none missed or failed; estimate 0.99 mm, last report 20.2 mm
1, last report	0 placed, 32 failed, 17 missed; 46.8 mm from the nearest product
2, at rest	0 placed, 60 failed on speed at 0.3005 m/s; position 0.96 mm
2, move with it	30 placed, none failed; 0.0027 m/s; 0.61 mm rms
2, move with it, feedforward off	30 placed, 15 failed, 1 hit; 16.1 mm rms
3, assumed speed, first 20 s	9 of 11 placed, none failed
3, assumed speed, 60 s	10 placed, 40 failed after the change, 14 missed; 99.7 and 120.4 mm
3, encoder	24 of 25 placed, none failed or missed
4, last report, $L = 0.30$ s	0 placed, 42 failed, 23 missed; 84.7 and 108.0 mm
4, last report, $L = 0$	0 placed, 42 failed; 55.0 and 63.0 mm
4, encoder, $L = 0.30$ s	24 of 25 placed, none failed; 0.99 mm
5, 50 % lost	27 placed, 3 missed, all 3 never seen; 1.28 mm
5, none lost	27 placed, 3 missed; 1.07 mm
6, arm only	24 placed, 11 hits; 47.1 and 24.3 N m; 6.20 mm rms, 43.7 mm worst
6, arm + payload	24 placed, 0 hits; 0.26 mm rms, 0.75 mm worst
6, arm only, $\alpha_{\max}$ 25 and 50	24 placed, 0 hits; 4.93 mm rms
7, $a_{carry} = 30$	0 placed, 24 of 25 dropped
7, $a_{carry} = 15$	4 placed, 20 dropped
7, $a_{carry} = 14$	24 placed, 0 dropped; peak 1.445 N
7, $a_{carry} = 12$	24 placed, 0 dropped; peak 1.230 N
8, oldest first	22 placed, 8 missed, 16 failed; 145.6 mm
8, feasible first	25 placed, 5 missed, none failed
2, 120 min^{-1}, move with it, 90 s	105 placed, 70.0 min^{-1}; 46 of 161 missed
9, as shipped	0 placed, 224 failed, 100 missed, 1 hit
9, move with it	0 placed, 35 dropped, 158 failed
9, and $a_{carry} = 8$	35 placed, 34 of them before 60 s; 158 failed
9, and encoder	98 placed, 2 missed, 0 failed, 0 dropped, 0 hits; 1.97 mm rms
9, and arm + payload	98 placed, 2 missed, 0, 0, 0; 0.25 mm rms; ends at 180.2 s

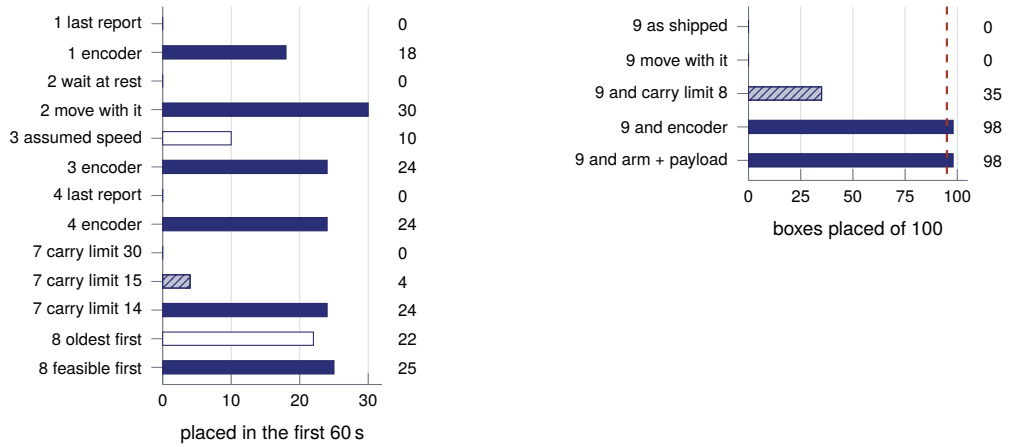

Figure 10.4 The products placed in Table 10.4 drawn as bars, labeled by case and setting. Left, the first minute of cases 1 to 4, 7 and 8: white bars are the setting that fails and solid bars the one that works, and the hatched bar is the carry limit equal to the bound the grip suggests. Right, the production run as each change is added to the one before, with the 95 boxes it asks for as a dashed line and the hatched bar placing only part of the run. Notice how far a single setting moves each count, and that the run is met only once the encoder estimator joins the first two changes.

and all 32 grasps fail. The latency contributes 15 mm of that, and the rest is the belt moving while the arm travels and the gripper closes, as Equation (10.9) predicts.

Case 4 raises the latency to 300 ms. Aimed at the last report, the gripper closes 84.7 mm from the nearest product and none of 42 grasps take (Table 10.4), and Figure 10.5 shows one of them closing

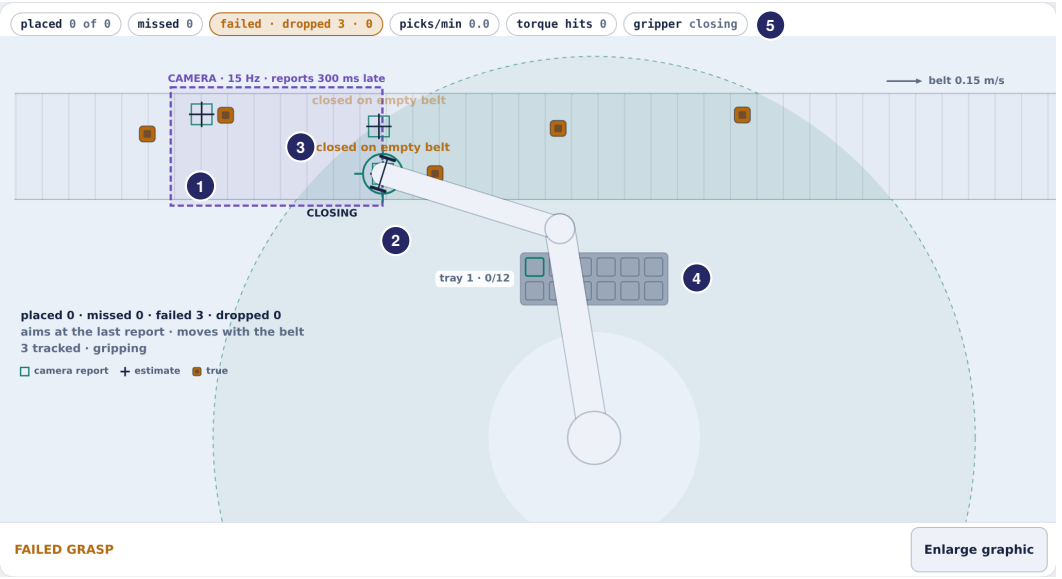

Figure 10.5 The picking application in scenario 4, Camera delay, cropped to the plant view, at $t = 8.20$ s with the settings the scenario loads, so the controller aims at the camera's last report. (1) The 15 Hz camera reports 300 ms late. (2) The gripper is closing inside the ring that marks its aim, over the last report of the product it was sent for, while that product has ridden 98.7 mm further along the 0.15 m/s belt. (3) Both floating labels read closed on empty belt. (4) The tray waits empty, and (5) the badges count three failed grasps, nothing placed and no torque-limit hits.

on empty belt. The application's grasp error is the distance to the nearest product on the belt, and the product each grasp was sent for is 108.0 mm away on average. With the latency set to zero the two distances fall to 55.0 and 63.0 mm and still no grasp takes. The distance to the product sent for falls by 45.0 mm, which is the belt speed times the latency, while the distance to the nearest product falls by only 29.7 mm, and the fit of Figure 10.2 extends the same arithmetic over eleven latencies. With the encoder estimator at 300 ms, 24 of 25 products are placed with no failed grasp and a position error of 0.99 mm, because the encoder count at the report's time stamp comes from the stored history.

🕮 **Pause.** Try Exercise 10.1 before reading on.

10.6.2 Speed at the Grasp

Case 2 runs the belt at 0.30 m/s with the arm still arriving early and waiting at rest. The gripper closes 0.96 mm from each product on average, and all 60 grasps in the first minute fail on relative speed, the product sliding under the gripper at 0.3005 m/s against the 0.15 m/s tolerance. On the slow belt of case 1 the same wait left 0.0994 m/s, inside the limit. Moving with the belt, the relative speed at the grasp falls to 0.0027 m/s and 30 products are placed with no failure. Switching the torque feedforward off in this case raises the tracking error from 0.61 to 16.1 mm rms, costs 15 failed grasps, 11 of them on position and 4 on speed, and brings one torque-limit hit.

🕮 **Pause.** Try Exercise 10.2 before reading on.

10.6.3 A Typed-In Speed Against the Encoder

Case 3 runs the belt at 0.20 m/s with the controller assuming 0.20 m/s, and at 20 s the line speeds up to 0.30 m/s. All 9 products that reach the arm in the first 20 s are placed. After the change the estimate falls behind by the speed error times the age of each report, and by 60 s 40 grasps have failed, closing

99.7 mm from the nearest product and 120.4 mm from the one each was sent for, with 14 products missed. The encoder estimator measures what the belt did rather than what it was told, and over the same minute it places 24 of the 25 products that arrive with no failed grasp.

10.6.4 Lost Frames

Case 5 loses half of all frames on a 0.25 m/s belt, where a product crosses the field of view in 1.6 s, and the camera goes dark for 3 s at 27.9 s. Each product seen is reported 12.0 times on average, against 23.5 with no frames lost. Twenty-seven products are placed and three are missed, and all three crossed the field of view in the dark and were never reported. No grasp fails, and the estimate stays 1.28 mm from its product on average against 1.07 mm with no frames lost, where the count is again 27 placed and 3 missed. Over six other seeds the misses are always the products never reported, and at one seed losing half the frames costs a single miss that the full stream avoids. A product seen once is carried forward by the encoder for as long as it rides the belt, and a product never seen cannot be picked by any estimator.

> ✦ DID YOU KNOW? **Half the Frames, the Same Count**
>
> From the simulator: in case 5, losing half of all camera frames costs no product. The arm places 27 and misses 3 either way, and the estimate moves only from 1.07 mm to 1.28 mm from its product (Table 10.4). Once a product has been seen, the belt encoder carries it.

10.6.5 Payload and Grip

Case 6 changes the product to a 1.5 kg case held by a 60 N cup, with the controller's arm model still knowing only the arm. The planner believes the carries fit inside the drives and they do not. The commands reach 47.1 N m at the shoulder and 24.3 N m at the elbow against limits of 40 and 17 N m, the torque limit is hit 11 times, and the tip lags its reference by 6.20 mm rms, 43.7 mm at worst. With the payload in the model there are no hits and the tracking error is 0.26 mm rms, 0.75 mm at worst. Lowering the acceleration limits to 25 and 50 rad/s^2 instead also stops the hits, at six of seven seeds, while the lag stays at 4.93 mm rms because the feedforward still leaves the case out. Case 7 cuts the grip to 15 % of 10 N. At the default carry limit of 30 m/s^2, 24 of the 25 products that arrive are dropped and the other is still on the belt. At 15 m/s^2, 20 are still dropped, and at 14.5 or 14 m/s^2 all 24 reach the tray. Over seven seeds a 15 m/s^2 limit drops 8 to 23 products and 14 m/s^2 drops none.

🕮 **Pause.** Try Exercise 10.3 before reading on.

10.6.6 Scheduling and the Throughput Ceiling

Case 8 sends a burst of twelve products at the minimum spacing into a stream of 20 a minute on a 0.25 m/s belt. Taking the oldest first, the arm spends 16 grasps on products it could no longer catch, closing 145.6 mm from the nearest one on average, and places 22 while 8 are missed. Feasible first skips what cannot be caught and places 25, with 5 missed and no failed grasp. Over seven seeds feasible first never fails a grasp while oldest first fails 8 to 18, and at one seed both place 31. Pushed with 120 arrivals a minute on a 0.30 m/s belt while moving with it, the arm places 105 products in 90 s, 70.0 min^{-1}, and misses 46 of the 161 that arrive. At that ceiling one pick takes 0.857 s, made up of 0.303 s of approach, 0.100 s of closing, 0.313 s of carry, 0.133 s of placing and release and 0.009 s of waiting.

10.6.7 The Production Run

The challenge sends 100 boxes of 0.6 kg at 36 a minute on a belt that speeds up from 0.22 to 0.30 m/s at 60 s, with 200 ms of latency, one frame in ten lost and the cup at 60 % of 10 N. It is met by placing at least 95 with no drop and no torque-limit hit. As shipped, waiting at rest with an assumed belt speed, it places none and fails 224 grasps and hits the torque limit once, and the run stops at 181.9 s when the last box is decided. The 68 failures before the speed change are on relative speed, and the

156 after it are on position while also moving too fast. Moving with the belt drops 35 boxes, because the full carry limit asks a 6 N grip for more than it holds. An 8 m/s^2 carry limit stops the drops and places 35, 34 of them before the speed change defeats the assumed speed. The encoder estimator brings the count to 98 of 100 with no failed grasp, and putting the payload in the arm model keeps 98 while cutting the tracking error from 1.97 to 0.25 mm rms over the run, which ends when the last box is decided at 180.2 s. The payload model also lowers the largest carry force from 5.46 to 4.89 N, that is from 9.10 to 8.15 m/s^2 against the 10 m/s^2 the grip allows. Over six other seeds the same four changes place 95 to 100 boxes, and the shipped settings place none.

10.7 Technology in Practice

The cell of Section 10.3 reads joint encoders, a belt encoder, a camera and a part-present switch, and it moves two joint drives and a gripper. Figure 10.6 draws its loops with real devices in the place of the schematic's symbols.

What Measures It

A robot synchronized with a moving conveyor is a conveyor tracking robot, and it needs frequent updates of the desired position on the conveyor. Line tracking is done either with the robot's base moving along the line or with the base fixed beside it, as here [18]. The camera finds the products, and the belt encoder says how far they have gone since. The camera of Equation (10.6), 15 Hz with a latency and 2 mm of noise, stands for such a vision system, and its lost frames for the frames a real one loses to its lighting. The first SCARA prototype measured each joint with an optical encoder [11], and the 17-bit counts of Table 10.1 stand for such encoders.

> **TECHNOLOGY** **Sensor: Machine Vision Camera**
>
> A vision system is one or more cameras interfaced to guide a robot to locate and identify a part. On a conveyor the camera finds each product and sends its offset to the robot controller, which combines it with the conveyor's speed to track the product until it comes within reach. One published line of this kind picked sachets from random places at 50 a minute to an accuracy of ±3 mm. Lighting is the usual trouble: moving parts, factory lights, the air and daylight from windows can upset a system that worked in the laboratory, and parallax and the match between camera and robot data call for a high mount and a calibration [18]. On a trend, a lighting fault looks like lost frames: the reports stop and the estimate rides on the belt encoder.

What Moves It

Each joint of that prototype was driven by a direct-current servo motor through a strain-wave (harmonic) gear reducer, and vibration in its servo motion led to a motor with tachogenerator feedback in the second [11]. The torque limits of 40 and 17 N m in Equation (10.1) stand for such drives, with the current loop and the gearbox taken as ideal (Section 10.3.5). The gripper of Equation (10.7), closing in 0.10 s and holding up to 10 N at full grip, stands for a vacuum cup, whose holding force must exceed the forces of the carry.

> **TECHNOLOGY** **Final Element: Vacuum Cup**
>
> A vacuum cup is a pneumatic pickup that holds a part by suction from a venturi or a vacuum pump, and it lets the part go when the vacuum supply is switched off. Cups are round or oval, from 12.7 mm to 101.6 mm across, and made of silicone, other rubbers or polyurethane. Choosing one takes experience and sometimes trial and error, and a hole in the part, a rough surface or a finish that must not be marked can rule a cup out. The force a grip must supply is the sum of the forces from mass, inertia and the Coriolis and centrifugal effects of the arm's motion, all of which change as the robot moves [18]. On a trend, a product lost in the carry is a part-present signal that drops out before the tray.

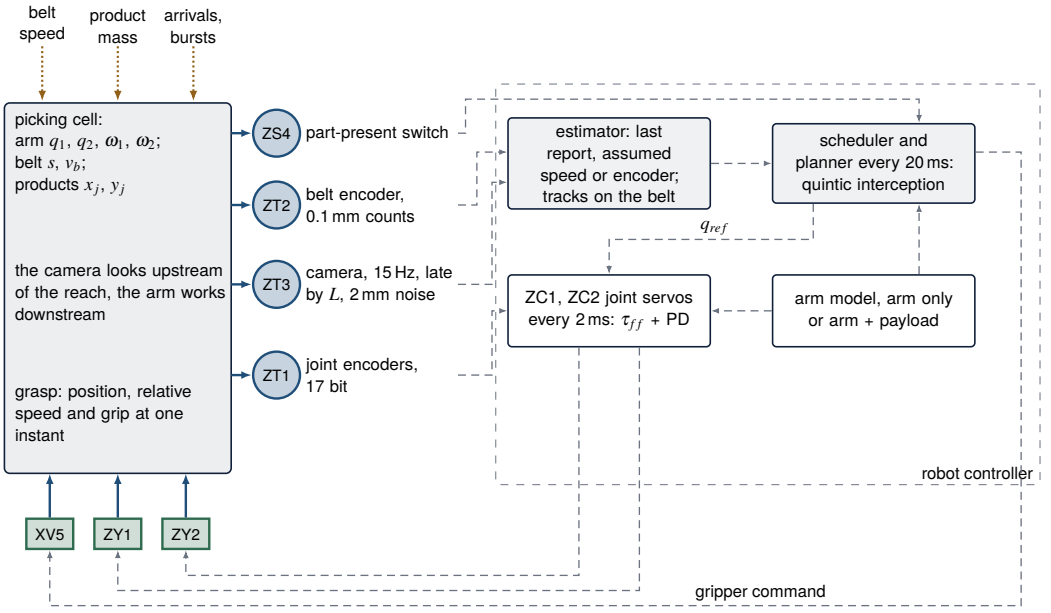

ZY1, ZY2: shoulder and elbow drives, a servo motor and gear reducer on each joint, limited to 40 and 17 N m. XV5: the vacuum valve of the gripper cup, which closes in 0.10 s and holds up to 10 N at full grip.

Figure 10.6 The loops of the chapter with real devices in place of the schematic's symbols. ZT1 counts the joint angles, ZT2 the belt travel, ZT3 is the camera and ZS4 the part-present switch. In the robot controller (dashed frame) the estimator feeds the scheduler and planner, which sets the joint references q_{ref} of the servos ZC1 and ZC2 and commands the gripper valve XV5. The servos drive ZY1 and ZY2 with torques computed from the arm model. Notice that the camera and the belt encoder meet only in the estimator, where the time stamp of each report lets the encoder count move it on.

Where the Controller Runs

The servo of this chapter runs every 2 ms, and the scheduler and the replanning every 20 ms. In discrete manufacturing a supervisor for each cell may coordinate a machine controller and a robot controller under the plant's supervisory loop, and such loops run with cycle times from milliseconds to minutes [19]. The first SCARA motions were computed by an 8-bit microcomputer, which shaped a motion curve with a deceleration interval twice its acceleration interval to prevent vibration [11]. The structure of this chapter lives in the robot controller: the computed-torque servo is its fast layer, and the estimator, the scheduler and the quintic planner sit in the layer above it.

What Is Modernizing It

Three-dimensional machine vision has lagged two-dimensional cameras because it needs light sources and calibration built into the sensor, and better auto-calibration routines are widening its use. Recognition based on geometry tolerates changes of lighting, focus and shape, and it is key to bin picking and tracking [18]. Lightweight arms with seven or more joints, often with force sensing in each joint, are now on the market, and robot hands are beginning to enter manufacturing [20]. Corrections that restore the rigid-body condition between camera and robot coordinates make the registration of the next segment more accurate [21].

IN PRACTICE **Registering the Camera to the Arm**

Before a vision-guided cell runs, the camera's coordinates are tied to the robot's. The common method measures the same points, called fiducials, in both frames and fits one rigid transformation between them. Noise and bias in those points then shift every target the robot is sent to. In NIST tests on two arms, corrections interpolated from fiducials near each target cut the root-mean-square target error to less than 20 % of its uncorrected value [21]. Lighting is settled on the line itself, because laboratory trials rarely match production [18]. The camera of Equation (10.6) reports in the arm's own coordinates with noise and no bias, so this step is absent from the model.

10.8 Limitations

The arm is planar and rigid with a point payload and ideal drives, so the vertical axis, link flexibility, the torque-speed curve and backlash are all absent. The camera and the encoder share a perfect clock, and recognition never fails inside the field of view except by losing a whole frame. The grasp is a threshold and a slip is total, which makes each failure mode clean but sharper than a real cup. The scheduler looks one product ahead and does not optimize an order of several or share the belt with another arm [8]. The numerical targets are the application's own design choices, not measured industrial performance.

One reporting detail affects how the numbers are read. The grasp error shown in the trends is the distance to the nearest product, which understates the miss when products are closely spaced, and this chapter reports the distance to the product each grasp was sent for beside it. The joint loops report no overshoot or settling time, because their reference moves a little at every sample and a step response is never defined.

OPEN CHALLENGE **Testing a Grasp on a Moving Part**

A NIST roadmap for robot dexterity lists grasping objects while in motion among the capabilities that need test methods. Its first test methods were to assume structured and static conditions, with no picking from a conveyor belt, and it notes that grasping cannot be separated from perceiving the object. Only in the longer term would tests cover cluttered, mixed parts that may be moving [20]. A moving grasp is hard to test because it couples the camera, the estimate, the arm and the gripper at one instant, as Equation (10.7) does. Progress would be a repeatable test that scores a whole picking cell against a stream of moving parts, as the production run of this chapter does for its model.

10.9 Conclusion

A picking robot acts on camera reports that are already old, and the simulator of this chapter makes that age, and what it costs, measurable. Aimed at the last report, a gripper on a 0.15 m/s belt closes behind the product it was sent for by $64.4 + 148.2L$ in mm, so removing the latency still leaves 63.0 mm, while moving each report on by the encoder travel since its time stamp places 24 of 25 products at every latency up to 0.5 s. Waiting at rest in the right place fails every grasp above the 0.15 m/s relative speed the gripper takes, and arriving at belt velocity places 30 of 32 at 0.30 m/s with 0.0027 m/s left at the grasp. A payload left out of the controller's model drives the torque commands to 47.1 and 24.3 N m against limits of 40 and 17, and a 1.5 N grip needs a carry limit of 14.5 m/s^2 or less rather than the 15 m/s^2 its holding force suggests. The arm and the grasp are idealized, and the clock is perfect. A fuller model would add the vertical axis, a torque-speed curve, clock offset between the camera and the encoder, recognition errors and a second arm sharing the belt. The application is freely available practice in meeting a moving object where it will be and at the speed it is going.

KEY POINTS

- A picking arm acts on camera reports that are already old, so it must meet each product where it will be and at the speed it is going.
- Aimed at the last report, a gripper still closes 63.0 mm behind its product with no latency, while the encoder estimate places 24 of 25 at every latency up to 0.5 s.
- Waiting at rest fails every grasp above the gripper's 0.15 m/s, and arriving at belt velocity places 30 of 32 at 0.30 m/s.
- A payload left out of the arm model drives the torques past their limits.

Where else this applies. The same age of a measurement is met in Volume 5, Chapter 13 and Volume 5, Chapter 7.

References

[1] S. Hutchinson, G. D. Hager, and P. I. Corke, "A tutorial on visual servo control," *IEEE Transactions on Robotics and Automation*, vol. 12, no. 5, pp. 651–670, 1996.

[2] P. K. Allen, A. Timcenko, B. Yoshimi, and P. Michelman, "Automated tracking and grasping of a moving object with a robotic hand-eye system," *IEEE Transactions on Robotics and Automation*, vol. 9, no. 2, pp. 152–165, 1993.

[3] T. D. Larsen, N. A. Andersen, O. Ravn, and N. K. Poulsen, "Incorporation of time delayed measurements in a discrete-time Kalman filter," in *Proceedings of the 37th IEEE Conference on Decision and Control*, vol. 4, 1998, pp. 3972–3977.

[4] E. A. Croft, R. G. Fenton, and B. Benhabib, "Optimal rendezvous-point selection for robotic interception of moving objects," *IEEE Transactions on Systems, Man and Cybernetics, Part B (Cybernetics)*, vol. 28, no. 2, pp. 192–204, 1998.

[5] D. Hujic, E. A. Croft, G. Zak, R. G. Fenton, J. K. Mills, and B. Benhabib, "The robotic interception of moving objects in industrial settings: strategy development and experiment," *IEEE/ASME Transactions on Mechatronics*, vol. 3, no. 3, pp. 225–239, 1998.

[6] J. S. Dai, "Robotics and automation for packaging in the confectionery industry," in *Robotics and Automation in the Food Industry*. Woodhead Publishing, 2013, pp. 401–419.

[7] R. Mattone, L. Adduci, and A. Wolf, "Online scheduling algorithms for improving performance of pick-and-place operations on a moving conveyor belt," in *Proceedings of the 1998 IEEE International Conference on Robotics and Automation*, vol. 3, 1998, pp. 2099–2105.

[8] H. I. Bozma and M. E. Kalalıoğlu, "Multirobot coordination in pick-and-place tasks on a moving conveyor," *Robotics and Computer-Integrated Manufacturing*, vol. 28, no. 4, pp. 530–538, 2012.

[9] C. G. Atkeson, C. H. An, and J. M. Hollerbach, "Estimation of inertial parameters of manipulator loads and links," *The International Journal of Robotics Research*, vol. 5, no. 3, pp. 101–119, 1986.

[10] F. Pierrot, C. Reynaud, and A. Fournier, "DELTA: a simple and efficient parallel robot," *Robotica*, vol. 8, no. 2, pp. 105–109, 1990.

[11] H. Makino, "Development of the SCARA," *Journal of Robotics and Mechatronics*, vol. 26, no. 1, pp. 5–8, 2014.

[12] J. Luh, M. Walker, and R. Paul, "Resolved-acceleration control of mechanical manipulators," *IEEE Transactions on Automatic Control*, vol. 25, no. 3, pp. 468–474, 1980.

[13] P. K. Khosla and T. Kanade, "Experimental evaluation of nonlinear feedback and feedforward control schemes for manipulators," *The International Journal of Robotics Research*, vol. 7, no. 1, pp. 18–28, 1988.

[14] G. J. Monkman, S. Hesse, R. Steinmann, and H. Schunk, *Robot Grippers*. Wiley-VCH, 2006.

[15] J. E. Bobrow, S. Dubowsky, and J. S. Gibson, "Time-optimal control of robotic manipulators along specified paths," *The International Journal of Robotics Research*, vol. 4, no. 3, pp. 3–17, 1985.

[16] T. Kröger and F. M. Wahl, "Online trajectory generation: basic concepts for instantaneous reactions to unforeseen events," *IEEE Transactions on Robotics*, vol. 26, no. 1, pp. 94–111, 2010.

[17] B. Siciliano, L. Sciavicco, L. Villani, and G. Oriolo, *Robotics: Modelling, Planning and Control*, ser. Advanced Textbooks in Control and Signal Processing. London: Springer, 2009.

[18] R. Bostelman and J. Falco, "Survey of industrial manipulation technologies for autonomous assembly applications," National Institute of Standards and Technology, Gaithersburg, MD, NIST Interagency/Internal Report 7844, 2012.

[19] 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.

[20] J. Falco, J. Marvel, and E. Messina, "A roadmap to progress measurement science in robot dexterity and manipulation," National Institute of Standards and Technology, Gaithersburg, MD, NIST Interagency/Internal Report 7993, 2014.

[21] M. Franaszek, G. S. Cheok, K. van Wyk, and J. A. Marvel, "Improving 3D vision-robot registration for assembly tasks," National Institute of Standards and Technology, Gaithersburg, MD, NIST Interagency/Internal Report 8300, 2020.

10.10 Exercises

10.1 Concept check · 2 min

Scenario 4 aims the gripper at the camera's last report while the arm moves with the belt at 0.15 m/s. The camera latency is set to zero and nothing else is changed. What happens to the grasps?

(a) They take, because a report with no latency describes the product as it is now

(b) They fail on relative speed, because the gripper waits at rest under the report

(c) They still fail, because the product moves on while the arm travels and the gripper closes

(d) They take for half the products, those reported in the frame just before the approach

10.2 Concept check · 2 min

In scenario 2 the belt runs at 0.30 m/s and the arm arrives early and waits at rest where each product will be. The gripper closes 0.96 mm from the product on average, yet all 60 grasps of the first minute fail. Why?

(a) The camera noise of 2 mm is larger than the position error the gripper can take

(b) The product's own acceleration at the catch exceeds what the holding force can follow

(c) The arm must reach belt speed while the gripper closes, and the drives hit their torque limit

(d) The product slides under the gripper at the belt speed, twice the relative speed it can take

10.3 Concept check · 1 min

In scenario 6 the controller's arm model knows only the arm, and the planner checks every move against 90 % of each drive. Why does the torque limit still get hit while a 1.5 kg case is carried?

(a) The planner computes the torque with the controller's arm model, which leaves the case out

(b) The torque share of 90 % leaves too little for the Coulomb friction of the joints

(c) The servo gains amplify the quantization of the joint encoders into torque spikes

(d) The catch conserves momentum, so every grasp sends a torque step through both joints

10.4 Theory and calculation · 6 min

Use the grasp tolerance $e_p = 8\,\text{mm}$ and the closing time $t_{close} = 0.10\,\text{s}$ of Table 10.2 and Equation (10.9).

(a) An instantaneous camera ($L = 0$) feeds a controller that aims at the report and does not predict the closing interval, so the shortest possible T_d is t_{close}. Find the fastest belt on which a grasp can still meet e_p. *Checkpoint:* 0.08 m/s.

(b) Scenario 4 runs at 0.15 m/s with $L = 0.30\,\text{s}$, and the time from delivery to grasp averages $T_d = 0.42\,\text{s}$. Find the distance from the gripper to the product it was sent for. *Checkpoint:* 108 mm.

(c) A camera with no latency replaces the old one. Find how much of that distance it removes and how much remains. *Checkpoint:* 45 mm removed, 63 mm remaining.

10.5 Theory and calculation · 6 min

At the ready pose ($q_2 = 62.6°$) the empty arm has $M_{11} = 0.9996\,\mathrm{kg\,m^2}$. The shoulder servo has $K_{p1} = 400\,\mathrm{N\,m/rad}$ and $K_{d1} = 25\,\mathrm{N\,m\,s/rad}$. Treat the shoulder as a single inertia M_{11} under the PD law with the feedforward exact, so that its error obeys $M_{11}\ddot{e} + K_{d1}\dot{e} + K_{p1}e = 0$.

(a) Find the natural frequency and the damping ratio of the shoulder loop. *Checkpoint:* 20.0 rad/s and 0.63.

(b) A point mass m_p at the tip adds $m_p r^2$ to M_{11}, where r is the tip's distance from the shoulder. With $l_1 = 0.40\,\mathrm{m}$ and $l_2 = 0.35\,\mathrm{m}$, find r at the ready pose and M_{11} with a 1.5 kg case held. *Checkpoint:* $r = 0.641\,\mathrm{m}$ and $1.617\,\mathrm{kg\,m^2}$.

(c) Find the natural frequency and damping ratio with the case, and compare the relative change of each. *Checkpoint:* 15.7 rad/s and 0.49.

10.6 Theory and calculation · 6 min

The encoder estimator averages the reports of one product in belt coordinates with the gain $k_n = \max(1/n, 0.25)$ after the nth report. Each report carries independent noise of $\sigma = 2\,\mathrm{mm}$ on each axis.

(a) For an exponential average $u \leftarrow u + k(u_{rep} - u)$ with a fixed k, the stationary variance is $k/(2-k)$ times the variance of one report. Find the standard deviation per axis for $k = 0.25$. *Checkpoint:* 0.756 mm.

(b) The mean length of a two-dimensional error whose axes are independent normals of standard deviation s is $s\sqrt{\pi/2}$. Find it and compare it with the 0.99 mm measured in case 1. *Checkpoint:* 0.95 mm.

(c) In scenario 5 the belt runs at 0.25 m/s past the 0.40 m field of view of a 15 Hz camera, and half the frames are lost. Find how many reports a product gets on average with all frames and with half of them. *Checkpoint:* 24 and 12.

10.7 App challenge · 8 min

Load scenario 4, Camera delay. It aims at the last report (Where the chocolate is set to Last report), moves with the belt, and the camera reports 300 ms late. Use the default seed and run each part for 60 s at 2×. Read the scorecard tiles Placed, Failed grasps and Missed at 60 s. Counts are exact at the default seed.

(a) Run the scenario as loaded. *Checkpoint:* 0 placed, 42 failed grasps, 23 missed.

(b) Reload the scenario with the reload button, set Camera latency to 0 ms before Run, and run again. Explain the result with Equation (10.9). *Checkpoint:* 0 placed, 42 failed grasps, 23 missed.

(c) Reload the scenario, set Where the chocolate is to Encoder before Run, keep the 300 ms latency, and run again. *Checkpoint:* 24 placed (Placed reads 24 of 24), no failed grasp, none missed.

10.8 App challenge · 8 min

Load scenario 7, Reduced grip. The cup holds 15 % of 10 N and each product weighs 0.10 kg. Use the default seed, set each value before Run, run for 60 s at 2×, and read the tiles Placed and Drops.

(a) Compute the largest tip acceleration the grip can hold. *Checkpoint:* $15\,\mathrm{m/s^2}$.

(b) Reload the scenario with the reload button, set Carry acceleration limit to $15.0\,\mathrm{m/s^2}$, and run. *Checkpoint:* 4 placed, 20 dropped.

(c) Reload, set the limit to $14.5\,\mathrm{m/s^2}$, and run. Explain why the bound of part (a) is not a usable setting. *Checkpoint:* 24 placed, none dropped.

10.9 App challenge · 9 min

Load the challenge, a production run of 100: 0.6 kg boxes at 36 a minute, a belt that speeds up from 0.22 to 0.30 m/s at 60 s and a grip of 60 % of 10 N. Use the default seed and run at 5× until the run stops itself, about 180 s of simulated time. Reload the scenario with the reload button before each part, make the listed changes before Run, and read the tiles Placed, Drops, Failed grasps and Tracking error. Counts are exact. Give the tracking error within 0.02 mm.

(a) Set Meet the chocolate to Move with it. *Checkpoint:* 0 placed, 35 dropped, 158 failed grasps.

(b) Also set Carry acceleration limit to 8.0 m/s^2. *Checkpoint:* 35 placed, none dropped, 158 failed grasps.

(c) Also set Where the chocolate is to Encoder. *Checkpoint:* 98 placed, none dropped, no failed grasp, 1.97 mm rms.

(d) Also set Controller's arm model to Arm + payload. *Checkpoint:* 98 placed, 0.25 mm rms.

10.10 Application area · 4 min

The chapter assumes that the camera and the belt encoder share one clock. On a real line the trigger time stamped on a frame can be offset from the encoder's clock.

(a) Find the position error a 10 ms offset adds on a 0.30 m/s belt. *Checkpoint:* 3 mm.

(b) The encoder estimator leaves about 1 mm of error at the grasp, and the grasp tolerance is 8 mm. Adding the two errors in the worst case, find the largest clock offset the cell can take at 0.30 m/s. *Checkpoint:* About 23 ms.

10.11 Application area · 3 min

Pushed with 120 arrivals a minute on a 0.30 m/s belt, the arm completes one pick in 0.857 s: 0.303 s of approach, 0.100 s of closing, 0.313 s of carry, 0.133 s of placing and release and 0.009 s of waiting.

(a) Find the ceiling in picks per minute. *Checkpoint:* 70.0 per minute.

(b) A line must clear 120 products a minute with no misses. Find the least number of identical arms, and name what the chapter's scheduler lacks for that layout. *Checkpoint:* Two arms.

10.12 Review the agent · 5 min

An AI assistant reads the scenario 6 results (a 1.5 kg case, Arm only model, 11 torque-limit hits and 6.20 mm rms of tracking error) and writes: "The servo is too soft for the heavy case. Double both position gains to 800 and 200 N m/rad so the joints hold their references more stiffly. The lag will shrink and the torque-limit hits will stop." Find the flaw and say how you would check it in the app.

CHAPTER 11

A Table Tennis Robot That Must Swing Before It Knows Where

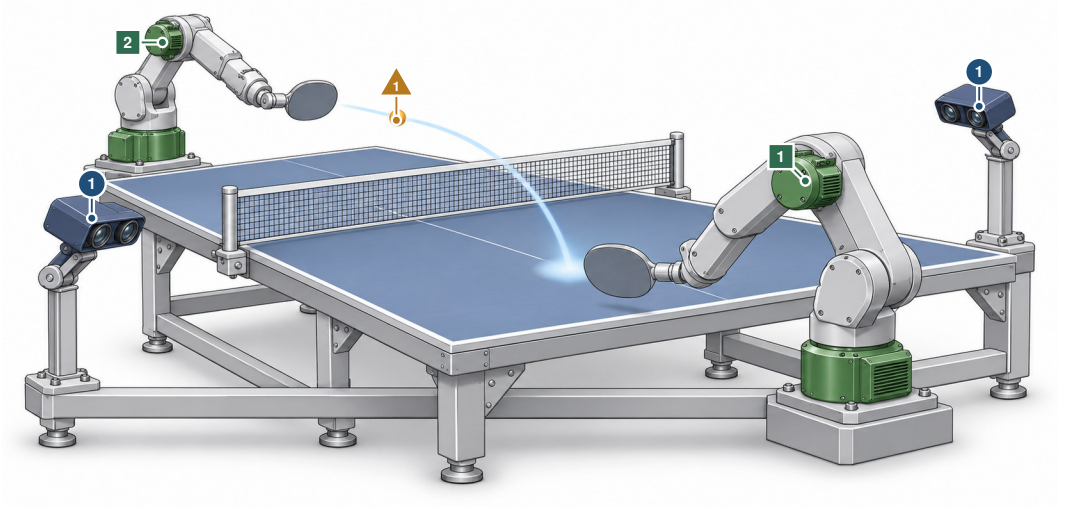

The controller moves 1 joint motors of the near arm, which swing the paddle 2 joint motors of the opponent arm

The controller reads 1 camera pairs, which report the ball a little late

What disturbs it 1 spin and air, which bend the ball's path

Notice that the near arm must start its swing toward the end of the faint arc, a prediction, before the ball has arrived there. Illustration, not to scale.

IN PLAIN TERMS A Robot That Must Swing Before It Knows

The main idea. A robot arm plays table tennis, and the ball crosses the table in about half a second. Cameras watch the ball, and a computer works out where it will meet the arm's paddle.

The challenge. Every extra look makes that guess better but leaves the arm less time to swing. The arm must stop watching at the right moment, and its guess must include the air and the spin that bend the ball's path.

Why it matters. An arm that waits too long, or forgets the air, swings at an empty spot and loses the point.

THREE WAYS TO READ THIS CHAPTER

Explore	11 min	the first page, the fact on ball speed in the Introduction, the history segment (Section 11.2), the try-it segments of cases 7 and 1 (Sections 11.4 and 11.6), the fact on the commit time, and the key points
Study	45 min	the whole chapter with its three pause points and the twelve exercises at its end
Practice	15 min	the first page, the measured cases and the game (Section 11.6), Technology in Practice (Section 11.7), Limitations (Section 11.8) and the open challenge

The control question A table-tennis ball crosses the table in about half a second, and every camera frame sharpens the prediction of where it will cross the hitting plane while leaving the arm less time to get there. When should the arm stop watching and commit, and what must its flight model, stroke and trajectory get right for the swing to arrive in time?

Learning objectives

- Compute the drag factor of the ball and explain why a predictor without drag commits close in place but wrong in time.
- Explain why the commit time has an optimum, and locate it from the returns and the time left in a commit sweep.
- Derive the paddle's normal speed from the impact law, and the rolling exit of a gripping bounce.
- Compare stepped, trapezoidal and minimum-jerk swings by their peak acceleration and by the lag of a drive tuned away from the shaping.
- Decide when a camera's rate and latency make an interception infeasible for any commit rule.

Prerequisites Sections 1.4.2 (arm kinematics), 1.4.4 (quintic and bang-bang profiles), 1.4.6 (the Kalman filter and the age of a report), 1.3 (loop action) and 1.6.4 (scenarios and the scorecard).

The app https://apmonitor.com/apps/table-tennis-robot/

See also Volume 5, Chapter 13 for an extended Kalman filter on late fixes, Volume 5, Chapter 12 for a state that is never measured and Volume 5, Chapter 3 for a measurement too slow for the loop. Volume 8, Chapter 15 uses the same six-axis arm to keep a ball bouncing.

11.1 Introduction

A ball served at 10 m/s and 3° from 0.30 m behind the far end of a table-tennis table reaches a hitting plane 0.30 m behind the near end line 0.504 s later in the simulator of this chapter, after a bounce 0.62 m from that end line. At that speed the air slows the ball by 12.6 m/s^2, more than gravity pulls it down, and a predictor that leaves the air out expects the ball after 0.373 s and puts the bounce at 0.28 m. A robot at the near end has to choose a hitting point inside that half second from camera frames that arrive late, and then bring a paddle there at the right speed within the limits of six joint drives. A 2014 promotional duel between a table-tennis professional and a six-axis industrial arm brought this problem to a general audience, and the problem is older and wider than the game.

DID YOU KNOW? Nine Meters a Second, and Faster

An amateur hits a table-tennis ball at up to 9 m/s, and a professional at three times that. A ball at amateur speed crosses the table in 0.4 s. It cannot carry a marker or paint, because either changes its weight or texture, so robots track it with passive cameras [1].

Two beliefs about control fail at this table. The first is that a controller should act once it knows. Every camera frame improves the prediction, and every frame spent watching is time the arm no longer has for its swing. Over twenty random serves to the full pipeline, a commit 50 ms after the serve returns 16 balls onto the table, a commit 200 to 275 ms after it returns 19, and a commit at 350 ms returns 10, with 243 ms left for a swing that needs a quarter of a second or more. The second belief concerns opponents. The tuned arm on the far side was designed with a declared blind spot, a commit fixed 150 ms after the student's strike and a spin it assumes rather than estimates, in the expectation that spin and wide placement would beat it. Measured over ten games against the suggested settings, its spin assumption decides none of them, and the points it loses come from its clock-frozen plan and from the ends of its reach.

The control problem belongs to a class, interception under a sensing delay with a decision deadline. A robot that catches a thrown ball predicts a free flight and closes its hand where the ball will be [2], and a robot table-tennis player chooses where and when to strike from a predicted trajectory before

it plans the stroke [3, 4]. In each case the estimate improves with time while the time left to act shrinks, so the moment of commitment is a design variable with an optimum, and the actuator's speed and acceleration limits set how much time the action needs. The picking arm of Chapter 10, which meets a product on a moving belt from camera reports that arrive late, a flying shear that must meet a moving strip and a batch cut at a predicted end point have the same shape. The flight itself is the other half of the problem, because a light ball's drag and Magnus lift move its bounce by decimeters [5, 6], drag alone by 34 cm for the serve above, and its spin changes how it leaves the table [7, 8].

The published robot players are each built around one platform, and how much the flight model, the commit time and the joint limits contribute is seen, when it is seen, inside a complete system [3, 4, 9–11]. A student who watches a robot miss cannot tell whether the prediction, the timing or the arm was at fault. This chapter takes those causes apart on a simulated table with a six-axis arm at each end, switching one on at a time, and counts the returns each one costs. The flight model, the gripping bounce, the extended Kalman filter [12], damped least-squares inverse kinematics [13, 14] and minimum-jerk trajectories [15] are standard methods. The measurements are of one model at one game seed. Among them are a commit-time optimum 319 to 390 ms before the hit, an estimator with a spin state that reads 33.5 rev/s of a 40 rev/s backspin before the bounce while every predictor without one misses by 31 to 37 cm, and the finding that the tuned opponent's declared blind spot is not where it loses. The simulator is a 33-state model with every impact located inside its integration step, the shared six-axis arm of the series and a measurement boundary that keeps both arms' controllers on their own camera frames, and the chapter states it completely enough to build again. Each of the nine cases and the game has a completion rule, and no setting of the pipeline wins case 9.

Section 11.2 follows robot table tennis from its first players. Section 11.3 gives the table, the ball, the two arms and the cameras. Section 11.4 spends the half second twice, on knowing where the ball will be and on moving there, and Section 11.5 describes the controllers as coded. Section 11.6 reports the nine cases and the game as measured, Sections 11.8 and 11.9 state the limitations and the conclusions, Section 11.7 names the devices of real robot players, and Section 11.10 turns the same cases into exercises.

11.2 Robots at the Table

Table tennis became a robotics benchmark because it packs perception, prediction and motion into a fraction of a second on a table of fixed size. A robot ping-pong player documented in a 1988 monograph framed the game as an experiment in real-time intelligent control, in which vision, the prediction of the ball and the planning of each return had to be finished in the time the ball takes to cross the table [3]. A low-cost player built around a personal computer was reported in 2003 [9], and a learning approach to robotic table tennis appeared in 2005 [10]. Striking movements that a robot learns to select and generalize were reported in 2013 [4], and in 2022 a muscular robot learned to play from scratch [11]. Across these decades the decomposition stayed recognizable, in which the ball is predicted, a hitting point and time are chosen, and the stroke is planned and executed.

> ⧖ HISTORY **From a Challenge of 1983 to Learned Play**
>
> Early work on robot table tennis is traced to a challenge of 1983: a competition whose simplified rules allowed for the technology of the time [16]. For most of the years since, robot players have predicted the ball from a dynamics model, either fitted or learned, chosen a virtual hitting point for the paddle, and then generated the joint trajectories to reach it. A newer line of research drops the hitting point and the ball model and controls the joints at a high rate with policies learned by trial or from demonstrations [1]. A player built that way, trained in simulation and moved to the real arm without further training, has since played full competitive matches against people it had never met [16].

The physics of the ball was worked out alongside. A review of sports-ball aerodynamics described how drag and the lift of a spinning ball depend on the flow around a sphere [5]. Fitted to recorded

Table 11.1 Signals by Role. Ranges Are Those the Interface Offers

Symbol	Role	Default	Range
$\mathbf{p}$, $\mathbf{v}$	states, ball position, velocity	per serve	—
ω	state, ball spin, constant in flight	per serve	—
q_i, $\dot{q}_i$	states, six joints and rates per arm	ready pose	ranges
$\hat{\mathbf{p}}_h$, t_h	CV, predicted hitting point and time	—	—
$q_{c,i}$	MV, joint commands	per plan	ranges
y_p, z_p	MV in Manual, paddle across, height	0, 20 cm	± 70, -5 to 50 cm
$\mathbf{z}_k$	measured, camera frame of $\mathbf{p}$	200 Hz	50 to 200 Hz
σ	estimate, $\hat{\mathbf{p}}_h$ uncertainty	—	—
$\mathbf{w}$	DV, gust across the hall	0	± 3 m/s
—	DV, worn ball, change in drag	0 %	± 30 %
—	DV, base drive derated to	100 %	40 to 100 %
—	DV, hidden-spin serve, 45 rev/s	—	event
—	DV, 50 ms camera dropout, net cord	—	events
—	outputs, score, returns, rallies	0	—

flights of the table-tennis ball, a drag coefficient of 0.54 ± 0.074 and a lift coefficient of 0.069 ± 0.029 were identified [6], and a rebound model with friction fitted to the same group's measurements gave a table a normal restitution of 0.93 and a sliding friction coefficient of 0.25 [8]. Measurements of bouncing balls show the contact gripping and slipping in turn, which a rigid grip law idealizes [7]. The laws of table tennis fix the ball at 40 mm and 2.7 g and ask a table to return a ball dropped from 30 cm to about 23 cm (Law 2.1.3).

The arm draws on manipulator kinematics and trajectory planning [17], which Sections 1.4.2 and 1.4.4 summarize. Damped least squares made inverse kinematics robust near singular configurations in 1986 [13, 14], a minimum-jerk profile was proposed in 1985 as a model of smooth reaching [15], and the Kalman filter of 1960 [12], in its extended form, is the usual estimator for a flight whose model is nonlinear. A robot catching thrown balls showed the same predict-and-intercept structure outside the game [2]. The simulator combines these pieces in a form small enough to teach with, and each one can be switched on or off.

11.3 The Table and Two Arms as Modeled

The world frame has x along the table from the near end line, where the student's arm plays, to the far end line at $x = L = 2.74$ m, y across it, positive to the student's left, with the center line at $y = 0$ and the side lines at ± 0.7625 m, and z up from the playing surface, as shown in Figure 11.1. The net stands over $x = 1.37$ m, 0.1525 m high, its posts 0.1525 m outside the side lines. Each arm stands on a fixed base 0.85 m behind its end line, 0.05 m below the playing surface, and strikes on a hitting plane $x = x_h$, 0.30 m behind its end line. Two cameras watch the ball for each arm. The variables appear with their roles in Table 11.1. The constants of the ball, the arms and the cameras and every threshold were chosen for teaching, as the model dialog and the Instructions page state.

11.3.1 Flight

The ball is a sphere of radius $r = 20$ mm and mass $m = 2.7$ g. Its position, velocity and spin obey Equation (11.1).

$$\begin{aligned} \dot{\mathbf{p}} &= \mathbf{v}, \qquad \dot{\omega} = \mathbf{0} \\ \dot{\mathbf{v}} &= -g\,\mathbf{e}_z - k_D\,|\mathbf{v}-\mathbf{w}|\,(\mathbf{v}-\mathbf{w}) + k_M\,\omega \times (\mathbf{v}-\mathbf{w}) \\ k_D &= \frac{\rho\, C_D\, \pi r^2}{2m} \end{aligned} \tag{11.1}$$

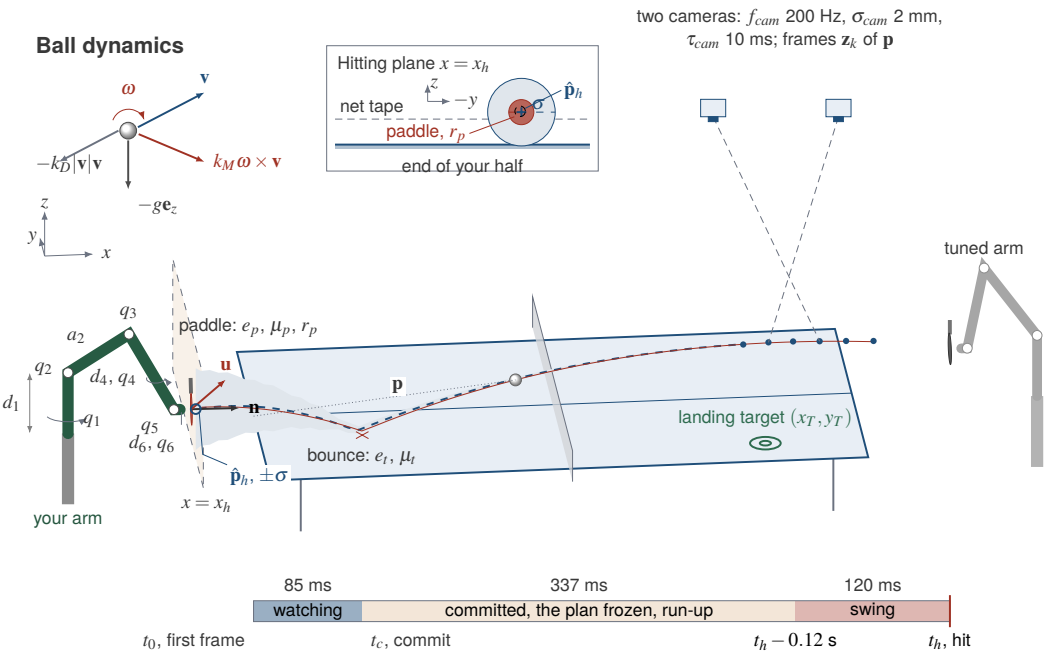

Figure 11.1 The table, the two six-axis arms and the student's pipeline at the commit on the first serve of the challenge with the suggested settings, drawn in the Table view's orthographic projection from the engine's state. The ball at $\mathbf{p}$ with velocity $\mathbf{v}$ and spin ω feels gravity $-g\,\mathbf{e}_z$, drag $-k_D|\mathbf{v}|\mathbf{v}$ and Magnus lift $k_M\,\omega \times \mathbf{v}$, and bounces with restitution e_t and grip μ_t. The cameras deliver frames $\mathbf{z}_k$ (dots) at f_{cam}, late by τ_{cam} with noise σ_{cam}. The dashed curve is the prediction at the commit and the shaded band its $\pm\sigma$, which reaches 202 mm at the predicted hitting point $\hat{\mathbf{p}}_h$ on the plane $x = x_h$, seen nearly edge on. The near arm, with joints q_1 to q_6 and lengths d_1, a_2, d_4, d_6, holds its paddle (restitution e_p, grip μ_p, radius r_p) at the planned hitting pose with normal $\mathbf{n}$ and velocity $\mathbf{u}$ that send the ball to the landing target. The upper-left inset separates the ball-force terms (schematic directions). The upper-middle inset draws the plane $x = x_h$ from behind the arm, with the paddle on $\hat{\mathbf{p}}_h$ inside its σ circle and the true crossing as a dashed ring. The bar is the Plan view's time budget for this ball, from the first frame t_0 to the commit t_c, the start of the swing $t_h - 0.12\,\mathrm{s}$ and the hit t_h. Not to scale across the panels.

where $\mathbf{w} = [0,\ w_y,\ 0]^{\mathsf{T}}$ is a gust across the hall, $\rho = 1.2\,\mathrm{kg/m^3}$ and $C_D = 0.45$, which give $k_D = 0.1257\,\mathrm{m^{-1}}$, and $k_M = 4.5 \times 10^{-4}\,\mathrm{m/rad}$. A worn ball scales k_D by the disturbance slider, and neither arm knows. At 5 m/s, 10 and 15 the drag deceleration is 3.1, 12.6 and 28.3 $\mathrm{m/s^2}$, against 3.8, 15.1 and 33.9 $\mathrm{m/s^2}$ for the fitted $C_D = 0.54$ [6], and the terminal speed is 8.83 m/s. The fitted lift coefficient corresponds to a Magnus factor of about 1.0×10^{-3} m/rad, so the teaching value is less than half the measured one. Spin keeps its value through a flight.

11.3.2 Bounce and Strike

Both impacts follow one law, written relative to a surface of unit normal $\mathbf{n}$, toward the ball, moving at $\mathbf{u}$, as shown in Equation (11.2).

$$\begin{aligned} v_n &= (\mathbf{v}-\mathbf{u})\cdot\mathbf{n}, \qquad J_n = -(1+e)\,v_n \\ \mathbf{s} &= (\mathbf{v}-\mathbf{u}) - v_n\mathbf{n} - r\,\boldsymbol{\omega}\times\mathbf{n} \\ \mathbf{J}_t &= \begin{cases} -\frac{\alpha}{1+\alpha}\,\mathbf{s} & \frac{\alpha}{1+\alpha}|\mathbf{s}| \le \mu J_n \\ -\mu J_n\,\mathbf{s}/|\mathbf{s}| & \text{otherwise} \end{cases} \\ \mathbf{v}^+ &= \mathbf{v} + J_n\mathbf{n} + \mathbf{J}_t, \qquad \boldsymbol{\omega}^+ = \boldsymbol{\omega} - \frac{\mathbf{n}\times\mathbf{J}_t}{\alpha r} \end{aligned} \tag{11.2}$$

where $\mathbf{s}$ is the slip of the contact point, $\alpha = I/(mr^2) = 2/3$ for a thin hollow sphere, so the gripping impulse is $-\frac{2}{5}\mathbf{s}$, and the impulses are per unit mass. A gripping ball leaves rolling on the surface, and a ball whose grip would need more than μ times the normal impulse slides. The spin about the normal is kept, and angular momentum about the contact point is conserved either way, which one of the application's checks confirms to 10^{-12}. The table has $\mathbf{n} = \mathbf{e}_z$, $\mathbf{u} = \mathbf{0}$, $e_t = 0.905$ and $\mu_t = 1.0$. With this drag, $e_t = 0.905$ returns a ball dropped from 30 cm to the 23 cm the laws ask. The grip is a teaching value chosen so that bounces grip, since the 10 m/s serve at 3° needs $\mu_t \ge 0.76$ to grip, while on the measured 0.25 [8] it would slide and reach the hitting plane 68 ms sooner. The paddle has the normal and velocity of the arm's tool at the contact, $e_p = 0.7$ and $\mu_p = 1.2$ for the rubber, both teaching values, and it strikes only a ball whose center passes within $r_p + 0.3r$ of the paddle's center, $r_p = 75$ mm. The normal part of Equation (11.2) gives $v_n^+ = -e_p(v_n - u_n) + u_n$, so turning a 6 m/s ball into a 10 m/s one takes a paddle moving at 3.41 m/s along its normal, and a paddle brushing across the ball puts spin on it through $\mathbf{J}_t$.

11.3.3 The Six-Axis Arms

Each arm is the shared servo-level six-axis arm of the simulator framework, `shared/src/framework/core/arm.js`, the same arm that holds the paddle in Volume 8, Chapter 15. Its chain is an anthropomorphic arm with a spherical wrist on a base turned by ψ ($\psi = 0$ for the near arm and π for the far one), as shown in Equation (11.3).

$$\begin{aligned} R &= R_z(\psi+q_1)\,R_y(q_2+q_3)\,R_z(q_4)\,R_y(q_5)\,R_z(q_6) \\ \mathbf{p}_t &= \mathbf{b} + d_1\mathbf{e}_z + a_2\,R_z(\psi+q_1)R_y(q_2)\,\mathbf{e}_z \\ &\quad + d_4\,R_z(\psi+q_1)R_y(q_2+q_3)\,\mathbf{e}_z + d_6\,R\,\mathbf{e}_z \\ \mathbf{n} &= R\,\mathbf{e}_z \end{aligned} \tag{11.3}$$

where $\mathbf{b}$ is the base, $\mathbf{p}_t$ the paddle's center and $\mathbf{n}$ its normal, the flange's z axis, with $d_1 = 0.30$ m, $a_2 = 0.42$ m, $d_4 = 0.40$ m and $d_6 = 0.08$ m, a reach of 0.90 m. The paddle's center and normal make a five-number task, and the roll q_6 about the normal is left free. On its hitting plane at 20 cm above the table the arm reaches 0.68 m across, and 0.62 to 0.68 m over the heights and face tilts it plays.

Each joint is a drive at the servo level, with no inertia matrix, Coriolis terms or torque limits, as shown in Equation (11.4).

$$\begin{aligned} \ddot{q}_i &= \mathrm{sat}_{\ddot{q}_{max,i}}\left(\omega_j^2(q_{c,i}(t) - q_i) - 2\zeta_j\omega_j\dot{q}_i\right) \\ |\dot{q}_i| &\le \dot{q}_{max,i} \end{aligned} \tag{11.4}$$

where $\omega_j = 2\pi f_j$ and ζ_j are the drives' bandwidth and damping, 15 Hz and 1 by default, and the command $q_{c,i}(t)$ is the controller's latest value carried along its rate between 5 ms updates, as a drive interpolates its setpoint. A drive at its speed limit gets no further acceleration in that direction. The base, shoulder and elbow are limited to 360 °/s and 3000 °/s^2, and the three wrist axes to twice that, a

departure from a single limit for all six that the wrist needs to turn the paddle within a quarter-second swing. At a lever of 0.5 m the base limits correspond to 3.14 m/s and 26.2 m/s^2 at the paddle. A joint's range is a soft limit. Moving toward an end at $\dot{q}_i$ with d_{end} of range left, the drive brakes at $\min(1.25\,\dot{q}_i^2/(2d_{end}),\ \ddot{q}_{max,i})$ once $\dot{q}_i^2/(2d_{end})$ reaches $0.8\,\ddot{q}_{max,i}$, whatever the command, and the plant keeps a hard clamp at the range as a backstop. Over every case and both challenge settings no joint arrives at a range end faster than 6.9 °/s.

11.3.4 Cameras and Events

Each arm has its own cameras, a stereo pair modeled as one measurement of the ball's center, as shown in Equation (11.5).

$$\mathbf{z}_k = \mathbf{p}(t_k) + \sigma_{cam}\,\mathbf{n}_k, \qquad t_k = k/f_{cam} \tag{11.5}$$

where frame k is delivered at $t_k + \tau_{cam}$ and $\mathbf{n}_k$ holds three unit normal deviates drawn from that arm's own seeded stream, three per frame whatever the noise level, $f_{cam} = 200$ Hz, $\tau_{cam} = 10$ ms and $\sigma_{cam} = 2$ mm. The student's cameras are adjustable down to 50 Hz and up to 150 ms, while the tuned arm's stay at the defaults. Spin is never measured. A dropout loses every frame exposed while it lasts, and no frame is taken while the ball is not in play.

The table, the paddles and the net act between integration steps. When the ball's surface crosses the table or a paddle face inside a step, the ball and the paddle are carried linearly back to the contact instant, the law of Equation (11.2) is applied there, and the ball flies the rest of the step on the true model. A ball through the net's plane below the tape is dead, and one that passes within one radius over it clips the tape, keeps 45 % of its speed along the table and pops up. The referee scores a strike or a serve as a shot that must bounce once on the receiver's half. A second bounce there or a ball that passes the receiver is a missed return, and a first bounce on the striker's own half, a ball past the receiver's end line without a bounce or a ball off the side is out. A game goes to eleven points, won by two, with the serve changing every two points and every point from ten all. Serves are launched rather than struck, from 0.10 m behind the server's end line, 0.25 m up (the ball machine's from 0.30 m behind the far end line), aimed on the true model at a seeded point 0.30 to 0.75 m from the receiver's end line and within 0.40 m of the center line, at a seeded speed of mean 8.5 m/s and standard deviation 0.6 m/s, with up to 20 rev/s of topspin or backspin. The serve counts as the server's shot, and the rule that a serve bounces on both halves is left out.

11.3.5 Parameters and Numerics

Table 11.2 gives every parameter. The 33 states, the ball's nine and each arm's twelve, are advanced by fourth-order Runge-Kutta steps of 0.5 ms, in which a 10 m/s ball moves 5 mm. Both arms' controllers and cameras run every 5 ms and the historian logs at 100 Hz, on the engine of Section 1.6.1. Within a step the controller runs first, then the integration, then the speed and range clamps, the impacts, the referee and the serve clock. At 0.5 ms the integrated step overshoots a contact by at most 1.8 mm at the table and 2.9 mm at a paddle before the back-interpolation, halving at 0.25 ms, and the planning serve's bounce lands at the same 0.61954 m at 0.5, 0.25, 0.1 and 0.01 ms. The 66.3 s game of the suggested settings plays in about 13 s at five times real time. Every serve, camera frame and opponent target comes from seeded streams, so a case replays identically after Reset, from the list after another case, and when opened first.

11.3.6 Assumptions

A rigid ball with constant aerodynamics. The drag and Magnus factors do not change with speed or spin, and the spin does not decay in flight. The fitted factors are larger [6], so real air matters more than it does here.

Instantaneous impacts with a grip law. A real contact has compliance, grips and slips in turn [7], and loses restitution with speed, and a real paddle's rubber, sponge and blade bend. The table's grip

Table 11.2 Model Parameters. A Is Adjustable in the Model Dialog, H Hidden, C Fixed in the Code

Symbol	Meaning	Value	
Ball and table			
r, m	ball radius, mass	20 mm, 2.7 g	C
k_D	drag factor, C_D 0.45	0.1257 m^{-1}	A
k_M	Magnus factor	4.5×10^{-4} m/rad	A
e_t, μ_t	table restitution, grip	0.905, 1.0	A
e_p, μ_p	rubber restitution, grip	0.7, 1.2	A
r_p	paddle radius	75 mm	H
L, W, h_n	table length, width, net height	2.74, 1.525, 0.1525 m	C
Arms			
d_1, a_2, d_4, d_6	chain lengths	0.30, 0.42, 0.40, 0.08 m	C
b	base, behind the end line, below the top	0.85, 0.05 m	C
x_h	hitting plane behind the end line	0.30 m	C
$\dot{q}_{max}$	speed limit, axes 1 to 3 (wrist $\times 2$)	360 °/s	A
$\ddot{q}_{max}$	acceleration limit, axes 1 to 3 (wrist $\times 2$)	3000 °/s^2	A
—	ranges, axes 1 to 6 ($\pm$°)	170, 100, 150, 200, 125, 200	C
f_0, ζ_0	drive tuning the commands assume	15 Hz, 1	C
Cameras			
f_{cam}, τ_{cam}	frame rate, latency	200 Hz, 10 ms	A
σ_{cam}	noise per axis	2 mm	A
Tuned opponent and serves			
$t_{c,opp}$	fixed commit after the strike	150 ms	A
b_{spin}	spin assumed until its bounce	1 (yes)	A
y_{opp}	target width, 0.65 of it used	0.55 m	H
v_s, σ_v	serve speed, spread	8.5, 0.6 m/s	A
ω_s, y_s	largest serve spin, widest target	20 rev/s, 0.40 m	A
—	game seed	1	A
Numerics			
Δt, T_c	integration step, control period	0.5, 5 ms	C
—	historian	100 Hz	C

is four times the measured value, which keeps a gripping bounce for spinning balls but makes fast, low bounces grip that would slide.

Servo-level arms. The kinematics are exact, but there is no inertia, torque limit or flexibility, and a real arm's limits change with its pose and payload. The paddle's strike uses the velocity of its center and leaves out its rotation at the contact.

Ideal cameras apart from noise, rate and latency. Calibration error, motion blur and occlusion are absent, and every frame of a ball in play is taken.

Fixed bases and a simplified game. Neither arm moves its base, so a ball short enough to bounce twice before the hitting plane is out of play. Serves are launched by a machine at the server's paddle, and lets and edge balls are left out.

11.4 Half a Second, Spent Twice

11.4.1 Where the Ball Will Be

The flight model decides the prediction more than the frames do. Launched at 10 m/s and 3° from 0.30 m behind the far end line with no spin and the net removed, the ball bounces 0.62 m from the near end line after 0.283 s and reaches the hitting plane at 0.504 s. Without drag it would bounce at 0.28 m and arrive at 0.373 s, a difference of 34 cm along the table and 131 ms in time. Spin moves the bounce as well. With 50 rev/s of backspin the lift holds the ball up and the bounce moves to 0.43 m, and with 50 rev/s of topspin it moves to 0.78 m. Through the grip of the bounce, 40 rev/s of topspin brings the ball to the plane at 0.439 s and 40 rev/s of backspin at 0.532 s, 93 ms apart.

A fit or a filter anchored to the frames can carry the ball forward only with the forces its model has, as shown in Figure 11.2a. On the planning serve, after the 30 frames of the first 150 ms, the median distance from the predicted hitting point to the ball at the predicted instant is 296 to 329 mm without drag and 10.5 mm to 10.9 mm with it. The no-drag error is mostly timing. Its predicted point lies

4.4 to 4.7 cm from the true crossing, but it expects the ball 73 to 81 ms early, and the paddle swings through an empty plane. Adding a spin state to a spinless ball costs accuracy at first, because the extended Kalman filter (EKF) must also explain the frames with a spin it does not yet know, and its miss after 30 frames is 40 mm.

Spin reverses the ranking, as shown in Figure 11.2b. On case 7's serve, 8.5 m/s with 40 rev/s of backspin, the drag-only fit and EKF stay 30 to 32 cm off until the bounce, the no-drag models are 37 cm off by then, and only the EKF on the drag-and-Magnus model sees the lift bend the flight. At frame 68, where case 7 commits, 14 ms before the bounce, it misses by 22 mm. The fit restarts its window after the bounce and catches up, to 23 mm by frame 80, but by then the plan is frozen.

> **TRY IT IN THE APP** **7 · The Hidden Spin**
>
> Run the case as loaded, with the Fit estimator. Every serve carries 40 rev/s of backspin, and the fit returns none of the ten. Press Reset, set Estimator to EKF and run again. The filter reads the spin from the bend of the flight before the bounce, and all ten come back.

11.4.2 When to Stop Watching

Information arrives at 200 Hz and the deadline does not move. An arm whose swing needs a quarter of a second or more inside its joint limits must commit at least that long before the hit, and waiting buys frames that narrow the prediction it commits to. Case 3 measures the trade over 20 random serves with the full pipeline, as shown in Figure 11.2c. Committing 50 to 150 ms after the serve, 539 to 444 ms before the hit, returns 16 or 17 and leaves the hit time uncertain by 31 to 39 ms rms. Committing 200 to 275 ms after it, 390 to 319 ms before the hit, returns 19. From 350 ms, with 243 ms left, the returns fall to 10 to 12 although the hit time is known to about 11 to 23 ms rms, because the swing no longer fits. The in-plane error at the commit falls only from 38 to 14 mm across the sweep, so the cost of committing early is mostly in the timing.

Pause. Try Exercise 11.2 before reading on.

11.4.3 What the Arm Can Do in the Time Left

The stroke must be delivered at an instant, with a position, an orientation and a velocity, by drives with limits. A minimum-jerk move of 0.3 m in 0.25 s needs 27.7 m/s^2 at its peak, 1.44 times the 19.2 m/s^2 of a bang-bang move (Equation (1.5)), and on the arm it asks 68 % of a speed limit and 100 % of the base's acceleration limit and ends 0.1 mm from its pose. The same move with the joint commands stepped to the pose ends 150 mm off. A move of 0.5 m in 0.25 s needs 32 m/s^2 and 4.0 m/s even on a bang-bang profile, and on the arm its minimum-jerk version asks 144 % of an acceleration limit and ends 116 mm off. The forward stroke is limited by the elbow to a paddle speed of about 2 m/s, so a drive planned for 7.5 m/s leaves at 6.35 m/s.

11.5 The Pipeline as Coded

Every 5 ms each arm takes the frame due, updates its tracker, predicts where the ball will cross its hitting plane, re-plans the stroke and the arm's motion until it commits, and writes six joint commands. The controllers read only their own camera frames and their own joint angles, with three declared exceptions. The referee's dead-ball signal ends a track, the wide placement aims away from the far paddle's true position, and the tuned arm takes its spin assumption after its own serve from the serve it launched. The truth, where the ball really crossed, is recorded for the display and the score and logged in columns prefixed `truth_`. No loop in this application is a PID loop. The velocity-form routine of Section 1.6.2 plays no part here, nothing integrates an error, so there is no windup to guard against, and the only limits are the drives' own.

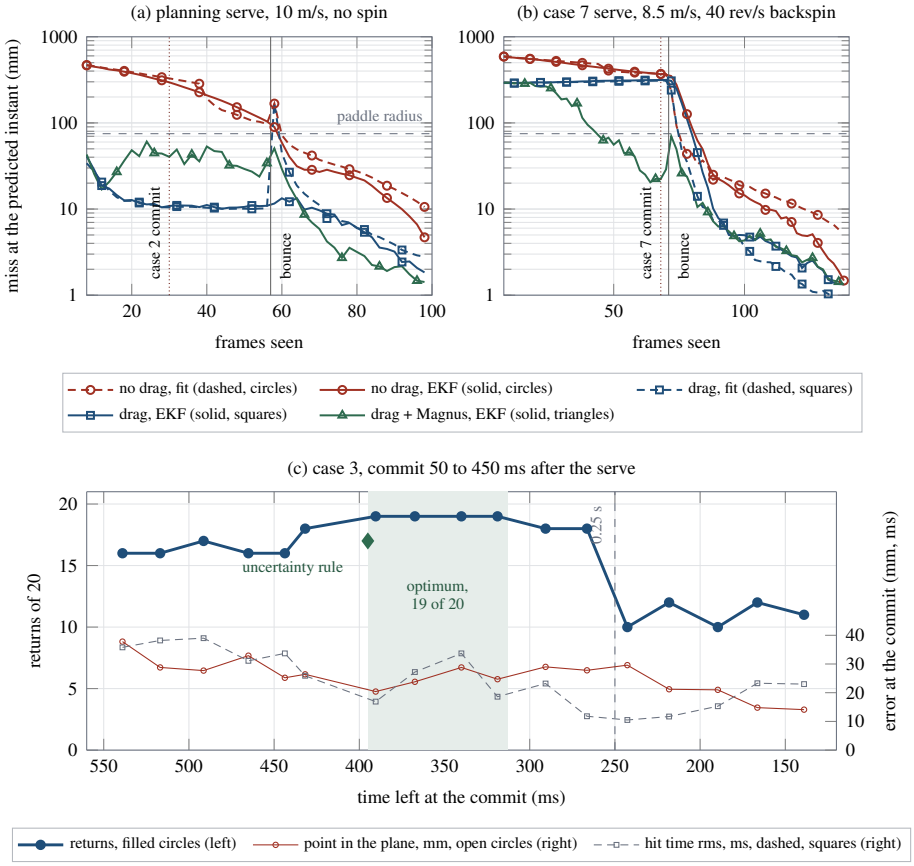

Figure 11.2 Prediction and the commit, from the application's own predictor and engine. (a) and (b) The distance from the predicted hitting point to the true ball at the predicted instant against the camera frames seen, the median of 20 noise seeds, for the least-squares fit (dashed) and the EKF (solid) on the no-drag (circles), drag (squares) and drag-and-Magnus (triangles) models. The fit has no spin state, so the fit on the drag-and-Magnus model equals the fit on drag and is not drawn. The horizontal line is the paddle's radius. (c) Case 3 over 20 random serves with the commit fixed 50 to 450 ms after the serve in 25 ms steps, with the returns that land (left axis, filled circles) and, at the commit, the in-plane error of the committed point (open circles) and the rms error of the hit time (dashed, squares) (right axis), against the time left before the hit. The band marks the optimum, 19 of 20, and the diamond the uncertainty rule.

11.5.1 Tracking and Prediction

A flight starts when a frame lands more than 3 cm from where the track expects it, at a strike or a serve, and the next frame initializes position and velocity from the two. The fit then solves for the position and velocity at the latest frame by one Gauss-Newton pass per frame over at most 80 frames since the flight began or since the bounce its model predicts, with the model's own curvature carried back from the estimate, and it has no spin. The EKF [12], the nonlinear form of the variance-chosen filter of Section 1.4.6, carries $\hat{\mathbf{x}} = [\hat{\mathbf{p}},\ \hat{\mathbf{v}},\ \hat{\omega}]$ on the chosen model of Equation (11.1) with $\mathbf{w} = \mathbf{0}$, its

own bounce, and the update of Equation (11.6).

$$\begin{aligned} P^- &= FPF^\mathsf{T} + Q, \qquad F = I + A\Delta + \tfrac{1}{2}(A\Delta)^2_{\mathbf{p}} \\ K &= P^- H^\mathsf{T}(HP^- H^\mathsf{T} + \sigma^2_{cam} I)^{-1}, \qquad H = [I \;\; 0 \;\; 0] \\ \hat{\mathbf{x}} &= \hat{\mathbf{x}}^- + K(\mathbf{z}_k - \hat{\mathbf{p}}^-), \qquad P = P^- - KHP^- \end{aligned} \tag{11.6}$$

where A is the flight's Jacobian, the squared term is kept only in the position rows, the propagation runs in steps of at most 5 ms, and Q adds $q_v\Delta$ to each velocity variance, $q_v = 0.01\,\mathrm{m}^2/\mathrm{s}^3$, and $100\,\mathrm{rad}^2/\mathrm{s}^3$ times Δ to each spin variance. When the model estimates spin, it starts at zero with a standard deviation of 120 rad/s and enters the estimate only through the Magnus curvature and the bounce, and otherwise it is held at zero. At a bounce the covariance passes through the bounce map's Jacobian, taken by finite differences at the contact state, and each velocity variance gains $(0.08\,\mathrm{m/s})^2$ for the bounce model's own error. The prediction rolls the estimate forward in steps of 4 ms, carrying the covariance every 20 ms, through the bounce on the arm's half, to the hitting plane x_h, or to planes 0.10 and 0.20 m nearer the table when the ball would pass x_h less than 3 cm above the surface. The uncertainty σ is the larger of the across and height standard deviations at the plane.

11.5.2 The Stroke, Inverted From the Landing

The stroke is a boundary-value problem. A landing target on the far half, the middle or 0.55 m off it away from the far paddle, fixes the outgoing velocity, found by Newton's method on the elevation and heading of rollouts of the arm's own flight model to 5 mm, with the ball passing at least 5 cm over the tape. The normal and velocity of the paddle then follow from Equation (11.2), as shown in Equation (11.7).

$$\begin{aligned} u_n &= \frac{v_n^{out} + e_p\, v_n^{in}}{1 + e_p}, \qquad \mathbf{u} = u_n\mathbf{n} + b\,\hat{\mathbf{t}} \\ \mathbf{n} &\parallel \mathbf{v}^{out} - \mathbf{v}^{in} - \mathbf{J}_t \end{aligned} \tag{11.7}$$

where $\hat{\mathbf{t}}$ is the face's upward direction and b the brush, iterated until the normal moves less than 10^{-7}, with the outgoing spin fed back into the aim over two or three passes. A target out of range raises the speed by 12 % at most twice, and a ball that would not clear the tape lowers it by 12 % at most four times. The push leaves at 6.0 m/s aimed 0.55 m from the far end line, the drive at 7.5 m/s aimed 0.40 m from it and the spin stroke at 7.0 m/s with $b = 2.4\,\mathrm{m/s}$, aimed 0.45 m from it. The paddle's center goes to $\hat{\mathbf{p}}_h - r\mathbf{n}$.

11.5.3 Inverse Kinematics and the Three Trajectories

Damped least squares [13, 14] solves the five-number task, as shown in Equation (11.8).

$$\Delta\mathbf{q} = J_t^\mathsf{T}\left(J_t J_t^\mathsf{T} + \lambda^2 I\right)^{-1} \begin{bmatrix} \mathbf{p}^* - \mathbf{p}_t \\ L(\mathbf{n} \times \mathbf{n}^*) \end{bmatrix} \tag{11.8}$$

where J_t is the geometric Jacobian with its angular rows projected across the normal and scaled by $L = 0.3\,\mathrm{m/rad}$, $\lambda = 0.02\,\mathrm{m}$, the step is capped at 0.5 rad and the joints are clamped into range after each one. Warm-started every 5 ms it takes about one iteration and 1 µs, and a pose counts as reached within 1 cm and 0.05 rad. The joint rates at the hit come from a regularized least squares that gives the paddle $\mathbf{u}$ with its normal held. The *PD step* sends the hitting pose to the drives unshaped. The *trapezoid* waits at the ready pose, then moves the paddle on a straight line at $25\,\mathrm{m/s}^2$ and arrives at rest. *Min-jerk* plans a quintic per joint, Equation (11.9), first to a run-up pose 0.12 m behind the hitting pose reached 0.22 s before the hit, then at the commit from the present reference to the hitting pose and its joint rates, then a 0.16 s follow-through and a 0.6 s recovery.

$$q_i(\tau) = \sum_{k=0}^{5} c_{ik}\tau^k, \qquad (q_i, \dot{q}_i, \ddot{q}_i)\big|_{\tau=T} = (q_{h,i},\ \dot{q}_{h,i},\ 0) \tag{11.9}$$

where the six coefficients also match the reference's position, rate and acceleration at $\tau = 0$, and T is the time left to the hit. A swing asking more than a joint limit at any of 25 samples is slowed, the stroke's speed scaled by 0.85 up to seven times. The reference r reaches the drives through the inverse of the servo tuned as commissioned, as shown in Equation (11.10).

$$q_c = r + \frac{2\zeta_0}{\omega_0}\dot{r} + \frac{\ddot{r}}{\omega_0^2}, \qquad \omega_0 = 2\pi \cdot 15\,\text{Hz}, \quad \zeta_0 = 1 \tag{11.10}$$

so a drive tuned as assumed follows the reference exactly in its linear range, and a drive detuned from it lags the plan. Each drive of Equation (11.4) is reverse acting in the convention of Section 1.3 ($a = +1$), because its acceleration rises when the angle falls short of its command. A ball predicted 30 cm to the student's left turns the base joint left, in the positive direction.

11.5.4 Commit Rules

Until the commit the plan is re-made every 5 ms from the latest prediction, once at least eight frames have arrived. *Time* commits a fixed time after the flight's first frame, *Frames* after a number of frames and *Uncertainty* when $\sigma \le \sigma_c$, 10 mm by default, or when the time left falls to $\max(0.12\,\text{s},\ T_{sw})$. For the minimum-jerk trajectory T_{sw} is the shortest swing, found by bisection between 0.06 and 0.5 s, whose quintic stays within 95 % of every limit, timed at the speed the stroke will get. For the drive and the spin stroke a long swing winds back before it goes forward, so the bound often sits at its 0.5 s cap. At the commit the plan freezes and the swing is planned in the time left, so a late commit leaves too little time and the swing is slowed.

11.5.5 Manual Mode and the Opponents

In Manual the paddle stays 0.45 m behind the end line, facing the table, where the two sliders or the arrow keys put it, through the same inverse kinematics, and on a phone a tap on the table moves it. The steady server returns every ball it reaches with a push within 0.15 m of the middle of the student's half, 0.40 to 0.60 m from the end line. The tuned arm has the student's best pipeline, the drag-and-Magnus EKF, minimum-jerk swings and 15 Hz drives, plays a drive stroke on every return at a seeded target 0.25 to 0.55 m from the student's end line and within 0.36 m of the center line, and returns to the center between strokes. Its blind spot, written in the model dialog, has two parts. It commits a fixed 150 ms after the student's strike whatever its uncertainty, and until the ball has bounced on its half it holds the spin that a flat return of its own shot would carry, 40 % of the incoming spin, without estimating it.

11.6 Nine Cases and a Game, Measured

Table 11.3 lists the cases and Table 11.4 what they returned. The numbers below come from the application's engine stepped outside the browser exactly as the page steps it, at game seed 1, and any setting a sentence changes was set before the run began. A return counts when the ball lands on the far half. The application passes its 36 verification checks. Figure 11.3 draws the returns of Table 11.4 as bars.

11.6.1 By Hand, and the Flight Model

Case 1 hands the paddle to two sliders against twelve machine serves at 7.0 to 8.0 m/s, which take 0.62 to 0.75 s to reach the hitting plane. Left where it starts, the paddle returns none and the first point is lost 1.85 s into the run, and the status line says how far, and to which side, the ball passed the paddle. Switched to Auto before Run, the arm returns all twelve. Case 2 serves the planning ball ten times to a no-drag fit that commits 150 ms after the serve. At the commit its hitting point is 3.6 cm from the true crossing, but it expects the ball 71 ms early, so each swing passes before the ball arrives and all ten returns go out. With drag the point is 0.9 cm off and the time 1 ms, and all ten land 0.6 cm from their targets. Drag and spin also returns all ten, 3.2 cm from the targets, because the fit has no spin state and only the stroke's aim allows for spin.

Table 11.3 The Nine Cases and the Challenge

#	Configuration	Teaching point
1	Manual, 12 machine serves near 7.5 m/s	a hand is too late
2	no-drag fit, commit 150 ms, 10 planning serves	the flight model
3	full pipeline, commit 50 ms, 20 random serves	the commit time
4	PD step, serves 0.30 m either side	the trajectory
5	full pipeline, drives at 3 Hz	the drive tuning
6	drag model, push, 12 planning serves	the stroke
7	fit, 10 serves with 40 rev/s backspin	a hidden state
8	push against the steady server	a rally in a loop
9	one 50 Hz camera, 100 ms late, 13 m/s drives	honestly infeasible
C	a game to eleven against the tuned arm	everything at once

Table 11.4 Outcomes Read From the Running Engine

Case and condition	Result
1, sliders untouched / Auto	0 of 12, first point lost at 1.85 s / 12 of 12
2, no drag / drag / drag + spin	0, 10, 10 of 10; at the commit 3.6, 0.9, 0.9 cm off, hit time 71, 1, 1 ms rms
3, commit 50 / 200 / 450 ms	16, 19, 11 of 20; 539, 390, 139 ms left
3, uncertainty rule	17 of 20; commits 191 ms in, 395 ms before the hit
4, PD step / trapezoid / min-jerk	0, 0, 10 of 10; 10, 6, 10 struck
5, drives 3 / 5 / 8 / 15 Hz	5, 9, 12, 10 of 12 land; 7 net at 3 Hz
6, push / drive / spin	12 of 12 each; 1.0, 0.5, 4.7 cm from target
7, fit / EKF / EKF on drag	0, 10, 0 of 10; EKF reads 33.5 rev/s of 40
8, push / drive, 120 s	40, 43 land; 21.1, 22.5 cm; longest rally 12, 14
9, slow camera, 144 designs	0 of 10 in every one
9, 200 Hz, 10 ms / 100 Hz, 50 ms	10 of 10 at 2.7 cm / 1 of 10
C, shipped / suggested	0–11 in 28.2 s / 11–9 in 66.3 s

Pause. Try Exercise 11.1 before reading on.

TRY IT IN THE APP **1 · By Hand**

Load the case, press Run and leave Paddle across and Paddle height where they start. Watch each ball pass the paddle: none of the twelve serves comes back, the first point is lost 1.85 s into the run, and the status line says how far, and to which side, the ball passed. Press Reset, switch to Auto and run again. The arm returns all twelve.

11.6.2 The Commit, the Trajectory and the Drives

Case 3 is the sweep of Figure 11.2c. The uncertainty rule returns 17 of 20, committing 191 ms after the serve on average and 395 ms before the hit, and all 20 of its commits come from the latest-swing bound, because σ never falls to 10 mm before the bounce. The spin about the line of flight leaves no trace in the frames until the bounce, so σ stays near 20 cm, and the rule is in practice a deadline. In case 4, with serves 0.30 m to either side, stepped commands strike all ten balls and return none, because the drives saturate, overshoot and arrive with no stroke speed. Moved 0.3 m across in 0.25 s on the arm alone, stepped commands hold the base, elbow and forearm-roll drives on their acceleration limits for 240, 239 and 243 ms of the 250 ms. The trapezoid strikes six, arrives at rest and returns none, while the minimum-jerk swing returns all ten, its drives at a limit for 7 % of the committed time against 52 % for the step. Case 5 detunes the drives to 3 Hz while the commands are still shaped for 15 Hz. The drives lag by about $2\zeta/\omega = 0.106$ s, every ball is struck late with the face tilted, seven of twelve go into the net and five land 54.5 cm from their targets on average. At 5, 8 and 15 Hz, 9, 12 and 10 land.

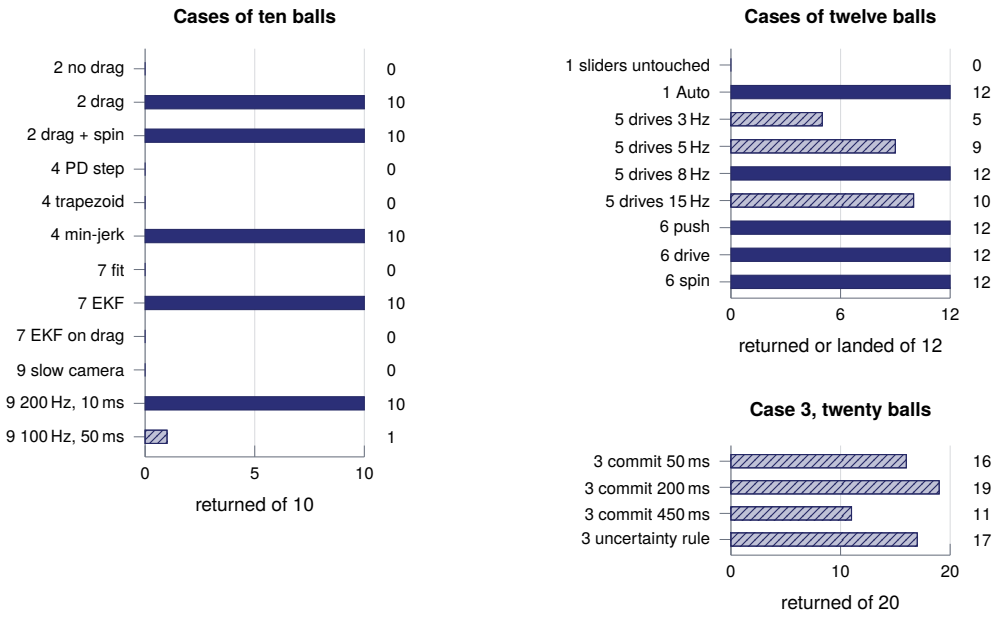

Figure 11.3 The returns of Table 11.4 drawn as bars, by case and setting, in three panels by the number of balls each case plays. Solid bars return every ball, hatched bars some of them, and a setting with no bar returns none; case 5 counts the balls that land. Notice that each case turns on one part of the pipeline, the flight model in case 2, the trajectory in case 4 and the spin state in case 7, and that in case 3 neither the earliest nor the latest commit returns the most.

> ✦ **DID YOU KNOW? Knowing More, Returning Less**
>
> From the simulator: in case 3, a commit 200 ms after the serve returns 19 of 20 balls with 390 ms left before the hit. Waiting until 450 ms sharpens the prediction but returns only 11, because 139 ms is too little time for the swing (Table 11.4).

11.6.3 The Stroke and a Hidden Spin

Case 6 returns the planning serve with the drag model. The push leaves at 5.79 m/s and lands 1.0 cm from its target on average, the drive at 6.35 m/s and 0.5 cm, and the drive placed wide 1.0 cm, so the planner's model is the flight's. The ball arrives with 35.1 rev/s of topspin picked up rolling off the table, and both flat strokes send it back with 14 to 15 rev/s of backspin. The spin stroke moves the paddle at 2.87 m/s, brushing up across the ball, leaves 1.7 rev/s of backspin and lands 4.7 cm off, because the drag model leaves out the lift of the spin it puts on. Case 7 serves 40 rev/s of backspin at 8.5 m/s. The fit returns none of ten and the EKF on the drag model none, while the EKF on the drag-and-Magnus model returns all ten, having read 33.5 rev/s of backspin from the curvature of the flight at the first serve's commit, 356 ms before the hit. The bounce makes the backspin visible to every model, because the contact slides and the ball leaves the table with 2.18 m/s of horizontal speed instead of the 3.66 m/s of the same serve without spin, but it comes too late.

🕮 **Pause.** Try Exercise 11.3 before reading on.

11.6.4 A Rally, and a Camera Too Slow

Case 8 plays the steady server for 120 s. With the push 40 of the student's returns land, 21.1 cm from their targets on average, and the longest rally is 12 strikes. The server fails to return 18 balls and the student 3, and 10 of the server's failures are pushes that bounce twice before its hitting plane,

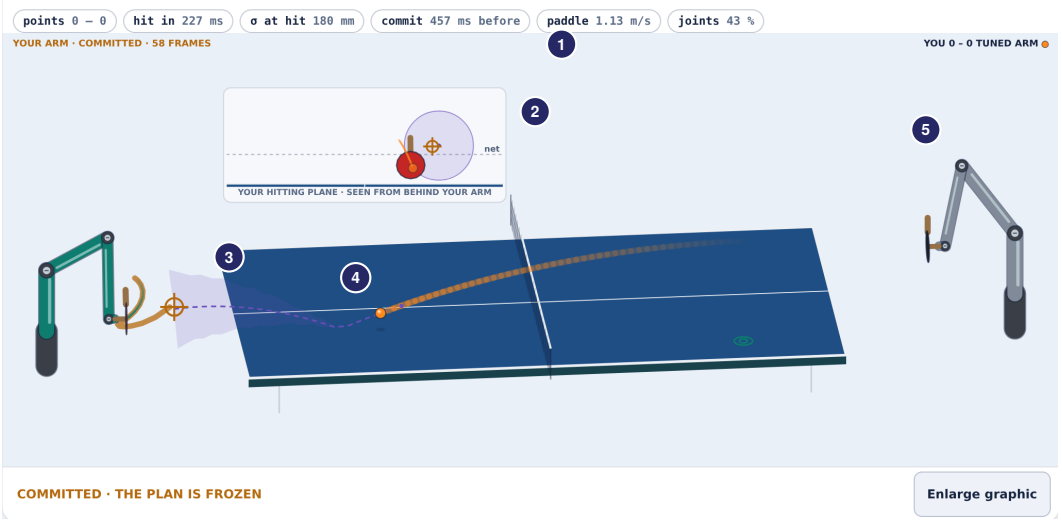

Figure 11.4 The student's arm committed on the tuned arm's first serve, 1.20 s into the challenge played with the suggested settings, on the Table view, cropped to the plant view. The tuned arm (5) served at 0.90 s, 8.50 m/s with 17.1 rev/s of topspin, and the ball (4) is about to bounce. The student's arm committed at 0.985 s, 85 ms after the first frame and 457 ms before the planned hit, and has 58 frames. The badges (1) read the score 0–0, the hit in 227 ms, σ at the hit 180 mm, the commit 457 ms before, the paddle at 1.13 m/s and the joints at 43 % of their limits. The dashed predicted path (3) widens inside its band toward the hitting-point marker (amber in the app), and the panel (2) above the near half draws the hitting plane from behind the arm with the 180 mm circle around the prediction.

which an arm on a fixed base cannot reach. With the drive 43 land, 22.5 cm from their targets, and the server fails 10 times. Over ten game seeds the push's longest rally in 150 s runs from 8 to 25 strikes, so a rally's length is a matter of chance and the case is scored on landings. Case 9 gives the arm one camera at 50 Hz with 100 ms of latency against 13 m/s drives. The drive reaches the plane 351 ms after it leaves, the eighth frame is exposed 140 ms into the flight and delivered 100 ms later, so the first plan comes 110 ms before the ball arrives. Every one of 144 designs, 16 commit settings by three strokes by three trajectories, returns none of ten. The 200 Hz, 10 ms camera returns all ten, 200 Hz with 50 ms of latency all ten at 17.0 cm, 100 Hz with 50 ms one, and 50 Hz with 20 ms none.

11.6.5 The Game, and Where the Tuned Arm Loses

The challenge is a game to eleven, won by two, against the tuned arm at the default game seed, with the model and the opponent as given, and the score is the points. The configurations are compared in Table 11.5. As shipped, with the no-drag fit, a commit 100 ms after the opponent's strike, the PD step and the push, the student loses 0–11 in 28.2 s and wins none of ten seeds, and the status line names the first loss as the ball 71 ms later than predicted with a joint at its limit. The drag model and the minimum-jerk swing alone win all ten seeds. The suggested settings, the drag-and-Magnus EKF, the uncertainty commit, minimum-jerk swings and the spin stroke placed wide, win 11–9 in 66.3 s and all ten seeds, and 20 of their 21 commits come from the latest-swing bound, 316 to 499 ms before the hit. The state of that game 1.20 s in is shown in Figure 11.4. The status line names the cause of every point the student's arm loses, a prediction off and by how much, a ball late or early, a late paddle, a drive at its limit or a pose out of reach. Play ten seeds runs the student's settings over game seeds 1 to 10 in the background and reports the win rate beside the measured benchmark of the shipped

Table 11.5 The Challenge Against the Tuned Arm, Game Seeds 1 to 10

Configuration	Seed 1	Won
shipped (no-drag fit, commit 100 ms, PD step, push)	0–11	0 of 10
shipped with drag and min-jerk	11–5	10 of 10
suggested (drag + Magnus EKF, uncertainty, min-jerk, spin wide)	11–9	10 of 10
suggested, placed in the middle	11–3	9 of 10
suggested, commit fixed at 150 ms	4–11	5 of 10
suggested, opponent estimates spin all the way	11–5	10 of 10
suggested, opponent commits at 200 ms / 300 ms	11–7 / 11–7	10 / 9 of 10

and suggested settings, and the challenge stops crediting a game as soon as any model parameter or disturbance moves from its published value.

The tuned arm was designed to be beaten at its declared blind spot, and the expectation was that a heavy spin stroke or a wide placement would do it. Measured, the two parts of the blind spot are not equal. With its spin assumption removed in the model dialog, the suggested settings still win all ten seeds, so the assumed spin decides nothing. Of the 110 points the tuned arm loses in the suggested settings' ten games, 52 come from its plan, frozen 150 ms after the student's strike, putting its paddle more than a paddle's radius from where the ball crossed, 34 from wide balls whose hitting pose sits at the end of its reach or a joint range, 13 from its own returns going out, 6 from swings beyond its joint limits, 3 from balls that bounce twice before its plane, and one each from a missing commit and a joint at its range end. Against pushes to the middle, the shipped settings with drag and minimum-jerk swings, 40 of the 113 points it loses are balls that bounce twice before its hitting plane. Moving its commit to 200 ms or 300 ms still loses 10 and 9 seeds of ten, so the fixed clock is not the only weakness. The arm that loses is weak where its geometry is, a reach that ends about 0.65 m off the middle and a hitting plane behind its end line, as much as where its model is. The gap between the design intent and the measurement is itself the finding. A good controller is beaten where it actually fails, which a match statistic finds and a design note only suggests. The student's own commit matters as much. The suggested pipeline made to commit at a fixed 150 ms wins only five of ten, and of the 47 points the student's arm loses in the ten suggested games, 15 are committed points more than 7.5 cm off and 10 are short balls.

The game's secondary counts depend on the step. At 0.25 ms the shipped settings still lose 0–11 and the suggested settings win, 11–5 instead of 11–9, with the same 0 and 10 of ten seeds, while the case-3 returns move by one or two serves at a setting and case 5's count at 5 Hz moves from 9 to 11. The verdicts and the optimum survive the halving, and the counts are quoted at the shipped step.

11.7 Technology in Practice

Each arm of Section 11.3 reads a camera pair and its own joint angles, and it moves six joint drives. Figure 11.5 draws the student's loop with real devices in the place of the schematic's symbols.

What Measures It

A robot player measures the ball with high-speed cameras alone, because the ball itself cannot be instrumented. The two cameras of a pair are synchronized so that each detection can be triangulated, and a neural detector can work on the raw sensor image, so that no time is spent converting it to color [1]. A player driven by artificial muscles read the ball's position from a color-based camera system at 180 Hz [11]. The stereo pair of Equation (11.5), with 200 Hz frames 10 ms late and 2 mm of noise, stands for such a system, and its dropout event for the frames a real one loses. One research player kept its returns until its cameras were slowed to 50 frames a second or delayed by 150 ms, when its swings came too late and its balls went into the net [1], a limit of the kind that case 9 measures.

TECHNOLOGY Sensor: High-Speed Stereo Cameras

A research player tracked the ball with two synchronized cameras with global shutters, a 4 ms exposure and 125 frames a second, mounted about 2 m above the play area on opposite sides of the table, and it triangulated each pair of detections into a position. The wide lenses that let both cameras see the whole play area made calibration harder and added uncertainty to the triangulation. When processing falls behind the frames, latency grows until frames are dropped [1], and a ball that climbs out of the field of view is lost to the robot [16]. On a trend either fault shows as a gap in the camera track, while the prediction runs on its model alone.

What Moves It

The final elements are the joint drives of a motor-driven industrial arm, which follow the references of the arm's own controller [1]. The drives of Equation (11.4), 15 Hz servos with speed and acceleration limits and soft range ends, stand for such drives. The inverse servo of Equation (11.10) plays the part of the identification a real arm needs before its motion matches the plan, and a drive detuned from that model lags, as case 5 shows.

TECHNOLOGY Final Element: Joint Servo Drives

Each joint of an industrial arm is a servo drive that follows a reference from the arm's own controller. In one table-tennis system the arm took a position and a speed target for each joint and returned joint feedback at 248 Hz, and its joints turned at up to 420 or 600 °/s. The controller passes the reference through a low-pass filter and gains of its own, so the arm was identified before its motion matched the simulator. Its safety layer stops the arm when it detects a collision, predicts a joint past its range, or finds a joint commanded so fast that it reaches a torque limit, and the stop freezes the robot [1]. On a trend a safety stop is a joint that halts in mid-swing and stays there.

Where the Controller Runs

The controllers of this chapter run every 5 ms on drives integrated every 0.5 ms. A research player ran its learned policies at 100 Hz in most experiments, while the arm's controller exchanged commands and feedback at 248 Hz through low-level control threads kept apart from the learning code [1]. A later player split the work in two layers: skill policies sent joint velocity commands at 50 Hz, and a high-level controller chose a skill once for every stroke of the opponent [16]. Here the tracker, the predictor, the commit rule and the stroke planner form the upper layer, and the shaped commands and the drives the lower one.

IN PRACTICE Measuring the Delays Before Tuning

A fast vision loop is commissioned by measuring its delays first. One team logged time stamps through its whole software stack, measured the mean and the spread of each component's latency, and gave the same delays to its simulator. Its cameras went on opposite sides of the table, which cut the bias of the triangulation [1]. The arm and its cell fall under ISO 10218, whose Part 1 covers the industrial robot and Part 2 the robot system and its integration [18]. The camera of Equation (11.5) has one fixed latency, so the model skips this step.

What Is Modernizing It

Policies trained in simulation now move to the real arm without further training, once the simulator carries the measured delays of each component and the balls that real play produces [1, 16]. A parallel simulation that runs beside the real arm serves as a safety layer and keeps it out of states defined as undesirable [1]. Arms driven by pneumatic artificial muscles are a second route. Each muscle contracts as its air pressure rises and can only pull, so two act against each other on one joint. They

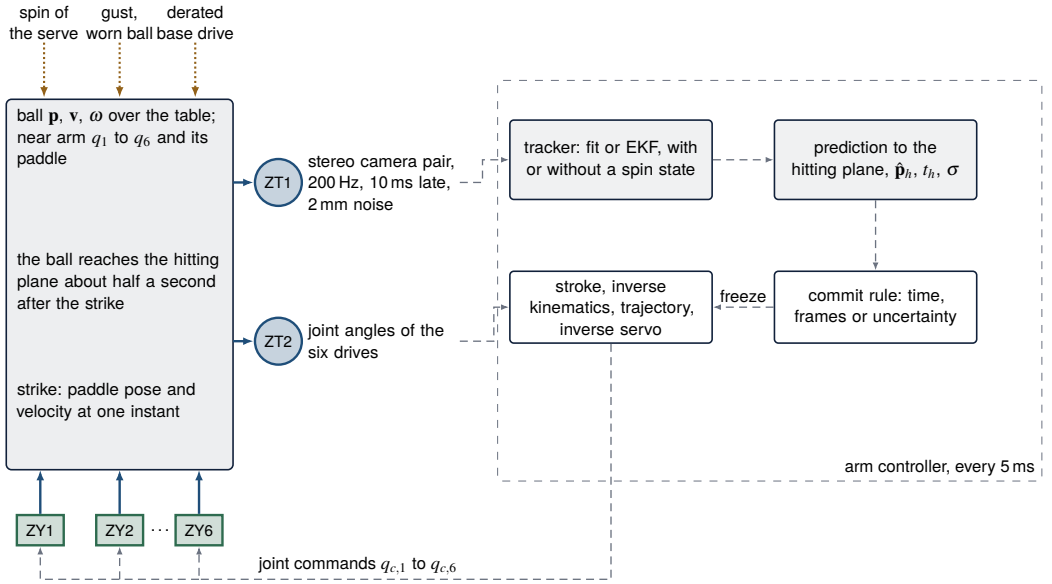

ZY1 to ZY6: joint servo drives, 15 Hz with damping 1, base, shoulder and elbow limited to 360 °/s and 3000 °/s^2, the three wrist axes to twice that.

Figure 11.5 The student's loop with real devices in place of the schematic's symbols. ZT1 is the stereo camera pair and ZT2 the joint angles. In the arm controller (dashed frame) the tracker feeds the prediction to the hitting plane, and the commit rule freezes the plan that the stroke, the inverse kinematics and the trajectory turn into six joint commands for the drives ZY1 to ZY6. Notice that the camera reaches the drives only through the commit: after it, new frames no longer change the swing.

give high accelerations, and their softness keeps the arm inside a safe region while it learns, but they are nonlinear, change with temperature and wear, and show hysteresis [11].

11.8 Limitations

The ball's drag and Magnus factors are teaching values smaller than fitted ones, its spin does not decay, and its impacts are instantaneous, with a grip law in place of a compliant contact and a table grip four times a measured one. The arms are servo-level, with exact kinematics and limited drives but no dynamics, fixed bases and hitting planes behind their end lines, and their forward stroke is limited by the elbow to about 2 m/s of paddle speed, so its drives leave at about 6.4 m/s. The cameras have no calibration error, blur or occlusion, serves are launched rather than struck, and the serve rule, lets and edge balls are left out.

> **OPEN CHALLENGE** **Fast Balls and Spin Nobody Sees**
>
> A learned robot player that beat every beginner it met still lost every match against the most advanced players. Its latency of about 100 ms limited the decisions it could make on a fast ball, and it could not read extreme spin, which opponents exploited [16]. Spin is hard because the cameras see only the ball's position, and spin shows only in a bend of the flight and in the bounce, as case 7 of this chapter shows. Progress would be a robot that reads heavy spin and answers fast balls well enough to beat advanced players. Its builders name predictive models of the ball and faster links between sensors and actuators as routes [16].

The tuned opponent is a teaching opponent, the student's pipeline with a declared blind spot and safe targets, and not a model of any player or machine. Its losses are counted by a set of rules applied to its own flight record, and a point with several faults is charged to the first that applies. The

uncertainty rule's latest-swing bound is conservative, so it commits up to 0.5 s before the hit for the drive and the spin stroke. A game's counts are sensitive to the integration step and the seed, which is why the challenge reports ten seeds beside the default one.

11.9 Conclusion

A table-tennis robot has about half a second to watch, decide and swing, and this simulator makes each part of that budget measurable on one table with two six-axis arms. Drag moves the planning serve's arrival by 131 ms and its bounce by 34 cm, so a no-drag predictor commits 3.6 cm from the true crossing and 71 ms early and returns nothing, while a drag model returns all ten serves within a centimeter of their targets. Committing 200 to 275 ms after the serve returns 19 of 20, earlier commits return 16 or 17 and later ones as few as 10. A minimum-jerk swing reaches a pose 0.3 m away in a quarter second to 0.1 mm, where stepped commands end 150 mm off, and only an estimator with a spin state reads a 40 rev/s backspin before the bounce. A slow camera defeats all 144 designs tried. The tuned opponent loses where its reach ends and its clock freezes its plan, not where its assumed spin was expected to fail it. The model is idealized in its ball, its contacts, its arms and its cameras. A fuller model would add spin decay and a compliant bounce, rigid-body arm dynamics, a moving base, struck serves and a second opponent whose blind spot lies in its model rather than its geometry. The simulator is a game in which knowing more always costs time to move.

KEY POINTS

- A robot player has about half a second to watch, decide and swing, and every look at the ball costs time to move.
- Drag moves the planning serve's arrival by 131 ms, so a no-drag predictor commits 3.6 cm from the true crossing but 71 ms early and returns nothing.
- Committing 200 to 275 ms after the serve returns 19 of 20, while earlier commits return 16 or 17 and later ones as few as 10.
- Only an estimator with a spin state reads a 40 rev/s backspin before the bounce.

Where else this applies. The same choice of when to act on a late measurement is met in Chapter 10, and the same six-axis arm keeps a ball bouncing in Volume 8, Chapter 15.

References

[1] D. D'Ambrosio, J. Abelian, S. Abeyruwan, M. Ahn, A. Bewley, J. Boyd, K. Choromanski, O. Cortes, E. Coumans, T. Ding *et al.*, "Robotic table tennis: A case study into a high speed learning system," in *Robotics: Science and Systems XIX*. Robotics: Science and Systems Foundation, 2023.

[2] B. Hove and J.-J. E. Slotine, "Experiments in robotic catching," in *Proceedings of the 1991 American Control Conference*. IEEE, 1991, pp. 380–386.

[3] R. L. Andersson, *A Robot Ping-Pong Player: Experiment in Real-Time Intelligent Control*, ser. The MIT Press Series in Artificial Intelligence. Cambridge, MA: MIT Press, 1988.

[4] K. Mülling, J. Kober, O. Kroemer, and J. Peters, "Learning to select and generalize striking movements in robot table tennis," *The International Journal of Robotics Research*, vol. 32, no. 3, pp. 263–279, 2013.

[5] R. D. Mehta, "Aerodynamics of sports balls," *Annual Review of Fluid Mechanics*, vol. 17, no. 1, pp. 151–189, 1985.

[6] J. Nonomura, A. Nakashima, and Y. Hayakawa, "Analysis of effects of rebounds and aerodynamics for trajectory of table tennis ball," in *Proceedings of SICE Annual Conference 2010*. IEEE, 2010, pp. 1567–1572.

[7] R. Cross, "Grip-slip behavior of a bouncing ball," *American Journal of Physics*, vol. 70, no. 11, pp. 1093–1102, 2002.

[8] A. Nakashima, Y. Ogawa, Y. Kobayashi, and Y. Hayakawa, "Modeling of rebound phenomenon of a rigid ball with friction and elastic effects," in *Proceedings of the 2010 American Control Conference*. IEEE, 2010, pp. 1410–1415.

[9] L. Acosta, J. J. Rodrigo, J. A. Mendez, G. N. Marichal, and M. Sigut, "Ping-Pong player prototype – A PC-based, low-cost, ping-pong robot," *IEEE Robotics & Automation Magazine*, vol. 10, no. 4, pp. 44–52, 2003.

[10] M. Matsushima, T. Hashimoto, M. Takeuchi, and F. Miyazaki, "A learning approach to robotic table tennis," *IEEE Transactions on Robotics*, vol. 21, no. 4, pp. 767–771, 2005.

[11] D. Büchler, S. Guist, R. Calandra, V. Berenz, B. Schölkopf, and J. Peters, "Learning to play table tennis from scratch using muscular robots," *IEEE Transactions on Robotics*, vol. 38, no. 6, pp. 3850–3860, 2022.

[12] R. E. Kalman, "A new approach to linear filtering and prediction problems," *Journal of Basic Engineering*, vol. 82, no. 1, pp. 35–45, 1960.

[13] C. Wampler, "Manipulator inverse kinematic solutions based on vector formulations and damped least-squares methods," *IEEE Transactions on Systems, Man, and Cybernetics*, vol. 16, no. 1, pp. 93–101, 1986.

[14] Y. Nakamura and H. Hanafusa, "Inverse kinematic solutions with singularity robustness for robot manipulator control," *Journal of Dynamic Systems, Measurement, and Control*, vol. 108, no. 3, pp. 163–171, 1986.

[15] T. Flash and N. Hogan, "The coordination of arm movements: an experimentally confirmed mathematical model," *The Journal of Neuroscience*, vol. 5, no. 7, pp. 1688–1703, 1985.

[16] D. B. D'Ambrosio, S. Abeyruwan, L. Graesser, A. Iscen, H. Ben Amor, A. Bewley, B. J. Reed, K. Reymann, L. Takayama, Y. Tassa *et al.*, "Achieving human level competitive robot table tennis," in *2025 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2025, pp. 74–82.

[17] B. Siciliano, L. Sciavicco, L. Villani, and G. Oriolo, *Robotics: Modelling, Planning and Control*, ser. Advanced Textbooks in Control and Signal Processing. London: Springer, 2009.

[18] R. Bostelman and J. Falco, "Survey of industrial manipulation technologies for autonomous assembly applications," National Institute of Standards and Technology, Gaithersburg, MD, NIST Interagency/Internal Report 7844, 2012.

11.10 Exercises

11.1 ○ Concept check · 2 min

In scenario 2, The predictor, the no-drag fit commits 150 ms after each serve. At the commit its hitting point lies 3.6 cm from where the ball really crosses the hitting plane, and all ten returns go out. Which explanation is right?

(a) The point is too far off across the plane for a paddle of 75 mm radius to meet the ball

(b) The commit leaves too little time for the swing, so the paddle reaches the plane after the ball

(c) The point is close, but the fit expects the ball 71 ms early and the swing passes before it

(d) The camera noise of 2 mm per axis scatters the fit, so the aim of each stroke is random

11.2 ○ Concept check · 2 min

In scenario 3, The commit time, a commit 350 ms after the serve returns 10 of 20 serves, against 19 of 20 for a commit at 200 ms, although the later commit knows the hit time to about 11 ms rms. What is the main cause?

(a) The frames after the bounce confuse the estimator, so the later prediction is worse

(b) The serves are faster at the later setting, because the serve speeds depend on the commit

(c) The tracker stops taking frames once it commits, so the paddle aims at a stale estimate

(d) About 243 ms remain, and the swing no longer fits inside the joint limits

11.3 ○ Concept check · 2 min

In scenario 7, The hidden spin, every serve carries 40 rev/s of backspin. The EKF on the drag-and-Magnus model returns all ten serves, while the same EKF on the drag model returns none. Why?

(a) Only the drag-and-Magnus model has a spin state, which the bending of the flight can inform

(b) The drag model overstates the air drag, so its predicted ball arrives later than the real one

(c) The drag-and-Magnus model reads the spin directly from the camera frames of the ball

(d) The drag model has no bounce, so it predicts the ball to pass straight through the table

11.4 Theory and calculation · 6 min

The ball has radius $r = 20\,\text{mm}$ and mass $m = 2.7\,\text{g}$. The app uses $\rho = 1.2\,\text{kg/m}^3$ and $C_D = 0.45$ in Equation (11.1), with $g = 9.81\,\text{m/s}^2$.

(a) Compute the drag factor k_D. *Checkpoint:* $0.1257\,\text{m}^{-1}$.

(b) Compute the drag deceleration of a ball moving at 10 m/s with no wind, and the terminal speed of a ball falling with no spin. *Checkpoint:* $12.6\,\text{m/s}^2$ and 8.83 m/s.

(c) Repeat (a) and the deceleration of (b) for the fitted $C_D = 0.54$, and say in which direction the teaching value biases the app against real air. *Checkpoint:* $0.1508\,\text{m}^{-1}$ and $15.1\,\text{m/s}^2$.

11.5 Theory and calculation · 7 min

Use the impact law of Equation (11.2) for a thin hollow sphere, $\alpha = 2/3$, with the table's $e_t = 0.905$ and $\mu_t = 1.0$ and the rubber's $e_p = 0.7$. Speeds are magnitudes along the direction stated.

(a) A ball arrives at the paddle at 6 m/s along the paddle's normal. What normal speed must the paddle have for the ball to leave at 10 m/s? *Checkpoint:* 3.41 m/s.

(b) A spinless ball reaches the table moving 6.10 m/s along the table and 2.21 m/s down. Show that the bounce grips, and find the speed along the table and the topspin with which the ball leaves. *Checkpoint:* 3.66 m/s and 183 rad/s, or 29.1 rev/s.

(c) The same serve with 40 rev/s of backspin reaches the table at 6.14 m/s along the table and 2.08 m/s down. Show that the bounce slides, and find the speed along the table after it. *Checkpoint:* 2.18 m/s.

11.6 Theory and calculation · 7 min

The paddle must move 0.3 m in 0.25 s, from rest to rest.

(a) Compute the peak acceleration of a minimum-jerk (quintic) move and of a bang-bang move, and their ratio. Equation (1.5) gives the peaks. *Checkpoint:* $27.7\,\text{m/s}^2$, $19.2\,\text{m/s}^2$ and 1.44.

(b) The base joint is limited to $360\,^\circ\text{/s}$ and $3000\,^\circ\text{/s}^2$. Convert these limits to paddle speed and acceleration at a lever of 0.5 m, and compare them with (a). *Checkpoint:* 3.14 m/s and $26.2\,\text{m/s}^2$.

(c) A drive obeys Equation (11.4) in its linear range with $\zeta = 1$ and $f_j = 3\,\text{Hz}$, but its commands are shaped by Equation (11.10) for 15 Hz. For a reference that ramps at a constant rate, find how far in time the joint lags the reference once the transient has passed. *Checkpoint:* 84.9 ms.

11.7 App challenge · 8 min

Load scenario 3, The commit time, and use the default seed. Each run ends by itself after 20 serves, about 50 s of simulated time, so run at 2x. Read the Returns on the table tile when the run ends. The counts are exact.

(a) Run the scenario as shipped, committing 50 ms after each serve. *Checkpoint:* 16 of 20.

(b) Press the reload button, set Commit time to 250 ms and run again. *Checkpoint:* 19 of 20.

(c) Press the reload button, set Commit time to 350 ms and run again. *Checkpoint:* 10 of 20.

Explain the shape of the three results with the time left at the commit that Figure 11.2c gives for each setting.

11.8 App challenge · 7 min

Load scenario 7, The hidden spin, and use the default seed. Each run ends by itself after ten serves, about 22 to 26 s of simulated time. Read the Returns on the table tile when the run ends.

(a) Run the scenario as shipped, with the Fit estimator. *Checkpoint:* 0 of 10.

(b) Press the reload button, set Estimator to EKF and run again. Export the CSV file and find the first row whose `phase` column reads COMMITTED. Read `est_topspin_rps` and `truth_ball_topspin_rps` in that row. Backspin is negative. *Checkpoint:* 10 of 10. At 1.25 s the estimate reads −33.5 rev/s, within 0.5 rev/s, against a true −40 rev/s.

(c) Press the reload button, set Estimator to EKF and Predictor model to Drag, and run again. *Checkpoint:* 0 of 10.

11.9 App challenge · 8 min

Load the challenge, a game to eleven against the tuned arm, and use the default seed. A game ends by itself, so run at 2x. Read the Points won tile and the clock when the game ends.

(a) Run the challenge as shipped. *Checkpoint:* 0–11, lost at 28.2 s, within 0.1 s.

(b) Press the reload button, press Reveal suggested tuning and run again. *Checkpoint:* 11–9, won at 66.3 s, within 0.1 s.

(c) Keep the suggested settings, but set Commit rule to Time and Commit time to 150 ms. Press Reset and run again. *Checkpoint:* 4–11, lost at 44.3 s, within 0.1 s.

Which of the tuned arm's weaknesses does the student's arm share in (c)?

11.10 Application area · 3 min

The app's table grip $\mu_t = 1.0$ is four times the 0.25 measured on a real table. What would change for the 10 m/s planning serve on a real table, and what should a robot builder do before trusting a bounce model?

11.11 Application area · 4 min

A 13 m/s drive reaches the hitting plane 351 ms after it leaves the far paddle. The predictor plans after eight frames, frame k is exposed k/f_{cam} after the launch with the first at zero, and each frame arrives τ_{cam} late. A swing needs a quarter of a second or more. For the cameras (50 Hz, 100 ms), (100 Hz, 50 ms) and (200 Hz, 10 ms), find when the first plan can be made and how much time is left. Which camera would you specify?

11.12 Review the agent · 5 min

A team uses the Uncertainty commit rule with the default threshold of 10 mm. An AI assistant advises: "The threshold is loose. Lower Commit uncertainty to 2 mm so that the arm commits only once the hitting point is sharp. In scenario 3 this will raise the returns above 17 of 20." Find the flaw in the advice and say how you would check it in the app.

CHAPTER 12

Self-Regulating Arc and Distance-Keyed Seam Preview in a Robotic Welding Simulator

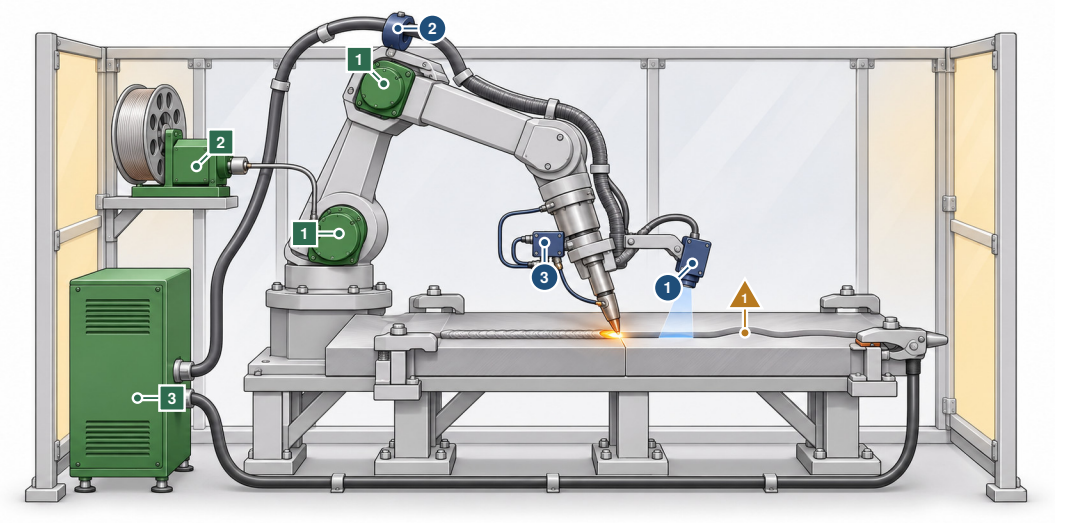

The controller moves 1 robot axes that move the torch up and across the seam 2 wire feeder 3 power source, its voltage set with the feed

The controller reads 1 laser that reads the groove ahead of the torch 2 current sensor on the welding cable 3 voltage sensor at the torch

What disturbs it 1 a seam that bends, rises or opens ahead of the torch

Notice that the laser reads the groove ahead of the arc, so each reading belongs to a place the torch has not yet reached. Illustration, not to scale.

IN PLAIN TERMS **A Welding Arc That Sets Its Own Length**

The main idea. A robot joins two steel plates by melting a wire into the joint with an electric spark called an arc. If the torch dips closer, more current melts the wire faster, and the arc grows back to its length with no controller.

The challenge. A laser reads the joint just ahead of the torch, so each reading is saved for the moment the torch reaches that spot. A reading saved for the wrong spot steers the torch off the joint, and nothing in the reading shows it.

Why it matters. A bead that misses an edge of the joint makes a weak weld.

THREE WAYS TO READ THIS CHAPTER

Explore	11 min	the first page, the history segment (Section 12.2), the fact on the source slope (Section 12.3), the try-it segments of cases 1 and 4 (Sections 12.4 and 12.5), the fact on adaptive fill (Section 12.6) 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 12.6), Technology in Practice (Section 12.7), Limitations (Section 12.8) and the open challenges

The control question A constant-voltage source and a constant-speed wire feeder hold the arc length of a robotic weld with no controller between them, and a laser reads the seam a fixed distance ahead of the arc. How does the arc correct itself, what does its current reveal about the torch height, and how must a reading taken ahead of the torch be filed so that it is used at the place it describes?

Learning objectives

- Linearize the melting balance to find the time constant of arc self-regulation and the current's sensitivity to the contact-tip distance at steady state and at the weave frequency.
- Explain why a constant-current source removes the self-regulating loop and how an arc-voltage-controlled feed restores it.
- Compute the lateral error that a misconfigured look-ahead produces on a bending seam, and predict where along the bend it is largest.
- Check a travel speed and a wire feed against the fusion, undercut and underfill limits of the bead correlations.
- Configure tracking, adaptive fill and the procedure for a seam whose gap and position vary, and justify each choice by the defect it prevents.

Prerequisites Sections 1.4.3 (position loops around velocity loops), 1.4.5 and 1.4.6 (following a path and the age of a report), 1.6.2 (velocity-form PID and loop action) and 1.6.4 (running and reading a scenario).

The app https://apmonitor.com/apps/robot-arc-welding/

See also Volume 3, Chapter 6 doses a coagulant from a measured raw water turbidity, a feedforward bounded by the plant as adaptive fill is. Volume 5, Chapter 10 controls with a psychrometric soft sensor, and Volume 5, Chapter 7 follows material by volume past a delayed analyzer, as the preview buffer files readings by distance.

12.1 Introduction

A welding robot replays a path that was taught on one part, and the part in front of it is never quite that part. Robots in industrial welding operations are common, although programming them and mastering the process they carry are still far from routine [1]. In gas metal arc welding the torch feeds a consumable wire through a contact tip into an arc, a power source and a wire feeder set the process between them, and the pool that the arc leaves behind has to fuse both plates, fill the groove and freeze without a defect [2]. Successful robotic arc welding was argued early to depend on sensory feedback, with optical sensing and sensing through the arc itself as the two prevalent approaches [3]. In the simulated cell of this chapter the procedure feeds 8 m/min of 1.2 mm wire at a contact-tip-to-work distance (CTWD) of 15 mm and travels at 10 mm/s, which carries 252.1 A at 27.05 V and puts 681 J/mm into a 6 mm plate.

The plant has two properties that a block diagram hides. The first is that the arc length has no controller. A constant-voltage source and a constant-speed feeder hold it between them, because a shorter arc draws more current, more current melts the wire faster than it arrives, and the extension burns back until the arc is long again. When the plate steps up 3 mm under the torch, the model's current rises from 252.1 to 316.5 A within 17.5 ms and the arc is 90 % restored 97.0 ms after the step, while the robot's 50 Hz controller does nothing at all. A constant-current source has no such loop, and the same step leaves its arc 0.99 mm long until an arc-voltage-controlled feeder brings it back, 90 % of the way in 834.5 ms. The simulator exists to overturn one belief, that every variable a plant holds steady is held there by a controller. The 20.5 A that the constant-voltage arc keeps after the step is also a measurement, and through-the-arc sensing reads the torch height from it [3, 4].

The second property is that the laser seam sensor and the arc are separated by distance, not by time. A line scanner 40 mm ahead of the arc measures the groove's position, root gap and height, and its reading becomes useful when the torch arrives at that place, 4 s later at 10 mm/s and never if the robot stops. The controller therefore files each reading under the travel distance it describes and asks for the reading at the torch's own position. A reading filed under the wrong distance describes the wrong place. With the look-ahead configured at 20 mm on a laser mounted at 40 mm, the torch

runs up to 1.45 mm off a seam that bends 3.5 mm to either side every 300 mm, and 372 mm of the 1200 mm seam lose fusion at a groove edge.

None of this is peculiar to welding. A variable inferred from others through a model is an inferential or soft sensor, and soft sensors built from process data are a subject of their own in the process industries [5]. The psychrometric soft sensor of Volume 5, Chapter 10 is one such inference, corrected by a delayed laboratory. A controller that sees its command a fixed interval ahead is a preview controller, a class with its own optimal theory [6] and an established application in the active suspension of vehicles [7], which Volume 7, Chapter 11 treats with road-preview feedforward. The weld adds three complications to those textbook forms. Its preview lead is a distance fixed by a bracket, its part moves between the reading and the torch's arrival, and its feedforward from the measured gap runs out of room where extra metal brings more heat than the joint can take. None of the three survives the reduction of the torch to one sensor, one actuator and a delay.

The welding literature treats arc sensing, seam tracking and adaptive filling as sensor and process technology [8, 9], and a control course meets preview and inferential control as block diagrams, if at all [10], so a student rarely sees an arc correct itself faster than any controller could, or a correct laser reading used at the wrong place. This chapter shows both properties on a simulated robotic arc welding cell whose log keeps the true seam, CTWD, arc and bead apart from what the laser, the current and the voltage report. The melting-rate model, arc sensing, seam tracking, preview control and adaptive fill follow published practice. Three measured results carry the chapter: a linearization of the model's own arc that predicts the recovery time of self-regulation to within 9 % and orders it across source slopes, the rise with frequency of the current's sensitivity to the distance, which sets the calibration a weaving tracker carries, and the law by which a misconfigured look-ahead becomes a lateral error. The model has 18 states at a 0.5 ms step, with a preview buffer keyed by distance, a bead scored every millimeter, a speed guard that reports an infeasible target, and wire and deposit ledgers that close, all given in enough detail to rebuild it. Eight cases and a graded seam separate a physical loop from a controlled one, an inference from a measurement, and a reading from the place it describes.

Section 12.2 reviews how arc sensing and seam vision arose. Section 12.3 gives the circuit, the wire, the axes, the part and the bead as modeled, and Section 12.4 linearizes the arc's own loop. Section 12.5 describes the controller as built, Section 12.6 reports the eight cases and the graded seam, Section 12.7 names the devices of real welding cells, and Sections 12.8 and 12.9 state the limitations and the conclusions. The exercises of Section 12.10 run the cases in the application.

12.2 Arc Sensing, Seam Vision and the Adaptive Cell

The gas metal arc process feeds its electrode continuously, so the power source and the feeder share a single task, which is to melt the wire exactly as fast as it arrives. The process, its power sources, its monitoring and control, and its automation are treated together in one text [2], and its modeling, sensing and automatic control in another [11]. The electrode extension, the length of wire between the contact tip and the arc, is where the balance is struck. Dynamic models of the extension combine the melting rate, the temperature-dependent resistivity of the wire and the arc voltage, derived from mass continuity and energy relations for use in adaptive control [12]. A later model of the extension driven by the process voltage reproduced the frequency response of the real process, was accurate to about ± 0.6 mm against high-speed video, and was offered for through-the-arc sensing and arc length control [13]. The drops that the arc carries across have been modeled in their own right [14].

Robots came before the sensors that could guide them. A 1983 review of robotic arc welding categorized the sensor requirements by design requirements, joint imperfections, weld shape deviations and process characteristics, and covered the optical and through-the-arc approaches [3]. The history segment below follows the visually guided robot and the arc sensor that came next, and the simulated cell carries the same combination as that robot, a sensor ahead of the torch, a lateral tolerance of half a millimeter and a procedure scheduled on the gap.

Table 12.1 Variables of the Cell by Role, Initialized for Case 1

Variable	Role	Value at reset	Range
CTWD	CV of HC, true	15.00 mm SP	10–25
V_{tip}	CV of AVC (CC only)	27.01 V, SP 27.0	18–36
e_b	CV of TC, true	0 mm	y_{SP} ±3
bead flags per mm	quality, true	sound	5 defects
Δz, Δy	MVs, corrections	0 mm, 0 mm	±10
WFS_{cmd}	MV, feeder	8.0 m/min	1.5–18
V_{oc}	setting, synergic trim	31.8 V	20–42
v_{SP}	setting, travel	10 mm/s	2–20
y_{off}, a, A_b	DVs, seam	0, 0, 0 mm	±5, ±10, 0–4
g_0, profile	DV, root gap	1.0 mm, constant	0–3
z_r, z_{st}	DVs, plate height	0, 3 mm	±5, ±4
k_d	DV, thermal drift	0	0–2
spatter, window	DVs, laser	0 %, clean	0–100
gas, cast δ_c	DVs, shielding, wire	100 %, 0 mm	0–120, ±1.5
i, l_s, ϕ	states, arc	251.9 A, 10.73 mm, 0.5	—
WFS, s, v	states, feed, travel	8.0 m/min, 0, 0	—
y, $\dot{y}$, z, $\dot{z}$	states, axes	0, 0, $z_p + 15$ mm, 0	—
i_f, V_f	states, transmitters	251.9 A, 27.01 V	—
E_p, e_1, e_b	states, pool, bead	0, 0, 0	—
E_{arc}, L_{fed}, L_{mlt}	states, ledgers	0, 0, 0	—

Ranges in the units of the value column, a in mm per m, k_d in mm per 100 kJ, z_r over 100 to 700 mm and z_{st} at 60 mm.

> **HISTORY** **A Robot That Saw the Seam, 1985**
>
> Conventional robots of the early 1980s struggled to weld sheet steel pressings 1 to 2 mm thick. Wear of tools and fixtures, springback and thermal distortion put the torch as much as ±3.0 mm from the seam. In 1985 a model-based visual feedback system guided a metal inert gas welding robot on such parts. It detected gap width, standoff error and lateral error in lap, T and butt joints, corrected the torch to within ±0.5 mm, and used the measured seam width to change the welding parameters during the weld [15]. Sensing through the arc itself grew beside it. An arc sensor for butt joints, published in 1991, read the torch position from the welding current and overcame that signal's fluctuation by fitting it to a curve inversely proportional to the tip-to-workpiece distance [4].

Seam tracking for robotic welding has since been reviewed as a family of sensors and feature-extraction methods for the start and end of the weld, its edges, the joint width and the path [8]. Adaptive welding answers distortion, varying seam preparation and a changing gap during the weld by adjusting the travel speed and the torch orientation, and the wire feed, arc length and process dynamics through synergic settings [9]. Humping and discontinuous beads often limit the travel speed, and with it the productivity, of gas metal arc welding, and dimensionless process maps give the limiting speed as a function of the process parameters [16]. Welding distortion, which moves the joint while it is being welded, has been predicted together with the residual stress for a thin-plate butt-welded joint [17].

12.3 The Cell as Modeled

The simulated cell welds a single-V butt joint 1.2 m long in 6 mm steel plate in one pass, with 1.2 mm wire in a spray arc at the procedure of Section 12.1. The robot is reduced to three Cartesian servo axes, travel along the taught straight path and lateral and height corrections across it, and the part is where it is rather than where it was taught. The arrangement, the instruments and every symbol of the equations are shown in Figure 12.1, and Table 12.1 classifies them as controlled, manipulated, disturbance and state variables.

Eighteen states are integrated, the arc current and electrode extension, the droplet phase, the wire feed behind the feeder, the travel distance and speed, the lateral and height positions and velocities,

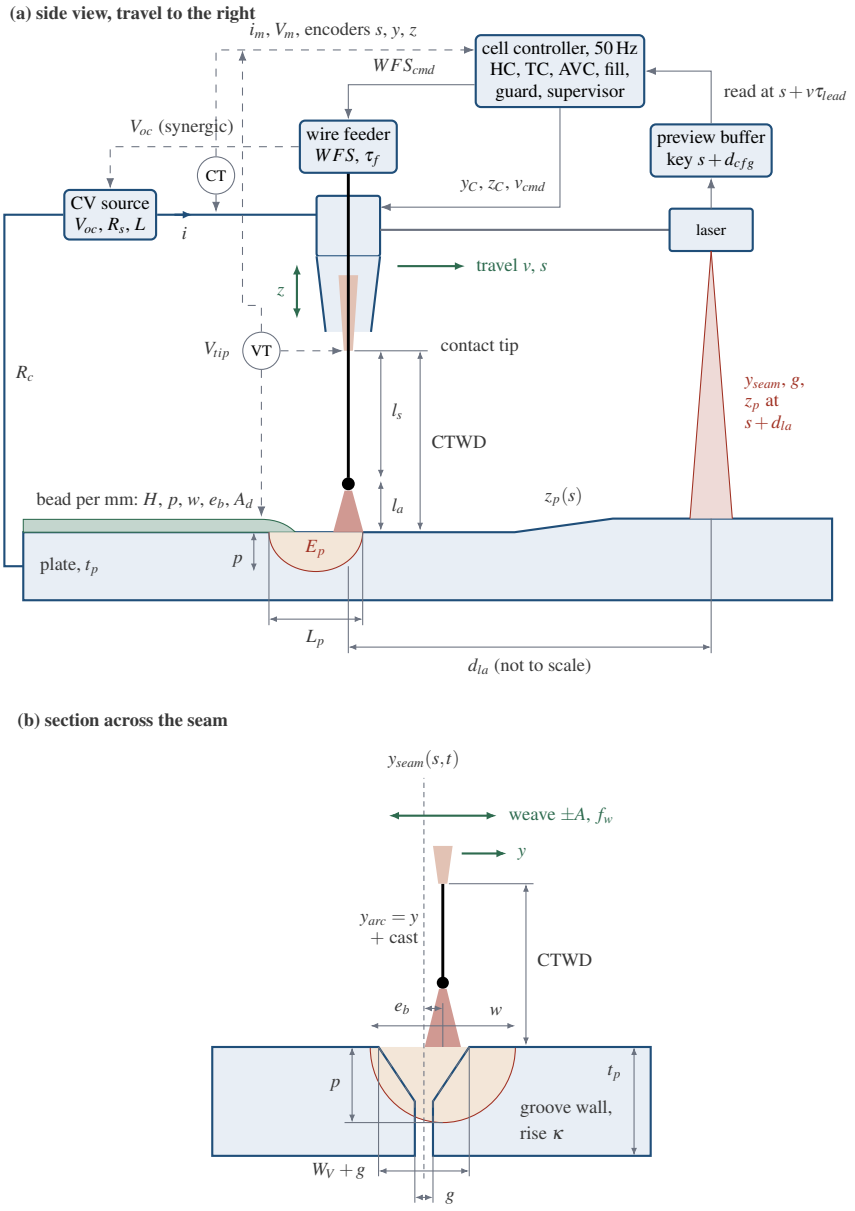

Figure 12.1 The welding cell as modeled. In (a), with travel to the right, the source (V_{oc}, R_s, L) drives i through the cable and tip (R_c), the extension l_s and the arc l_a, which span the CTWD, the feeder pushes wire at WFS behind τ_f, the pool of length L_p holds E_p and penetrates p, and the laser reads y_{seam}, g and z_p a distance d_{la} ahead for the buffer to file under $s + d_{cfg}$. CT and VT are the current and torch-voltage transmitters. In (b), drawn below (a) and across the seam at the arc, the V-groove of top width $W_V + g$ in plate of thickness t_p has walls the arc sees rising by κ per unit offset, and the bead of width w sits at $y_{seam} + e_b$. Distances in (a) are not to scale, and e_b in (b) is exaggerated.

the current and voltage transmitters, the pool energy, two stages of the bead's offset from the seam, and three ledgers of arc energy, wire fed and wire melted, all on one step and in SI units.

12.3.1 The Circuit and the Arc

The source's line, an open-circuit voltage less a slope, drives the current through an inductance, the cable and contact tip, the electrode extension and the arc, as shown in Equation (12.1).

$$\begin{aligned} L\frac{\mathrm{d}i}{\mathrm{d}t} &= V_{oc} - (R_s + R_c)\, i - V_{tip} \\ V_{tip} &= V_{arc} + r_w\, l_s\, i \\ V_{arc} &= V_0 + E_a \max(l_a - \Delta l_d,\ l_{min}) + R_a\, i \end{aligned} \tag{12.1}$$

where i is the arc current, V_{oc} the open-circuit voltage, R_s the source slope, R_c the cable and contact-tip resistance, V_{tip} the voltage a voltmeter reads between the contact tip and the work, r_w the hot resistance of the wire per unit length, V_0 the anode and cathode falls together, E_a the column gradient, l_a the arc length, Δl_d the shortening by the hanging droplet, $l_{min} = 0.1\,\mathrm{mm}$ and R_a the arc's own resistance. With the arc switched off the current decays as $\mathrm{d}i/\mathrm{d}t = -i/1\,\mathrm{ms}$. At the default 31.8 V, a slope of 1.5 V per 100 A and $L = 0.3\,\mathrm{mH}$ the arc settles at 251.9 A and 27.01 V with $l_s = 10.73\,\mathrm{mm}$, and the circuit's own time constant L/R_Σ is 6.3 ms, with $R_\Sigma = R_s + R_c + R_a + r_w l_s = 47.7\,\mathrm{m\Omega}$.

> ✦ DID YOU KNOW? **Two Volts for Every Hundred Amperes**
>
> A constant-voltage welding source is not quite constant: its voltage falls by about 2 V for each 100 A it delivers. The designer picks that slope so that a short arc gets more power and a long one less, which holds the arc length [18]. The source of this chapter falls 1.5 V per 100 A.

12.3.2 The Wire and Its Melting Rate

The feeder pushes the wire forward while the arc and the extension's own resistance melt it back, as shown in Equation (12.2).

$$\begin{aligned} \frac{\mathrm{d}l_s}{\mathrm{d}t} &= WFS - MR, \qquad MR = k_1\, i + k_2\, l_s\, i^2 \\ l_a &= CTWD - l_s \\ \tau_f \frac{\mathrm{d}WFS}{\mathrm{d}t} &= WFS_{cmd} - WFS \end{aligned} \tag{12.2}$$

where MR is the melting rate in length of wire per second while the arc burns, the first term is the arc's heat, the second is the resistive preheating of the extension, and the application attributes the two-term form to Lesnewich. The balance $MR = WFS$ defines the steady arc, which is found at each reset by bisection on the circuit's residual, bracketed between the current at which the extension would fill the whole CTWD and the one that would burn it back to the tip.

12.3.3 Constant Current and the Droplet

The comparison source regulates the current itself inside a compliance voltage, as shown in Equation (12.3).

$$\frac{\mathrm{d}i}{\mathrm{d}t} = \min\left(\frac{I_{set} - i}{\tau_{cc}},\ \frac{V_{max} - R_c\, i - V_{tip}}{L}\right) \tag{12.3}$$

where the setting I_{set}, the regulator time $\tau_{cc} = 2\,\mathrm{ms}$ and the compliance $V_{max} = 50\,\mathrm{V}$ define the source. A droplet grows on the wire tip over one period of a phase ϕ and shortens the arc while it hangs, as

shown in Equation (12.4).

$$\begin{aligned} \frac{\mathrm{d}\phi}{\mathrm{d}t} &= f_d(i), \qquad f_d = \mathrm{clamp}\Big(f_T\Big(\frac{i}{i_T}\Big)^3,\ 20,\ 500\,\mathrm{Hz}\Big) \\ d &= \Big(\frac{6MRA_w}{\pi f_d}\Big)^{1/3}, \qquad \Delta l_d = \kappa_d F\,\phi\,d \end{aligned} \tag{12.4}$$

where A_w is the wire's cross-section, d the droplet diameter, $\kappa_d = 0.5$ the share of it that shortens the arc, and F a size factor. When ϕ reaches 1 the droplet detaches, ϕ drops back by one and F is drawn as $\mathrm{clamp}(1 + 0.15\,n,\ 0.5,\ 1.6)$ with n standard normal. At the default arc the model sheds 225.6 droplets a second, about nine physics steps per droplet, and the ripple they leave is the sawtooth on the scope.

12.3.4 Travel and the Torch Axes

Travel follows its command through a first-order response with an acceleration limit, and the two correction axes are position loops around velocity loops (Section 1.4.3), as shown in Equation (12.5).

$$\begin{aligned} \frac{\mathrm{d}s}{\mathrm{d}t} &= v, \qquad \frac{\mathrm{d}v}{\mathrm{d}t} = \mathrm{sat}_{a_s}\Big(\frac{v_{cmd} - v}{\tau_s}\Big) \\ \ddot{y} &= \mathrm{sat}_{a_y}\big(\omega_{v,y}\big[\mathrm{sat}_{v_y}\big(\omega_{p,y}(y_{ref} - y)\big) - \dot{y}\big]\big) \\ y_{ref} &= y_C + A\sin\big(\varphi_w + 2\pi f_w(t - t_w)\big) \\ \ddot{z} &= \mathrm{sat}_{a_z}\big(\omega_{v,z}\big[\mathrm{sat}_{v_z}\big(\omega_{p,z}(z_C - z)\big) - \dot{z}\big]\big) \end{aligned} \tag{12.5}$$

where $\mathrm{sat}_m(x)$ limits x to $\pm m$, y_C and z_C are the commanded lateral and height positions, A and f_w the weave amplitude and frequency, and φ_w and t_w the phase and time of the last control update, which keep the weave continuous when its frequency changes. The application's own check drives both axes through a 10 mm step and finds that neither exceeds its rate or acceleration limit while the lateral axis reaches its rate limit.

12.3.5 The Part

The seam, the gap and the plate height are functions of the travel distance, and the seam also moves with the heat put into the part, as shown in Equation (12.6).

$$\begin{aligned} y_{seam} &= y_{off} + a\,s + A_b\sin\frac{2\pi s}{\lambda_b} + k_d\,\frac{E_{arc}(t)}{100\,\mathrm{kJ}}\,\frac{s}{L_s} \\ z_p &= z_r\,\rho(s) + z_{st}\,[s \geq 60\,\mathrm{mm}] \\ y_{arc} &= y + \delta_c\,l_s/11\,\mathrm{mm} \\ CTWD &= z - z_p - \kappa\min\big(|y_{arc} - y_{seam}|,\ (g + W_V)/2\big) \end{aligned} \tag{12.6}$$

where $\lambda_b = 300\,\mathrm{mm}$, $L_s = 1.2\,\mathrm{m}$ is the seam length, $\rho(s)$ rises linearly from 0 to 1 between 100 and 700 mm, δ_c is the wire cast that bends the tip off the torch axis, and the root gap $g(s)$ is the nominal g_0 plus one of three piecewise-linear profiles, constant, opening by 1.1 mm between 150 and 750 mm, or varying, at $-0.2\,\mathrm{mm}$ up to 250 mm, $+1.2\,\mathrm{mm}$ from 550 to 700 mm and $+0.2\,\mathrm{mm}$ beyond 950 mm. The part is clamped at the start of the seam and bows as it heats, so its free end moves k_d millimeters for every 100 kJ of arc energy delivered, and every point of the seam moves with it, welded or not. The groove walls rise by κ per unit of lateral offset as the arc sees them, which is what a weave reads, and both κ and the drift law are illustrative.

12.3.6 The Pool, the Bead and Its Record

The pool integrates the arc's power and loses it as it moves on, and the bead's offset is the arc's offset averaged over the pool's passage, as shown in Equation (12.7).

$$\begin{aligned} \frac{\mathrm{d}E_p}{\mathrm{d}t} &= V_{tip}\,i - E_p\,\frac{\max(v, v_c)}{L_p}, & H &= \frac{E_p}{L_p} \\ \frac{\mathrm{d}e_1}{\mathrm{d}t} &= r_b\,(y_{arc} - y_{seam} - e_1), & \frac{\mathrm{d}e_b}{\mathrm{d}t} &= r_b\,(e_1 - e_b) \end{aligned} \tag{12.7}$$

where the pool holds the energy E_p over a length $L_p = 8\,\mathrm{mm}$, $v_c = 4\,\mathrm{mm/s}$ stands for conduction into the plate, H is the heat input per unit length, and $r_b = \max(v, 8\,\mathrm{mm/s})/(L_p/4)$, about $L_p/2v$ over the two stages. A moving pool settles at $H = V_{tip}\,i/v$, the heat input the application's check confirms to within 0.5 %, and a stationary one keeps heating toward $V_{tip}\,i/v_c$, which is why the start dwells and the crater fill is short. The bead follows the correlations of Equation (12.8).

$$\begin{aligned} p &= a_p H^{b_p}\,\frac{1 + \beta_g g}{1 + \beta_A A}, & w &= a_w H^{b_w} + k_w A \\ v_{max} &= v_u \Big(\frac{H}{675\,\mathrm{J/mm}}\Big)^{e_u}, & A_n &= A_V + t_p\,g \end{aligned} \tag{12.8}$$

where p is the penetration, w the width, v_{max} the travel speed above which the bead undercuts and A_n the metal a millimeter of groove needs to be filled flush. A millimeter lacks fusion when $p < p_{min}$ or when $w/2 - |e_b| < (g + W_V)/2 + m_e$, burns through when $p > t_p$, undercuts when $v > v_{max}(H)$, underfills when its deposit $A_d < A_n$, and is porous when the shielding coverage $\chi = f_{gas} \min\big(1, (18\,\mathrm{mm}/CTWD)^{1/2}\big)$ falls below χ_{min}, with f_{gas} the gas flow as a share of the set flow. At the default 681 J/mm and a 1 mm gap the correlations give $p = 4.23\,\mathrm{mm}$, $w = 8.03\,\mathrm{mm}$ and $v_{max} = 13.1\,\mathrm{mm/s}$, the deposit $\eta A_w WFS/v$ is $14.6\,\mathrm{mm}^2$ per millimeter against a need of $12.0\,\mathrm{mm}^2$, and the heat input window between lack of fusion and burn-through runs from 499 to 1123 J/mm. Every coefficient is a teaching value chosen so that the default procedure sits inside its window. None of it is a qualified procedure, and the real limits come from a procedure qualification record under ISO 15614-1 or AWS D1.1 and from the acceptance levels of ISO 5817, which the application names.

Each millimeter is classified when the arc leaves it, from the pool's heat input, the gap at the millimeter's center, the weave amplitude, the bead offset, the travel speed, the metal deposited into that millimeter and the shielding. Deposit is booked as η times the wire melted, spread over the pool length ahead of the arc during the start dwell and behind it during the crater fill. When the arc goes out, at the end of the seam or on a stop, the last pool length is recorded again with a crater depression $A_{cr} = 4\,\mathrm{mm}^2$ removed from it, so a crater that is not filled underfills. An extension shorter than 0.5 mm ends the weld as a burn-back, an arc longer than 12 mm breaks, a wire driven 1 mm into the pool stubs out, and an arc shorter than 0.3 mm after the droplet is flagged as stubbing.

The current and voltage transmitters are first-order front ends, $\tau_m\,\mathrm{d}i_f/\mathrm{d}t = i - i_f$ and $\tau_m\,\mathrm{d}V_f/\mathrm{d}t = V_{tip} - V_f$ with $\tau_m = 5\,\mathrm{ms}$, and the three ledgers integrate $V_{tip}\,i$, WFS and MR, so that the wire balance and the deposit close to 10^{-9} relative in the application's own check. The constants are collected in Table 12.2.

12.3.7 Initialization and Numerics

Every scenario opens with the arc struck on the start of the seam and the pool cold. The steady arc of the scenario's source and feed at the CTWD setpoint sets the current, the extension, the feed and both transmitters, the droplet is at mid-growth, the torch stands at $z_p(0)$ plus the setpoint, and the travel speed and the pool energy are zero. Before the arc strikes, the laser has already scanned the first 40 mm at 0.5 mm spacing, and with laser tracking the torch starts on the seam it found. Each 0.5 ms

Table 12.2 Model Constants and Where They Live. Adjustable Constants Have a Slider Under Show Model, Hidden Ones Have None, and Fixed Ones Exist Only in the Source

Symbol	Quantity	Value	Access
k_1	arc melting coefficient	0.300 mm/(s A)	adjustable
k_2	resistive melting coefficient	8.48×10^{-5} 1/(s A^2)	adjustable
V_0, E_a	arc falls, column gradient	15.0 V, 1.2 V/mm	adjustable
R_a, R_c	arc, cable and contact tip	22 mΩ, 4 mΩ	adjustable
r_w	hot wire per unit length	0.62 mΩ/mm	adjustable
κ_d	droplet shortening share	0.5	adjustable
i_T, f_T, σ_d	spray transition, scatter	220 A, 150 Hz, 0.15	hidden
τ_{cc}, V_{max}	CC regulator, compliance	2 ms, 50 V	hidden
τ_f, τ_s	feeder, travel lags	30 ms, 50 ms	adjustable
a_s	travel acceleration limit	200 mm/s^2	adjustable
$\omega_{p,y}$, $\omega_{v,y}$	lateral axis gains	36, 70 rad/s	adjustable
v_y, a_y	lateral limits	50 mm/s, 1.0 m/s^2	adjustable
$\omega_{p,z}$, $\omega_{v,z}$	height axis gains	25, 50 rad/s	adjustable
v_z, a_z	height limits	40 mm/s, 1.0 m/s^2	adjustable
L_s, t_p	seam length, plate	1.2 m, 6.0 mm	adjustable
A_V, W_V, κ	groove area, width, wall rise	6 mm^2, 4 mm, 0.6	adjustable
d_{la}	laser as mounted	40 mm	adjustable
L_p, v_c	pool length, conduction	8 mm, 4 mm/s	adj., fixed
a_p, b_p	penetration	0.04184, 0.70	adjustable
β_g, β_A	gap and weave on p	0.05, 0.05 mm^{-1}	adjustable
a_w, b_w, k_w	width	0.3079, 0.5, 1.0	adjustable
p_{min}, m_e	fusion, edge margin	3.4 mm, 0.3 mm	adjustable
v_u, e_u	undercut speed	13 mm/s at 675 J/mm, 0.8	adjustable
χ_{min}, η, A_{cr}	shielding, deposition, crater	0.6, 0.97, 4 mm^2	adjustable
σ_L, p_{drop}	laser noise, dropouts	0.05 mm, 1 %	adjustable
σ_I, σ_V	current, voltage noise	1.5 A, 0.15 V	adjustable
τ_m	transmitter front ends	5 ms	hidden
l_{min}, λ_b	arc floor, bend wavelength	0.1 mm, 300 mm	fixed

step of the model is one fourth-order Runge-Kutta step, because the circuit's 6.3 ms and the droplet period of about 4.4 ms are the fastest dynamics, while the controller and the laser run every 20 ms and the historian logs a row every 20 ms. Halving the step moves the peak and settled current and the arc length at 100 ms and at steady state by at most 0.66 %, against the application's 1.5 % limit. The engine, the speed selector and the seeding are those of Section 1.6.1. Here the speed selector ranges from a quarter of real time to twice it, a 1.2 m seam at 10 mm/s takes 121.6 s, at a quarter of real time the scope resolves single droplets, and the droplet sizes, the sensor noise and the laser dropouts come from streams seeded at reset from seed 20261102.

12.3.8 Assumptions

The arc is a lumped melting rate with a linear arc characteristic, whose coefficients the application describes as typical of 1.2 mm steel wire in an argon-rich gas, not fitted to one wire. The extension has one hot resistance per millimeter, where dynamic models of the extension carry a temperature-dependent resistivity [12]. Transfer is spray only, a droplet sawtooth on the arc length rather than a dynamic model of drop detachment [14], so short-circuit and globular transfer are absent and a wire driven into the pool is an event rather than a transfer mode. This breaks down below the 220 A transition and in any procedure that relies on short circuits.

The sources are ideal characteristics, a line with a slope and an inductance and a 2 ms current regulator inside a 50 V compliance, so the waveform control of an inverter source is outside the model, and the synergic line is represented only as a source voltage that follows the scheduled feed. The robot is three Cartesian axes with exact encoders, so arm kinematics, torch angles, path interpolation and the robot's own sensor interface delay are outside the model.

The joint is one single-V butt joint welded in one pass, the groove walls act on the arc through a linear rise, and the distortion is a rotation about the start of the seam in proportion to the arc energy,

whereas the distortion of a real thin-plate butt joint is predicted together with its residual stress [17]. The pool is one lumped energy of fixed length, the start and the crater are approximations over that length, and heat does not carry over from one seam to the next. The laser measures points of the groove with Gaussian noise and random dropouts and does none of the seam feature extraction that the tracking literature reviews [8], so reflections, smoke and tack welds do not exist for it. The bead correlations and defect rules are illustrative throughout.

> **OPEN CHALLENGE** **Narrow Gaps Hide From Both Sensors**
>
> A narrow joint defeats both sensors of this chapter. Arc sensors have been studied since the 1980s and sold for T and V grooves, but in narrow gap welding their arc stability and sensing accuracy are not satisfactory [19]. A laser stripe finds little to bend at a gap under 0.2 mm, and methods that add lighting or images of the joint ask for strict illumination of the surface [20]. Progress would be one sensor that finds the three-dimensional position and the surface of a narrow joint on any finish, near the arc.

12.4 The Arc's Own Loop

The circuit's time constant, 6.3 ms, is short against the recovery of the arc, so the self-regulation can be read from the model with the inductance and the droplet set aside. The current is then algebraic, and linearizing Equations (12.1) and (12.2) about a steady arc gives a first-order loop in the extension, as shown in Equation (12.9).

$$\begin{aligned} \frac{\mathrm{d}\,\delta l_s}{\mathrm{d}t} &= -\lambda\,\delta l_s + \mu\,\delta CTWD \\ \lambda &= \left(k_1 + 2k_2\,l_s\,i\right)\frac{E_a - r_w\,i}{R_\Sigma} + k_2\,i^2 \\ \mu &= \left(k_1 + 2k_2\,l_s\,i\right)\frac{E_a}{R_\Sigma} \end{aligned} \tag{12.9}$$

where the first term of λ is the arc's share of the loop, a longer extension shortening the arc and raising the current by $(E_a - r_w i)/R_\Sigma$ per unit length, and the second is the extension's own resistive preheating. At the default arc the two shares are 16.6 and 5.38 s^{-1}, so $\lambda = 22.0\,\mathrm{s}^{-1}$ and the time constant of self-regulation is $\tau_{sr} = 1/\lambda = 45.5$ ms. A first-order loop reaches 90 % in $2.303\,\tau_{sr} = 104.7$ ms, and the full nonlinear model restores 90 % of the 3 mm step of case 1 in 97.0 ms, as shown in Figure 12.2. The droplet sizes are random, and eight other seeds give 96.5 to 100 ms. The source slope enters through R_Σ. With no slope and 28.0 V open circuit, which keeps the same weld, the linearization gives 77.7 ms against a measured 66.5 ms, and at 3 V per 100 A and 35.6 V it gives 127.8 ms against 120.0 ms. One likely reason the 3 mm step recovers faster than the small-signal estimate in all three is that the melting rate grows with the square of the current, which overshoots to 336.1, 316.5 and 304.0 A.

> **TRY IT IN THE APP** **1 · The Arc Regulates Itself**
>
> Load the case and press Run on the Arc view. About 7 s in, the torch meets the 3 mm plate step. Watch the current on the scope jump toward 316 A and decay while the side view shows the arc growing back, with no loop acting in MANUAL. Press Reset, set Power source to Constant current on the Power source card and run again: the arc stays short.

The same linearization gives the current's response to the distance, as shown in Equation (12.10).

$$G(j\omega) = \frac{\delta i}{\delta CTWD} = -\frac{E_a}{R_\Sigma}\,\frac{j\omega + k_2\,i^2}{j\omega + \lambda} \tag{12.10}$$

A change too fast for the extension to follow meets the full $E_a/R_\Sigma = 25.2$ A per millimeter, while at steady state the extension burns back and leaves 6.16 A per millimeter, close to the 6.12 A per

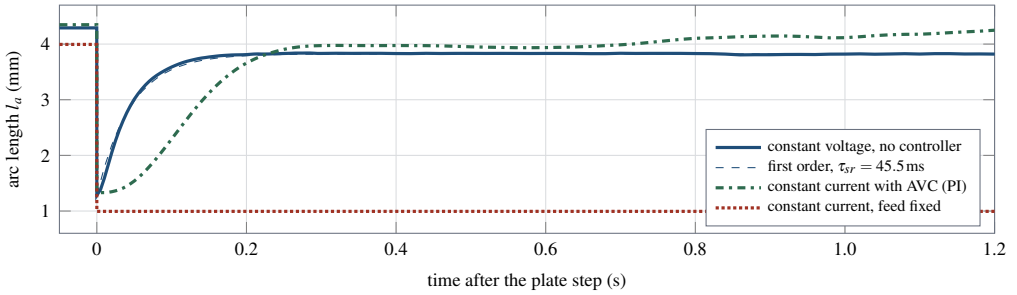

Figure 12.2 Arc length after the 3 mm plate step of case 1 at 6.95 s, for three sources on the same procedure. The constant-voltage arc (thick solid) regulates itself with no controller and follows the first-order response of Equation (12.9) (thin dashed) closely. The constant-current arc with a fixed feed (thick dotted) stays 3 mm short, and the arc-voltage-controlled feed (thick dash-dotted) restores it through a PI loop.

millimeter that the controller computes by bisection with the droplet at mid-growth. For the 3 mm step of case 1 the steady solution gives 251.9 A at 15 mm and 272.3 A at 12 mm, and the engine settles 20.5 A higher. Between the two limits lies the weave. At 3 Hz the magnitude of Equation (12.10) is 17.0 A per millimeter, 2.77 times its steady value, and a weaving tracker that divides by the steady sensitivity has to carry a calibration of that size. The application carries 2.5, measured at its default 1.5 mm and 3 Hz weave. At 2 and 4 Hz Equation (12.10) changes by factors of 0.80 and 1.14 relative to 3 Hz, and the calibrated tracker reads a known 0.3 mm offset as 0.16 to 0.23 mm at 2 Hz and 0.38 to 0.40 mm at 4 Hz, larger departures than the linear model because the lateral axis and the averaging near each extreme likely add their own dependence on frequency.

A constant-current source removes the loop. The melting balance at a fixed 250 A fixes the extension at 11.01 mm whatever the distance, so the arc takes the whole step, from 3.99 to 0.99 mm, and the torch voltage falls from 26.67 to 23.07 V. The arc-voltage-controlled feed (AVC) closes the loop around the voltage instead, on a plant that a manual feed step measures at −3.30 V per m/min with 63 % of the response in 227 ms. With the 1 s integral time the loop recovers 85 % of the length it holds 1.5 to 2.3 s after the step by 0.3 s and creeps the rest, reaching 90 % of it in 834.5 ms with the feed at 6.90 m/min. The creep is slow enough for the noise to move that time, which eight other seeds put between 858 and 996 ms. A counterfactual run without the integral term, which the 5 s upper end of the integral-time slider does not reach, recovers 70 % of the same length by 0.3 s and then holds the arc near 3.2 mm, 1.2 mm shorter. A larger inductance slows the current rather than the arc. At 1 mH the current peaks 39.0 ms after the step at 305.1 A, against 17.5 ms and 316.5 A at 0.3 mH, and the arc is 90 % restored in 93.5 ms.

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

12.5 Implemented Controller

The cell controller executes every $\Delta t_c = 20$ ms, forty physics steps, in the order of Section 1.6.1, with the droplets, events and bead record booked after each integration. No loop reads the part or the arc directly.

12.5.1 Instruments

The encoders report s, y and z exactly, the feeder tachometer reports WFS, and the current and torch voltage are sampled from the transmitter states with Gaussian noise of 1.5 A and 0.15 V. The laser returns the groove's center, root gap and height at $s + d_{la}$ every 20 ms, each with 0.05 mm of noise, and loses 1 % of its readings plus the spatter share, or all of them while its window is covered. The gas flow switch carries 1 % noise.

Table 12.3 Loops, Directions and Shipped Settings. All Execute Every 20 ms

Loop and law	Measurement	Actuator	Setting
HC, PI, $\gamma = +1$	$\widehat{CTWD}$ or laser	Δz	$K_c = 0.5$, $\tau_I = 0.25$ s
AVC, PI, $\gamma = -1$	V_m, 40 ms filter	*WFS*	$K_c = 0.5$, $\tau_I = 1$ s
TC, laser, Eq. (12.12)	$\hat{y}_{seam}(s + v\tau_{lead})$	Δy	$K_t = 0.5$
TC, weave, Eq. (12.14)	$\hat{\Delta}$ per cycle	Δy	$K_t = 0.5$
trim, integral	$\hat{\Delta}$ per cycle	y_{trim}	$K_{trim} = 0.2$
fill, Eq. (12.15)	$\hat{g}$ at the torch	*WFS* or v	$A_r = 2\,\text{mm}^2$

Δz, Δy within ± 10 mm, *WFS* 1.5 to 18 m/min. Gains in mm per mm (HC), m/min per V (AVC) and per update (K_t, K_{trim}).

12.5.2 The PI Algorithm

The height loop (HC) and the AVC loop call the framework's one velocity-form routine, Equations (1.19) to (1.22) of Section 1.6.2, executed every Δt_c. This chapter writes the controller action as $\gamma = \pm 1$, because a is the seam angle here. The height loop has $\gamma = +1$ and is reverse acting, since raising the torch lengthens the CTWD. The AVC loop has $\gamma = -1$ and is direct acting, since more feed shortens the arc and lowers the torch voltage. In both loops the output bias is zero, the tracking time T_t equals τ_I, and $\tau_D = 0$. The running sum is never clamped, and back-calculation corrects only the integral. In MANUAL, and whenever a loop is not live, it tracks so that its output follows the operator's station, its own last output or, for the AVC on a constant-voltage source, the scheduled feed, and a return to automatic is bumpless. The loops and their shipped settings are collected in Table 12.3.

12.5.3 Height Through the Arc

The height loop holds the CTWD at its 15 mm setpoint by adding Δz to the taught height, reading the CTWD from the laser's buffered groove height at the torch or through the arc. The through-the-arc estimate inverts the steady melting balance of Equation (12.2) and the arc voltage of Equation (12.1) from the measured current, voltage and feed, as shown in Equation (12.11).

$$\begin{aligned} \hat{I}_s &= \frac{WFS - k_1 \bar{\imath}}{k_2 \bar{\imath}^2} \\ \widehat{CTWD} &= \hat{l}_s + \frac{\bar{V} - V_0 - R_a \bar{\imath} - r_w \hat{l}_s \bar{\imath}}{E_a} + \frac{\kappa_d d}{2} \end{aligned} \tag{12.11}$$

where $\bar{\imath}$ and $\bar{V}$ are the samples averaged by a 0.15 s first-order filter that removes the droplet ripple, and d is the droplet diameter at $\bar{\imath}$ and the feed. The coefficients are the controller's own copy of the model, taken at reset, so a parameter changed during a run biases the estimate as a wire change would. On the flat start of case 2, between 3 and 10 s with the height loop closed through it, the estimate reads the CTWD 0.042 mm short with a scatter of 0.065 mm (1σ), while eight other seeds put that bias between 0.005 and 0.044 mm. It is right only once the arc has settled. After the 3 mm step of case 1 it reads 12.08 mm at the scan 90 ms after the step and dips to 11.87 mm, and its scatter last carries it more than 0.15 mm from 12 at 1.45 s, a time that eight other seeds spread from 0.07 to 3.0 s. The loop holds its output while the torch fills the crater, while the laser it reads is blind, and when the arc is on a constant-current source, whose current carries no height.

12.5.4 The Preview Buffer and Laser Tracking

The laser line falls $d_{la} = 40$ mm ahead of the arc, but the buffer knows nothing about the bracket. The controller files each reading under a key made of the torch's encoder position when the reading was taken and the look-ahead it believes the sensor sits at, and it asks for the reading at the torch's own

position, as shown in Equation (12.12).

$$\begin{aligned} k_j &= s_j + d_{cfg}, \qquad \hat{y}_{seam} = \mathscr{B}\big(s + v\,\tau_{lead}\big) \\ \Delta y_k &= \Delta y_{k-1} + K_t\big(\hat{y}_{seam} + y_{SP} + y_{trim} - \Delta y_{k-1}\big) \end{aligned} \tag{12.12}$$

where s_j is the torch position at reading j, d_{cfg} the configured look-ahead, $\mathscr{B}(q)$ the buffer's value at key q, $\tau_{lead} = 30\,\text{ms}$ a lag compensation for the lateral axis, y_{SP} an offset setpoint, and y_{trim} the through-the-arc trim when it is on. The buffer holds 6000 readings in a ring. It interpolates between valid readings at most 1.5 mm apart, holds the last valid reading for up to 10 mm beyond it, and then reports the laser blind. Section 1.4.6 corrects a report for its age in time. Here a reading belongs to a place along the seam instead, and its age in time does not matter. The tracking law is a first-order approach with an equivalent time constant of 28.9 ms at $K_t = 0.5$, and the gap and the plate height are read at the torch itself, without the lead.

A wrong d_{cfg} files every reading at the wrong place, and nothing in the readings can reveal it. The torch at s receives the reading taken when it stood at $s - d_{cfg}$, which describes the seam at $s + \Delta d$ with $\Delta d = d_{la} - d_{cfg}$, so on a bend of amplitude A_b the lateral error is $y_{seam}(s + \Delta d) - y_{seam}(s)$, whose largest value is $2A_b|\sin(\pi\Delta d/\lambda_b)|$, as shown in Equation (12.13).

$$e_{max} = 2A_b\left|\sin\frac{\pi\,(d_{la} - d_{cfg})}{\lambda_b}\right| \tag{12.13}$$

The error peaks where the seam crosses the taught path at its steepest and vanishes at the crests, where the seam runs parallel to the path, and a bead that must reach both groove edges loses fusion wherever $|e_b|$ exceeds $w/2 - (g + W_V)/2 - m_e$, 1.22 mm at the default procedure.

> **TRY IT IN THE APP** **4 · The Look-Ahead Set Wrong**
>
> Run the case on the Seam view. The amber readings end 20 mm short of the laser line, and the bead leaves the groove where the seam crosses the taught path. Press Reset, set Look-ahead distance (configured) to 40 mm on the TC card and run again. The worst offset falls from 1.451 to 0.067 mm.

12.5.5 Tracking Through the Weave

With a weave, the arc approaches one groove wall at each extreme. The controller averages the measured current while the encoder puts the torch more than $0.7A$ from the weave center on either side and, once per weave cycle, converts the difference into the torch's offset from the groove center, as shown in Equation (12.14).

$$\hat{\Delta} = \frac{\bar{\imath}_R - \bar{\imath}_L}{2\,\kappa G_0\,c}, \qquad \Delta y \leftarrow \Delta y - K_t\big(\hat{\Delta} - y_{SP}\big) \tag{12.14}$$

where $\bar{\imath}_R$ and $\bar{\imath}_L$ are the mean currents near the right and left extremes, G_0 is the steady sensitivity of the current to the CTWD that the controller recomputes from its model every cycle, and $c = 2.5$ is the calibration of Section 12.4. A torch to the right of the groove center sees the right wall sooner, a shorter distance and a higher current, and is moved left. In the laser with arc trim mode the laser law of Equation (12.12) runs every 20 ms and the trim integrates $y_{trim} \leftarrow y_{trim} - K_{trim}\hat{\Delta}$ once per cycle. When the laser is blind and the torch is weaving, the fallback, if it is on, switches the lateral correction to Equation (12.14). Without a weave there is nothing to fall back on, and the correction holds.

12.5.6 Adaptive Fill and the Synergic Line

The fill law schedules the feed or the speed on the gap the laser measured at the torch, as shown in Equation (12.15).

$$\begin{aligned} A_t &= A_V + t_p\,\hat{g} + A_r \\ \text{by feed:}\quad WFS &= \max\left(WFS_{SP},\ A_t\,v/(\eta A_w)\right) \\ \text{by speed:}\quad v &= \min\left(v_{SP},\ \eta A_w WFS_{SP}/A_t\right) \\ V_{oc,cmd} &= V_{oc} + k_{syn}\left(WFS - WFS_{SP}\right) \end{aligned} \tag{12.15}$$

where A_r is the reinforcement target above the groove, $k_{syn} = 1.2\,\mathrm{V}$ per m/min is the synergic slope, and the feed is clamped to 1.5 to 18 m/min. By feed, wire is added only when the gap needs more than the procedure deposits, so a narrow gap keeps the procedure's heat. By speed, the torch slows down at the procedure's feed. Both add heat to add metal, which is the limit case 5 finds. The fill law is feedforward from a measured disturbance, like the coagulant dose computed from the raw water turbidity in Volume 3, Chapter 6, and like any feedforward it is bounded by what the plant can take.

12.5.7 The Window Guard and the Supervisor

The window guard runs every tenth control step, or whenever the scheduled speed changes, and asks the bead correlations whether the asked speed is sound at the measured gap with a margin of 0.05 on every rule, using the steady arc at the CTWD setpoint and the fill law's feed. If it is not, the guard searches down in steps of 0.5 mm/s to 3 mm/s, bisects ten times for the fastest sound speed and holds the travel there. It also tests 33 feeds from 1.5 to 18 m/min, each at its synergic voltage, and reports the asked speed as infeasible when none of them welds it.

The supervisor starts the arc with the torch still for a 0.9 s dwell, welds, brakes into the end of the seam on $v_{cmd} = \min\left(v_{sched}, \sqrt{2 \cdot 0.9\,a_s\,(L_s - s - 0.2\,\mathrm{mm})}\right)$, with v_{sched} the speed after the fill law and the guard, fills the crater for 0.5 s at half the procedure's feed, no more than 5 m/min, and switches the arc off. It marks the laser blind when the buffer has no usable reading, and it stops the weld when the flow switch reads below 70 % for 0.3 s with the gas stop on.

12.5.8 Plant, Measured and Internal Models

The equations of Section 12.3 are the simulation model. The application carries no identified control model, and manual steps on the flat plate of case 2 measure what the loops see, each read as the stepped run less an otherwise identical unstepped one. A 1 mm step of the height station moves the true CTWD by 1.000 mm, 63 % of the way 51 ms after the scan that applies it, and the through-the-arc estimate by 0.998 mm in 101 ms, while the filtered current falls by 5.9 A. The AVC plant is the −3.30 V per m/min of Section 12.4. The controller's internal models are the inversion of Equation (12.11) and its sensitivity G_0, the calibration c and the wall rise κ of Equation (12.14), the configured look-ahead d_{cfg}, the groove need and deposit of Equation (12.15), and the bead correlations and steady arc inside the guard. None of them models the part's drift.

12.6 Measured Cases

Table 12.4 names the eight cases and the graded seam. Each case starts from the procedure of Section 12.1, with constant voltage, no height control, no tracking, weave, fill or guard, the start dwell, the crater fill and the gas stop, and differs from it only in what its row lists, except that cases 4 to 8 track the seam with the laser and cases 6 and 7 also weave 1.5 mm. Every number quoted below was measured by driving the engine from a script at the default seed, with each change made after the case is selected and before the run, as a learner makes it, which gives the same run as the change carried through a Reset (Section 1.6.4). Times are in simulated seconds, and the bead record of each run is summarized in Table 12.5. Positions along the seam are the lower edge of the first millimeter concerned. Worst offsets are partly set by the laser and sensor noise, and where a comparison is within its spread over eight other seeds the text says so. Figure 12.3 draws the bead record as bars.

Table 12.4 Scenarios, With the Disturbance Each Imposes

#	Scenario and disturbance	Teaching point
1	3 mm plate step at 60 mm, MANUAL	a loop with no controller
2	plate rises 4 mm, 100 to 700 mm	the current as a sensor
3	seam 2.5 mm off the taught path	preview places the torch
4	look-ahead set to 20 mm, bend 3.5 mm	a reading used early
5	gap opens 1.0 to 2.1 mm	adaptive fill, its limit
6	drift 1.0 mm per 100 kJ, weave	a preview 4 s old
7	laser blind 12 to 30 s, bend 2.5 mm	hold, fall back, stop
8	travel 10 to 14 mm/s at 20 s	the undercut boundary
9	gap 0.8–2.2 mm, offset 1 mm, 12.5 mm/s	graded seam

Table 12.5 Bead Record of Each Case and Variant Over the Whole 1200 mm Seam. LoF Is Lack of Fusion, BT Burn-Through, UC Undercut, UF Underfill and P Porosity

Case and configuration	Defect (mm)	Kind	Worst $\lvert e_b \rvert$ (mm)
2 height off, through the arc, laser	0	—	0.000
3 tracking off	1200	LoF	2.500
3 laser	0	—	0.041
3 through the arc, 1.5 mm weave	10	LoF	2.541
5 fill off	810	UF	0.041
5 fill by feed, by speed	0	—	0.051
5 by feed, gap opening to 2.5 mm	536	BT	0.041
6 laser	0	—	0.380
6 tracking off	610	LoF	8.429
6 laser with trim, through the arc	0	—	0.154, 0.233
7 hold the last reading	63	LoF	2.605
7 fall back to the arc	0	—	0.527
7 fall back, no weave	103	LoF	2.515
7 gas loss at 45 s, stop off	823	LoF, P	2.605
8 14 mm/s as shipped	1008	LoF, UC, UF	0.028
8 14.5 m/min at 39.6 V, or the guard	0	—	0.028, 0.025

12.6.1 The Arc Regulates Itself

Case 1 welds in MANUAL at 8 m/min and 31.8 V. Over the second before 6 s the arc carries 252.1 A at 27.05 V with a 4.29 mm arc, and over 0.3 s the scope shows the current between 250.7 and 254.7 A and the torch voltage between 26.56 and 27.41 V. The torch meets the step at 6.95 s. The arc is suddenly 1.28 mm long, the current peaks at 316.5 A after 17.5 ms, and the arc is back to 3.01 mm after 50 ms, 3.59 mm after 100 ms and 3.82 mm once settled, 90 % restored at 97.0 ms, with the current settled at 272.6 A and the torch voltage at 26.66 V. The run 100 ms after the step is shown in Figure 12.4. The laser met the same step 4.0 s earlier, 40 mm at 10 mm/s, and the buffered height at the torch changes on the torch's arrival, as the application's own check confirms.

12.6.2 The Current as a Height Sensor

Case 2 raises the plate 4 mm between 100 and 700 mm. With the height loop off the CTWD falls from 15.00 to 11.00 mm, and the filtered current rises from 251.2 A at 5 s to 274.9 A at 60 s, where the CTWD is 11.73 mm, and to 280.5 A on the raised plate, about 7 A per millimeter. The pool's heat input rises from 681 to 743 J/mm and the penetration from 4.23 to 4.49 mm. The bead stays sound over all 1200 mm, but it is welded 4 mm closer than its procedure. Taken through the arc, the height loop reads that current as a distance and raises the torch with the plate, holding the CTWD between

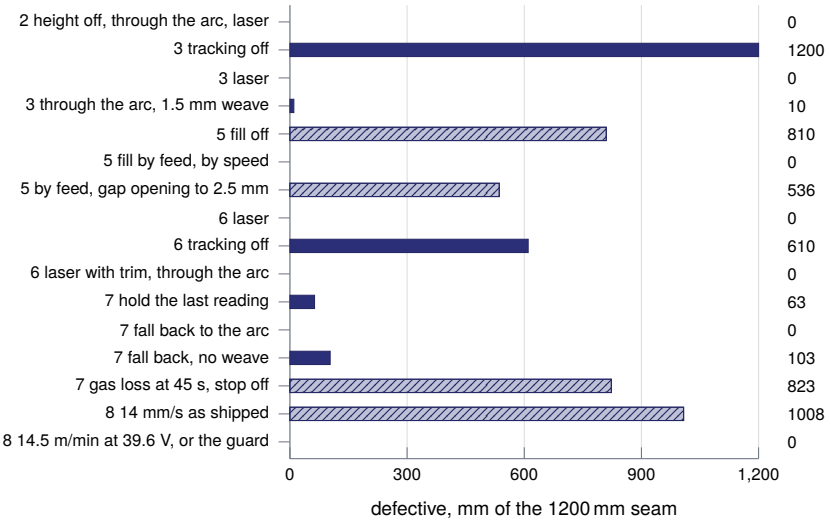

Figure 12.3 The bead record of Table 12.5 drawn as bars, labeled by case and configuration: the millimeters of the 1200 mm seam recorded defective. Solid bars lack fusion only, hatched bars carry underfill, burn-through, porosity or undercut, and a row with no bar welds the whole seam sound. Notice that every case has a configuration with no bar, and that the hatched bar of case 5 is the fill law itself, pushed to a wider gap.

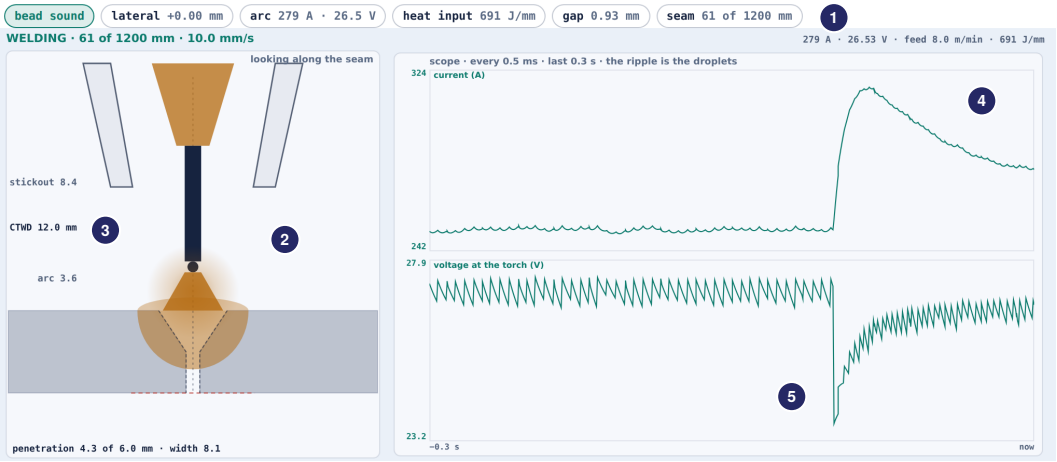

Figure 12.4 Scenario 1, The arc regulates itself, in MANUAL on the Arc view, cropped to the plant view, at $t = 7.05$ s, 0.10 s after the torch met the 3 mm plate step, with no control changed. (1) The badges read a sound bead, a lateral offset of +0.00 mm, 279 A at 26.5 V, 691 J/mm, a laser gap of 0.93 mm and 61 mm of 1200 mm welded at 10.0 mm/s. (2) The side view looking along the seam shows a 3.6 mm arc over a pool 8.1 mm wide that penetrates 4.3 mm of the 6.0 mm plate, (3) with a CTWD of 12.0 mm and a stickout of 8.4 mm. (4) The 0.5 ms scope over the last 0.3 s shows the current leaving its droplet ripple near 252 A toward 316 A and decaying, (5) while the torch voltage drops by more than 3 V and recovers with the droplet sawtooth.

14.87 and 15.13 mm and the current between 250.6 and 253.3 A. The laser's groove height holds it between 14.93 and 15.04 mm. Closed on the 3 mm step of case 1 instead, the through-the-arc loop

brings the CTWD back within 0.5 mm of its setpoint for good after 1.01 s, ten times the 97.0 ms in which the arc restored its own length.

The inference is only as good as the controller's copy of the model. Raising k_2 by 10 %, to 9.30×10^{-5} 1/(s A^2), at 5 s of case 2 with the height read through the arc lowers the true CTWD from 15.05 to 14.22 mm (means over 3 to 5 and 7 to 9 s). The plant now melts the wire faster than the controller's copy predicts, so the inversion of Equation (12.11) reads the CTWD long and the loop lowers the torch. The bias stays until a Reset, which keeps the new coefficient and recopies it into the controller's model, and the CTWD then holds 15.04 mm.

12.6.3 Preview and the Misfiled Reading

Case 3 sets the part 2.5 mm off the taught path. Without tracking, the 8.1 mm bead ends 4.1 mm to the seam's side of the torch while the groove's far edge lies at 5.0 mm, so every millimeter lacks fusion, and the arc on the wall top draws 260.3 to 263.0 A. With the laser, the seam found by the scan before the arc struck places the torch on it and the buffered readings keep it there, with a worst bead offset of 0.041 mm (0.027 to 0.121 mm over eight other seeds). Tracking through the arc with a 1.5 mm weave has to find the seam from its own current, and 10 mm at the start lack fusion while it pulls in from 2.5 mm. With the laser but no start dwell, the first 10 mm lack fusion, nine of them also undercut, the first millimeter frozen from a pool at 117 J/mm.

Case 4 bends the seam 3.5 mm to either side every 300 mm and configures the look-ahead at 20 mm. The torch follows the seam 20 mm further on, the worst offset is 1.451 mm against the 1.455 mm of Equation (12.13), and 372 mm lack fusion in runs of about 48 mm every 150 mm, centered where the seam crosses the path at its steepest. The Seam view shows the fault before any bead does, because the readings, drawn where they were filed, end 20 mm short of the laser line. Swept over configured look-aheads of 30, 50 and 60 mm, the worst offset is 0.728, 0.823 and 1.488 mm against 0.732, 0.732 and 1.455 mm from Equation (12.13), and only 60 mm loses fusion as well, over 397 mm. Filing 20 mm too far ahead errs slightly more than filing 20 mm too near on every seed tried, likely because a reading used late adds to the residual lag of the axis and the pool while one used early offsets part of it, but the 30 and 50 mm pair lies within the noise, which spreads the 30 mm result from 0.721 to 0.821 mm over eight other seeds. With the look-ahead at 40 mm the worst offset is 0.067 mm, and removing the 30 ms lag compensation moves it only to 0.073 mm, inside the 0.042 to 0.122 mm that other seeds give with it. A slow through-the-arc trim does not rescue a wrong look-ahead on this seam, and with a 1.5 mm weave the worst offset becomes 1.501 mm, although the wider bead then covers both edges.

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

12.6.4 Adaptive Fill and Its Limit

Case 5 opens the gap from 1.0 to 2.1 mm between 150 and 750 mm. At 8 m/min and 10 mm/s the procedure deposits 14.6 mm^2 per millimeter, enough for a gap of 1.44 mm, and without the fill law the bead underfills from 389 mm to the end, 810 mm in all. Filling by feed raises the feed with the measured gap, to 11.18 m/min at 900 mm with the source at 35.65 V, and every millimeter is sound, with the pool's heat input peaking at 1029 J/mm and the wire used rising from 16.18 to 20.11 m. Filling by speed also welds the whole seam sound, slowing to 7.00 mm/s at 900 mm, but the seam takes 150.9 s instead of 121.6 s. Neither law has room to spare. With the nominal gap at 1.4 mm, so that it opens to 2.5 mm, filling by feed burns through from 663 mm, 536 mm in all, and filling by speed from 675 mm. The extra metal brings more heat than the open root can take.

> ✦ DID YOU KNOW? **More Metal, Then a Hole**
>
> From the simulator: in case 5, filling by feed welds the whole seam sound as the gap opens to 2.1 mm. Opened to 2.5 mm, the same law burns through 536 mm of the seam, while no fill at all leaves 810 mm underfilled (Table 12.5). Extra wire brings extra heat.

12.6.5 A Part That Moves, and a Blind Laser

Case 6 lets the part bow by 1.0 mm per 100 kJ while the torch weaves 1.5 mm, and the 836 kJ of the seam move the far end 8.36 mm. Without tracking the bead ends 8.43 mm off the seam and 610 mm lack fusion from 577 mm. With the laser, each reading is 4 s old when the torch arrives, so the bead drifts up to 0.380 mm off, sound but most of a 0.5 mm tolerance. The through-the-arc trim reads the groove at the torch and brings the worst offset to 0.154 mm, and the weave alone gives 0.233 mm. The preview is early but old, the arc current fresh but late, and the two together beat either alone on every seed tried. The wafer stage of Volume 14, Chapter 4 pairs a sensor ahead of its slit with a trailing trim for the same reason.

🕮 **Pause.** Try Exercise 12.3 before reading on.

Case 7 covers the laser's window from 12 to 30 s on a seam bending 2.5 mm. The buffer's 40 mm of readings and the 10 mm hold keep the torch sighted until the scan at 16.96 s, and the first new reading reaches it at 33.98 s. Holding the last reading keeps the torch straight while the seam curves away, and 63 mm lack fusion between 195 and 260 mm (57 to 64 mm over eight other seeds) with the bead 2.60 mm off. Falling back to the weave keeps every millimeter sound at a worst offset of 0.527 mm, and without a weave the fallback has nothing to read and 103 mm lack fusion. A shielding loss fired at 45 s trips the flow switch at the scan 0.28 s later and the supervisor stops the weld at 443 mm, leaving 8 mm of porous, unfilled crater, while with the stop off 760 mm are porous.

12.6.6 The Speed Window and an Honest No

Case 8 raises the travel speed from 10 to 14 mm/s at 20 s, with the feed and the voltage unchanged. The heat input falls to 486 J/mm, below the 499 J/mm the fusion rule needs at this gap, the undercut speed falls to 10.0 mm/s, and the deposit falls to 10.4 mm^2 against a need of 12, so 1008 mm are defective from 191 mm, all of them underfilled, 1001 mm undercut and 980 mm short of fusion. Raising the feed to 14.5 m/min at 39.6 V at the speed change welds every millimeter sound at 14 mm/s, in 92.7 s and with 20.17 m of wire. The window guard instead holds the speed at 11.28 mm/s, where undercut is the binding rule at its 5 % margin, and the seam takes 110.1 s. With adaptive fill by feed as well, the guard holds the speed between 11.6 and 12.4 mm/s, a cap that follows the noise on the gap it measures, and the seam takes 105.0 s. Asked for 18 mm/s, the guard finds no feed between 1.5 and 18 m/min on the synergic line that welds it, reports the speed infeasible and holds 11.28 mm/s. At a 1.0 mm gap the fastest sound speed inside the feeder's range is 17.1 mm/s, at 18.0 m/min and 465 A.

12.6.7 The Graded Seam

The challenge welds 1.2 m at 12.5 mm/s over a gap varying between 0.8 and 2.2 mm, on a part set 1 mm off its path whose free end drifts 0.22 mm per 100 kJ, and asks for every millimeter in specification, the bead within ± 0.5 mm of the seam and an average of at least 12 mm/s over the arc-on time, dwell and crater included, as compared in Table 12.6.

As shipped, with no tracking, no fill and the default procedure, 1198 mm of the 1200 mm are defective and the bead runs up to 2.48 mm off the seam. The suggested configuration welds every millimeter sound with a worst offset of 0.133 mm, 0.37 mm inside the band (0.072 to 0.121 mm over eight other seeds), while the far end drifts 2.04 mm, the feed rises to 14.35 m/min at the widest gap and the pool's heat input peaks at 962 J/mm, for 925 kJ of arc energy. Each element pays for itself. Without the fill law 503 mm underfill where the gap is wide, and with the default 8 m/min at 31.8 V the narrow stretch undercuts over 331 mm, because at 12.5 mm/s that procedure's 544 J/mm puts the undercut speed at 10.9 mm/s. Without tracking, the offset and the drift take the bead off the groove

Table 12.6 The Graded Seam, Shipped, Suggested and With One Element of the Suggestion Removed. Every Run Averages 12.30 mm/s

Configuration	Defect (mm)	Kind	Worst (mm)	Wire (m)
shipped	1198	LoF, UC, UF	2.480	12.98
suggested[a]	0	—	0.133	18.08
suggested, fill off	503	UF	0.133	16.22
suggested, 8 m/min at 31.8 V	331	UC	0.131	17.18
suggested, tracking off	867	LoF	3.071	18.08

[a] laser tracking, fill by feed, 10 m/min at 34.2 V, 12.5 mm/s, $A_r = 1\,\text{mm}^2$

edge from 333 mm. Every run finishes at 12.30 mm/s, so quality and precision alone decide the target, which the application judges on the step that fills the crater, at 97.58 s.

12.7 Technology in Practice

The cell of Section 12.3 reads a laser, the current, the voltage, the feed and the encoders, and it moves two axes, the feed and the source voltage. Figure 12.5 draws its loops with real devices in the place of the schematic's symbols.

What Measures It

Through-the-arc sensing needs only current and voltage sensors, with the wire feed speed sometimes added, and its sensors sit well away from the weld. In one laboratory cell a Hall-effect transducer in the welding circuit measured up to 600 A within 1 %. The arc voltage cannot be read alone: across the source it includes the drops along the cables, at the contact tube and along the extension, and moving the leads to the torch and the work removes only the cable drop [21]. The transmitters of Section 12.5, with 5 ms front ends and noise of 1.5 A and 0.15 V, stand for such sensors. A laser triangulation unit on the torch gives the distance to the work and the joint as well, at a cost in price, size and intrusiveness [21]. An arc sensor, by contrast, coincides with the arc and resists the arc's light [19]. The laser of Section 12.5, with 0.05 mm of noise and 1 % of readings lost, stands for it.

> **TECHNOLOGY** **Sensor: Laser Stripe Seam Sensor**
>
> A laser seam sensor projects a stripe of laser light across the joint, a camera images it, and optical triangulation turns the stripe's deformation into the joint's position in three dimensions. The method is common for joints with large grooves, its images are strongly disturbed by the arc, and commercial sensors work at more than 100 mm from the part. It fails where the joint gives the stripe nothing to bend: for a butt joint narrower than 0.2 mm the deformation almost disappears [20]. On a trend a failing sensor shows as missing or jumping readings, which the model's dropouts, noise and covered window stand for.

What Moves It

Two kinds of final element act on the weld. The robot's axes carry the torch and apply the height and lateral corrections, and the feeder and the source set the arc. Modern pulsed sources hold the arc length themselves, comparing the arc voltage with a reference and adjusting the pulse frequency and the current [19]. The feeder lag τ_f of 30 ms and the synergic slope of Equation (12.15) stand for the feeder and the source.

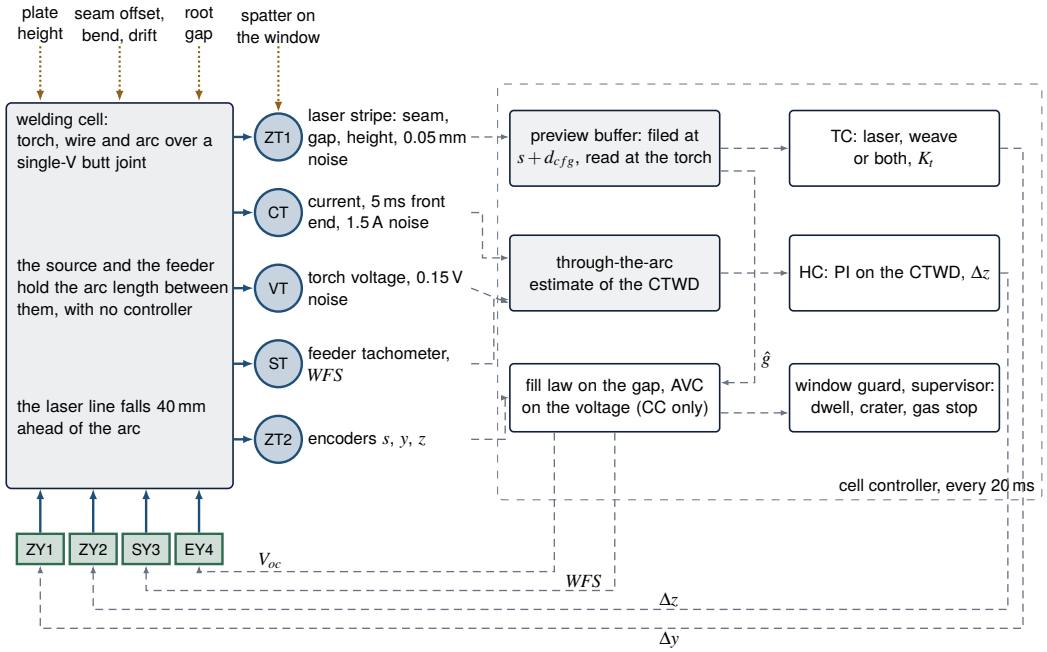

ZY1, ZY2: lateral and height servo axes of the robot, corrections within ±10 mm. SY3: wire feeder, 30 ms lag, 1.5 to 18 m/min. EY4: power source, open-circuit voltage on the synergic line. The travel axis, slowed by the fill law or the guard, is not drawn.

Figure 12.5 The loops of the chapter with real devices in place of the schematic's symbols. ZT1 is the laser stripe sensor ahead of the torch, CT and VT the current and torch-voltage transmitters, ST the feeder tachometer and ZT2 the axis encoders. In the cell controller (dashed frame) the preview buffer feeds the tracking law TC and the fill law, the through-the-arc estimate feeds the height loop HC, and the AVC loop acts on the feed only on a constant-current source. The controller moves the lateral and height axes ZY1 and ZY2, the wire feeder SY3 and the source voltage EY4. Notice that the arc length itself has no loop in the controller: the source and the feeder close it inside the cell.

TECHNOLOGY **Final Element: Power Source and Wire Feeder**

The source and the feeder set the arc between them. A constant-voltage source with a fixed feed holds the arc length by self-regulation, the most common method, while a constant-current source needs a feed motor driven by the process voltage [21]. A pulsed source typically needs a feedback loop in the feeder to follow changes of the distance [18]. In one laboratory cell the current regulator followed its command up to 500 Hz and the capstan feeder only up to 7 Hz [21]. Debris in the contact tube makes the feed hesitate and then stop [18], and a wire that stops lets the arc burn back toward the tip.

Where the Controller Runs

The cell controller of this chapter runs every 20 ms, forty physics steps, while the arc's own loop needs no controller at all. In discrete manufacturing a supervisor for each cell may coordinate a machine controller and a robot controller, and such loops run with cycle times from milliseconds to minutes [22]. An adaptive welding cell joins a six-axis robot, a welding power source and a camera through a process controller that combines the three, and it adjusts the travel speed, the torch orientation and the synergic settings of the source [9]. The preview buffer, the fill law and the window guard of this chapter sit in that cell layer, above the axis servos and beside the source's own regulator.

> **IN PRACTICE** **The Contact Tube Wears Out**
>
> The contact tube is a copper-alloy part with a bore 0.2 to 0.5 mm larger than the wire, and it wears in service. As it wears, the point where current passes into the wire moves, the length of wire carrying current changes, and the arc grows unstable. That instability sets the interval at which the tube is replaced. A welder compensates by moving the gun and spots a worn tube by sight and sound, so an automated cell needs a monitor: in tests at NIST the power of the circuit voltage between 0 and 4 Hz correlated with the wear [18]. The model's contact tip keeps a fixed resistance.

What Is Modernizing It

Cameras now watch the weld pool itself: one adaptive cell measures the pool's width and the position of its front in real time and turns their changes into an adjustment of the process [9]. Deep learning is applied to the monitoring of welding and to predicting weld quality, and a laser vision sensor with deep neural networks has predicted the internal bead shape and the tensile strength of a joint from its external bead shape and the welding conditions [23].

12.8 Limitations

Pulsed, short-circuit and globular transfer, inverter waveform control, arm kinematics and image processing are outside the model, as its assumptions state, and the bead correlations, defect rules, wall rise and drift law are illustrative teaching values rather than a qualified procedure or a distortion model. The through-the-arc lateral calibration is fixed at 3 Hz and reads a 0.3 mm offset 0.54 to 0.78 times at 2 Hz and 1.25 to 1.32 times at 4 Hz. The through-the-arc height is an inversion of the steady arc and is biased and noisy during a transient. The part, the gap and the noise come from one seed, so a case started the same way repeats exactly, which makes comparisons exact and removes the variation a cell sees from part to part, while worst offsets that differ by less than their spread over other seeds do not rank two configurations.

> **OPEN CHALLENGE** **Welding a Joint That Changes Under the Torch**
>
> Distortion, varying seam preparation and a gap that changes during welding remain a major challenge for automated gas metal arc welding. A human welder adjusts speed and torch angle on the fly, while a robot needs sensors to detect the change and actuators to answer it [9]. Models learned from data need data of high quality, and data on unstable welds are scarce in low-volume work on complex paths, where a misjudged weld is costly. Progress would be a cell that adapts from little data, as a framework that adds physical laws to a learned model begins to do [24].

12.9 Conclusion

This chapter studied a robotic arc welding cell that holds its arc length without a controller and reads its seam with a sensor that leads the arc in distance rather than time, in a simulator that logs the true arc, seam and bead beside the cell's measurements. A 3 mm plate step drives the current to 316.5 A and the arc recovers 90 % of its length in 97.0 ms, against 104.7 ms from a linearization of the model, while a constant-current arc stays 0.99 mm long until its voltage-controlled feed restores it in 834.5 ms. The 20.5 A the arc keeps is a height signal that holds the CTWD within 0.13 mm of its setpoint as the plate rises 4 mm, and at a 3 Hz weave the current's sensitivity to the distance is 2.77 times its steady value, close to the 2.5 that the weaving tracker carries. A look-ahead configured 20 mm short turns a 3.5 mm bend into a 1.451 mm error against the 1.455 mm of the geometric law and 372 mm without fusion. The graded seam improves from 1198 mm defective to none at 12.30 mm/s. The models of transfer, bead and distortion are simplified and their limits illustrative. A fuller model would add short-circuit and pulsed transfer, arm kinematics with a sensor interface delay, a laser with image processing, and a bead model fitted to a qualified procedure. The welding

cell gives a student, in a browser, a process that closes its own loop and a sensor whose readings belong to places rather than to moments.

KEY POINTS

- A welding arc holds its own length: a shorter arc draws more current, which melts the wire faster until the arc is long again.
- After a 3 mm plate step the arc recovers 90 % of its length in 97.0 ms, while a constant-current arc stays 0.99 mm long.
- The 20.5 A the arc keeps is a height signal that holds the CTWD within 0.13 mm of its setpoint as the plate rises 4 mm.
- A look-ahead configured 20 mm short turns a 3.5 mm bend into a 1.451 mm error and 372 mm without fusion.
- On a wider gap, the extra metal of adaptive fill brings extra heat.

Where else this applies. Readings filed by place also steer the road-preview suspension of Volume 7, Chapter 11, and a feedforward bounded by the plant also doses the coagulant of Volume 3, Chapter 6.

References

[1] J. N. Pires, A. Loureiro, T. Godinho, P. Ferreira, B. Fernando, and J. Morgado, "Welding robots," *IEEE Robotics & Automation Magazine*, vol. 10, no. 2, pp. 45–55, 2003.

[2] J. Norrish, *Advanced Welding Processes*. Woodhead Publishing, 2006.

[3] G. E. Cook, "Robotic arc welding: Research in sensory feedback control," *IEEE Transactions on Industrial Electronics*, vol. IE-30, no. 3, pp. 252–268, 1983.

[4] J.-W. Kim and S.-J. Na, "A study on arc sensor algorithm for weld seam tracking in gas metal arc welding of butt joints," *Proceedings of the Institution of Mechanical Engineers, Part B: Journal of Engineering Manufacture*, vol. 205, no. 4, pp. 247–255, 1991.

[5] P. Kadlec, B. Gabrys, and S. Strandt, "Data-driven soft sensors in the process industry," *Computers & Chemical Engineering*, vol. 33, no. 4, pp. 795–814, 2009.

[6] M. Tomizuka, "Optimal continuous finite preview problem," *IEEE Transactions on Automatic Control*, vol. 20, no. 3, pp. 362–365, 1975.

[7] A. Hać, "Optimal linear preview control of active vehicle suspension," *Vehicle System Dynamics*, vol. 21, no. 1, pp. 167–195, 1992.

[8] A. Rout, B. B. V. L. Deepak, and B. B. Biswal, "Advances in weld seam tracking techniques for robotic welding: A review," *Robotics and Computer-Integrated Manufacturing*, vol. 56, pp. 12–37, 2019.

[9] A. Biber, R. Sharma, and U. Reisgen, "Robotic welding system for adaptive process control in gas metal arc welding," *Welding in the World*, vol. 68, no. 9, pp. 2311–2320, 2024.

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

[11] D. S. Naidu, S. Ozcelik, and K. L. Moore, *Modeling, Sensing and Control of Gas Metal Arc Welding*. Elsevier, 2003.

[12] Z. Bingul, G. E. Cook, and A. M. Strauss, "Dynamic model for electrode melting rate in gas metal arc welding process," *Science and Technology of Welding and Joining*, vol. 6, no. 1, pp. 41–50, 2001.

[13] Z. Bingul and G. E. Cook, "A real-time prediction model of electrode extension for GMAW," *IEEE/ASME Transactions on Mechatronics*, vol. 11, no. 1, pp. 47–54, 2006.

[14] L. A. Jones, T. W. Eagar, and J. H. Lang, "A dynamic model of drops detaching from a gas metal arc welding electrode," *Journal of Physics D: Applied Physics*, vol. 31, no. 1, pp. 107–123, 1998.

[15] W. F. Clocksin, J. S. E. Bromley, P. G. Davey, A. R. Vidler, and C. G. Morgan, "An implementation of model-based visual feedback for robot arc welding of thin sheet steel," *The International Journal of Robotics Research*, vol. 4, no. 1, pp. 13–26, 1985.

[16] T. C. Nguyen, D. C. Weckman, and D. A. Johnson, "Predicting onset of high speed gas metal arc weld bead defects using dimensional analysis techniques," *Science and Technology of Welding and Joining*, vol. 12, no. 7, pp. 634–648, 2007.

[17] D. Deng and H. Murakawa, "Prediction of welding distortion and residual stress in a thin plate butt-welded joint," *Computational Materials Science*, vol. 43, no. 2, pp. 353–365, 2008.

[18] M. A. Mornis, T. P. Quinn, T. A. Siewert, and J. P. H. Steele, "Sensing of contact tube wear in gas metal arc welding," National Institute of Standards and Technology, Boulder, CO, NIST Interagency/Internal Report 3996, 1992.

[19] W. Liu, L. Li, Y. Hong, and J. Yue, "Linear mathematical model for seam tracking with an arc sensor in P-GMAW processes," *Sensors*, vol. 17, no. 3, p. 591, 2017.

[20] B. Xue, B. Chang, G. Peng, Y. Gao, Z. Tian, D. Du, and G. Wang, "A vision based detection method for narrow butt joints and a robotic seam tracking system," *Sensors*, vol. 19, no. 5, p. 1144, 2019.

[21] R. B. Madigan, T. P. Quinn, and T. A. Siewert, "Control of gas-metal-arc welding using arc-light sensing," National Institute of Standards and Technology, Boulder, CO, NIST Interagency/Internal Report 5037, 1995.

[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] K. Lee, I. Hwang, Y.-M. Kim, H. Lee, M. Kang, and J. Yu, "Real-time weld quality prediction using a laser vision sensor in a lap fillet joint during gas metal arc welding," *Sensors*, vol. 20, no. 6, p. 1625, 2020.

[24] J. Liu, F. Jiang, S. Tashiro, S. Chen, and M. Tanaka, "A physics-informed and data-driven framework for robotic welding in manufacturing," *Nature Communications*, vol. 16, p. 4807, 2025.

12.10 Exercises

12.1 ⌕ Concept check · 2 min

In case 1 the torch meets a 3 mm plate step in MANUAL on the constant-voltage source, with the feed fixed at 8 m/min and no loop acting. The arc is suddenly 1.28 mm long and is 90 % restored 97.0 ms later. What restores it?

(a) The height loop raises the torch, because the current has risen above its reference

(b) The shorter arc draws more current, which melts the wire faster than it is fed

(c) The feeder slows the wire, because the synergic line follows the falling torch voltage

(d) The droplets detach more often at the higher current, and each detachment lengthens the arc

12.2 ⌕ Concept check · 1 min

In case 4 the laser sits 40 mm ahead of the arc, and the controller files its readings as if it sat 20 mm ahead. The seam bends 3.5 mm to either side every 300 mm. Where along the bend is the bead farthest from the seam?

(a) At the crests of the bend, where the seam changes direction fastest

(b) Equally everywhere, since every reading is used the same 20 mm early

(c) Where the seam crosses the taught path, at its steepest slope

(d) Only near the start, until the buffer has filled with readings

12.3 ⌕ Concept check · 2 min

In case 6 the part bows as it heats, and its far end moves 8.36 mm over the weld while the torch weaves 1.5 mm. Laser tracking alone leaves a worst bead offset of 0.380 mm, tracking through the arc alone 0.233 mm, and the laser with the through-the-arc trim 0.154 mm. Which statement explains the ranking?

(a) The laser's noise is larger than the current's, so any arc reading improves its estimate

(b) The trim corrects the configured look-ahead, which the drift has moved off the mounted distance

(c) The weave widens the bead, so the same lateral error counts for less at the groove edges

(d) Each laser reading is 4 s old when the torch arrives, and the trim reads the groove at the torch

12.4 Theory and calculation · 9 min

At the default arc the current is $i = 251.9\,\mathrm{A}$, the extension $l_s = 10.73\,\mathrm{mm}$ and the circuit resistance $R_\Sigma = 47.7\,\mathrm{m\Omega}$. The melting coefficients are $k_1 = 0.300\,\mathrm{mm/(s\,A)}$ and $k_2 = 8.48 \times 10^{-5}\,1/(\mathrm{s\,A^2})$, the column gradient is $E_a = 1.2\,\mathrm{V/mm}$ and the hot wire has $r_w = 0.62\,\mathrm{m\Omega/mm}$.

(a) Compute the arc's share and the extension's share of λ in Equation (12.9), and λ itself. *Checkpoint:* $16.6\,\mathrm{s^{-1}}$ and $5.38\,\mathrm{s^{-1}}$, so $\lambda = 22.0\,\mathrm{s^{-1}}$.

(b) Compute the time constant of self-regulation and the time to 90 % of a step. The full model restores 90 % of the case 1 step in 97.0 ms. Give one reason the model is faster than the linearization. *Checkpoint:* $\tau_{sr} = 45.5\,\mathrm{ms}$ and $2.303\,\tau_{sr} = 104.7\,\mathrm{ms}$.

(c) From Equation (12.10), compute the current's sensitivity to the CTWD at steady state and at a 3 Hz weave, and their ratio. Compare the ratio with the calibration $c = 2.5$ of Equation (12.14). *Checkpoint:* 6.16 A and 17.0 A per millimeter, a ratio of 2.77.

12.5 Theory and calculation · 6 min

Case 4 bends the seam with amplitude $A_b = 3.5\,\mathrm{mm}$ and wavelength $\lambda_b = 300\,\mathrm{mm}$, and the laser is mounted at $d_{la} = 40\,\mathrm{mm}$. The default bead is $w = 8.03\,\mathrm{mm}$ wide, the root gap is 1.0 mm, the groove's top width is $W_V + g$ with $W_V = 4\,\mathrm{mm}$, and the edge margin is $m_e = 0.3\,\mathrm{mm}$.

(a) Use Equation (12.13) to compute the worst lateral error for configured look-aheads of 20, 30, 50 and 60 mm. *Checkpoint:* 1.455, 0.732, 0.732 and 1.455 mm.

(b) Compute the largest bead offset that still fuses both groove edges, and say which of the four configurations the law predicts will lose fusion. *Checkpoint:* 1.22 mm. The 20 and 60 mm configurations.

(c) How far may the configured look-ahead be from the mounted one before the law predicts lack of fusion? Give the answer to 0.1 mm. *Checkpoint:* 16.7 mm, so from about 23.3 to 56.7 mm.

12.6 Theory and calculation · 7 min

In case 8 the travel speed rises from 10 to 14 mm/s with the feed at 8 m/min and the source unchanged. At 10 mm/s the pool's heat input is 681 J/mm. The correlations are $p = a_p H^{b_p}(1 + \beta_g g)$ with $a_p = 0.04184$, $b_p = 0.70$, $\beta_g = 0.05\,\mathrm{mm^{-1}}$ and no weave, a fusion limit $p_{min} = 3.4\,\mathrm{mm}$, an undercut speed $v_{max} = v_u(H/675\,\mathrm{J/mm})^{e_u}$ with $v_u = 13\,\mathrm{mm/s}$ and $e_u = 0.8$, and a deposit $\eta A_w WFS/v$ with $\eta = 0.97$ and 1.2 mm wire. The root gap is 1.0 mm, so the groove needs $12.0\,\mathrm{mm^2}$ per millimeter.

(a) Assume the arc's power does not change with the speed, and compute the heat input at 14 mm/s. *Checkpoint:* 486 J/mm.

(b) Compute the penetration and the undercut speed at that heat input. *Checkpoint:* 3.34 mm and 10.0 mm/s.

(c) Compute the deposit per millimeter at 14 mm/s, and list every defect the bead rules predict. *Checkpoint:* $10.4\,\mathrm{mm^2}$. Lack of fusion, undercut and underfill.

12.7 App challenge · 8 min

Load scenario 1, The arc regulates itself. It runs in MANUAL on the Arc view at the default seed, and the torch meets the 3 mm plate step at 6.95 s. Run each part at 1×, export the CSV, and read the columns `arc_length_mm`, `current_A` and `wfs_mpm` in the rows named. Read the arc length to ±0.05 mm, the current to ±3 A, since the droplets ripple it, and the feed to ±0.05 m/min.

(a) Run the scenario as loaded to 8 s. Read the arc length and the current at 6.90 s and at 7.90 s, and say how much of the 3 mm step the arc does not recover. *Checkpoint:* 4.29 mm and 252 A, then 3.81 mm and 272 A; the arc stays about 0.48 mm shorter.

(b) Reload the scenario with the reload button, set Power source to Constant current before Run, stay in MANUAL, and run to 8 s. Read the arc length at 6.90 s and at 7.90 s. *Checkpoint:* 3.99 mm and 0.99 mm.

(c) Reload, set Power source to Constant current, switch the header to Auto before Run so the AVC loop sets the feed, and run to 9 s. Read the arc length and the feed at 8.90 s, each to ± 0.1, since the slow integral leaves both still moving. *Checkpoint:* 4.44 mm and 6.88 m/min.

12.8 App challenge · 9 min

Load scenario 4, The look-ahead set wrong. It tracks with the laser, which is mounted 40 mm ahead of the arc, and the look-ahead is configured at 20 mm. Use the default seed and run at $2\times$ until the run stops itself when the crater is filled, about 122 s of simulated time. Reload the scenario with the reload button before each part, and make the change before Run. Read the scorecard tiles Worst lateral error and Defects. Worst lateral error is exact to the 0.01 mm the tile shows at the default seed.

(a) Run the scenario as loaded. *Checkpoint:* 1.45 mm; Defects 372 mm, lack of fusion.

(b) Set Look-ahead distance (configured) to 30 mm and run. *Checkpoint:* 0.73 mm; Defects 0 mm.

(c) Set Look-ahead distance (configured) to 60 mm and run. Compare the three results with Equation (12.13), and say which error the law underestimates. *Checkpoint:* 1.49 mm; Defects 397 mm, lack of fusion.

12.9 App challenge · 8 min

Load the challenge, “a 1.2 m seam at 12 mm/s”. It welds at 12.5 mm/s over a gap that varies between 0.8 and 2.2 mm, on a part set 1 mm off its path that drifts as it heats. Use the default seed and run at $2\times$ until the run stops itself when the crater is filled, about 98 s of simulated time. Reload the scenario with the reload button before each part, make every change before Run, and read the tiles Defects, Worst lateral error, Average speed and Wire used.

(a) Set Seam tracking to Laser, Adaptive fill On with Fill by on Wire feed, Wire feed speed (CV) to 10.0 m/min, Source voltage (open circuit) to 34.2 V and Reinforcement target to 1.0 mm^2, and run. *Checkpoint:* Defects 0 mm, 0.13 mm, 12.3 mm/s, 18.1 m; the challenge target is met.

(b) Make the same changes but leave the feed at 8.0 m/min and the source at 31.8 V, and run. *Checkpoint:* Defects 331 mm, undercut; 0.13 mm, 12.3 mm/s, 17.2 m.

(c) Explain part (b) with the undercut correlation of Equation (12.8), taking the heat input of the 8 m/min procedure at 12.5 mm/s as 681 J/mm scaled by the speed. *Checkpoint:* 545 J/mm, an undercut speed of 10.9 mm/s.

12.10 Application area · 3 min

A shielding-gas loss in case 7 at 45 s trips the flow switch 0.28 s later, and the supervisor stops the weld at 443 mm, leaving 8 mm of porous, unfilled crater. With the gas stop off, 760 mm of the seam are porous.

(a) Which setting should a production cell use, and what does the repair cost in each case? *Checkpoint:* Stop on gas loss; a crater to repair against most of the seam.

(b) Why does the monitor wait for the flow to stay below 70 % for 0.3 s instead of stopping at the first low reading? *Checkpoint:* The flow switch is noisy, and a short dip does not spoil the shield.

12.11 Application area · 3 min

A technician replaces a bent laser bracket with a longer one that mounts the scanner 60 mm ahead of the arc, and nobody changes the configured look-ahead of 40 mm. The cell is checked by welding a straight test coupon whose seam is offset 2 mm from the taught path, and it passes.

(a) What does Equation (12.13) predict on the coupon, and on production parts that bend like case 4? *Checkpoint:* No error on the straight coupon; 1.455 mm on the bend.

(b) How can the fault be found before a bead shows it? *Checkpoint:* Compare where the readings are filed with where the laser line falls, or weld a known feature.

12.12 Review the agent · 5 min

An AI assistant reviews case 5 and writes: “Fill by wire feed welds the whole seam sound as the gap opens from 1.0 to 2.1 mm, because the law adds exactly the wire the measured gap needs. The law has no limit except the feeder’s 18 m/min, so the same settings will weld parts whose gap opens to 2.5 mm. Set the Root gap slider to 1.4 mm and expect a sound seam.” Find the flaw, and say how you would check the claim in the application.

CHAPTER 13

Robotic Peg-in-Hole Simulator for Compliance, Jamming and Force Control Through Contact

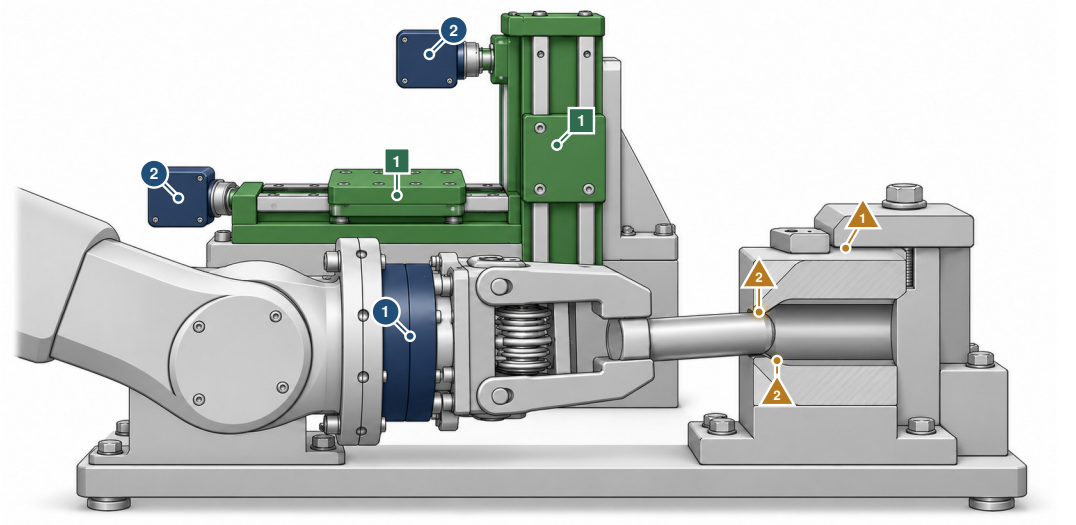

The controller moves 1 robot axes that move the peg across and into the hole

The controller reads 1 wrist force sensor, read through a lag 2 position encoders on the axes

What disturbs it 1 fixture shifted or tilted from where the program aims 2 friction where the peg touches the hole

Notice the tilted peg touching the bore at two amber points, the two-point contact that can jam it, with the insertion drawn lying on its side. Illustration, not to scale.

IN PLAIN TERMS **Pushing a Peg Into a Hole**

The main idea. A robot arm puts a pin into a hole that is only a hair wider than the pin. Once the pin touches the part, the robot can no longer choose exactly where the pin goes, only how hard it pushes.

The challenge. A robot that holds its path very stiffly turns a tiny miss into a force large enough to break the part. A springy hold lets the hole guide the pin in. But a tilted pin can still wedge, and a late force reading can hide hard knocks.

Why it matters. Shafts, pins and connectors all go together this way, so these choices decide whether a part is made or scrapped.

THREE WAYS TO READ THIS CHAPTER

Explore	11 min	the first page, the fact on the first touch (Section 13.1), the history segment (Section 13.2), the try-it segments of cases 2 and 7 (Section 13.5), the fact on the hundred parts (Section 13.6) and the key points
Study	45 min	the whole chapter with its three pause points and the twelve exercises at its end
Practice	15 min	the first page, the measured cases (Section 13.6), Technology in Practice (Section 13.7), Limitations (Section 13.8) and the open challenges

The control question Once a peg touches its part, the part fixes some of the peg's coordinates and the controller can only decide how hard to push along them. How do stiffness, friction, landing speed and a late force signal each turn a position error of a few tenths of a millimeter into a contact force, a jam or a damaged part?

Learning objectives

- Compute the fixture offset a stiff position servo can absorb on a chamfer, from the chamfer force factor and the damage limit.
- Estimate the force of a landing from the effective mass of the contact corner and the approach speed.
- Predict with a moment balance the depth at which a tilted peg jams, and choose the rotational stiffness that frees it.
- Explain why the delay margin of the insertion force loop is set by the sensor lag and the filter, and why a calm force reading can hide a hammering press.
- Decide where a guarded move must slow down, and where the compliance center belongs.

Prerequisites Sections 1.4.8 (impedance, landing force and the delay of a force loop), 1.6.2 (velocity-form PID), 1.3 (loop action) and 1.6.4 (scenarios, scorecard and CSV file).

The app https://apmonitor.com/apps/robot-force-assembly/

See also Volume 7, Chapter 13 for a contact process whose stability boundary is set by a delay, and Volume 7, Chapter 4 for a force that friction caps at the contact patch.

13.1 Introduction

An industrial robot repeats a programmed path to a small fraction of a millimeter, which is enough for welding, painting and transfer, where nothing on the path pushes back. Assembly ends in contact instead. A shaft entering a bearing, a pin entering a connector and a cartridge entering a housing are versions of one task, a peg entering a hole whose clearance is small beside the uncertainty in where the hole is, because fixture tolerance, part variation and robot calibration all add to the position error [1, 2]. Once the peg touches the part, the part's surfaces fix some of the peg's coordinates, and along those directions the controller can prescribe only a force [3]. Hybrid position and force control assigns each task direction to one of the two, and impedance control prescribes the relation between them [4–6]. Section 1.4.8 sets out the impedance law, the landing force and the delay margin of a force loop in general. This chapter applies them to one peg.

The belief this simulator is meant to overturn is that a more precise position controller makes a better assembly robot. After contact, precision becomes force. The robot's own servo in the simulator holds the tip with 200 N/mm, so each micrometer by which the program and the part disagree costs 0.2 N. With the fixture shifted 0.4 mm the servo damages the part and the contact force climbs to 136.1 N, while a programmable spring of 3 N/mm puts the same peg into the same part with a largest contact force of 13.7 N. Stiffness is the rate of exchange between position error and force, and once the peg is in contact it is a design choice rather than a mark of quality.

The class of problem is general. A grinding wheel on a casting, a polishing pad under a regulated normal force, a friction stir tool plunged into a plate, the foot of a walking robot meeting the ground and a surgical instrument meeting tissue all pass from free motion, where position control is enough, into constrained motion, where force has to be controlled or at least bounded [2, 6]. Volume 9, Chapter 3 plunges such a tool under a force limit, and Chapter 14 closes a force loop on tissue. Three difficulties recur across the class, and a block diagram that treats the environment as a linear spring shows none of them. Geometry and friction can jam a part that more force will not free [1]. An impact is decided by the moving mass and the speed at contact [7], and in the simulator it is over within a millisecond, before any feedback can act. A force loop is bounded by the dynamics that lie between the actuator, the sensor and the environment [8, 9]. The simulator produces each difficulty on purpose

and draws the contact forces the controller never reads beside the lagged reading it does act on, so that a jam, an impact and a hammering press are measured rather than inferred.

> ✦ DID YOU KNOW? **The First Touch Is the Hardest**
>
> The force spike of the first impact between tool and part is often three to five times larger than the forces of the rest of an assembly [10]. In NIST tests of two arms stepping to a contact force, the overshoot stayed well under 10 N along the horizontal axes and was typically well over 100 N along the vertical one [11].

This chapter attaches a number to each of these difficulties, with a peg-in-hole simulator in which stiffness, friction, landing speed and sensing can each be changed alone and their consequence read as a contact force, a jam or a damaged part. Impedance control, force loops closed around a position reference and quasi-static jamming analysis each have a long record [1, 5, 12], and the chapter uses them as established methods. Section 13.2 traces force control in assembly. Section 13.3 sets down a planar contact model with six penalty sites and a friction law that holds a stuck contact without creep, with every relation and constant needed to rebuild it. Section 13.4 gives the three controllers, which read only a lagged, noisy and optionally delayed wrist sensor. Section 13.5 checks four short estimates against the nonlinear simulator, for the offset at which the stiff servo damages the part, for the force of a landing, for the depth at which a jammed peg stops and for the delay a force loop can tolerate. Section 13.6 measures the eight cases and the challenge, in which each failure is met on its own. Section 13.7 names the devices of real assembly cells, Section 13.8 states the limitations, Section 13.9 draws the results together, and the exercises of Section 13.10 work through the cases.

13.2 Force Control and Compliant Assembly

Automated mechanical assembly was studied in the 1970s as a problem of using force and position information together while parts are in contact. The history segment below follows its first two answers, force feedback and the passive remote center compliance.

> ⧖ HISTORY **Two Answers to the Peg, 1977 and 1982**
>
> In the 1970s, automating mechanical assembly meant using force and position information together while parts touched. One answer, published in 1977, used force and position sensors to make small corrections in the paths of parts in contact. It showed that such force feedback strategies must be posed as control problems, with stability bounds derived from simple linear actuator models and sampled-data control [8]. The other answer was passive. A quasi-static analysis of 1982 set out the geometric and force equilibrium conditions for mating rigid parts, explained the action of the remote center compliance and gave guidelines for choosing its parameters [1]. Such a device lets the held part give way to forces and torques, so that it adjusts to small lateral and angular errors [10].

Programmable compliance followed within a few years. A formal account of compliant motion separated the constraints that the task geometry imposes from those the controller remains free to choose [3]. Hybrid position and force control combined readings from a force-sensing wrist with joint positions to meet position and force trajectories in a coordinate system tied to the task [4]. Impedance control then argued that control of position or of force alone is inadequate once a manipulator is coupled to its environment, and that a manipulator handling inertial objects must behave as an impedance [5]. A review in 1987 traced these lines, compared the impedance and hybrid strategies built on them, and discussed the stability problem they share [2], while dynamic models of the robot, the sensor and the environment were introduced to explain where that problem comes from [9].

A second arrangement closes the force loops outside the robot's own position control. The position loops decouple and linearize the arm and present the external controller with a plant that is easy to model, and the effect of a passive compliance on task speed and disturbance rejection was analyzed in the same framework [12]. Explicit force control strategies were then compared in theory and

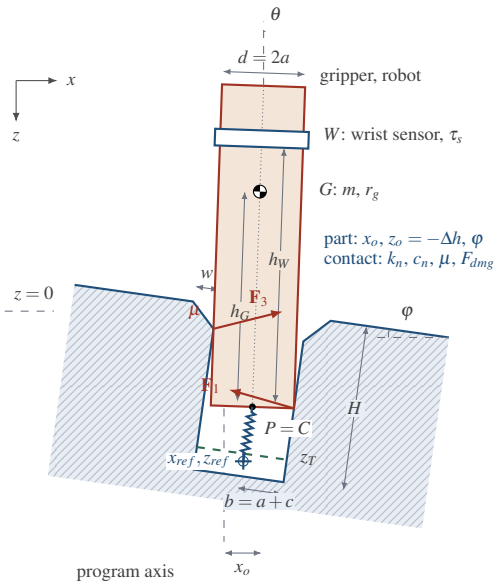

Figure 13.1 Section through the peg and the chamfered hole, drawn in two-point contact and not to scale. The part sits at offset x_o, height $z_o = -\Delta h$ and tilt φ from the program's nominal hole, and its fixture frame (X, Z) has its origin at the hole's top center. The peg of width $d = 2a$ leans θ from the vertical, its tip P coincides with the compliance center C at the default $h_C = 0$, and G and the wrist sensor W lie h_G and h_W above it. The impedance spring pulls C toward the reference (x_{ref}, z_{ref}), and the contact forces $\mathbf{F}_1$ at the right corner and $\mathbf{F}_3$ at the left entry edge carry friction μ. The hole of half-width $b = a + c$ has a chamfer of width w, depth H and the target depth z_T.

in experiment [13]. When the position error exceeds what a chamfer can capture, the hole has to be found before it can be entered, and search strategies for peg-in-hole assembly under position uncertainty address that case [14]. Two later threads frame the damage limit and the landing speed of this simulator. Robots that work beside people made the force of an unintended impact a safety requirement, measured with automotive crash-test methods on a lightweight arm [7], a budget that Chapter 15 keeps beside a person, and the force control capability of robots is now benchmarked by experiment [11].

13.3 Peg, Part and Contact Model

The simulated cell is a vertical section through a peg 12 mm wide held by the robot above a part with a chamfered hole, as shown in Figure 13.1. The hole leaves 50 µm of clearance on each side, its mouth carries a 45° chamfer 0.8 mm wide, it is 14 mm deep, and the insertion is complete at 10 mm. The wrist, the gripper and the peg move as one rigid body of 2 kg whose center of mass G lies 90 mm above the peg tip P. The part sits in a fixture whose pose the controller does not know, made of a lateral offset x_o, a tilt φ about the hole's top center and a height error Δh, while the program aims at a nominal hole at $x = 0$ with its top face at $z = 0$. The model is written in SI units and displayed in millimeters, micrometers and degrees.

Table 13.1 classifies the variables. Fifteen states are integrated, the position x_G, z_G and tilt θ of the center of mass and their rates v_x, v_z, ω, one tangential friction spring b_i at each of six contact sites, and three lagged components of the wrist sensor reading. The depth coordinate z points down into the hole, and θ is positive when the top of the peg leans toward $+x$.

Table 13.1 Variables, With Nominal Values of the Default Cell

Symbol	Role	Nominal	Range
F_{ins}	CV, insertion force reading	SP 12 N	SP 2–40
F_x	CV, side force reading	SP 0 N	SP ± 10
z_{ref}	MV, z command	path	-5 to 13 mm
x_{ref}	MV, x command	0 mm	± 3 mm
F_{cx}, F_{cz}, M_c	wrench at C	—	± 300 N, ± 20 N m
x_o	DV, fixture offset	0 mm	± 2
φ	DV, fixture tilt	0°	± 3
μ	DV, friction coefficient	0.20	0.05–0.8
Δh	DV, part height error	0 mm	-1 to 2.4
x_G, z_G, θ	states, pose of G	tip 3 mm up	—
v_x, v_z, ω	states, rates	0	—
$b_0 \dots b_5$	states, friction springs	0	$k_t \lvert b_i \rvert \le \mu F_{n,i}$
$F_{s,x}, F_{s,z}, M_s$	states, lagged sensor	0	—
Z	output, true depth	—	target 10 mm
$\max_i \lvert \mathbf{F}_i \rvert$	output, largest contact	—	damage > 60 N

13.3.1 Rigid Body

The end effector obeys the Newton-Euler equations written at G, where the inertia is diagonal, as shown in Equation (13.1).

$$\begin{aligned} m\dot{v}_x &= F_{cx} + \textstyle\sum_i F_{ix} \\ m\dot{v}_z &= F_{cz} + mg + \textstyle\sum_i F_{iz} \\ mr_g^2\,\dot{\omega} &= M_c + (h_C - h_G)\,(F_{cz}\sin\theta + F_{cx}\cos\theta) \\ &\quad + \textstyle\sum_i \left[(p_{ix} - x_G)F_{iz} - (p_{iz} - z_G)F_{ix}\right] \end{aligned} \tag{13.1}$$

where (F_{cx}, F_{cz}, M_c) is the wrench the controller applies at the compliance center C, a height h_C above the tip, held constant over each control period, $\mathbf{F}_i$ is the contact force at site i acting at $\mathbf{p}_i$, r_g is the radius of gyration about G and g is gravity. The unit vectors $\mathbf{u} = (\sin\theta, -\cos\theta)$ up the peg and $\mathbf{r} = (\cos\theta, \sin\theta)$ across it place the tip at $P = G - h_G\mathbf{u}$, so that $x = x_G - h_G\sin\theta$ and $z = z_G + h_G\cos\theta$. Because the center of mass sits high, a sideways force at the tip also turns the peg, and the tip presents an effective lateral mass of $m/(1 + h_G^2/r_g^2) = 0.47$ kg to a sudden push.

13.3.2 Contact Sites

The part's profile in the fixture frame is the top face, two 45° chamfers, two walls at $|X| = b = a + c$ and a bottom at $Z = H$, where $a = d/2$ is the peg's half-width and c the clearance per side. World points map into that frame through Equation (13.2).

$$\begin{aligned} X &= (x - x_o)\cos\varphi + (z - z_o)\sin\varphi \\ Z &= -(x - x_o)\sin\varphi + (z - z_o)\cos\varphi \end{aligned} \tag{13.2}$$

with $z_o = -\Delta h$, so that Z evaluated at the tip is the true insertion depth. Contact is evaluated at six sites. Sites 0 and 1 are the peg's bottom corners $P \mp a\mathbf{r}$, tested against the whole profile. Sites 2 to 5 are the part's convex edges, the tops of the chamfers at $X = \mp(b + w)$, $Z = 0$ and the entry edges at $X = \mp b$, $Z = w$, tested against the peg's bottom face and its two sides. At each site the penetration δ_i is the distance to the nearest surface, $\mathbf{n}_i$ is the unit normal along which the part pushes the peg, and $\mathbf{t}_i = (-n_{iz}, n_{ix})$ is the tangent. One-point contact, two-point contact and jamming are not flags. They are what the six sites do with the geometry and the friction.

13.3.3 Normal Force and Friction

The normal force is a unilateral spring and damper whose damping fades in over the first δ_r of penetration, as shown in Equation (13.3).

$$F_{n,i} = \max\Big(0,\ k_n\delta_i + c_n\dot{\delta}_i \min(1,\ \delta_i/\delta_r)\Big) \tag{13.3}$$

in which $\dot{\delta}_i$ is the rate at which the peg's material point closes on the surface along $\mathbf{n}_i$, $k_n = 5\,\mathrm{N/\mu m}$ lumps the local compliance of the part, the fixture and the peg, $c_n = 1500\,\mathrm{N\,s/m}$ and $\delta_r = 0.25\,\mu\mathrm{m}$. The maximum keeps the contact from pulling and the ramp keeps the force continuous at touchdown, two faults of a plain spring and dashpot that a damping term growing with penetration also removes [15, 16]. The ramp is a quarter of the 1 µm that a 5 N press penetrates.

Friction is an elastic-plastic Coulomb law, as shown in Equation (13.4).

$$\begin{aligned} &\dot{b}_i = v_{t,i}, \qquad F_{t,i} = -\mathrm{clamp}(k_t b_i + c_t v_{t,i},\ \pm\mu F_{n,i}) \\ &b_i \leftarrow \mathrm{clamp}(b_i,\ \pm\mu F_{n,i}/k_t) \quad \text{after each step} \end{aligned} \tag{13.4}$$

in which $v_{t,i}$ is the slip velocity along $\mathbf{t}_i$, $k_t = k_n/2$, $c_t = c_n/2$ and μ is one coefficient for sticking and sliding alike. The spring b_i is zero whenever its site is out of contact. A stuck site holds any tangential load below $\mu F_{n,i}$ without creeping, and a sliding one has its spring returned to the friction cone after every step. Representing true sticking is the known difficulty in simulating friction near zero velocity [17], and a tangential spring with a plastic limit belongs to the elastoplastic family of friction models [18]. A law regularized around zero slip would let a loaded contact creep, and a wedged peg would slowly ooze through instead of jamming. The contact forces sum to the wrench of Equation (13.1), and their moments are also taken about the sensor point W, a height h_W above the tip, and about the tip itself.

13.3.4 Wrist Sensor, Encoders and Damage

The wrist sensor is a first-order lag on the contact wrench at W, and the controller samples it through noise, an optional delay and its own filter, as shown in Equation (13.5).

$$\begin{aligned} \tau_s\,\dot{\mathbf{F}}_s &= \mathbf{F}_{con} - \mathbf{F}_s \\ \mathbf{F}_{meas}(t) &= \mathscr{L}_{\tau_f}[\mathbf{F}_s(t-\tau_d) + \mathbf{n}_F] \end{aligned} \tag{13.5}$$

where $\mathbf{F}_{con}$ is the contact wrench, $\tau_s = 2\,\mathrm{ms}$ stands for the sensor's bandwidth and its anti-aliasing filter, $\mathbf{n}_F$ is Gaussian noise of 0.15 N on each force and 0.003 N m on the moment, τ_d is a transport delay read from a ring buffer by linear interpolation, and $\mathscr{L}_{\tau_f}$ is the controller's first-order filter of $\tau_f = 3\,\mathrm{ms}$, stepped by the exact hold expression. The sensor is taken as perfectly tared, so it reports the contact wrench without the weight or the inertia of the gripper. The joint encoders give the tip's x and z with 0.5 µm of noise and θ with 5 µrad, and the controller forms rates from a first difference filtered at 0.5 ms. The controller never reads the contact forces or the fixture pose.

The part is damaged at the first physics step on which any single contact carries a resultant $\sqrt{F_n^2 + F_t^2}$ above $F_{dmg} = 60\,\mathrm{N}$. An attempt ends inserted when the controller has declared it done and the true depth is within 0.1 mm of the target, and otherwise damaged, stuck inside the hole, or missed when the peg never finds it. In the challenge a new part is loaded after each attempt, with its fixture offset and tilt drawn uniformly from an envelope. A fixture change takes effect only while the peg is clear of the part, with no site carrying load and the tip at least 0.3 mm above the surface.

13.3.5 Parameters and Numerics

The plant parameters are listed in Table 13.2. A fixed 0.1 ms fourth-order Runge-Kutta step of the engine of Section 1.6.1, the framework's h, integrates Equation (13.1) with the friction springs and the sensor lags, and the plastic update of Equation (13.4) and the damage check follow every step.

Table 13.2 Plant Parameters. Adjustable Ones Are Sliders in the Model Panel, Hidden Ones Are Never Displayed, Code Ones Are Fixed in the Program

Symbol	Meaning	Value	Class
m	wrist, gripper and peg	2.0 kg	adjustable
r_g	radius of gyration about G	50 mm	hidden
h_G, h_W	G and sensor above the tip	90, 110 mm	hidden
F_{max}	force limit at the tool	300 N	adjustable
M_{max}	moment limit at the tool	20 N m	hidden
d, L	peg width, length below gripper	12, 50 mm	hidden
c	clearance per side	50 µm	adjustable
w	chamfer width at 45°	0.8 mm	adjustable
H	hole depth	14 mm	hidden
z_T	insertion target	10 mm	adjustable
k_n	contact stiffness	5 N/µm	adjustable
c_n	contact damping	1500 N s/m	adjustable
δ_r	damping ramp	0.25 µm	code
k_t, c_t	friction spring and damper	$k_n/2$, $c_n/2$	code
F_{dmg}	damage limit, one contact	60 N	adjustable
τ_s	wrist sensor lag	2 ms	adjustable
τ_d	force signal delay	0 ms	adjustable
σ_F	force noise, one sigma	0.15 N	adjustable
σ_x	encoder noise at the tool	0.5 µm	hidden
g	gravity	9.81 m/s^2	hidden

The stiffest contact mode, the 0.47 kg tip on 5 N/µm, rings near 3.3×10^3 rad/s, so the step resolves each period about nineteen times. Halving the step, with the control period held at 0.3 ms, changes the largest force of the angular-error case by 0.0003 % and leaves its insertion time unchanged. Every run starts with the tip at rest 3 mm above the nominal surface and the weight carried, the historian logs at 100 Hz, and generators seeded from 20261021 supply the noise and, separately, the challenge draws, so a run replays exactly.

13.3.6 Assumptions

The section is planar, with one lateral axis, the depth and one tilt. A round peg in a round hole adds a second tilt axis and contact along the rim, which changes the jamming forces by a geometric factor.

The wrist, the gripper and the peg are one rigid body pushed by the commanded wrench. Arm dynamics, joint flexibility, backlash and structural resonance are absent, although these modes between the motor and the tool are what limit force control on a real arm [9].

Contact is a penalty spring of one stiffness that lumps the part, the fixture and the peg. Penetration stays under 12 µm while every contact is below the damage limit, less than a quarter of the clearance, so the geometry is not distorted where it matters, but burrs, waviness and local yielding are not modeled.

One friction coefficient serves for sticking and sliding and it does not wear. Real surfaces often stick harder than they slide, which would add a stick-slip chatter the model does not show.

Damage is a threshold on the largest single contact force at one physics step. Real damage also depends on how long a load lasts, how often it recurs and where it lands.

The wrist sensor is perfectly tared and the fixture is rigid. The part does not move in the fixture, and the fixture moves only between insertions.

13.4 Three Controllers on One Wrench Law

Control executes on every third physics step, every $\Delta t_c = 0.3$ ms, and is handed that interval as its sample time. The structure common to the three controllers is shown in Figure 13.2.

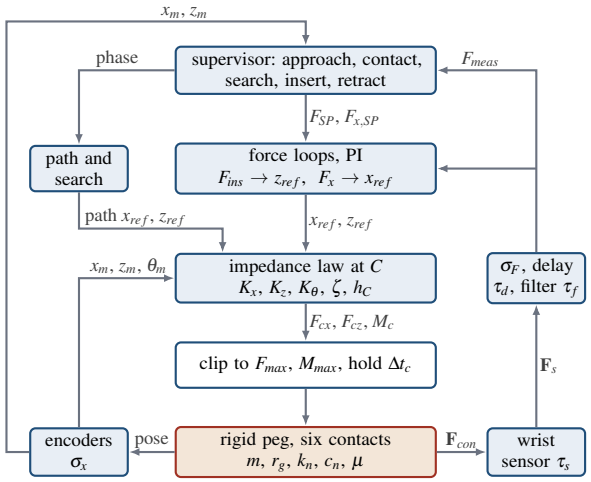

Figure 13.2 Signal flow of the controllers. The supervisor sets the phase and the force setpoints, the force loops or the programmed path write the references, and one impedance law turns them into the wrench at C. The controller reads only the encoders and the lagged, noisy and optionally delayed wrist sensor, never the contact forces.

13.4.1 The Cartesian Wrench Law

Every strategy ends in one law applied at the compliance center, as shown in Equation (13.6).

$$\begin{aligned} F_{cx} &= K_x\left(x_{ref,C} - x_C\right) + D_x\left(\dot{x}_{ref} - \dot{x}_C\right) \\ F_{cz} &= K_z\left(z_{ref,C} - z_C\right) + D_z\left(\dot{z}_{ref} - \dot{z}_C\right) - mg \\ M_c &= K_\theta\left(\theta_{ref} - \theta_m\right) - D_\theta\,\omega_m \end{aligned} \tag{13.6}$$

with $D_x = 2\zeta\sqrt{K_x m}$, $D_z = 2\zeta\sqrt{K_z m}$ and $D_\theta = 2\zeta\sqrt{K_\theta I_C}$, where $I_C = m[r_g^2 + (h_G - h_C)^2]$. The measured position of C is $x_C = x_m + h_C \sin\theta_m$ and $z_C = z_m - h_C\cos\theta_m$, the reference is carried up to C the same way from the programmed tip position, $\theta_{ref} = 0$ is the robot's vertical, and the rates of C include the $\omega_m h_C$ terms. Each axis is the impedance law of Equation (1.14), written at C. Each force is clipped to $\pm F_{max}$, the moment to $\pm M_{max}$, and all three are held for the control period. Position and force are never demanded on the same axis at once, because a force loop that owns an axis does so by writing that axis's reference.

The term $-mg$ carries the payload's weight as a force at C but not its moment about C. The weight acting $h_G - h_C$ above C therefore adds a moment $mg(h_G - h_C)\sin\theta$ in the direction of the lean, which with C at the tip is a negative rotational stiffness of $mgh_G = 1.77\,\mathrm{N\,m/rad}$. The contacts feel $K_\theta - 1.77\,\mathrm{N\,m/rad}$, which is 8.2 at the default of 10 and 1.2 at 3. Below 1.77 N m/rad the peg would be unstable when held in free space, which is why the rotational stiffness has a lower limit of 2 N m/rad, where the contacts feel 0.23 N m/rad.

13.4.2 Stiff and Impedance Programs

The controller settings are listed in Table 13.3. The stiff program is the robot's own position servo, the law of Equation (13.6) with $K_x = K_z = 200\,\mathrm{N/mm}$, $K_\theta = 500\,\mathrm{N\,m/rad}$, $\zeta = 0.7$ and C at the tip, and it follows the programmed path to the target depth without reading the force sensor. Its damping, sized for the full 2 kg, acts on a tip whose lateral effective mass is 0.47 kg, so the servo is stable only while the control period stays short. In model runs with the control period changed, a 0.1 mm step of the lateral command overshoots by 4.8 % at the executed 0.3 ms, 5.6 % at 0.5 ms and 11.8 % at 0.6 ms, and at 0.7 ms it diverges, reaching 411 % within 0.3 s. The effective mass is the cause. With

Table 13.3 Controller Settings. Level B Is Shown to a Beginner, A Only in Advanced Mode, and C Marks a Fixed Constant of the Code

Symbol	Meaning	Default	Range	L
K_x, K_z	lateral, axial stiffness	3, 5 N/mm	0.2–100, 1–50	B
K_θ	rotational stiffness	10 N m/rad	2–200	B
ζ	damping ratio	0.7	0.2–2	A
h_C	compliance center height	0 mm	0–110	A
v_{app}	approach speed	10 mm/s	1–50	B
h_{slow}	slow-down height	1 mm	0–3	B
F_{SP}, $F_{x,SP}$	force setpoints	12, 0 N	2–40, ±10	B
$K_{c,F}$, $\tau_{I,F}$	insertion loop	0.2 mm/N, 0.02 s	0–0.5, 0.005–1	B
$K_{c,x}$, $\tau_{I,x}$	lateral loop	0.03 mm/N, 0.1 s	0–0.5, 0.01–2	B
F_{lim}	overload limit	40 N	5–150	B
τ_f	force signal filter	3 ms	0–30	A
F_c, F_s	contact, search force	2, 5 N	0.5–10, 1–20	A
A	search amplitude	1 mm	0–2.5	A
n_r	retries allowed	2	0–5	A
K, K_θ	stiff servo	200 N/mm, 500 N m/rad	—	C
v_{fast}	transfer speed	50 mm/s	—	C
v_{ret}, v_{lat}	retract, lateral speed	25, 20 mm/s	—	C
$\dot{z}_{max}$	rate limit on z_{ref}	25 mm/s	—	C
T_{sw}	search period, 3 per search	0.8 s	—	C

the center of mass lowered to 50 mm above the tip, which raises that mass to 1.00 kg, the same step overshoots 6.6 % at 1 ms and 11.8 % at 1.1 ms, and it diverges only at 1.2 ms.

The impedance program follows the same path through the springs of Table 13.3 and does not read the force sensor either. Its path is programmed F_{SP}/K_z beyond the target, to 12.4 mm at the defaults, so that the axial spring pushes with the 12 N setpoint when the tip reaches 10 mm. Both programs command the reference down at v_{fast} until it is h_{slow} above the nominal surface and at v_{app} after that, with the reference speed fed forward. Both declare the insertion done when the measured tip comes within 20 µm of the target, and both give up once the reference has rested at the end of its path for 1 s.

13.4.3 The Contact-Aware Sequence

The sequence runs on the same impedance law and reads the wrist sensor. It approaches like the impedance program, with its path continuing to 13 mm, until the measured insertion force exceeds the contact threshold F_c, and it records the contact height. It then presses with the insertion loop at the search force F_s for 0.12 s, and if the peg has not dropped 0.3 mm below the contact height in that time, it searches. The search sweeps the lateral reference sinusoidally about its value at contact with a period of 0.8 s and an amplitude that grows to A over the first 1.2 s, and it ends either when the peg drops 0.3 mm or, after three periods, as a failed alignment. The insertion phase regulates the insertion force to F_{SP} and the side force to $F_{x,SP} = 0$, so that the hole centers the peg. The sequence retracts to 1 mm above the contact height at v_{ret} and tries again, up to n_r times, when the measured force magnitude exceeds F_{lim}, when the peg gains less than 0.1 mm in 0.4 s while inserting, or when a search fails, and any attempt is abandoned after 6 s.

13.4.4 The Force Loops

Both force loops use the shared velocity-form PID of Section 1.6.2, Equations (1.19) to (1.22), here with proportional and integral action only, $\tau_D = 0$. The output u of each loop is the reference it owns, sampled every Δt_c. Back-calculation acts on the integral term alone with the tracking time constant $T_t = \tau_I$, and the clamp and the rate limit apply to the emitted output, never to the running sum of the proportional increments. When a loop is not in control it tracks, with its integral set so that its output equals the reference in use, and the entry into force control is bumpless by construction.

The insertion loop has $a = +1$ and writes z_{ref} within −5 to 13 mm at no more than 25 mm/s, with $K_{c,F} = 0.2\,\text{mm/N}$ and $\tau_{I,F} = 0.02\,\text{s}$. It is reverse acting in the convention of Section 1.3, because a

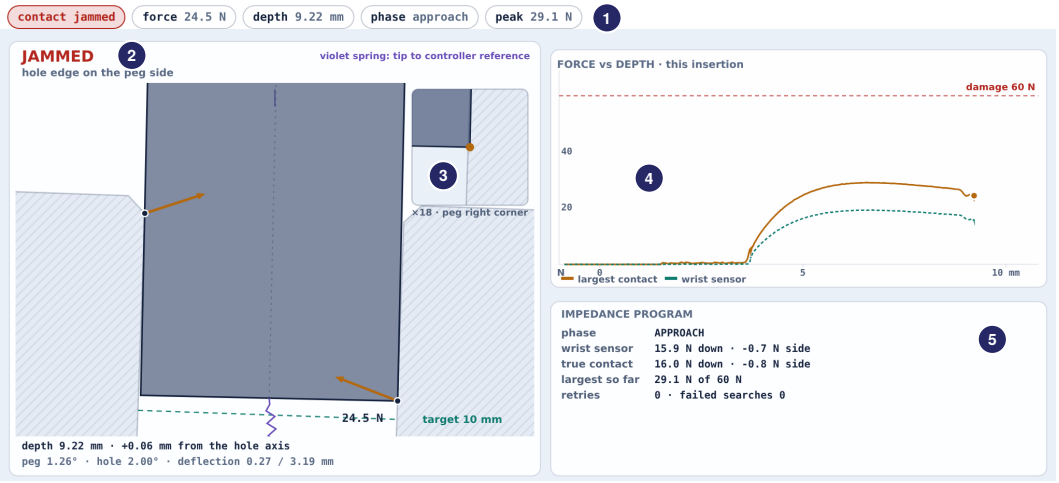

Figure 13.3 Scenario 4, Friction increase, as named in the dropdown, at $t = 2.00\,\mathrm{s}$ of a run of the impedance program with nothing changed, played at 0.25 times real speed, cropped to the plant view. The badge row (1) reads the contact jammed and a peak of 29.1 N. The cutaway (2) reads JAMMED, with force arrows at the left entry edge on the side of the peg and at the peg's right corner on the far wall, where the largest contact carries 24.5 N and the magnifier (3) looks. The zigzag spring below the peg runs 3.19 mm down to the controller's reference, and the peg leans 1.26° in a hole tilted 2.00°, 9.22 mm deep. The force against depth plot (4) rises from about 3.7 mm to a broad plateau near 29 N, with the wrist reading below it. The program panel (5) reads the insertion force at 15.9 N against 16.0 N true and a side force of –0.7 N, with the largest force so far 29.1 N of 60.

reading below the setpoint moves the reference deeper, which raises the force. The lateral loop has $a = -1$, direct acting, and writes x_{ref} within $\pm 3\,\mathrm{mm}$, with $K_{c,x} = 0.03\,\mathrm{mm/N}$ and $\tau_{I,x} = 0.1\,\mathrm{s}$, so it moves the reference in the direction of the measured side force and yields to whichever wall pushes the peg. While a force loop owns an axis, the reference speed fed forward on that axis is zero.

In manual mode the operator sets the z and x commands directly, with jogs of 0.1 and 0.05 mm, and they become the references of Equation (13.6) at the stiffness of the selected strategy while both force loops track. A jog of 0.1 mm against the part with the stiff servo is about 20 N of contact force.

Three models are kept apart. The simulation model is that of Section 13.3. The controllers carry no model of the plant beyond the payload weight, the nominal hole at $x = 0$ and $z = 0$ and the programmed target, and nothing is identified from data. The estimates of Section 13.5 are quasi-static balances and a linearized force loop that the application never uses.

13.5 What Sets the Contact Force

Four short estimates built from the constants of Tables 13.2 and 13.3 account for most of what the scenarios show, each checked against the simulator, and Figure 13.3 shows the instant the third examines.

13.5.1 Stiffness and the Chamfer

A peg whose corner lands on the chamfer $x_o - c$ outside the hole wall enters only after the chamfer has pushed it that far sideways, which costs $K_x(x_o - c)$ against the lateral spring. A corner sliding down a 45° chamfer with friction μ receives a force whose vertical part is $(1+\mu)/(1-\mu)$ times its

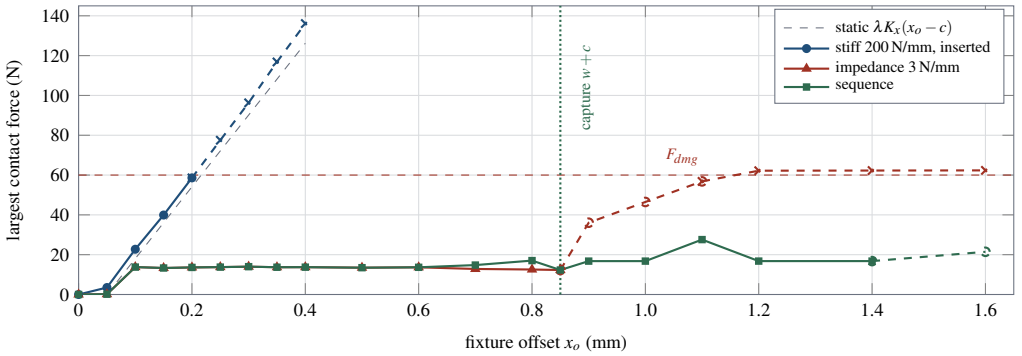

Figure 13.4 Largest contact force against fixture offset for the three controllers at their defaults, from case 2 with the offset changed. Filled circles are the stiff servo, filled triangles the impedance program and filled squares the sequence, all on solid lines while the part is inserted. On the dashed continuations, crosses are damaged parts and open circles missed holes. The thin dashed line is the static estimate. The stiff servo follows the static estimate of Equation (13.7) and passes the 60 N limit between 0.2 and 0.25 mm, the impedance program is set by the landing until the capture range ends, and the search carries the sequence to 1.4 mm.

lateral part, so the contact carries λ times the lateral force it delivers, as shown in Equation (13.7).

$$\lambda = \frac{\sqrt{(1+\mu)^2+(1-\mu)^2}}{1-\mu} = \frac{\sqrt{2}\sqrt{1+\mu^2}}{1-\mu} \tag{13.7}$$

At $\mu = 0.2$ the factor is $\lambda = 1.80$ with a vertical to lateral ratio of 1.50, and the damage instant of case 2 matches both, with 60.5 N at the left corner of which 33.6 N is lateral and 50.3 N vertical. The stiff servo can therefore absorb only $F_{dmg}/(\lambda K_x) = 0.17$ mm of sideways correction, and it damages the part once $x_o > c + F_{dmg}/(\lambda K_x) = 0.22$ mm. In case 2 the peg meets the chamfer 0.45 mm down and is damaged 0.62 mm down, after that same correction of 0.17 mm. The measured largest force is 22.7 N at 0.1 mm, 58.6 N at 0.2 mm and 77.5 N, with the part damaged, at 0.25 mm, against the static 18.0, 54.1 and 72.1 N, and its slope between 0.1 and 0.4 mm is 378 N/mm against $\lambda K_x = 361$ N/mm, as shown in Figure 13.4. The 3 N/mm spring needs 1.05 N for the correction at 0.4 mm, so its largest force is the landing, which stays between 12.3 and 14.0 N for every offset from 0.1 mm to the capture range of $w + c = 0.85$ mm.

> **TRY IT IN THE APP** **2 · Lateral Offset**
>
> Load the case and press Run. Watch the force arrow on the chamfer and the largest contact force: the stiff servo damages the part, and the force climbs to 136.1 N before the run stops itself. Press Reset, set Controller to Impedance before pressing Run, and the same part goes in with a largest force of 13.7 N.

Pause. Try Exercise 13.1 before reading on.

13.5.2 The Landing

A landing is a moving mass stopped by the part's stiffness, and it is over in about a millisecond. The corner that meets the chamfer presents an effective mass along the chamfer normal given by Equation (13.8).

$$m_n = \left[\frac{1}{m} + \frac{(\mathbf{r}_{GP} \times \mathbf{n})^2}{m r_g^2}\right]^{-1}$$
$$F_{peak} \approx v \cos 45^\circ \sqrt{k_n m_n} \tag{13.8}$$

where $\mathbf{r}_{GP}$ runs from G to the landing corner and v is the landing speed. The corner's 0.83 kg stopped without damping by 5 N/µm gives 1.44 N per mm/s. A least-squares line through the origin fits the simulator at 1.35 N per mm/s over landing speeds from 2.9 to 59.6 mm/s, from 4.24 N at the slowest to 80.2 N at the fastest, and the first impact of the stiff servo peaks at 13.3 N from 9.9 mm/s on the same line, because nothing in any of the three controllers moves in a millisecond. In case 5 the force stays above half its 13.5 N peak for 0.7 ms, while the 2 ms sensor and the 3 ms filter pass at most 2.6 N to the controller.

🕮 **Pause.** Try Exercise 13.2 before reading on.

13.5.3 The Jam

In two-point contact the clearance forces the peg to lean $\theta = \varphi - 2c/\ell$, where $\ell = Z - w$ is its engagement below the entry edge. The moment that the rotational spring and the weight leave unbalanced, $(K_\theta - mgh_G)\theta$, has to come from the two side contacts as a couple of normal forces N about ℓ apart, and their friction, with the small vertical share of the tilted normals, resists the push of the axial spring. The peg stops at the first depth where Equation (13.9) holds.

$$\begin{aligned} 2\mu N + N(\sin\varphi - \sin\theta) &\geq K_z\left(z_{goal} - Z\right) \\ N = \frac{(K_\theta - mgh_G)\,\theta}{\ell}, \qquad \theta &= \varphi - \frac{2c}{\ell} \end{aligned} \tag{13.9}$$

where $z_{goal} = z_T + F_{SP}/K_z = 12.4$ mm is the end of the impedance path. For case 4 the balance places the jam at 9.19 mm against the 9.22 mm measured, and at μ of 0.30, 0.40 and 0.50 at 9.73, 8.61 and 7.40 mm against 9.75, 8.66 and 7.52 mm. Over a grid of 36 runs, K_θ from 3 to 20 N m/rad and μ from 0.2 to 0.8, it separates the 21 jams from the 15 insertions without an error. Its depth is within 0.43 mm for K_θ up to 10 N m/rad and within 0.92 mm at 20, where larger forces penetrate up to 9.4 µm and the peg leans 0.18° less than the clearance forces, the $2\delta/\ell$ that the penetration allows.

At the instant of Figure 13.3, 2.00 s into case 4, the balance can be read off the contacts. The peg leans 1.26° at 9.22 mm, the spring and the weight leave 0.181 N m about the tip, and over the 8.5 mm between the contacts that gives $N \approx 21.3$ N against the measured 23.1 and 21.8 N. Both sites sit on their friction cones, with tangential forces of 8.04 and 7.61 N, ratios of 0.348 and 0.349 at $\mu = 0.35$, and their friction of 15.6 N together with 0.3 N from the tilted normals balances the 15.9 N that the axial spring's 3.19 mm deflection supplies. Equation (13.9) points to two remedies the cases test, a lower friction coefficient and a softer rotational spring, and case 4 adds a third, keeping the compliance center at the tip, where the lateral spring cannot turn the peg.

🕮 **Pause.** Try Exercise 13.3 before reading on.

13.5.4 The Delay Margin of the Force Loop

While the sequence presses, the contact is a thousand times stiffer than the axial spring, so the contact force follows the reference through K_z and the loop gain is $K_{c,F}K_z$. With the sensor lag, the filter and the delay, the open loop is Equation (13.10).

$$L(s) = K_{c,F}K_z\left(1 + \frac{1}{\tau_{I,F}s}\right)\frac{\mathrm{e}^{-\tau_d s}}{(\tau_s s + 1)\left(\tau_f s + 1\right)} \tag{13.10}$$

At the default $K_{c,F} = 0.2$ mm/N the gain $K_{c,F}K_z$ is 1.0. With the contact mode and the sampling the code applies (the wrench held over each 0.3 ms period, the filter updated from the current sample and the integral carried from the previous scan) included, the crossover is 116 rad/s and the phase margin without delay is 124°, so the loop tolerates 18.7 ms of delay and would oscillate with a period of 54 ms at that limit. The margin of a purely integral force loop, Equation (1.16), does not apply here. With $K_{c,F}K_z$ at 1.0 the proportional term alone keeps the magnitude of the controller and spring at one or more at every frequency, so the sensor lag and the filter, not the integral gain, set the crossover.

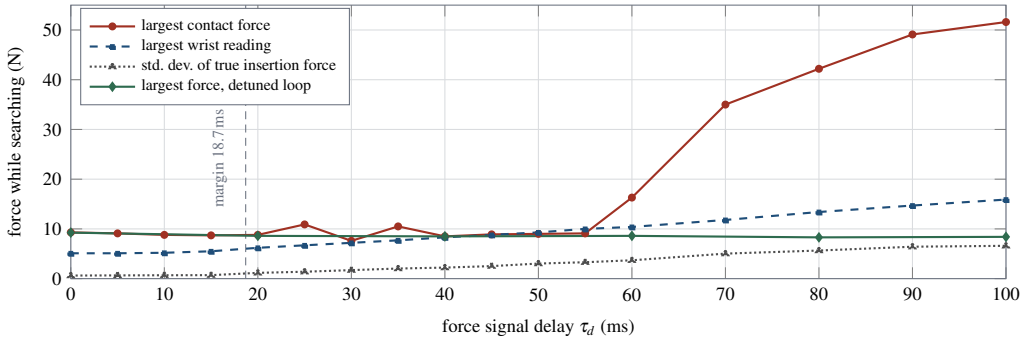

Figure 13.5 The searching press of case 7 against the force signal delay. The default loop starts to oscillate near the 18.7 ms margin of Equation (13.10), the rate limit holds the swing below the search force until about 55 ms, and beyond it the landings grow while the wrist reading stays near a third of them. Filled circles on a solid line are the largest contact force, filled squares on a dashed line the largest wrist reading, triangles on a dotted line the standard deviation of the true insertion force, and diamonds the largest force of the detuned loop, 0.1 mm/N and 0.05 s, which is unaffected.

Halving the gain to 0.1 mm/N and lengthening the integral time to 0.05 s moves the crossover to 11.5 rad/s with 117° of margin, which tolerates 177 ms.

The simulator bears the margin out. In a counterfactual run with the reference rate limit removed from a copy of the controller, the standard deviation of the true insertion force while searching rises from 0.69 N at a delay of 14 ms to 2.66 N at 19 ms, and at 20 ms the press loses contact, lifting off 19 times and landing at up to 43.1 N. As shipped, the 25 mm/s limit on z_{ref} bounds the same oscillation. It starts at the same delay, 0.70 N at 15 ms and 1.16 N at 20 ms with a period of 58 ms, and grows only as fast as the limit lets the reference swing, to 3.03 N at 50 ms. Lift-offs begin at 55 ms, where the swing exceeds the 5 N search force, and at 80 ms the landings reach 42.2 N while the largest reading is 13.4 N, as shown in Figure 13.5. The detuned loop holds the deviation between 0.62 and 0.71 N, and the peg in contact 99 % of the time, out to 100 ms. A delay inside the loop also sets the chatter boundary of the milling cut in Volume 7, Chapter 13, where each tooth cuts the surface the previous tooth left.

Table 13.4 Scenarios and Their Presets

#	Scenario	Preset
1	Perfect alignment	stiff servo
2	Lateral offset	stiff servo, $x_o = 0.4$ mm
3	Angular error	impedance, $\varphi = 2°$
4	Friction increase	as case 3, $\mu = 0.35$
5	Fragile component	impedance, $x_o = 0.5$ mm, $F_{dmg} = 10$ N
6	Excessive stiffness	impedance, $x_o = 0.5$ mm, $K_x = 40$ N/mm
7	Delayed force signal	sequence, $x_o = 1.2$ mm, $\tau_d = 80$ ms
8	Early contact	sequence, $\Delta h = 2$ mm, $x_o = 0.3$ mm
C	Challenge	stiff servo, 100 parts, $\mu = 0.35$, offsets within ± 1.2 mm and tilts within $\pm 2°$

Table 13.5 Measured Results. Forces Are the Largest Single Contact Force

Case and setting	Measured
1, aligned	in at 1.14 s, no contact
1, x_o 0.1 / 0.2 / 0.3 mm	22.7 / 58.6 N / damaged, 96.3 N
2, stiff	damaged at 0.223 s, 136.1 N
2, impedance / sequence	13.7 N, in at 1.14 / 0.63 s
2, 1.2 mm, impedance / sequence	damaged, 62.2 N / in, 1.59 s
2, 1.2 mm, sequence, $K_x = 1$ N/mm	missed, 2 failed searches
3, $K_\theta = 10$	in at 1.33 s, 27.8 N, 1.32°
3, $K_\theta = 50$ / stiff	damaged 4.86 / 4.25 mm in, 67.1 / 117.2 N
4, $K_\theta = 10$	stuck at 9.22 mm, 29.1 N
4, K_θ 3 / 5	in at 1.20 / 1.28 s, 10.3 / 14.7 N
4, $h_C = 110$ mm	stuck at 4.96 mm, 56.3 N
5, v_{app} 10 / 7 / 5 / 3 mm/s	13.5 (damaged) / 9.4 / 7.0 / 4.2 N
5, insertion time at 7 / 5 / 3 mm/s	1.62 / 2.25 / 3.72 s
6, K_x 40 / 20 / 10 N/mm	96.7 (damaged) / 18.3 / 13.5 N
7, $\tau_d = 80$ ms	42.2 N true, 13.4 N read, 60 % in contact
7, 80 ms, 0.1 mm/N, 0.05 s	8.3 N, 99 %, in at 1.72 s
7, $\tau_d = 0$	9.3 N, 99 %, in at 1.59 s
8, as loaded	damaged at 38 ms, 80.2 N
8, $h_{slow} = 2.5$ mm	12.5 N, in at 0.63 s
C, stiff / impedance / sequence	5 / 66 / 87 of 100 inserted
C, sequence, $K_\theta = 3$	100 of 100, 30.2 N, 157.9 s

TRY IT IN THE APP **7 · Delayed Force Signal**

Press Run and watch both traces of the force-versus-depth plot while the sequence presses and searches. The landings reach 42.2 N while the wrist reading never passes 13.4 N. Press Reset, set Force loop gain to 0.1 mm/N and Force loop integral time to 0.05 s before pressing Run, and the press stays in contact 99 % of the time.

13.6 Scenarios and Measured Results

Table 13.4 lists the cases, which build from case 1 to the challenge, with the preset each loads. Each case is credited, as Section 1.6.4 describes, when the learner makes the attempts its briefing names, with every attempt judged at its end, whether the part is inserted, stuck, missed or damaged. The challenge is judged on the case as given, and changing a model parameter or a disturbance, or firing an extra disturbance button, voids the attempt. The values of Table 13.5 were read from headless runs of the shipped application at its default seed, with the scenario loaded as shipped and any named change made before the run. Figure 13.6 draws the largest contact forces of that table as bars.

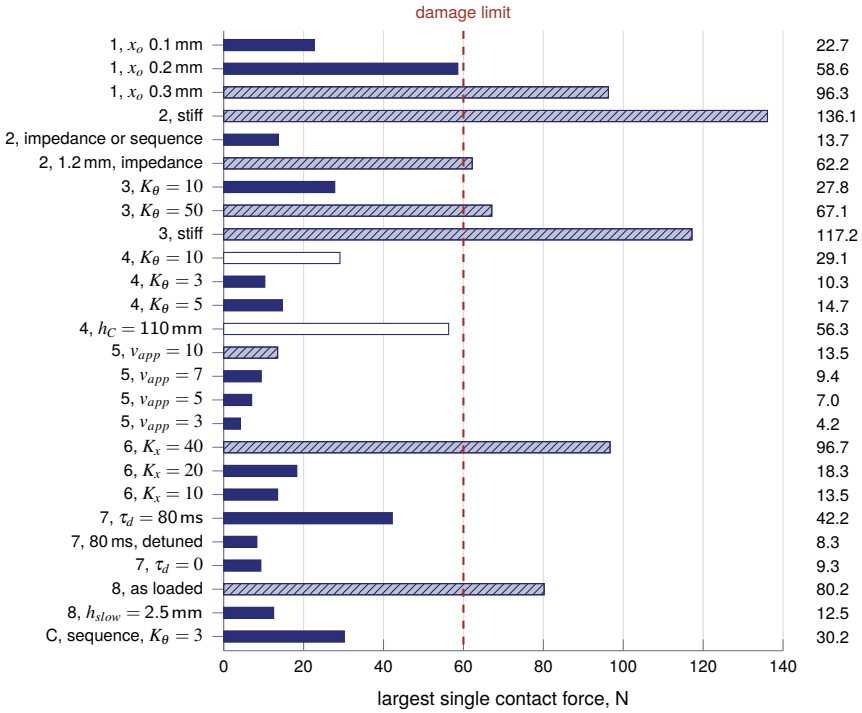

Figure 13.6 The largest contact forces of Table 13.5 drawn as bars, by case and setting (K_θ in N m/rad, K_x in N/mm, v_{app} in mm/s). Hatched bars damaged the part, open bars left it stuck, and solid bars did neither. The dashed line is the 60 N damage limit, and the part of case 5 cracks above 10 N. Notice that every case has a setting whose bar ends far left of the line, and that the bars past it belong to stiff settings, to the fast landing of case 8 and to a hole beyond the chamfer's reach.

13.6.1 Position Control and Its Tolerance

Aligned, the stiff servo puts the peg in at 1.14 s without touching the part, and shifted 0.3 mm it damages the part 0.72 mm down the chamfer as the force reaches 96.3 N. Case 2 sets the offset at 0.4 mm. The part is damaged 0.223 s into the run, and the force keeps climbing to 136.1 N before the run stops. The impedance program inserts the same part at 1.14 s with 13.7 N, and the sequence at 0.63 s, because it detects the contact and regulates the insertion force instead of creeping down at the approach speed. Below 0.35 mm the sequence never reads more than its 2 N threshold, keeps to the path and takes the same 1.14 s.

Beyond the capture range the order reverses. At 1.2 mm the impedance program keeps pushing on the top face until the peg cocks 2.2° onto the edge of the chamfer and the part is damaged at 62.2 N, while the sequence detects the contact, sweeps until the peg drops into the chamfer and inserts it at 1.59 s with 16.8 N. The search reaches offsets of 1.4 mm and misses at 1.6 mm. A lateral spring of 1 N/mm is too soft to search, because the 1 N of friction under the 5 N search force takes a millimeter of deflection, the whole amplitude, and the sequence gives up after two failed searches at the 6 s attempt limit.

13.6.2 Tilt, Friction and the Rotational Spring

Case 3 tilts the fixture 2°. At 10 N m/rad the contacts turn the peg, and it reaches the target at 1.33 s with 27.8 N, leaning 1.32° in the 2.00° hole. The diametral clearance of 0.1 mm over the 9.2 mm between the entry edge and the tip allows 0.62° of the 0.68° difference, and contact penetrations of up to 5.4 µm take up the rest. At 50 N m/rad the same turn leaves 5.9 times the unbalanced moment

for the contacts to supply, and the part is damaged 4.86 mm in with 67.1 N, and the stiff servo, at 500 N m/rad, damages it 4.25 mm in with 117.2 N.

Case 4 dries the parts to $\mu = 0.35$ and changes nothing else. The right corner first meets the tilted wall 1.54 mm down, the second contact forms at 3.7 mm, and the largest force peaks at 29.1 N at 6.6 mm before easing toward the jam at 9.22 mm. The impedance reference reaches the end of its path at 1.38 s, the scene reads JAMMED from 1.80 s, and the program gives up at 2.38 s. At 3 N m/rad the peg goes in at 1.20 s with 10.3 N, and at 5 N m/rad with 14.7 N, while moving the compliance center to the wrist, 110 mm up, jams it at 4.96 mm with 56.3 N.

> ⚠ **OPEN CHALLENGE** **Built-In Compliance or Programmed Compliance**
>
> A cell can take up position error with a mechanism, such as a compliant wrist, or with an active compliant controller. Large position or force uncertainties can be accommodated by the controller, while smaller ones can be removed by making the mechanism more compliant. When to tune the controller and when to change the mechanism is still an open question of peg-in-hole research [19]. Progress would be a design rule that splits a cell's tolerance budget between the hardware and the control law before either is built.

13.6.3 The Force Budget

Case 5 lands the impedance program 0.5 mm off on a part whose latch cracks above 10 N. At the default 10 mm/s the landing reaches 13.5 N and cracks it 0.35 mm down the chamfer on the first step over the limit, while the controller's reading never exceeds 2.6 N. The peak falls in proportion to the speed, to 9.4 N at 7 mm/s, 7.0 N at 5 mm/s and 4.2 N at 3 mm/s, and the insertion time rises to 1.62, 2.25 and 3.72 s, so the force budget rather than the geometry sets the cycle time. Case 6 raises the lateral stiffness to 40 N/mm at the same offset. The spring pulls the peg back toward the programmed x against the hole's left wall as it enters, the contact forces turn it, and 5.15 mm in it leans 1.57° with its left corner digging into the wall past the limit. The largest force reaches 96.7 N before the peg reaches the target. At 20 N/mm the peak is 18.3 N and at 10 N/mm it is 13.5 N, which is the landing itself.

13.6.4 A Late Signal and an Early Contact

Case 7 gives the sequence a hole 1.2 mm off, so that it must press on the top face and search, and delays the force signal by 80 ms. With no delay the press keeps at least one site loaded above 0.05 N, the count the application logs, 99 % of the searching time, its largest force is 9.3 N, and the part is in at 1.59 s. With the delay the loop oscillates as Section 13.5 predicts. The peg is in contact only 60 % of the searching time and lifts off 12 times, the landings reach 42.2 N, and the controller's reading never passes 13.4 N, so the 40 N overload never trips. The part still goes in at 1.72 s, only because 42.2 N is under this part's limit. Detuned to 0.1 mm/N and 0.05 s, the same delay leaves the press in contact 99 % of the time with 8.3 N, and the part is in at 1.72 s.

Case 8 raises the part 2 mm, with an offset of 0.3 mm that the chamfer can catch. The slow-down still begins 1 mm above the nominal surface, so the peg meets the chamfer 38 ms into the run at 59.6 mm/s, faster than the 50 mm/s transfer because the damped spring overshoots in speed while the tool catches up with its reference. The first step over the limit carries 70.0 N, the peak is 80.2 N, and the wrist sensor has registered 7.4 N by the time the run stops. Detecting contact cannot help, because the impact is over before the signal arrives, and the impedance program, which reads no sensor at all, is damaged at the same 80.2 N. A slow-down height of 2.5 mm lands the peg at 9.5 mm/s with 12.5 N, and the part is in at 0.63 s.

13.6.5 One Hundred Parts

The challenge loads 100 parts in a row, each in a fixture with its own offset within ±1.2 mm and tilt within ±2°, on dry parts at $\mu = 0.35$, and it is met when at least 95 are inserted with no damage event. The shipped stiff program inserts 5, damages 93 and jams 2, with a largest contact force of 69.4 N, 90

parts short of the target, and the set ends after 58.2 s. The impedance program at its defaults inserts 66, with 21 missing holes beyond the capture range, 8 jammed and 5 damaged. The sequence at its defaults inserts 87, and the 13 that jam all have tilts of 1.71° or more, while every tilt it inserts is 1.69° or less, and its 36 retries repeat the motion without curing them. With $K_\theta = 3\,\mathrm{N\,m/rad}$ the sequence inserts all 100 with no damage event and no retry, with a largest contact force of 30.2 N over the whole set and a mean insertion time of 1.01 s, and the set takes 157.9 s. The last part is done at 157.7 s, and the run pauses 0.2 s later.

The four tallies are scored on the same 100 parts: the part draws come from a random stream of their own, seeded from the scenario, so the sensor noise an attempt consumes cannot change the next part. Over ten further seeds, with the four configurations again meeting the same parts on each, the stiff program inserts 2 to 6, the impedance program 59 to 70, the default sequence 80 to 88, and the suggested configuration 100 on every seed with no damage and a largest force of 28.3 to 30.1 N.

> **DID YOU KNOW? The Most Precise Program Scores Worst**
>
> From the simulator: on the same 100 parts of the challenge, the stiff servo inserts 5, the impedance program 66 and the contact-aware sequence 87. The sequence with a rotational spring of 3 N m/rad inserts all 100 (Table 13.5).

13.7 Technology in Practice

The cell of Section 13.3 reads a wrist force sensor and the joint encoders, and it moves the peg through one wrench law. Figure 13.7 draws its loops with real devices in the place of the schematic's symbols.

What Measures It

Force-based robotic assembly exists in industry mainly as force and torque sensors at the end of the arm, sold by robot makers and integrators, often with no performance specification beyond the sensing itself [11]. Sensors that read the deformation of a compliant body with strain gauges, with capacitive electrodes or with optical fibers have all been developed [20]. The wrist sensor of Section 13.3, a 2 ms lag with 0.15 N of noise ahead of the controller's 3 ms filter, stands for such a sensor, and the encoders with 0.5 µm of noise stand for the joint encoders of the arm.

> **TECHNOLOGY Sensor: Six-Axis Wrist Force Sensor**
>
> Force-based assembly cells mostly sense contact with a force and torque sensor at the end of the arm [11]. A conventional six-axis sensor measures three forces and three moments from the deformation of a single compliant structure. Strain gauges read that deformation with high accuracy, but they are sensitive to electromagnetic interference and often show cross-axis coupling, in which a load in one direction appears in the others [20]. On a trend, coupling shows as a side force that rises and falls with the push. An offset left in the reading makes the arm jerk when it switches from position to force control [10]. The model's sensor, perfectly tared, has neither fault.

What Moves It

The robot's own position loops do most of the work. Force loops closed around them see a plant that the position loops have decoupled and linearized, and the role of a passive compliance in the speed of the task and in its rejection of disturbances has been analyzed [12]. Impedance control is typically used for fast contact tasks such as polishing, metal removal and some assemblies [10]. The clipped wrench of Equation (13.6), held over each control period, stands for the arm's drives.

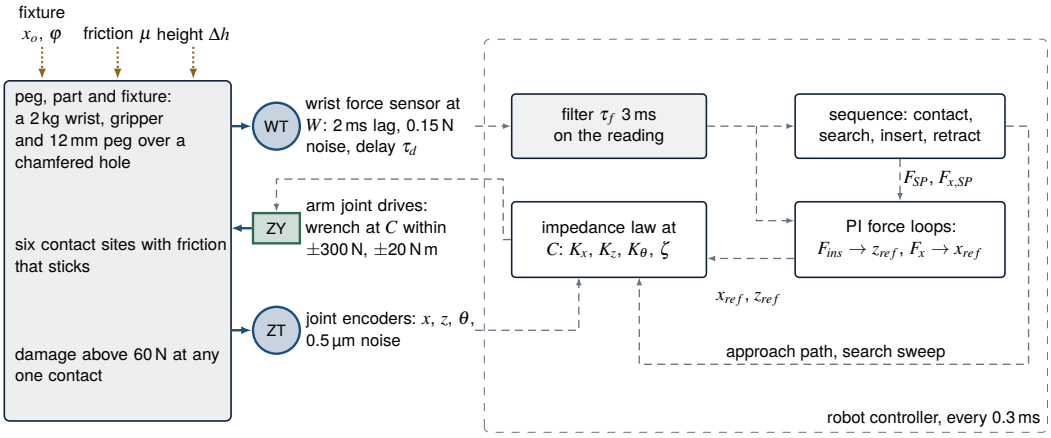

Figure 13.7 The insertion loops of the chapter with real devices in place of the schematic's symbols. The wrist force sensor WT and the joint encoders ZT feed a robot controller that runs every 0.3 ms. The sequence sets the force setpoints, the PI force loops write the references, and one impedance law at the compliance center C turns them into the wrench the arm's joint drives ZY apply. Notice that the arm itself is the final element, and that every force the controller acts on passes through the sensor's lag and the filter first.

> **TECHNOLOGY** **Final Element: The Arm as a Programmable Spring**
>
> The final element is the arm itself, driven by its joint servos. Active stiffness control turns the robot into a programmable spring, where a passive tool would carry real springs and dampers in line with the gripper. Some arms infer the applied force from the torques or the currents at each joint instead of reading a wrist sensor [10], and arms built to work beside people resolve the force at the tool from torque sensors in the joints [11]. A badly damped arm resonates, and on a trend it hops along the approach direction instead of settling [10]. The wrench law of Equation (13.6) stands for these drives.

Where the Controller Runs

In discrete manufacturing a supervisor for each cell may coordinate a machine controller and a robot controller, and supervisory and lower-level loops run with cycle times from milliseconds to minutes [21]. The controllers of this chapter act every 0.3 ms, at the fast end of that range. In a real cell the impedance law and the two force loops would sit in the robot controller, beside its position servos, and the contact-aware sequence in the robot's program. That sequence is an event-based search strategy, which moves the peg around the expected hole, randomly or in a raster or a spiral, until the peg can be pushed in [10]. The sequence of this chapter sweeps a sinusoid.

> **IN PRACTICE** **Testing Force Control Against a Load Cell**
>
> A force-controlled cell is tested against an independent load cell, not the robot's own sensor, whose reading carries its own accuracy, latency and tooling offsets [10]. In one NIST test method the robot steps to a target force from free space, 32 times in each direction tested. The peak overshoot, which indicates the impact force, the settling time and the steady-state error are recorded. Assembly itself is judged by its success rate, its time and the largest and average force applied to the parts [11]. The simulator draws the true contact force beside the reading in the same spirit.

What Is Modernizing It

Arms built to work beside people now carry torque sensors in their joints and resolve the force at the tool from them, which lets them both assemble by force and react to a collision with a person [11].

Cameras join the force sensor: vision gives the hole's position from afar but is blocked in narrow spaces, so one strategy fuses images, the measured contact force and the joint positions by Bayesian probability, which needs no large training set, and judges the result by force and position together [22]. Learning methods form the third development. Strategies learned from human demonstrations and from interaction with the environment by reinforcement learning are being compared with the contact models and compliant control that this chapter uses [19]. Test methods with independent load cells and standard artifacts, including a peg-in-hole artifact, are meant to help users select a robot for force-based work [11].

13.8 Limitations

The numbers are the simulator's own design choices, a 60 N damage limit, a 5 N/µm contact and a 2 kg end effector, and no claim about a real cell's forces, cycle time or yield can rest on them. The model is planar, with one rigid body and no arm dynamics, and the assumptions of Section 13.3.6 state where each simplification breaks.

A single physics step over the limit counts as damage, so an impact of 0.1 ms weighs the same as a sustained press. Retries repeat the motion that failed, which is why the 13 jams of the default sequence in the challenge survive 36 retries. The weight's moment about the compliance center is not compensated, which makes the rotational stiffness the contacts feel 1.77 N m/rad lower than the slider reads. The loop tiles of the scorecard score the insertion force against a setpoint of zero whenever no force loop is in control, so for the stiff and impedance programs they integrate the force the sensor reported rather than a regulation error. Robot safety functions and their certification are outside what the simulation claims to represent.

> ⚠ **OPEN CHALLENGE** **Fast or Reliable, Rarely Both**
>
> Force control has been researched for decades, yet industry has been slow to adopt it for assembly. The time to implement it and a steep learning curve weigh against it, and reliable solutions are not fast, while fast solutions are not reliable. A robot can be programmed to be an expert at one assembly, but generalizing that expertise across assemblies is difficult [10]. Learned strategies promise to generalize, but their low sample efficiency still limits them on real robots [19]. Progress would be a cell that inserts a new part quickly and reliably without being tuned for it.

13.9 Conclusion

This chapter studied a robotic peg-in-hole simulator in which the moment of contact changes what the controller can command. A planar 2 kg end effector carries a 12 mm peg into a hole with 50 µm of clearance per side, through six penalty contacts whose friction sticks and a wrist sensor the controller reads only after its lag. The robot's 200 N/mm servo tolerates a fixture offset of about 0.22 mm, as the chamfer factor of 1.80 predicts, and damages the part at 0.4 mm with the force reaching 136.1 N, while a 3 N/mm spring inserts it at 13.7 N. Landings cost 1.35 N per mm/s of landing speed whatever the controller, from 4.2 N at 3 mm/s to 80.2 N for a part met during the transfer. A 2° hole on dry parts jams 9.22 mm deep, where a moment balance predicts 9.19, and inserts at 10.3 N once the rotational stiffness drops to 3 N m/rad. The force loop loses its phase margin at 18.7 ms of delay, its rate limit postpones lift-off to 55 ms, and at 80 ms it hammers the part at 42.2 N while its reading stays under 13.4 N. Across 100 seeded parts the stiff program inserts 5 and the sequence with a soft wrist inserts 100 with a largest force of 30.2 N, and 100 on each of ten further seeds. The model omits arm dynamics, the second tilt axis and duration-dependent damage. A fuller model would add a round peg, adaptive retries and a damage-energy criterion. Within those limits the simulator lets each failure of control through a contact transition be met alone and measured.

KEY POINTS

- Once the peg touches the part, a stiff robot turns every small position error into force.
- The 200 N/mm servo tolerates a fixture offset of about 0.22 mm, while a 3 N/mm spring inserts the part shifted 0.4 mm at 13.7 N.
- A landing is over within a millisecond, before any feedback can act, so the speed at contact sets its force.
- A tilted peg on dry parts jams where a moment balance predicts, and a softer rotational spring frees it.
- With a delayed force signal the press hammers the part while the controller's reading stays calm.

Where else this applies. A delay inside the loop also sets the chatter boundary of the milling cut in Volume 7, Chapter 13, and Chapter 14 closes a force loop on tissue.

References

[1] D. E. Whitney, "Quasi-static assembly of compliantly supported rigid parts," *Journal of Dynamic Systems, Measurement, and Control*, vol. 104, no. 1, pp. 65–77, 1982.

[2] ——, "Historical perspective and state of the art in robot force control," *The International Journal of Robotics Research*, vol. 6, no. 1, pp. 3–14, 1987.

[3] M. T. Mason, "Compliance and force control for computer controlled manipulators," *IEEE Transactions on Systems, Man, and Cybernetics*, vol. 11, no. 6, pp. 418–432, 1981.

[4] M. H. Raibert and J. J. Craig, "Hybrid position/force control of manipulators," *Journal of Dynamic Systems, Measurement, and Control*, vol. 103, no. 2, pp. 126–133, 1981.

[5] N. Hogan, "Impedance control: An approach to manipulation: Part I—theory," *Journal of Dynamic Systems, Measurement, and Control*, vol. 107, no. 1, pp. 1–7, 1985.

[6] L. Villani and J. De Schutter, "Force control," in *Springer Handbook of Robotics*, B. Siciliano and O. Khatib, Eds. Cham: Springer International Publishing, 2016, pp. 195–220.

[7] S. Haddadin, A. Albu-Schäffer, and G. Hirzinger, "Requirements for safe robots: Measurements, analysis and new insights," *The International Journal of Robotics Research*, vol. 28, no. 11–12, pp. 1507–1527, 2009.

[8] D. E. Whitney, "Force feedback control of manipulator fine motions," *Journal of Dynamic Systems, Measurement, and Control*, vol. 99, no. 2, pp. 91–97, 1977.

[9] S. Eppinger and W. Seering, "Introduction to dynamic models for robot force control," *IEEE Control Systems Magazine*, vol. 7, no. 2, pp. 48–52, 1987.

[10] J. Marvel and J. Falco, "Best practices and performance metrics using force control for robotic assembly," National Institute of Standards and Technology, Gaithersburg, MD, NIST Interagency/Internal Report 7901, 2012.

[11] J. Falco, J. Marvel, R. Norcross, and K. Van Wyk, "Benchmarking robot force control capabilities: Experimental results," National Institute of Standards and Technology, NIST IR 8097, 2016.

[12] J. De Schutter and H. Van Brussel, "Compliant robot motion II. A control approach based on external control loops," *The International Journal of Robotics Research*, vol. 7, no. 4, pp. 18–33, 1988.

[13] R. Volpe and P. Khosla, "A theoretical and experimental investigation of explicit force control strategies for manipulators," *IEEE Transactions on Automatic Control*, vol. 38, no. 11, pp. 1634–1650, 1993.

[14] S. R. Chhatpar and M. S. Branicky, "Search strategies for peg-in-hole assemblies with position uncertainty," in *Proceedings 2001 IEEE/RSJ International Conference on Intelligent Robots and Systems*, vol. 3. Maui, HI, USA: IEEE, 2001, pp. 1465–1470.

[15] K. H. Hunt and F. R. E. Crossley, "Coefficient of restitution interpreted as damping in vibroimpact," *Journal of Applied Mechanics*, vol. 42, no. 2, pp. 440–445, 1975.

[16] G. Gilardi and I. Sharf, "Literature survey of contact dynamics modelling," *Mechanism and Machine Theory*, vol. 37, no. 10, pp. 1213–1239, 2002.

[17] D. Karnopp, "Computer simulation of stick-slip friction in mechanical dynamic systems," *Journal of Dynamic Systems, Measurement, and Control*, vol. 107, no. 1, pp. 100–103, 1985.

[18] P. Dupont, V. Hayward, B. Armstrong, and F. Altpeter, "Single state elastoplastic friction models," *IEEE Transactions on Automatic Control*, vol. 47, no. 5, pp. 787–792, 2002.

[19] J. Xu, Z. Hou, Z. Liu, and H. Qiao, "Compare contact model-based control and contact model-free learning: A survey of robotic peg-in-hole assembly strategies," arXiv preprint arXiv:1904.05240, 2019.

[20] S. Y. Lee, J. Y. Sim, D.-Y. Seok, Y. B. Kim, J. Shim, U. Kim, and H. R. Choi, "A redundant-sensing-based six-axis force/torque sensor enabling compactness and high sensitivity," *Sensors*, vol. 26, no. 3, p. 871, 2026.

[21] 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.

[22] L. Qin, H. Wang, Y. Yuan, and S. Qin, "Multi-sensor perception strategy to enhance autonomy of robotic operation for uncertain peg-in-hole task," *Sensors*, vol. 21, no. 11, p. 3818, 2021.

13.10 Exercises

13.1 Concept check · 2 min

In scenario 2 the fixture sits 0.4 mm from where the program aims. The stiff servo damages the part with a contact force that climbs to 136.1 N, while the impedance program with its 3 N/mm lateral spring inserts the same peg with a largest force of 13.7 N. Which statement explains the difference?

(a) The chamfer must push the peg sideways by the offset less the clearance, and that costs force in proportion to the lateral stiffness

(b) The impedance program reads the wrist sensor and slows its approach as soon as the contact force starts to rise

(c) The stiff servo approaches the part faster, so its landing on the chamfer is harder than that of the impedance program

(d) The stiff servo's damping is sized for the full mass, so its loop becomes unstable once the peg touches the chamfer

13.2 Concept check · 2 min

In scenario 8 the part is 2 mm taller than programmed, and the contact sequence, which reads the wrist sensor and retracts above its force limit, damages it at 80.2 N. A learner lowers the force limit from 40 to 5 N before Run, and the part is damaged at the same force. Why?

(a) The force limit is checked only while the force loop inserts, and this contact arrives during the approach

(b) The impact is over in about a millisecond, while the sensor lag and the filter pass only a few newtons in that time

(c) The landing force is set by the lateral spring, and lowering the force limit leaves that spring unchanged

(d) The damage limit in this scenario is 10 N, below the force of any landing the program can make

13.3 Concept check · 1 min

In scenario 4 the impedance program jams a peg 9.22 mm deep in a hole tilted 2° on dry parts ($\mu = 0.35$). Which single change, made before Run, lets it insert the peg?

(a) Lower the approach speed to 3 mm/s, so the peg enters the tilted hole more gently

(b) Raise the rotational stiffness to 50 N m/rad, so the peg is held straighter in the hole

(c) Lower the rotational stiffness to 3 N m/rad, so the contacts need a smaller moment to turn the peg

(d) Move the compliance center to the wrist, 110 mm above the tip, as on a passive wrist

13.4 Theory and calculation · 6 min

The stiff servo holds the tip with $K_x = 200\,\mathrm{N/mm}$. The clearance is $c = 50\,\mu\mathrm{m}$ per side, the chamfer is at 45°, and a single contact above $F_{dmg} = 60\,\mathrm{N}$ damages the part.

(a) For $\mu = 0.20$, find the chamfer factor λ of Equation (13.7) and the ratio of the vertical to the lateral force on the corner. *Checkpoint:* $\lambda = 1.80$ and 1.50.

(b) Find the largest sideways correction the servo can make before the contact passes the damage limit, and the fixture offset at which that happens. *Checkpoint:* 0.17 mm and 0.22 mm.

(c) Repeat (b) for the dry parts of the challenge, $\mu = 0.35$, and find the lateral force the 3 N/mm impedance spring needs at an offset of 0.4 mm. *Checkpoint:* $\lambda = 2.31$, an offset of 0.18 mm, and 1.05 N.

13.5 Theory and calculation · 6 min

The end effector has $m = 2\,\text{kg}$, a radius of gyration $r_g = 50\,\text{mm}$ about its center of mass G, which lies $h_G = 90\,\text{mm}$ above the tip, and the contact stiffness is $k_n = 5\,\text{N}/\mu\text{m}$. The peg's corner is $a = 6\,\text{mm}$ from the peg axis and lands on a 45° chamfer whose normal points up and toward the hole.

(a) Find the effective lateral mass the tip presents to a sudden sideways push. *Checkpoint:* 0.47 kg.

(b) The lever arm of the chamfer normal about G is $(h_G - a)/\sqrt{2}$. Find the effective mass m_n of Equation (13.8) and the landing force per unit of approach speed. *Checkpoint:* 0.83 kg and 1.44 N per mm/s.

(c) The part of scenario 5 cracks above 10 N. Find the fastest safe approach speed from (b) and from the slope of 1.35 N per mm/s that the chapter fits to the simulator. *Checkpoint:* 6.9 mm/s and 7.4 mm/s.

13.6 Theory and calculation · 7 min

Scenario 4 tilts the hole $\varphi = 2°$, with clearance $c = 50\,\mu\text{m}$ per side, a chamfer of $w = 0.8\,\text{mm}$, $\mu = 0.35$, $K_z = 5\,\text{N}/\text{mm}$, $K_\theta = 10\,\text{N m}/\text{rad}$, $mgh_G = 1.77\,\text{N m}/\text{rad}$ and the end of the impedance path at $z_{goal} = 12.4\,\text{mm}$. Use Equation (13.9).

(a) At the depth $Z = 9.22\,\text{mm}$, find the engagement ℓ and the lean θ that the clearance forces. *Checkpoint:* $\ell = 8.42\,\text{mm}$ and $\theta = 1.32°$.

(b) Find the normal force N, the resisting force and the push of the axial spring, and say whether the peg stops there. *Checkpoint:* $N = 22.5\,\text{N}$, 16.0 N against 15.9 N, so it stops.

(c) Repeat (b) with $K_\theta = 3\,\text{N m}/\text{rad}$. *Checkpoint:* $N = 3.4\,\text{N}$, 2.4 N against 15.9 N, so it does not.

13.7 App challenge · 8 min

Load scenario 2, Lateral offset, which loads the stiff servo with the fixture 0.4 mm off. Use the default seed and let each run stop itself.

(a) Run the scenario as shipped. Read the Peak contact force and Damage events tiles, and the time of damage from the note. *Checkpoint:* 136.1 N, one damage event, at 0.223 s.

(b) Press Reset, set Controller to Impedance and run again. Read the Peak contact force and Insertion time tiles. *Checkpoint:* 13.7 N and 1.14 s.

(c) Press Reset, set Controller to Sequence and run again. Read the same tiles, and say why this run is faster. *Checkpoint:* 13.7 N and 0.63 s.

13.8 App challenge · 8 min

Load scenario 4, Friction increase. Use the default seed and let each run stop itself.

(a) Run the scenario as shipped. Read the depth at which the note reports the peg stuck, the Peak contact force tile, and from the exported CSV file the first time at which the column `contact_state` reads `JAMMED`. *Checkpoint:* 9.22 mm, 29.1 N, 1.81 s.

(b) Press Reset, set Rotational stiffness to 3 and run again. Read the Insertion time and Peak contact force tiles. *Checkpoint:* 1.20 s and 10.3 N.

(c) Press Reset, set Rotational stiffness to 5 and run again. Compare the three runs with the jam balance of Equation (13.9). *Checkpoint:* 1.28 s and 14.7 N.

13.9 App challenge · 8 min

Load scenario 7, Delayed force signal, in which the sequence presses and searches on a hole 1.2 mm off while its force reading arrives 80 ms late. Use the default seed, let each run stop itself and export the CSV file after each.

(a) Run the scenario as shipped. Read the Peak contact force, Insertion time and Retries tiles. In the file, find the largest value of `Fins_meas_N`, the reading the controller acts on. *Checkpoint:* 42.2 N, 1.72 s, 0 retries, and 13.3 N within 0.2 N.

(b) Press Reset, set Force loop gain to 0.1 mm/N and Force loop integral time to 0.05 s, and run again. Read the Peak contact force tile, and in the file the largest `contact_peak_N` over the rows whose `phase` reads `SEARCH`. *Checkpoint:* 16.8 N on the tile and 8.3 N while searching.

(c) Reload the scenario with the reload button, open Show Model, set Force signal delay to 0 ms and run again. Read the Insertion time tile. *Checkpoint:* 1.59 s.

13.10 Application area · 4 min

A cell engineer reviews the guarded move of scenario 8. The parts may stand up to 2 mm taller than programmed, the transfer runs at 50 mm/s, the approach at 10 mm/s, and a single contact above 60 N damages the part. Use the landing slope of 1.35 N per mm/s that the chapter fits to the simulator.

(a) Find the highest landing speed under the damage limit, and say whether any slow-down height below the height error can be made safe by a faster controller. *Checkpoint:* About 44 mm/s. No.

(b) The slow-down height is raised from 1 to 2.5 mm. Estimate the extra time per part when the part is at its nominal height, ignoring acceleration. *Checkpoint:* About 0.12 s.

13.11 Application area · 3 min

A supplier offers a passive compliant wrist whose center of compliance lies at its mounting flange, about 110 mm above the peg tip. The parts are those of scenario 4. What does the chapter predict for this wrist, where should the compliance center be instead, and which assumption of the model must be checked before the result is trusted for a round peg?

13.12 Review the agent · 5 min

An AI assistant is asked whether the insertion force loop of the sequence tolerates a 25 ms fieldbus delay and replies as follows. “The force follows the reference through the 5 N/mm axial spring, so the loop is an integrator of gain $g = K_z K_{c,F}/\tau_{I,F} = 5 \times 0.2/0.02 = 50\,\mathrm{s}^{-1}$. Its delay margin is $\pi/(2g) = 31$ ms, so the shipped tuning is safe up to about 30 ms and the 25 ms link needs no retuning.” Find the flaw and say how you would check it.

CHAPTER 14

Surgical Robot Simulator for Control Through a Fixed Port onto Moving Tissue

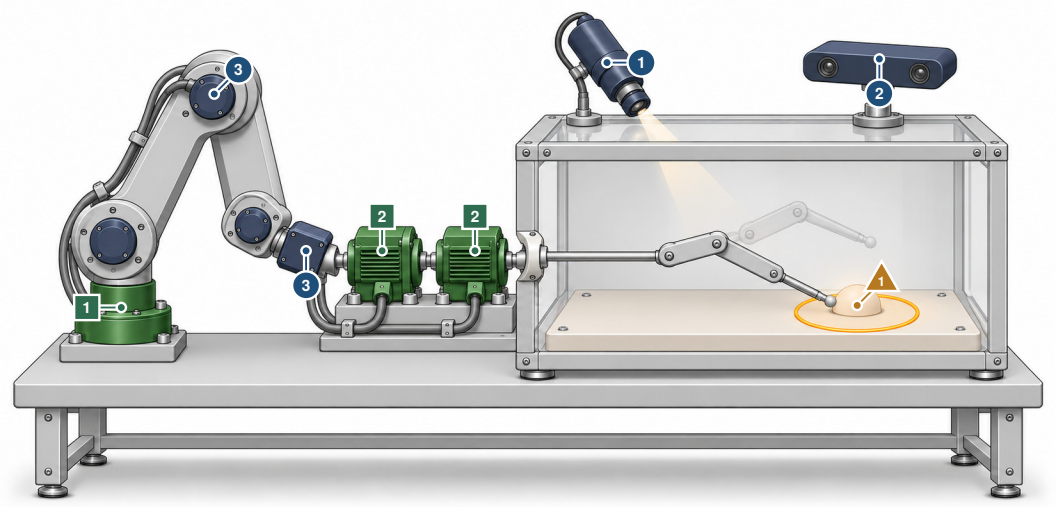

The controller moves 1 arm drives that pivot the shaft about the port and slide it in 2 drives that pull the wrist's cables

The controller reads 1 endoscope camera, the late picture the surgeon steers by
2 tissue tracker that reports the organ's motion late 3 joint encoders on the arm

What disturbs it 1 the organ's motion with each breath or heartbeat

Notice that the instrument pivots about the fixed port in the box wall, that no sensor reads the force at the tip, and that the target moves while the camera and the tracker see it late. Illustration, not to scale.

IN PLAIN TERMS A Robot Between the Hand and the Patient

The main idea. In keyhole surgery, long tools pass through small cuts in the belly. The tip then moves the opposite way to the hand. A surgical robot undoes that flip, and it can shrink and smooth the hand's motion before it reaches the tip.

The challenge. Each helper has a price. Smoothing out the shake of the hand adds a delay that the surgeon does not expect, and the organs keep moving with every breath, so the robot has to guess where they will be.

Why it matters. Knowing what each layer fixes and what it costs tells a designer which errors a robot can remove and which it cannot.

THREE WAYS TO READ THIS CHAPTER

⊘	**Explore**	12 min	the first page, the two history segments (Section 14.2), the fact on robotic surgery's spread (Section 14.5), the try-it segments of cases 3 and 8 and the fact on the error budget (Sections 14.5 and 14.5.6) 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 measured cases (Section 14.5), Technology in Practice (Section 14.6), Limitations (Section 14.7) and the open challenges

The control question A surgical robot inverts the fulcrum of the port, scales and filters the hand, holds tissue with a force loop on an estimate and forecasts the patient's breathing. How much of the tip error does each layer remove, and what does each one cost?

Learning objectives

- Compute the lever ratio of a hand-held instrument and the tip speed that a pivot rate limit allows at a given insertion.
- Explain how a tremor filter trades tremor at the tip against a lag that the surgeon's forward model does not expect.
- Predict the relaxation, creep and tear of a standard linear solid under a position hold and under a force loop on a biased estimate.
- Compare ventilator feedforward, an adaptive harmonic predictor and MPC on a breathing organ, and read a factorial error budget.
- Recognize a disturbance that no setting can reject, as on the beating heart.

Prerequisites Sections 1.6.2 (velocity form), 1.4.8 (force loop through a stiffness), 1.4.6 (report age), 1.4.9 (checking a monitor against the truth) and 1.4.10 (error budgets). Section 1.6.3 is needed only for the MPC option, and Chapter 13 is a gentler first force loop.

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

See also Volume 5, Chapter 8 for the delayed force channel this chapter leaves out, and Volume 5, Chapter 13 for late fixes applied at the instant they describe. Volume 7, Chapter 13 for a vibration that repeats every tooth, and Volume 7, Chapter 12 for a feedforward canceler built on an assumed model.

14.1 Introduction

A straight grasper held by its handle passes through a port in the abdominal wall with 120 mm of shaft inside the patient and 240 mm outside. When the hand moves 20 mm to the right, the tip moves 9.90 mm to the left. Pushed 80 mm further in, the same 20 mm of hand motion moves the tip 24.65 mm the other way. The port is a fulcrum, and its lever arms change every time the surgeon changes the depth. A surgical robot removes that reversal by solving, many times a second, for the pivots and the insertion that put the tip where the hand went. It then adds layers of its own: a scale and a filter between the hand and the command, a loop on a force it estimates rather than measures, a rule that keeps the tool out of a region on a map, and a forecast of where a breathing organ will be when the command reaches it. Each of these runs in the simulator of this chapter, and each has a cost that the simulator measures.

The plant corrects four beliefs that a block diagram of a teleoperator leaves in place. The first is that the robot copies the hand. It does not, and near the port the inversion that undoes the fulcrum asks the shaft for more rate than it has, so the tip falls up to 51.5 mm behind its command. The second is that more filtering is safer. A 2 Hz tremor filter cuts the tremor at the tip to 12 µm but adds 120 ms of lag that the surgeon's eye-hand loop does not expect, and three dissection strokes overshoot into an artery that a 6 Hz filter never touches. The third is that a force held below the tissue's tear force is safe. A force loop at 2.5 N on a flap that tears at 3 N lets the flap creep until its stretch tears it at 13.3 s. The fourth is that a faster sensor defeats a fast disturbance. Behind a tracker 40 ms late or 200 ms late, no option holds the tip within 1 mm of a coronary artery on a beating heart, and the best does 1.53 mm.

The problem belongs to a wider class, motion through a kinematic constraint with a human in the loop, acting on a load that relaxes and that moves with a disturbance whose period, or whose phase, is measured. A periodic exogenous signal is the object of repetitive control, which places a model of the period inside the loop [1]. Roll eccentricity in a strip mill, blade passing in a rotating machine, the tooth-passing vibration of the end mill of Volume 7, Chapter 13, a cyclic process with a recipe clock and the beating of a heart share that structure, and the surgical form has been treated with predictive control on the respiratory cycle [2] and on the heartbeat [3]. The fixed remote center is a holonomic

constraint on the shaft, built into the mechanism of the arms that hold laparoscopic instruments [4, 5], and it turns a hand's motion into joint rates that grow without bound as the tip nears the port.

However, the published motion compensation, virtual fixtures, tremor cancellation and autonomous suturing are each demonstrated on their own platform [2, 6–8], and a student cannot change one layer of a surgical robot at a time and see what each costs on the same seed. This chapter apportions how much of a surgical robot's tip error comes from the fulcrum, the console, the tissue and the patient's motion, and which control element removes which, using a browser simulator with eleven procedure excerpts and a challenge. The remote center, the standard linear solid, the three-phase needle force, harmonic prediction by recursive least squares, the forbidden-region fixture and supervised autonomy are all established [5, 7–10], and the chapter uses them as they stand. Its analysis is a two-level factorial error budget, measured on the running engine over five seeds, in which the surgeon's visuomotor delay costs nothing on a still mark and 86 % of the mean-square error on a breathing one, and an adaptive harmonic predictor removes 91 % of that interaction. A forecast taken on the ventilator's phase makes the tracker's latency almost irrelevant from 40 to 200 ms, and without compensation not one of 240 bites on a breathing kidney is kept, whatever else is removed. The model has 28 states, and its controllers read a measurement packet and never the simulator's truth. The surgeon is scripted with an internal forward model, and the recursive least-squares update is written in symmetric form because the textbook form diverged. Every case and the challenge carry a rule that awards credit when it is met (Section 1.6.4).

Section 14.2 traces the domain from stereotaxy to supervised suturing. Section 14.3 gives the instrument, the tissue, the patient's motion, the sensors and the scripted surgeon as simulated, and Section 14.4 gives the controllers as coded: the console, the virtual fixtures, the force loop, motion compensation and the autonomous suturing subtask. Section 14.5 reports the cases as measured, with the error budget in Section 14.5.6. Section 14.6 names the devices of real surgical robots, Section 14.7 states the limitations, Section 14.8 summarizes the results, and the exercises follow in Section 14.9.

14.2 From Stereotaxy to Supervised Suturing

Robots entered the operating room as positioners, and the history segment below follows the first of them and the fulcrum effect that laparoscopy put on the hand.

> ⧖ HISTORY **From Positioner to Fulcrum, 1988 to 1998**
>
> In 1988 a robot for stereotactic brain surgery guided by computed tomography was reported, and the report was about improving the robot's absolute positioning accuracy [11]. Laparoscopic surgery put a different burden on the hand, because the abdominal wall acts as a fulcrum on every instrument that passes through it. An ergonomic study of 1998 had novices make defined incisions under laparoscopic laboratory conditions. Those who worked with the image's vertical axis inverted made more incisions, and more correct ones, than those with a normal image, and the difference persisted over 50 trials. The study concluded that the fulcrum effect has a major detrimental influence on the acquisition of endoscopic skill [12].

The first robots inside the abdomen held the camera or an instrument for the surgeon. A telerobotic assistant for laparoscopy, reported in 1995, was built around a remote-center-of-motion mechanism so that the instrument pivoted about the port by construction [4]. By 2003 a review of computer-integrated surgery could describe a family of such systems, from robots that execute a plan made on images to assistants that work beside the surgeon, many of them on a remote center [5]. The console-based platforms in clinical use gave the surgeon no sense of the force at the tip, so force control stayed visual [13]. In 2001 a cholecystectomy was performed on a patient in Strasbourg by surgeons in New York over a dedicated network, a demonstration that the console and the patient could be an ocean apart [14]. The delay and passivity questions that long distance raises are the subject of Volume 5, Chapter 8 and are left out here.

The control literature turned next to what the hand cannot do. Adaptive cancellation of physiological tremor, a band near 10 Hz that a human cannot suppress, was demonstrated for microsurgery in 1998 [6]. Breathing and the heartbeat move the organs under the instruments by millimeters to centimeters, and predictive control with a repetitive model of the respiratory cycle was shown to cancel organ motion on a robot holding an instrument in a live animal [2]. The heart was harder, because its motion is faster, has many harmonics and changes from beat to beat, and algorithms for beating-heart surgery were built on a model of the heart's motion and a measured biological signal such as the electrocardiogram [3]. Rules imposed on the robot's own command, forbidden regions and guidance paths, were surveyed in 2014 under the names active constraints and virtual fixtures [7].

Autonomy came last, and the second history segment below follows its first two steps on the bowel. Over the same decade robotic surgery spread from urology to the common procedures of general surgery, which a statewide study of hospitals traced from 2012 to 2018 [15]. Outcome work set clinical limits that a control engineer can read as budgets. In partial nephrectomy, each minute of warm ischemia with the renal hilum clamped was associated with a higher risk to the kidney's function [16], which is the reason for the clock of cases 9 and 12. The simulator puts these steps on one robot so that each can be tried alone or with the others.

> ⧖ HISTORY **The Robot Takes the Needle, 2016 and 2022**
>
> In 2016 a robot performed supervised autonomous surgery on soft tissue in living pigs, in an open surgical setting. A three-dimensional and near-infrared fluorescent imaging system tracked the tissue, a program planned the sutures of an intestinal anastomosis, and the outcomes were better than those of expert surgeons and of robot-assisted surgery in the tissues tested [8]. In 2022 the task was done laparoscopically, in pigs followed for a week afterward. The robot detected the breathing motion of the tissue and told the operator to start a replanning step, and the operator selected among suture plans the robot had generated. On the scale of autonomy that runs from pure teleoperation, level 0, to full autonomy, level 5, a system that generates strategies and relies on a human to select or approve one is level 3, conditional autonomy [17].

14.3 Instrument, Tissue and Patient as Simulated

The simulated platform is one docked teleoperated robot with a console, a wristed instrument driven by the surgeon's hand, a third arm that grasps and retracts tissue, and an endoscope, all through fixed ports, as shown in Figure 14.1. The frame is the endoscope's, with x to the right of the image, y up the image toward the head and z into the patient, and the operative surface lies at $z = 0$ in the organ's own frame, which the organ's motion carries. Every tissue property, limit, motion, latency and time budget is a teaching value, both the model dialog and the Instructions page label them so, and no number applies to any patient, instrument or product. Table 14.1 classifies the signals.

14.3.1 The Instrument Through Its Port

Each instrument is a rigid shaft through its port $\mathbf{p}_{port}$, with six joints, the shaft's pitch θ and yaw ϕ about the port, the insertion d from the port to the wrist, the roll ψ about the shaft, and a two-axis wrist with pitch α on a link $l_{w1} = 5\,\text{mm}$ and yaw β on a jaw $l_{w2} = 10\,\text{mm}$. The forward kinematics are shown in Equation (14.1).

$$\begin{aligned} R_s &= R_0\,R_y(\theta)\,R_x(\phi), \qquad \mathbf{s} = R_s\,\mathbf{e}_z \\ \mathbf{p}_w &= \mathbf{p}_{port} + d\,\mathbf{s} \\ \mathbf{p}_{tip} &= \mathbf{p}_w + R_s R_z(\psi) R_x(\alpha)\left(l_{w1}\mathbf{e}_z + R_y(\beta)\,l_{w2}\mathbf{e}_z\right) \end{aligned} \tag{14.1}$$

where R_0 is a fixed rotation that points the shaft's zero at the operative field. The shaft passes through the port by construction, so the remote center is a property of the joint coordinates rather than a constraint the integrator must hold. The application's tests check it anyway, on the integrated states of every case, from two independently written expressions of the shaft, and the worst distance from the port to the shaft line is zero within the 0.05 mm bound.

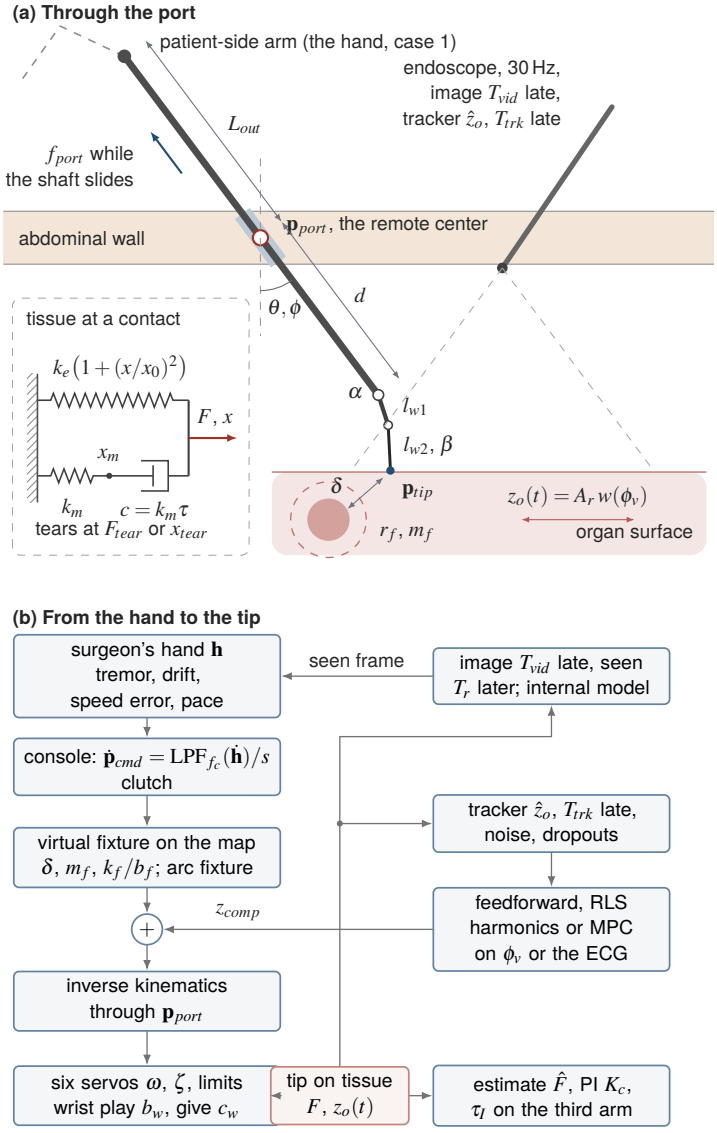

Figure 14.1 The instrument, the tissue and the control stack. (a) A section through the abdominal wall along the instrument, in the layout of the application's Port view. The shaft pivots by θ and ϕ about the port $\mathbf{p}_{port}$, a remote center that does not move, and slides by the insertion d. The wrist pitches by α on the link l_{w1} and yaws by β on the jaw l_{w2} to the tip $\mathbf{p}_{tip}$. Held by hand (case 1) the handle sits L_{out} outside the port. The organ moves by $z_o(t)$ on the ventilator's phase ϕ_v, a forbidden region of radius r_f and margin m_f lies at a distance δ from the tip, and the endoscope and the tracker see everything a latency late. The inset is the standard linear solid of every contact, with the equilibrium spring k_e, the Maxwell arm k_m and c with its internal state x_m, the stretch x and the force F. (b) The chain from the hand to the tip and the signals each controller reads. Not to scale.

Table 14.1 Signals by Role. Ranges Are Those the Interface Offers

Symbol	Role	Default	Range
$\mathbf{p}_{tip}$	CV, the tip in the task axis	—	—
$\mathbf{p}_{cmd}$	SP of the tip, from the console	—	—
$\hat{F}$	CV of the force loop (estimate)	—	—
F_{SP}	SP of the force loop	0.8 N	0.1 to 3.5 N
$\mathbf{h}$	MV, hand h_x, h_y, h_z (Manual)	0	±150 mm
x_s	MV, pull depth (Position)	4 mm	0 to 15 mm
$\mathbf{q}_c$	MV, six joint commands per arm	from IK	ranges
$\mathbf{q}$, $\dot{\mathbf{q}}$	states, 2×6 joints and rates	per case	ranges
x_m	states, three Maxwell arms	0	—
ϕ_h	state, the heart's phase	0	—
$\hat{z}_o$	measured, tracker report	100 ms late	40 to 200 ms
ϕ_v, t_R	measured, ventilator phase, ECG R wave	—	—
$\mathbf{q}$, $\dot{\mathbf{q}}$	measured, joint encoders	—	—
F, z_o	hidden, true force and organ	—	—
A_r	DV, respiration (peak to peak)	per case	0 to 20 mm
A_c, HR	DV, heart motion and rate	8 mm, 72 bpm	0 to 15 mm
—	DV, beat-to-beat variability	±5 %	0 to ±10 %
e_{reg}	DV, fixture registration error	0	±8 mm
—	DV, instrument age	0 %	0 to 100 %
f_{port}	DV, port friction	0.3 N	0 to 1 N
—	DV, tracker dropout (2 s)	—	event

The inverse takes a commanded tip, the jaw's commanded direction $\mathbf{u}$ and the roll, and returns the joints. The wrist angles follow in closed form from $\mathbf{u}$ in the shaft's frame, and the shaft's direction and insertion follow from where the wrist must then be. The two depend on each other only through the first wrist link, 5 mm against a shaft of 20 to 250 mm, so the fixed-point iteration of Equation (14.2) converges geometrically, at about the ratio of the link to the insertion.

$$\mathbf{s}^{(k+1)} = \frac{\mathbf{p}_2 - l_{w1}\,\mathbf{e}_1\big(\alpha(\mathbf{s}^{(k)})\big) - \mathbf{p}_{port}}{\big\|\mathbf{p}_2 - l_{w1}\,\mathbf{e}_1\big(\alpha(\mathbf{s}^{(k)})\big) - \mathbf{p}_{port}\big\|} \tag{14.2}$$

where $\mathbf{p}_2 = \mathbf{p}_{tip} - l_{w2}\mathbf{u}$ is the wrist's yaw joint and $\mathbf{e}_1$ is the direction of the first wrist link. Over 4000 random poses the worst case took 19 passes to 10^{-13} m, and the forward and inverse kinematics round-trip within 3×10^{-14} m. A sideways tip speed v needs a pivot rate of about v/d, so the pivots saturate as the tip comes back toward the port.

Case 1 has no robot. A straight instrument of 360 mm is held by its handle, which lies L_{out} outside the port on the line through the tip, and inside the port lie $L_{in} = 360\,\text{mm} - L_{out}$. A small lateral motion Δh of the handle moves the tip by Equation (14.3).

$$\Delta x_{tip} = -\frac{L_{in}}{L_{out}}\,\Delta h \tag{14.3}$$

At $L_{in} = 120$ mm the ratio is -0.500 and at 200 mm it is -1.250, the values the application's tests check. A 20 mm move gives 9.90 mm and 24.65 mm, a little less than the ratio because the handle swings on an arc about the port.

14.3.2 Joint Servos and the Cable-Driven Wrist

Every joint of both arms follows its command through the second-order servo of Equation (14.4), copied with its provenance from the six-axis arm of the shared framework, whose kinematic chain is a different structure.

$$\begin{aligned} \ddot{q}_i &= \mathrm{sat}_{a_{max,i}}\big(\omega^2(q_{c,i} - q_i) + 2\zeta\omega\,(\dot{q}_{ff,i} - \dot{q}_i)\big) \\ |\dot{q}_i| &\le v_{max,i}, \qquad q_{min,i} \le q_i \le q_{max,i} \end{aligned} \tag{14.4}$$

where $\omega = 2\pi \cdot 5\,\text{Hz}$, $\zeta = 0.8$, and the rate feedforward $\dot{q}_{ff,i}$ is the command's own difference over the 2 ms control period. The rate limits are 90 °/s on pitch and yaw, 0.2 m/s on insertion and 180 °/s on roll and the wrist, the acceleration limits are five times the rate limits per second, the shaft's angles stay within $\pm 70°$, the wrist's within $\pm 90°$ and the insertion between 20 and 250 mm. Approaching a range end, a drive starts to brake once stopping would need 80 % of its acceleration limit, and a hard stop after each step holds the range.

The wrist is driven by cables, so the jaw's delivered pitch differs from the servo's by a backlash and by a compliance under the load across the jaw, as shown in Equation (14.5).

$$\alpha_{tip} = \text{play}(\alpha, b_w) - \text{sign}(\dot{\alpha})\,\frac{c_w}{l_{w2}}\,F_t \tag{14.5}$$

where play is the play operator of half-width b_w, which grows linearly from 0.5° new to 3° at the end of the instrument's life, $c_w = 0.3\,\text{mm/N}$ is the give at the jaw's tip, and F_t is the needle's drag along its path. Both are applied once per control period as a kinematic offset, never inside the derivative function. The needle's point is 28 mm from the wrist's pivot, so one degree of play is about 0.5 mm at the point.

14.3.3 Tissue at the Contacts

Tissue is lumped at the points of contact. Each held or compressed tissue is a standard linear solid [9], an equilibrium spring that stiffens with stretch in parallel with a Maxwell arm, as shown in Equation (14.6).

$$\begin{aligned} F &= k_e\,x\big(1 + (x/x_0)^2\big) + k_m\,(x - x_m) \\ \dot{x}_m &= \frac{x - x_m}{\tau}, \qquad \tau = \frac{c}{k_m} \end{aligned} \tag{14.6}$$

where x is the stretch the instrument imposes and x_m the state of the Maxwell arm. Under a held stretch the force relaxes from the instantaneous stiffness $k_e + k_m$ toward k_e and loses the share $k_m/(k_e + k_m)$ of its linear part with time constant τ. Under a held force the stretch creeps with the longer time constant $\tau(k_e + k_m)/k_e$. A tissue tears when $F > F_{tear}$ or $x > x_{tear}$, tested after each step and latched. The bladder flap of case 4 has $k_e = 150\,\text{N/m}$, $k_m = 250\,\text{N/m}$, $\tau = 4\,\text{s}$, $x_0 = 15\,\text{mm}$, $F_{tear} = 3\,\text{N}$ and $x_{tear} = 10\,\text{mm}$, so a held stretch loses 0.625 of its linear force and a held force creeps with a time constant of 10.7 s. The flap's stretch is the third arm's travel from the grasp along the pull. The suture line under tension uses $k_e = 120$ and $k_m = 180\,\text{N/m}$, $\tau = 3\,\text{s}$, $x_0 = 6\,\text{mm}$ and $F_{tear} = 1.5\,\text{N}$. Bowel between the stapler's jaws uses $k_e = 300$ and $k_m = 600\,\text{N/m}$, $\tau = 15\,\text{s}$ and $x_0 = 2\,\text{mm}$, and the jaws close 3 mm of bowel to a gap of 1.5 mm in 3 s. The staple-line stress index is the force at firing over the force the same gap holds once fully relaxed.

A curved needle of radius $r_n = 7\,\text{mm}$, five eighths of a circle held 75° from its swage, follows a reduced form of the three phases of needle insertion [10]. The surface deflects ahead of the point up to $F_{pk} = 1.0\,\text{N}$ at 2 mm, the point punctures, and the shaft then drags $F_{fr} = 0.4\,\text{N}$. Any motion of the needle's center away from where it entered cuts the tissue sideways, as shown in Equation (14.7).

$$F_l = \frac{\|\Delta\mathbf{c}\|}{1/k_{lat} + c_w\left((r_n + l_{w2})/l_{w2}\right)^2} \tag{14.7}$$

where $\Delta\mathbf{c}$ is the center's departure and $k_{lat} = 600\,\text{N/m}$ is the tissue across the needle, in series with the wrist's give carried out to the needle's center. The bite tears when F_l exceeds the suture line's 1.5 N. A bite is scored on the larger of its entry and exit errors against the plan, six bites 3 mm apart and 3 mm deep. Vessels that must not be touched are capsules, a segment and a radius, pushing back at 400 N/m of penetration. A contact counts on entry and a force above 0.5 N is an injury, which latches and stops the run.

14.3.4 The Patient's Motion

The operative organ moves rigidly by $z_o(t)$ along a fixed direction, as shown in Equation (14.8).

$$\begin{aligned} z_{resp}(t) &= A_r\, w(\phi_v), \qquad \phi_v = 2\pi t / T_v \\ z_{card}(t) &= A_c\, \frac{c(\phi_h) - c_{lo}}{c_{span}}, \qquad \dot{\phi}_h = \frac{2\pi}{RR} \\ c(\phi) &= \textstyle\sum_{k=1}^{3} a_k \cos(k\phi + \vartheta_k) \end{aligned} \tag{14.8}$$

where the breath shape w rises over the first third of the cycle as a raised cosine (an inspiratory to expiratory ratio of 1:2) and returns over the rest as an exponential in the phase with rate 1.27 that reaches zero at the cycle's end, so w runs from 0 to 1 and A_r is the peak-to-peak motion. The kidney moves $A_r = 12\,\text{mm}$ every $T_v = 4\,\text{s}$, mostly along the image's y axis. The heart's harmonics have $a_k = (1, 0.45, 0.2)$ and $\vartheta_k = (0, 0.8, 1.9)$ rad, c_{lo} and c_{span} normalize the sum to a peak-to-peak of one, and $A_c = 8\,\text{mm}$ at 72 beats per minute. The interval RR of each beat is drawn uniformly within $\pm 5\,\%$ of the mean, and an R wave marks the start of each beat. The heart's phase is the one state of the motion.

14.3.5 What the Controllers Can Know

The endoscope captures a frame every $1/30$ s and shows it $T_{vid} = 80\,\text{ms}$ later. A tissue tracker reports the organ's displacement from the same frames with its own latency T_{trk} (100 ms by default), Gaussian noise of 0.2 mm per axis, and nothing while blood or smoke hides the surface. The tip force is not measured. The estimate is the drive's effort less the servo's own model, so it carries the friction the model does not know, as shown in Equation (14.9).

$$\begin{aligned} \hat{F} = \mathrm{LPF}_{20\,\text{ms}}\Big(& F + f_{port} \tanh\left(-\dot{d}/v_s\right) \\ & + f_{cab} \tanh\left(|\dot{\alpha}|/\omega_s\right) + n_F \Big) \end{aligned} \tag{14.9}$$

where $f_{port} = 0.3\,\text{N}$ is the port's Coulomb friction along the shaft, $f_{cab} = 0.15\,\text{N}$ the wrist cables', $v_s = 0.5\,\text{mm/s}$, $\omega_s = 0.05\,\text{rad/s}$ and n_F is Gaussian noise of 0.02 N, through a first-order filter of 20 ms. A shaft sliding outward reads high and one sliding inward reads low, and at rest the bias vanishes. Every controller reads a measurement packet with exactly fourteen fields, the frame shown now, the frame the surgeon perceives, the newest tracker report, the two force estimates, the joint encoders, the ventilator's phase, the last R wave and its interval, the stapler's jaw force and the occlusion flag. A test wraps the packet in a proxy and fails if a controller touches anything else, and the simulator's truth reaches only the scene and the 12 exported columns that begin with `truth_`.

14.3.6 The Scripted Surgeon

In Auto mode a scripted surgeon moves the hand toward the current target of a task script by a visuomotor loop on the frame shown $T_r = 180\,\text{ms}$ ago, itself 80 ms old. A position loop with 1 Hz of bandwidth through 0.26 s of dead time is unstable, by the delay margin of an integrating loop in Equation (1.16), so the surgeon carries an internal forward model of the tip motion they have intended, $\mathbf{I}(t)$, and judges the tip by Equation (14.10).

$$\begin{aligned} \hat{\mathbf{p}} &= \mathbf{p}_{seen}(t_{cap}) + \mathbf{I}(t) - \mathbf{I}(t_{cap}) \\ \dot{\mathbf{I}} &= \mathrm{sat}\big(K\,(\mathbf{p}_{target} - \hat{\mathbf{p}}),\; v_h\,\pi_s/s\big) \end{aligned} \tag{14.10}$$

where t_{cap} is the capture time of the frame the surgeon sees, $K = 2\pi\pi_s\ \text{s}^{-1}$ with the pace π_s (1 by default), s is the motion scale and $v_h = 0.15\,\text{m/s}$ is the hand's speed ceiling. The model covers the image's latency and the surgeon's own reaction and expects no lag from the robot. A slow tremor filter adds lag that the model does not know, and the surgeon pushes on toward a tip the image shows

Table 14.2 Model Parameters. A Is Adjustable in the Model Dialog, D a Disturbance, F Fixed in the Source

Symbol	Meaning	Value	
l_{w1}, l_{w2}	wrist link, jaw	5, 10 mm	F
$\omega/2\pi$, ζ	servo bandwidth, damping	5 Hz, 0.8	C
v_{max}	pitch and yaw, insertion, roll and wrist	90 °/s, 0.2 m/s, 180 °/s	F
c_w	wrist give at the jaw's tip	0.3 mm/N	A
b_w	wrist play, new to end of life	0.5 to 3°	D
k_e, k_m, τ	flap	150, 250 N/m, 4 s	F, A
x_0, F_{tear}, x_{tear}	flap	15 mm, 3 N, 10 mm	F, A, A
F_{exp}	tension the dissection needs (case 4)	0.7 N	F
—	suture line k_e, k_m, τ, F_{tear}	120, 180 N/m, 3 s, 1.5 N	F, A
—	bowel k_e, k_m, τ	300, 600 N/m, 15 s	F, A
F_{pk}, F_{fr}, k_{lat}	needle peak, drag, lateral	1.0 N, 0.4 N, 600 N/m	F, F, A
r_n	needle radius	7 mm	F
—	vessel contact stiffness, injury	400 N/m, 0.5 N	F
A_r, T_v	kidney motion, ventilator period	12 mm, 4 s	D
A_c, HR	epicardial motion, heart rate	8 mm, 72 bpm	D
T_{vid}, T_r	video latency, reaction	80 ms, 180 ms	A, F
T_{trk}	tracker latency	100 ms	D
—	tracker noise per axis	0.2 mm	F
f_{port}, f_{cab}	port, cable friction in $\hat{F}$	0.3 N, 0.15 N	D, F
K, v_h	surgeon gain, hand speed	2π s^{-1}, 0.15 m/s	F
—	tremor, drift, speed error	0.3 mm, 0.7 mm, 8 %	F
Δt	physics and control step	2 ms	F

C: a setting on a control card. The servo, the flap's relaxation time and both tear forces, the bowel's relaxation time and the lateral stiffness are on the model dialog. Every value is a teaching value.

short of its target. The hand then moves s times the intended tip motion with three imperfections, a tremor of 0.3 mm RMS from a resonator at 10 Hz with a damping ratio of 0.2 (a half-power band of 8 to 12 Hz), a drift of $0.7\,\text{mm}\,\pi_s$ RMS with a 1 s correlation time, and a speed error of $8\,\%\,\pi_s$ with a 0.1 s correlation time. A needle drive rotates the hand's wrist at $1.2\,\text{rad/s}\,\pi_s$ and makes 60 % of the translation the rotation needs. On a breathing organ the surgeon starts each drive in the quiet tail of expiration, read from the ventilator's phase. In Manual mode the hand follows the student's sliders or arrow keys at up to 0.15 m/s and carries the same tremor from the same generator. The surgeon's parameters are illustrative and describe no surgeon.

14.3.7 Parameters and Numerics

A fourth-order Runge–Kutta step of 2 ms integrates the physics (Section 1.6.1), the console, the inverse kinematics and the force loop run every 2 ms, the surgeon and the tracker at the 30 Hz frame, and model predictive control every 10 ms. The 28 continuous states are the six joints and rates of each arm, the three Maxwell arms and the heart's phase. Everything discrete happens after the integration, in this order, the joint stops, the hand-held lever of case 1, the wrist's offsets, the needle's puncture, exit, tear and hand-off, vessel contacts, incursions, the flap's tear and slip, the stapler, the R waves, the dropouts and the ischemia clock. Halving the step keeps the conclusions of cases 3, 6 and 8. Every random draw comes from its own seeded stream, so a case replays bit for bit after Reset. Headless, the model runs 24 000 steps per second with model predictive control on case 8 and 166 000 on case 1, and in Chrome the eight-minute challenge keeps 20 × real time while drawing. The parameters are listed in Table 14.2.

14.3.8 Assumptions

The joints are servo-level, with no rigid-body dynamics, so a contact force moves the instrument only through the wrist's give. A real arm's inertia and its own controller would add lag and coupling that matter faster than the 5 Hz servo. One instrument and the third arm are simulated, while the endoscope, the second instrument and the patient-side arms are drawn. The needle's hand-off to the second instrument and its return between bites are timed steps. The ports are fixed remote centers,

Table 14.3 Controller Settings, Defaults and Ranges. Cases Ship Their Own Values Where Stated

Symbol	Setting	Default	Range
s	motion scale, hand to tip	3	1, 2, 3, 5
f_c	tremor filter cutoff	6 Hz	1 to 30 Hz
—	clutch	engaged	released
π_s	scripted surgeon's pace	1.0	0.5 to 2.0
K_c	force loop gain	10 mm/s per N	1 to 50
τ_l	force loop reset time	2 s	0.2 to 20 s
F_{max}	force ceiling	2.0 N	0.2 to 5 N
—	force loop mode	Force	Position
k_f	fixture stiffness (adv.)	2000 N/m	200 to 5000
—	virtual fixture, arc fixture	off, on	on, off
—	motion compensation	none	ff, adaptive, MPC
N_h	harmonics	3	1 to 5
λ	forgetting factor	0.99	0.95 to 0.999
—	autonomy, map	off, tracked	run or halt, perfect
t_w	wait before firing	0 s	0 to 60 s

and the abdominal wall's breathing is drawn around them rather than applied, which ignores the force the wall puts on the cannula. Tissue is lumped at the contacts, so strain is not redistributed, a tear has no location and a needle does not steer. The needle is driven by pitching the jaw in the wrist's plane, where a real needle driver rolls it. The organ moves as a rigid body with no deformation, the tracker sees that displacement only, and the ECG gives the heart's phase from its R waves alone. There is no haptic channel to the hand, no bleeding and no smoke beyond the tracker's dropouts. Each simplification keeps the one-dimensional mechanism a lesson needs and leaves out what would decide a real operation.

14.4 The Controllers as Coded

Every 2 ms the controller builds the measurement packet, moves the hand, applies the console, the fixtures, the autonomy, the force loop and the compensation, and solves the inverse kinematics, in that order. Its settings are listed in Table 14.3.

14.4.1 The Console

The console differentiates the hand over the control period, filters the velocity with a second-order Butterworth low-pass and divides it by the scale, as shown in Equation (14.11).

$$\mathbf{p}_{cmd,k} = \mathbf{p}_{cmd,k-1} + \frac{\Delta t}{s}\,\mathrm{LPF}_{f_c}\!\left(\frac{\mathbf{h}_k - \mathbf{h}_{k-1}}{\Delta t}\right) \tag{14.11}$$

The filter is discretized by the bilinear transform with the cutoff prewarped, so its −3 dB point is exactly f_c, and it passes $1/\sqrt{1+(f/f_c)^4}$ of a component at f. At 6 Hz it passes 0.339 of a 10 Hz tremor and delays a 1 Hz motion by 37.9 ms, and at 2 Hz it passes 0.040 and delays the same motion by 120.3 ms (the group delay is 132.4 ms). Scaling the velocity rather than the position means that releasing the clutch, recentering the hand and engaging again does not move the tip. The needle's rotation follows the hand's wrist one to one through the same filter.

14.4.2 Virtual Fixtures

A forbidden region is a capsule of radius $r_f = 8$ mm around a structure, with a margin $m_f = 6$ mm, placed where the map says the structure is. The fixture acts on the command increment $\Delta\mathbf{p}$, as shown in Equation (14.12).

$$\begin{aligned} \Delta p_\perp &\leftarrow \Delta p_\perp \,\mathrm{clamp}(\delta/m_f, 0, 1) \\ \mathbf{p}_{cmd} &\leftarrow \mathbf{p}_{cmd} - \delta\,\hat{\mathbf{n}}\,\min\!\left(1, (k_f/b_f)\,\Delta t\right) \end{aligned} \tag{14.12}$$

where δ is the command's distance to the edge of the registered region, $\Delta p_{\perp}$ the part of the increment toward the region (scaled only when it is positive), the second line acts only while $\delta < 0$, $\hat{\mathbf{n}}$ is the outward normal, and $k_f / b_f = 2000/40 = 50\,\mathrm{s}^{-1}$. A registration error moves the region on the map, not in the patient, and the audit counts incursions against the true region, the check against the truth of Section 1.4.9. During a needle drive, the arc fixture replaces the command with the point that keeps the needle's center where the drive began, so the jaw turns about the needle's own center whatever the hand translates.

14.4.3 The Force Loop

The third arm's retraction in case 4 and the needle driver's suture pull run on the shared velocity-form law of Section 1.6.2, with only the emitted output clamped. In this loop's symbols it is Equation (14.13).

$$\begin{aligned} e_k &= F_{SP} - \hat{F}_k \\ v_k &= \mathrm{clamp}\big(K_c\, e_k + I_k,\ -v_{lim},\ v_{up}\big) \\ I_{k+1} &= I_k + \Big[\frac{K_c}{\tau_I}\, e_k + \frac{v_k - (K_c e_k + I_k)}{\tau_I}\Big]\Delta t \end{aligned} \tag{14.13}$$

where v is the arm's pull velocity along the retraction direction, $v_{lim} = 5\,\mathrm{mm/s}$, and v_{up} is 5 mm/s below the force ceiling and 0 above it, so the loop never pulls harder once the estimate exceeds the ceiling. The loop is reverse acting in the convention of Section 1.3 ($a = +1$), because a larger pull raises the force. The derivative is off, back-calculation anti-windup uses a tracking time constant equal to τ_I, and the sample period is 2 ms. Because the loop sets a pull velocity, its proportional term is the integral force law of Section 1.4.8 acting on the pull depth, here through tissue that relaxes, and Chapter 13 closes the same kind of loop on a measured force at a peg. In Position mode the arm approaches the station's depth at 5 mm/s while the law tracks it, so a switch in either direction is bumpless. The loop acts on the biased estimate of Equation (14.9), never on the true force. A suture pull is done when the estimate has held within 0.1 N of its setpoint for 0.8 s, or after 4 s.

14.4.4 Motion Compensation

Each compensation option adds an offset $\mathbf{z}_{comp}$ to the command so that the surgeon works in the organ's frame. Ventilator feedforward fits, per axis, a gain and an offset by least squares to the tracker's reports over the last breath, on a template that knows the ventilator's inspiratory to expiratory ratio but not the patient's exhalation (a raised cosine up and a raised cosine down, where the organ in fact returns exponentially). The adaptive harmonic predictor fits, per axis, a truncated Fourier series of the phase, as shown in Equation (14.14).

$$\hat{z}(\phi) = c + \sum_{k=1}^{N_h} \big(a_k \cos k\phi + b_k \sin k\phi\big) = \theta^{\mathsf{T}} \varphi(\phi) \tag{14.14}$$

Its coefficients are updated on each tracker report, at the phase the report was captured, the correction for report age of Section 1.4.6 applied to a phase. Chapter 10 moves each camera report on by the belt travel since its capture in the same spirit. The update is recursive least squares with forgetting written in symmetric form, as shown in Equation (14.15).

$$\begin{aligned} \mathbf{k} &= \frac{P\varphi}{\lambda + \varphi^{\mathsf{T}} P \varphi}, \qquad \theta \leftarrow \theta + \mathbf{k}\,(\hat{z}_o - \theta^{\mathsf{T}}\varphi) \\ P &\leftarrow \frac{1}{2\lambda}\Big[(P - \mathbf{k}\varphi^{\mathsf{T}} P) + (P - \mathbf{k}\varphi^{\mathsf{T}} P)^{\mathsf{T}}\Big] \end{aligned} \tag{14.15}$$

with $P_0 = 10^4 I$. With forgetting, the rounding error's antisymmetric part grows by $1/\lambda$ each report, and the textbook update lost positive definiteness after about 2500 reports at $\lambda = 0.99$ and diverged in the challenge. The forecast is evaluated at the phase of $t + 50\,\mathrm{ms}$, past the servo's lag, so the offset

lands where the organ will be when it reaches the tip. The ventilator's phase is exact. The heart's is taken from the ECG as $\phi = 2\pi \min\left(1.25, (t - t_R)/RR_{last}\right)$, so a beat longer or shorter than the last is a phase error that no report can reveal until the next R wave. A model is trusted after one whole cycle of reports and blends in over 0.5 s.

The model predictive option runs the shared linear model predictive controller of Section 1.6.3 on a double integrator per axis, the offset and its rate driven by an acceleration, every 10 ms over a 0.5 s horizon of 50 steps with graduated move blocking. It tracks the adaptive forecast with a weight of 1 on the offset, 10^{-6} on moves and 10^{-8} on the input (in millimeters), keeps the acceleration inside $0.8\,a_{max}\,d$ and the rate inside a soft band of $0.8\,v_{max}\,d$ (what the shaft's pivots can add at the present insertion), and integrates the planned acceleration between solves. On a failed solve the adaptive offset takes over bumplessly, and the option is labeled MPC only while the optimization ran. The supervisor's watchdog freezes every offset when the tracker has reported nothing valid for 0.1 s, lets it decay linearly to zero over 2 s, stops the models' learning and blends back onto the forecast over 0.5 s when the track returns.

14.4.5 Supervised Autonomous Suturing

In case 10 a planner, a tracking controller and a supervisor replace the hand. The planner places six bites 3 mm apart and 3 mm deep from the tracked tissue edges, each point carrying the edge finder's seeded error of 0.15 mm. The controller moves the jaw to the approach pose by a minimum-jerk blend $\sigma(\tau) = 10\tau^3 - 15\tau^4 + 6\tau^5$ over 1.2 s, the quintic of Section 1.4.4, rotates the needle about its own center through the bite over 1.6 s, waits 0.8 s for the hand-off, carries the suture to 4 mm above the exit, pulls it with the force loop to 0.5 N and releases. The supervisor halts on a lost track, an estimated force over its 1.6 N ceiling, tissue more than 2 mm from where the plan was made, or the student's Halt, and the motion brakes along its own path over 0.25 s. Freezing the command instead let the fast wrist overrun by about 7° and swing the point 3 mm sideways into the tissue. It resumes only on the student's Run, after replanning the remaining bites. A stage sequencer from the shared framework keeps the ledger, one stage per bite.

14.4.6 Three Models

Section 14.3 is the simulation model, and the controllers' internal models are much smaller. The inverse kinematics knows the chain exactly but not the wrist's play or give, the force loop has no model of the tissue, the predictor's model is a truncated Fourier series of the phase, the model predictive controller's is a double integrator, and the surgeon's is a forward model of their own motion with no robot lag. No model is identified from data. The simulated truth is logged beside every estimate for the audit. The latched injury stop and the out-of-view warning are judged on that truth, while every decision a controller makes is taken on the packet.

14.5 Twelve Cases, Measured

The eleven cases are excerpts of a few minutes from robot-assisted operations, the general surgical ones among the procedures whose move to the robot a statewide study traced [15], each at the step where one control element decides the outcome, and the challenge combines them, as listed in Table 14.4. Each result below was measured headlessly at the case's default seed with the scripts named in the application's build record, and the outcomes are collected in Table 14.5.

> **+ DID YOU KNOW? Robotic Surgery's Rapid Spread**
>
> In a statewide registry of 73 hospitals, the share of general surgery procedures done with a robot rose from 1.8 % in 2012 to 15.1 % in 2018. For inguinal hernia repair, the step of case 5, it rose from 0.7 % to 28.8 % [15].

Table 14.4 The Eleven Cases and the Challenge

#	Procedure and step	Control element
1	Cholecystectomy, the instrument by hand	the fulcrum, Equation (14.3)
2	Cholecystectomy, through the robot	inverse kinematics, rate limits near the port
3	Cholecystectomy, Calot's triangle	scaling against tremor, filter lag
4	Hysterectomy, holding the bladder flap	relaxation, creep, a force loop on an estimate
5	Inguinal hernia, the triangle of doom	a virtual fixture and its registration
6	Prostatectomy, the anastomosis	the needle's arc, the wrist's play and give
7	Colectomy, the stapler	the wait before firing
8	Partial nephrectomy, the breathing kidney	feedforward, adaptive, MPC, a dropout
9	Partial nephrectomy, the clock	pace against an ischemia budget
10	Small-bowel anastomosis, the robot sutures	a planner and a supervisor that halts
11	Coronary bypass on a beating heart	an infeasible rejection
12	Challenge, a renorrhaphy within 25 min	all of them on one seed

Table 14.5 Outcomes Read From the Running Engine, Default Seed

#	Setting	Outcome
2	tip at 29.4 mm insertion, 60 mm move	pitch at 90 °/s for 0.56 s, tip up to 51.5 mm behind
3	1:1, 30 Hz / 3:1, 6 Hz / 3:1, 2 Hz	6 / 0 / 3 artery contacts, tremor at the tip 111 / 31 / 12 µm
4	Position, 4 mm	peak 1.64 N, slip at 11.5 s, 0.64 N at 120 s
4	Force 0.8 N / 2.0 N / 2.5 N	creeps to 4.83 mm / 9.51 mm at 120 s / tears at 13.3 s
5	fixture off / on / +4 mm / −4 mm	1 / 0 / 2 / 0 incursions, 8.6 s blocked at −4 mm
6	new / end of life / no arc fixture	bites 0.29–0.73 / 0.69–1.60 and one missed / up to 1.47 N across
7	wait 0 / 15 / 30 / 45 s	stress index 2.24 / 1.46 / 1.17 / 1.06
8	none / ff / adaptive / MPC	3.29 / 2.19 / 1.10 / 1.13 mm RMS, 10–90 s
9	pace 0.6 / 1.0 / 1.4 / 1.8	over / 22.15 / 20.82 / over 25 min
10	tracked / perfect map	halt 62 ms after the dropout / no halt, bites 0.32–0.74 / 0.32–0.40 mm
11	none / ff / adaptive / MPC, 100 ms	3.86 / 3.22 / 1.71 / 1.57 mm RMS
12	shipped / suggested	clamp on at 25 min, 13 incursions, 9 tears / off at 22.1 min, none

14.5.1 The Fulcrum and Its Inversion

Held by hand in case 1, the instrument reverses and scales the hand's motion as Equation (14.3) says, and the hand's tremor goes through the same lever. At rest the tip carries 107 µm RMS of it in the 8 to 12 Hz band. Case 2 puts the same instrument on the robot at one to one with a 30 Hz filter, and a 20 mm hand move carries the command 20.5 mm with the tip within 1.53 mm of it, the tremor the filter lets through. With the tip brought back to 29.4 mm of insertion, a 60 mm hand move drives the shaft's pitch to its 90 °/s limit for 0.56 s, and the tip falls up to 51.5 mm behind its command before it catches up. At 105 mm of insertion the same limit allows 165 mm/s at the tip, and at 30 mm only 47 mm/s. The inversion is exact, and the limits are the joints'.

14.5.2 Scaling Against Lag

In case 3 the scripted surgeon clears the tissue in Calot's triangle with eleven strokes of 6 mm toward the cystic artery, each ending 1.45 mm from its wall. At one to one with a 30 Hz filter the hand's tremor and drift reach the tip whole, the artery is touched six times, the deepest 0.08 mm into the wall, and the strokes take 80.9 s because the tip will not settle. At 3:1 and 6 Hz nothing touches, the closest approach is 0.49 mm outside the wall, and the strokes take 35.9 s. At 3:1 and 2 Hz the tremor at the tip falls to 12 µm, but three strokes reach the artery, the deepest 0.95 mm in. The cause is the lag the surgeon's forward model does not contain. With the tremor, drift and speed error switched off, a single 6 mm stroke at 3:1 overshoots by 0.40 mm at 6 Hz and by 1.76 mm at 2 Hz, which is more than the 1.45 mm the stroke leaves. At 20 Hz the tip carries 50 µm of tremor against 31 µm at 6 Hz and still touches nothing. Over ten seeds the contacts are 2 to 7 at one to one, none at 3:1 with 6 Hz and 3 to 8 at 2 Hz, and none reaches the 0.5 N of an injury.

> **TRY IT IN THE APP** **3 · Cholecystectomy: Calot's Triangle**
>
> Press Run as shipped and count the artery contacts on the scorecard: six at one to one. Press Reset, set Motion scale to 3:1 and Tremor filter cutoff to 6 Hz, and run again: none. Press Reset and lower the cutoff to 2 Hz. The tremor at the tip falls to 12 µm, yet three strokes reach the artery.

Pause. Try Exercise 14.1 before reading on.

14.5.3 Tissue That Relaxes, Creeps and Tears

The third arm of case 4 pulls the bladder flap 4 mm and holds that position while the surgeon dissects, and the dissection needs 0.7 N of tension. The force peaks at 1.64 N, falls to 0.67 N at 15 s and settles at 0.64 N, what the equilibrium spring holds, so the flap slips at 11.5 s. The force loop at 0.8 N holds the estimate at its setpoint but walks the arm out as the tissue creeps, from 3.61 mm of stretch at 10 s to 4.83 mm at 120 s. While it walks the shaft slides outward, and at 10 s the estimate reads 0.807 N against a true 0.736 N, the port's friction acting on the loop. Early in the walk the true force sags just below the 0.7 N the dissection needs, and the scorecard counts one slip. From 60 to 120 s, with the arm nearly still, the estimate and the truth both average 0.800 N. At 2.5 N the stretch creeps past the 10 mm at which the flap tears, at 13.3 s, with the force 0.5 N below its tear force. At 2.0 N it creeps to 9.51 mm in 120 s and holds. A position hold fails slowly by relaxation and a force hold fails slowly by creep, and neither failure is visible in the force a loop regulates.

Pause. Try Exercise 14.2 before reading on.

The stapler of case 7 closes on the colon at 10 s and fires after the wait the student sets. Fired at once it meets 1.57 N, a stress index of 2.24 against the 0.70 N the same gap holds once relaxed. After 15, 30 and 45 s the index is 1.46, 1.17 and 1.06. The first relaxation time removes most of the excess, and every later second buys less.

14.5.4 Rules on the Command

In case 5 the surgeon raises the peritoneal flap along a path that passes 6 mm from the axis of the iliac vessels, inside the 8 mm region nothing may enter. With the fixture off the tip enters once, 2.44 mm deep. With it on the tip is turned aside at the edge, 8.02 mm from the axis. Registered 4 mm away from the dissection, the fixture admits the path, and the tip enters twice, 1.93 mm deep, with the fixture on. Registered 4 mm toward it, nothing enters, but the fixture holds the tip in safe tissue for 8.6 s, 11.79 mm from the axis. A fixture is as good as the map it is registered to, and its error appears either as an incursion or as safe motion blocked.

Pause. Try Exercise 14.3 before reading on.

The anastomosis of case 6 places six bites with the arc fixture on. With a new instrument the bites land 0.29 to 0.73 mm from their marks, and the needle carries at most 0.60 N across it from the wrist's give under the needle's drag. At the end of the instrument's life the 3° of play turns the needle about the wrist, the first bite's point never comes out of the tissue, and the rest land 0.69 to 1.60 mm off. Without the arc fixture a hand that turns a needle also pushes it, and the needle cuts sideways with up to 1.47 N across it, just under the 1.5 N that tears this tissue. At seeds 7 and 99 one bite tears.

14.5.5 The Patient's Motion

In case 8 the surgeon holds the tip on a mark on a kidney that moves 12 mm every 4 s, with a tracker 100 ms late that loses the kidney for 2 s at 60 s. From 10 to 90 s the hand alone holds the tip 3.29 mm RMS from the mark, ventilator feedforward 2.19 mm, the adaptive predictor 1.10 mm and model predictive control 1.13 mm. Feedforward knows the ventilator but not the exhalation. Three harmonics leave 0.35 mm of the breath's shape unexplained (five leave 0.20 mm), and model predictive control is no better than the predictor it tracks because at this insertion its rate and acceleration limits never bind, over 9000 solves with none failed. During the dropout every option freezes its offset and lets

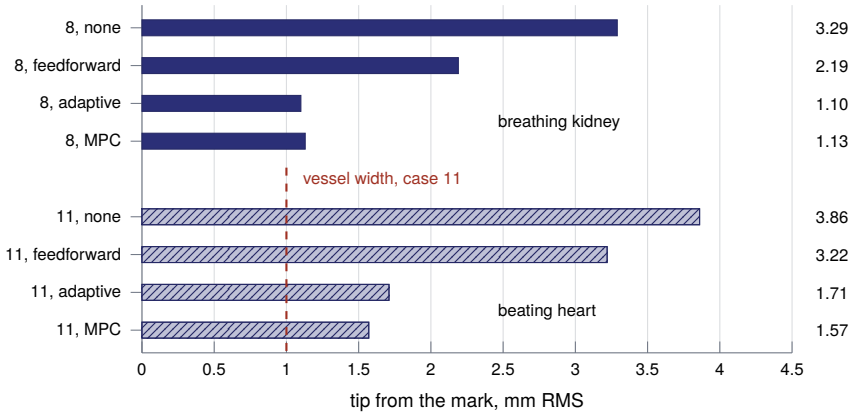

Figure 14.2 Distance of the tip from the mark under the four motion compensation options, from Table 14.5: the breathing kidney of case 8 (solid bars) and the beating heart of case 11 (hatched bars), default seed. Notice that learning the motion's shape cuts the kidney's error from 3.29 to 1.10 mm, while on the heart no option reaches the dashed width of the vessel.

it decay, and from 60 to 64 s the error is 3.43 to 3.48 mm RMS whatever the option, up to 6.87 mm. Figure 14.5 shows the organ one second into the dropout, with the compensation frozen.

> **TRY IT IN THE APP** **8 · Partial Nephrectomy: The Breathing Kidney**
>
> Load the case, keep Motion compensation at None and press Run. Watch the organ panel of the Console view: the hand chases the kidney late, and the tip holds 3.29 mm RMS from the mark. Press Reset, choose Adaptive harmonic and run again. The error falls to 1.10 mm, and at 60 s the offset freezes and decays while smoke hides the kidney.

The tracker's latency barely matters, as shown in Figure 14.3. From 40 to 200 ms the adaptive predictor on the kidney moves from 1.09 to 1.11 mm, because the forecast is evaluated on the ventilator's phase and the latency only delays which reports it has learned from. On the beating heart of case 11 the hand alone is 3.86 mm RMS from the vessel, feedforward 3.22 mm, the adaptive predictor 1.71 mm and model predictive control 1.57 mm, with the tracker 100 ms late. At 40 ms the best is 1.61 mm and at 200 ms 1.53 mm, and with no beat-to-beat variability model predictive control reaches 1.44 mm. No setting comes within the 1 mm width of the vessel, so the case is infeasible as declared. The motion repeats in 0.83 s, its third harmonic is at 3.6 Hz, close to the 5 Hz servo, and a beat 5 % longer than the last is a phase error of up to 18° that no forecast sees before the next R wave. Surgeons answer the heart with a stabilizer that holds a patch of it still, or by stopping it. Figure 14.2 draws the four options on both organs.

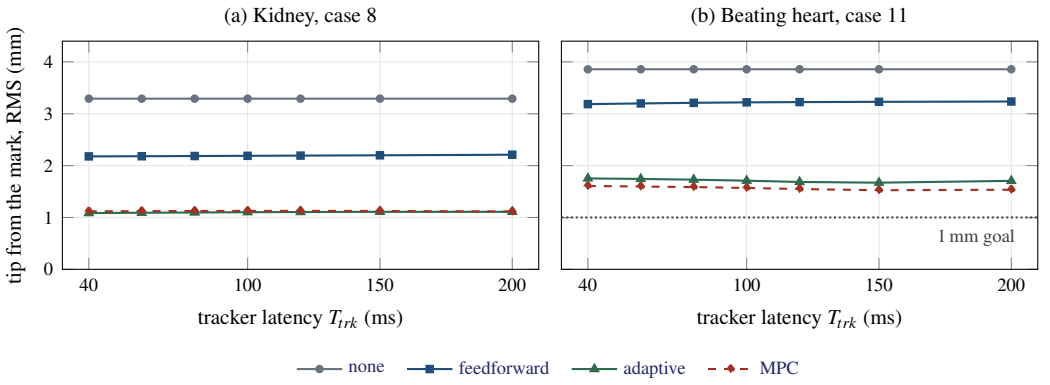

Figure 14.3 Tip distance from the mark against the tracker's latency, default seed. (a) The breathing kidney of case 8, RMS from 10 to 90 s with the 2 s dropout at 60 s. (b) The beating heart of case 11, RMS from 5 to 60 s, with the 1 mm width of the vessel dotted.

Table 14.6 Error Budget of the Hold on the Breathing Kidney, Case 8, Five Seeds. Mean-Square Error in mm^2

	No compensation		Adaptive	
Term	MSE	Share	MSE	Share
Motion alone (M)	1.735	17 %	0.237	18 %
Delay × motion (DM)	8.797	86 %	0.774	60 %
Hand alone (H)	0.057	1 %	0.057	4 %
Hand × delay (HD)	0.078	1 %	0.079	6 %
Instrument, all its terms	−0.449	−4 %	0.122	9 %
All other terms	0.016	0 %	0.022	2 %
Total (RMS)	10.234 (3.20 mm)		1.291 (1.14 mm)	

OPEN CHALLENGE **Holding a Tool on a Moving Organ**

Telerobotic systems that follow organs moved by breathing and the heartbeat have been developed for twenty years, and their potential for this work remains largely unexploited. Guidance by images alone suffers when landmarks are hidden, when the tissue deforms under the tool, and when the organ's motion makes the force fed back to the surgeon oscillate [18]. Progress would be a compensator that holds a tool on a beating heart within the width of a vessel, and keeps doing so while its tracker is briefly blind.

14.5.6 An Error Budget

Four sources of tip error were switched on and off in a two-level full factorial on the running engine, over the default seed and seeds 7, 99, 1 and 2. They are the hand (H, its tremor, drift and speed error, removed by setting all three to zero), the visuomotor delay (D, the 80 ms of video and the 180 ms of reaction, removed by setting both to zero, leaving the 30 Hz frame), the instrument (W, the wrist's give and half its life's play, removed to no give and the 0.5° of a new wrist), and the organ's motion (M, the 12 mm breath). The mean-square error pooled over the seeds is split exactly into a term for every subset of sources by Möbius inversion over the 16 cells, so a main effect is what a source costs alone and an interaction is what sources cost together beyond their sum. The decomposition of case 8's hold is listed in Table 14.6.

The delay costs nothing on a still mark (0.00 mm RMS with it alone), and the motion alone, with an instant visuomotor loop, costs 1.32 mm. Together they cost 3.25 mm, and their interaction is 86 % of the mean-square error. The surgeon's forward model covers the surgeon's own motion and not the organ's, so a target that moves is chased a quarter of a second late. The adaptive predictor cuts that

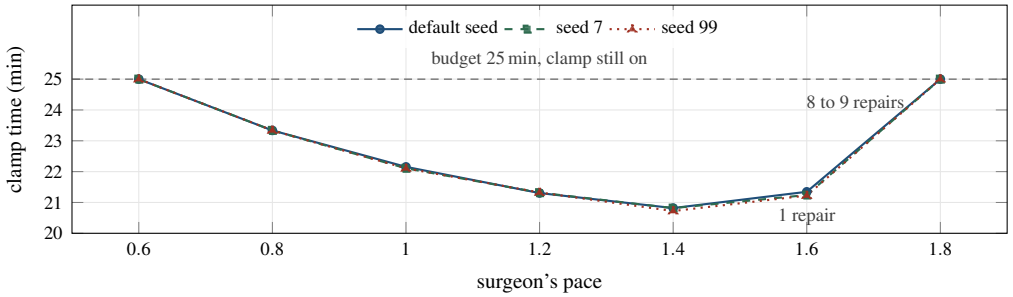

Figure 14.4 Warm ischemia time against the scripted surgeon's pace, case 9, at three seeds. A run still clamped at the 25 min budget is plotted at the budget. One repair appears at a pace of 1.6 and eight or nine at 1.8.

interaction from 8.797 to 0.774 mm^2, by 91 %, and the motion's own term from 1.735 to 0.237, by 86 %, while the hand's terms are unchanged at 0.057 and 0.078 mm^2, because scaling and filtering, not compensation, are what act on the hand. Over the five seeds the full hold is 3.18 to 3.22 mm RMS without compensation and 1.11 to 1.17 mm with it. With the instrument new, as case 8 ships, the pooled hold without compensation is 3.27 mm. In a hold the wrist carries no load, and its terms are small and of either sign.

> ✦ DID YOU KNOW? **A Delay That Costs Nothing Until the Organ Moves**
>
> From the simulator: on a still mark the surgeon's visuomotor delay costs nothing. On the breathing kidney its interaction with the motion is 86 % of the hold's mean-square error, and the adaptive predictor cuts that term from 8.797 to 0.774 mm^2 (Table 14.6).

The same factorial on case 6's six bites on still tissue, with the arc fixture on, gives a pooled 0.78 mm RMS, of which 0.37 mm remains with every source removed. Of the 0.612 mm^2, the instrument's four terms are 0.339 (55 %), the hand 0.072 and the delay 0.048, and with the half-worn wrist 3 to 5 of the 30 bites in each of its cells miss or tear. On the breathing kidney of case 9's renorrhaphy, with the clips and repairs switched off, no compensation leaves not one of the 240 bites placed by the eight cells with motion kept, whatever else is removed. With the adaptive predictor and every source present the pooled error is 1.19 mm with 3 of 30 bites missed or torn. There the terms are as large as the main effects and change sign. Removing the delay alone raises the pooled error to 1.86 mm with 12 bites lost, a mechanism the factorial did not isolate, so a single-source attribution does not hold for bites on a moving organ.

14.5.7 Pace Against the Clock

In case 9 the renal artery has been clamped for 17 of a 25 min budget, six renorrhaphy bites remain, each secured with a clip of 40 s, and a bite that tears, misses or lands more than 1.5 mm off costs 30 s and a new bite. The clamp time is shown against the pace in Figure 14.4. At 0.6 the clamp is still on at 25 min, at 1.0 it comes off at 22.15 min, at 1.4 at 20.82 min, and at 1.8 the surgeon makes nine repairs and overruns again. The optimum is interior at all three seeds, 20.72 to 20.82 min at a pace of 1.4. The loss index of the scorecard is a declared teaching curve, $e^{0.05(t-25)}$ with t in minutes, and the budget comes from the clinical finding that every minute of warm ischemia counts [16], not from any threshold that study states.

14.5.8 The Robot Sutures

In case 10 the robot closes an enterotomy by itself with the tracker 50 ms late. At bite 4 (23.42 s) the tracker loses the bowel, and the supervisor halts 62 ms later, braking along the path. When the student



Figure 14.5 The breathing kidney of case 8 with the adaptive harmonic predictor, at $t = 61.0$ s, one second into the tracker's 2 s dropout, on the Console view, cropped to the plant view. (1) The badge row, with the track lost and the compensation frozen. (2) The hand, scaled, against the command and the true tip (dashed). (3) The organ panel shows the true motion (dashed) leaving the tracker's last report while the frozen offset decays, and the tip is 6.31 mm from its mark. (4) The supervisor's state and the latencies it works with. (5) The status line.

sets the autonomy back to Run, three seconds later in the measured run, it replans and finishes the last bite at 47.0 s. The six bites land 0.32 to 0.74 mm from their marks, spaced within 0.19 mm of 3 mm. Each suture is pulled until the estimate reads 0.5 N, while the true tension is 0.28 N, the port's friction once more acting on the loop. The perfect map, a labeled teaching comparison, never halts and places every bite 0.32 to 0.40 mm off, the error that belongs to the instrument itself.

14.5.9 A Renorrhaphy Within the Budget

The challenge is a renorrhaphy on the breathing kidney with the hilum below the defect, an instrument at half its life, a 2 s dropout at 150 s and the clamp on for 17 of 25 minutes, and it ships at one to one, 30 Hz, no compensation, no fixture and a pace of 1.5. The target is the clamp off within 25 min with no hilum incursion, no tear, no stop and six kept bites within 1.5 mm. As shipped the surgeon chases the kidney, the bites land off and tear, the tip strays into the hilum's region 13 times, and after seven repairs the clamp is still on at 25 min, with 10 and 17 incursions at seeds 7 and 99. The suggested configuration (adaptive compensation, the fixture on, 3:1, 6 Hz and a pace of 1.2) takes the clamp off at 22.1 min with no incursion, no tear and the worst kept bite 1.46 mm off, and passes at seeds 7 and 99 at 22.1 and 23.2 min. At a pace of 1.0 the same settings overrun at seed 7, where two bites need repairs. The bite tolerance is 1.5 mm rather than 1 mm because the wrist's give, a half-worn wrist and the predictor's residual put most bites at the suggested settings beyond 1 mm.

14.6 Technology in Practice

The robot of Section 14.3 reads joint encoders, an endoscope, a tissue tracker and a force it estimates, and it moves the tip through joint servos and a cable-driven wrist. Figure 14.6 draws its loops with real devices in the place of the schematic's symbols.

What Measures It

Research systems find the organ by tracking markers or the tissue's surface in the endoscope's view, and they detect the breathing motion from it [17]. The tracker of Section 14.3, 100 ms late with 0.2 mm of noise and silent while blood or smoke hides the surface, stands for such a tracker. The force at the tip is harder to measure. Sensors at the tip face severe limits of size, geometry, cost, biocompatibility and sterilization, so the force would ideally be estimated without them. On a robot with low inertia and friction, the gap between the commanded and the actual position indicates the force, but the robot's own dynamic forces often mask the small forces of tissue [13]. A tension read near the motors carries the friction of the cables [19]. The estimate of Equation (14.9), biased by the port's and the cables' friction, stands for this kind of measurement.

> **TECHNOLOGY** **Sensor: Tracking Tissue in the Endoscope's Image**
>
> A tissue tracker finds the organ in the endoscope's images. In one research system a three-dimensional endoscope projects fringes onto the tissue to recover its surface, and a near-infrared camera follows biocompatible markers placed on it. The markers keep the track through blood and thin tissue, and the camera works 5 to 8 cm from the tissue, with the suturing tool withdrawn from the view so that it does not hide the surface [17]. A hidden landmark corrupts the measurement [18]. On a trend the report stops while the true organ moves on, as in case 8 at 60 s.

What Moves It

The instruments ride on an electromechanical arm under software control [20]. In the chapter the remote center is built into the joint coordinates. A research system of 2022 instead imposed the remote center as a constraint inside its kinematic solver, at the port of the laparoscopic tool [17]. The second-order servos of Equation (14.4), with their rate and acceleration limits, stand for the arm's joint drives.

> **TECHNOLOGY** **Final Element: The Cable-Driven Wristed Instrument**
>
> A dexterous surgical robot typically has seven degrees of freedom, including translations, rotations and the grip [13]. The wrist at the tip is driven by cables in sheaths, the most common drive of surgical tools for its small size, strength and flexibility. Friction between cable and sheath changes direction with the motor, so a tension reading shifts suddenly at a reversal before it settles, and a tension read near the motors carries that friction. A cable that sticks shows as an abnormal tension, and the robot can then stop itself [19]. The play and the give of Equation (14.5) stand for this drive in the model.

Where the Controller Runs

The surgeon's console and the patient-side arms are usually in the same operating room, where a survey puts the communication delay at 5 to 10 ms [18]. The controllers of this chapter act every 2 ms, and the model predictive option every 10 ms. The 2022 research system shows the same layers as the chapter: a task planner sends motion primitives to a trajectory generator, a kinematic solver turns them into joint trajectories, and closed-loop joint controllers follow them [17]. The console, the fixtures, the force loop and the compensation of Section 14.4 sit between the hand and that solver.

> **IN PRACTICE** **What the Regulator Asks of a Surgical Robot**
>
> In the United States, one rule for a software-controlled, table-mounted electromechanical system with instruments on an arm places it in Class II with special controls. The maker must run a training program for the device, distribute it only to facilities that maintain it, and report complaints and adverse events every year. Bench tests must show the device's motion accuracy and precision, instrument reliability and the human-device interface. Human factors validation must show safe use in an operating room, and the software needs verification, validation, hazard analysis and a review of its cybersecurity threats [20].

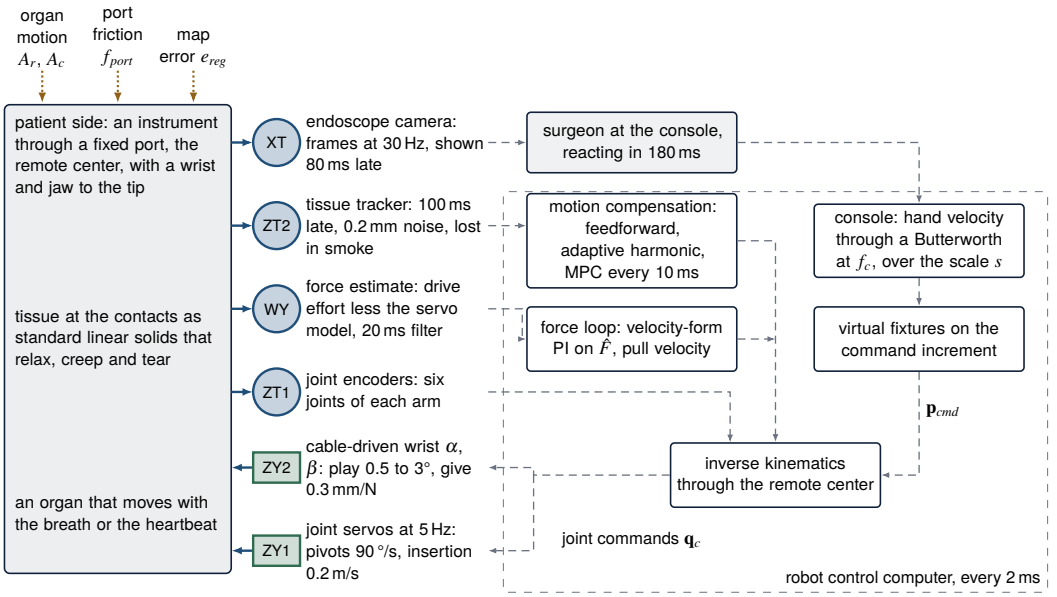

Figure 14.6 The loops of the chapter with real devices in place of the schematic's symbols. The surgeon sees the endoscope camera XT and moves the hand. A robot control computer acting every 2 ms scales and filters the hand in the console, bends the command at the virtual fixtures, adds the compensation learned from the tissue tracker ZT2, runs the force loop on the estimate WY, and solves the inverse kinematics for the joint servos ZY1 and the cable-driven wrist ZY2, read back by the encoders ZT1. Notice that the tip force has no sensor of its own, and that the surgeon closes the outer loop through a picture that is already late.

What Is Modernizing It

Autonomy is the first development. Research robots now plan and place sutures on soft tissue under supervision, and the operator approves or selects their plans [8, 17]. Sensing at the tip is the second. Fiber Bragg grating, capacitive and tactile sensors have been built for palpation, and a small tension sensor array at the distal end of a cable-driven robot is less affected by friction than one at the motors [19]. Motion compensation is the third. Telerobotic systems that synchronize the tool with an organ moved by breathing or the heartbeat have been studied for two decades, with position, force and impedance control schemes [18]. Security is the fourth. The United States rule for a mountable electromechanical surgical system asks for an assessment of the threats and vulnerabilities of its software as part of a cybersecurity review [20].

14.7 Limitations

The arms have exact kinematics and limited servos but no inertia, so a contact pushes the instrument back only through the wrist. Tissue is lumped at the contacts as standard linear solids with fixed tear limits, so a tear has no location, a needle cannot steer and the stretch of one bite does not load the next. The organ moves as a rigid body on a scripted shape, the ports are fixed, and there is no bleeding, smoke, port placement or haptic channel. The scripted surgeon is a forward-model controller with tremor and drift scaled by one pace, and it describes no person. Real surgeons adapt to a console, anticipate a breath and learn a worn instrument, and none of that is here. The error budget is measured at five seeds on three tasks and holds within those conditions, and the attribution on the moving kidney's bites was not resolved. Every tissue property, tolerance, ceiling, latency and time budget was chosen for teaching and carries no clinical meaning. Without a learning study, nothing is claimed

about what students learn, and the simulator is not a surgical trainer, clinical software or a decision aid.

OPEN CHALLENGE **Feeling the Tissue Through the Robot**

The lack of effective haptic feedback is often reported as a major limitation of teleoperated surgical robots. A perfectly transparent teleoperator would need perfect models of both robots, zero delays and perfect force sensing or estimation, and pushing toward transparency lets small errors and delays cause oscillations that surgery cannot accept. Force sensors at the tip must also meet limits on size, cost, biocompatibility and sterilization [13]. Progress would be an affordable, sterilizable instrument that senses its tip force, displayed to the hand without loss of stability.

14.8 Conclusion

A surgical robot realizes the hand's intent through a point that must not move, and this simulator measures what each layer of that realization costs. The inversion of the fulcrum is exact until the tip nears the port, where the pivots reach 90 °/s and the tip falls 51.5 mm behind. A 6 Hz filter at 3:1 keeps the hand off the artery that a 2 Hz filter overshoots into by 0.95 mm. A held position slips at 11.5 s, and a force loop at 2.5 N tears its flap by creep at 13.3 s. On the breathing kidney the adaptive predictor holds the tip at 1.10 mm RMS against 3.29 mm by hand, and on the heart no setting does better than 1.53 mm. The factorial budget places 86 % of the hold's mean-square error in the interaction of the surgeon's delay with the organ's motion and shows the predictor removing 91 % of it. The model is lumped, rigid and scripted. A fuller model would add arm dynamics, a deforming organ and a haptic channel with its stability limits, and only a study of students using the cases could say what they learn. The simulator is a surgical robot that can be taken apart one control layer at a time.

KEY POINTS

- A surgical robot carries out the hand's intent through a point that must not move.
- Inverting the fulcrum is exact until the tip nears the port, where the pivots saturate.
- At 3:1 a 6 Hz filter keeps the hand off the artery, and a 2 Hz filter overshoots into it.
- A position hold fails slowly by relaxation, and a force hold by creep.
- The surgeon's delay interacting with the organ's motion holds 86 % of the hold's mean-square error, and the predictor removes 91 % of it.

Where else this applies. A disturbance on a measured period also drives the tooth-passing vibration of Volume 7, Chapter 13.

References

[1] S. Hara, Y. Yamamoto, T. Omata, and M. Nakano, "Repetitive control system: A new type servo system for periodic exogenous signals," *IEEE Transactions on Automatic Control*, vol. 33, no. 7, pp. 659–668, 1988.

[2] R. Ginhoux, J. Gangloff, M. de Mathelin, L. Soler, M. M. A. Sanchez, and J. Marescaux, "Active filtering of physiological motion in robotized surgery using predictive control," *IEEE Transactions on Robotics*, vol. 21, no. 1, pp. 67–79, 2005.

[3] O. Bebek and M. C. Cavusoglu, "Intelligent control algorithms for robotic-assisted beating heart surgery," *IEEE Transactions on Robotics*, vol. 23, no. 3, pp. 468–480, 2007.

[4] R. H. Taylor, J. Funda, B. Eldridge, S. Gomory, K. Gruben, D. LaRose, M. Talamini, L. Kavoussi, and J. Anderson, "A telerobotic assistant for laparoscopic surgery," *IEEE Engineering in Medicine and Biology Magazine*, vol. 14, no. 3, pp. 279–288, 1995.

[5] R. H. Taylor and D. Stoianovici, "Medical robotics in computer-integrated surgery," *IEEE Transactions on Robotics and Automation*, vol. 19, no. 5, pp. 765–781, 2003.

[6] C. N. Riviere, R. S. Rader, and N. V. Thakor, "Adaptive cancelling of physiological tremor for improved precision in microsurgery," *IEEE Transactions on Biomedical Engineering*, vol. 45, no. 7, pp. 839–846, 1998.

[7] S. A. Bowyer, B. L. Davies, and F. Rodriguez y Baena, "Active constraints/virtual fixtures: A survey," *IEEE Transactions on Robotics*, vol. 30, no. 1, pp. 138–157, 2014.

[8] A. Shademan, R. S. Decker, J. D. Opfermann, S. Leonard, A. Krieger, and P. C. W. Kim, "Supervised autonomous robotic soft tissue surgery," *Science Translational Medicine*, vol. 8, no. 337, 2016.

[9] Y.-C. Fung, *Biomechanics*. New York, NY: Springer, 1993.

[10] A. M. Okamura, C. Simone, and M. D. O'Leary, "Force modeling for needle insertion into soft tissue," *IEEE Transactions on Biomedical Engineering*, vol. 51, no. 10, pp. 1707–1716, 2004.

[11] Y. S. Kwoh, J. Hou, E. A. Jonckheere, and S. Hayati, "A robot with improved absolute positioning accuracy for CT guided stereotactic brain surgery," *IEEE Transactions on Biomedical Engineering*, vol. 35, no. 2, pp. 153–160, 1988.

[12] A. G. Gallagher, N. McClure, J. McGuigan, K. Ritchie, and N. P. Sheehy, "An ergonomic analysis of the fulcrum effect in the acquisition of endoscopic skills," *Endoscopy*, vol. 30, no. 7, pp. 617–620, 1998.

[13] A. M. Okamura, "Haptic feedback in robot-assisted minimally invasive surgery," *Current Opinion in Urology*, vol. 19, no. 1, pp. 102–107, 2009.

[14] J. Marescaux, J. Leroy, M. Gagner, F. Rubino, D. Mutter, M. Vix, S. E. Butner, and M. K. Smith, "Transatlantic robot-assisted telesurgery," *Nature*, vol. 413, no. 6854, pp. 379–380, 2001.

[15] K. H. Sheetz, J. Claflin, and J. B. Dimick, "Trends in the adoption of robotic surgery for common surgical procedures," *JAMA Network Open*, vol. 3, no. 1, p. e1918911, 2020.

[16] R. H. Thompson, B. R. Lane, C. M. Lohse, B. C. Leibovich, A. Fergany, I. Frank, I. S. Gill, M. L. Blute, and S. C. Campbell, "Every minute counts when the renal hilum is clamped during partial nephrectomy," *European Urology*, vol. 58, no. 3, pp. 340–345, 2010.

[17] H. Saeidi, J. D. Opfermann, M. Kam, S. Wei, S. Leonard, M. H. Hsieh, J. U. Kang, and A. Krieger, "Autonomous robotic laparoscopic surgery for intestinal anastomosis," *Science Robotics*, vol. 7, no. 62, p. eabj2908, 2022.

[18] L. Cheng and M. Tavakoli, "COVID-19 pandemic spurs medical telerobotic systems: A survey of applications requiring physiological organ motion compensation," *Frontiers in Robotics and AI*, vol. 7, p. 594673, 2020.

[19] Z. Zhou, J. Yang, M. Runciman, J. Avery, Z. Sun, and G. Mylonas, "A tension sensor array for cable-driven surgical robots," *Sensors*, vol. 24, no. 10, p. 3156, 2024.

[20] U.S. Food and Drug Administration, "Mountable electromechanical surgical system for transluminal approaches," Code of Federal Regulations, Title 21, Section 878.4961, revised as of April 1, 2024, U.S. Government Publishing Office, govinfo.gov, 2024.

14.9 Exercises

14.1 ⌕ Concept check · 2 min

In case 3 at a motion scale of 3:1, lowering the tremor filter cutoff from 6 to 2 Hz cuts the tremor at the tip from 31 to 12 µm, yet three strokes now reach the cystic artery where none did at 6 Hz. What causes the contacts?

(a) The 2 Hz filter passes more of the 10 Hz tremor than the 6 Hz filter does

(b) The shaft's pivots reach their 90 °/s limit near the port during each stroke

(c) At a low cutoff the console no longer divides the hand's drift by the motion scale

(d) The filter's lag is missing from the surgeon's forward model, so the surgeon pushes on past the target

14.2 ⌕ Concept check · 1 min

In case 4 a force loop holds the bladder flap at 2.5 N, below the flap's 3 N tear force, and the flap tears at 13.3 s. Why?

(a) Under a held force the flap creeps, and its stretch passes the 10 mm stretch limit

(b) Port friction biases the estimate low, so the true force climbs past 3 N

(c) The integral term winds up, and the arm overshoots the setpoint past the tear force

(d) The force relaxes below the 0.7 N the dissection needs, and the slipping flap tears

14.3 Concept check · 2 min

In case 5 the virtual fixture is on, but its region is registered 4 mm off, toward the path of the dissection. What does the run show?

(a) The tip enters the true region twice, because the shifted map admits the path

(b) Nothing enters the true region, but the fixture holds the tip in safe tissue for several seconds

(c) The tip chatters at the edge of the region, because the fixture pushes it out at $50\,\mathrm{s}^{-1}$

(d) Nothing changes, because the fixture is judged against the true region and not the map

14.4 Theory and calculation · 5 min

In case 1 a straight instrument of 360 mm is held by its handle with 120 mm inside the port. In case 2 the robot holds the same kind of shaft, whose pitch and yaw are limited to 90 °/s.

(a) Use Equation (14.3) to find the tip motion for a 20 mm move of the hand, before and after the instrument is pushed in by 80 mm. *Checkpoint:* −10.0 mm and −25.0 mm.

(b) Find the sideways tip speed the pivot limit allows at an insertion of 105 mm and of 30 mm. *Checkpoint:* 165 mm/s and 47.1 mm/s.

(c) At one to one the hand moves at up to 0.15 m/s. Find the pivot rate that speed needs at 30 mm of insertion, as a multiple of the limit. *Checkpoint:* 5.0 rad/s, or 286 °/s, 3.2 times the limit.

14.5 Theory and calculation · 8 min

The console filters the hand's velocity with a second-order Butterworth low-pass of cutoff f_c. Its gain at a frequency f is $1/\sqrt{1+(f/f_c)^4}$, and its phase lag is $\varphi = \arctan\left(\sqrt{2}\,r/(1-r^2)\right)$ with $r = f/f_c < 1$, a delay of $\varphi/(2\pi f)$.

(a) Find the share of a 10 Hz tremor that passes at $f_c = 6$ and at 2 Hz, and the tremor left in the command at a scale of 3:1 if the hand's 0.3 mm RMS were all at 10 Hz. *Checkpoint:* 0.339 and 0.040, so 34 µm and 4.0 µm.

(b) Find the delay each filter adds to a 1 Hz motion. *Checkpoint:* 37.9 ms and 120.3 ms.

(c) The surgeon's loop is an integrator of gain $K = 2\pi\ \mathrm{s}^{-1}$. Use Equation (1.16), whose loop gain Kk_i is here the surgeon's K, to find its delay margin, and compare it with the 0.26 s the surgeon's image is late. *Checkpoint:* 0.25 s, less than 0.26 s.

(d) The forward model covers the 0.26 s and nothing else. How much more unmodeled delay does the 2 Hz filter add at 1 Hz than the 6 Hz filter? *Checkpoint:* 82.4 ms.

14.6 Theory and calculation · 8 min

The bladder flap of case 4 is a standard linear solid, Equation (14.6), with $k_e = 150\,\mathrm{N/m}$, $k_m = 250\,\mathrm{N/m}$, $\tau = 4\,\mathrm{s}$ and $x_0 = 15\,\mathrm{mm}$. It tears at 3 N or at 10 mm of stretch, and the dissection needs 0.7 N.

(a) The arm holds a stretch of 4 mm. Find the force at once, with $x_m = 0$, and the force once fully relaxed. Does the hold keep 0.7 N? *Checkpoint:* 1.64 N and 0.64 N, so no.

(b) Find the share of the linear force a held stretch loses and the time constant of creep under a held force. *Checkpoint:* 0.625 and 10.7 s.

(c) A force loop holds the true force at its setpoint once the arm is at rest. Find the relaxed stretch at 0.8 N and at 2.5 N. *Checkpoint:* 4.83 mm and 10.90 mm.

(d) Find the smallest setpoint at which the flap tears by creep, however long it takes. *Checkpoint:* 2.17 N.

14.7 App challenge · 8 min

Load scenario 4, Hysterectomy: holding the flap. It ships with the Force loop in Position (MAN) at a pull depth of 4.00 mm. Use the default seed, run each case for 120 s at a playback speed of 5, and export the CSV file after each run.

(a) Run as shipped. Read the Peak tissue force tile and the true force on the Flap force tile at 120 s, and find the first row in which column `slip` reads 1. *Checkpoint:* 1.64 N, 0.64 N and 11.5 s, within 0.05 s.

(b) Press Reset, set the Force loop to Force (AUTO) with the Force setpoint at 0.80 N, and run again. Read column `flap_stretch_m` at 10 and at 120 s. *Checkpoint:* 3.61 mm and 4.83 mm, within 0.02 mm.

(c) Press Reset, set the Force setpoint to 2.50 N, and run again. Find the first row in which column `tears` reads 1. *Checkpoint:* 13.3 s, within 0.1 s.

Compare (b) and (c) with the relaxed stretches of the flap calculation earlier in this set.

14.8 App challenge · 6 min

Load scenario 7, Colectomy: the stapler. The jaws close on the colon at 10 s, and the stapler fires after the Wait before firing, 0 s as shipped. Use the default seed and a playback speed of 5, and run each case until the stapler has fired and the run pauses.

(a) Run as shipped. Then press Reset, set the wait to 15 s and run. Then press Reset, set it to 45 s and run. Read the Staple-line stress tile each time. *Checkpoint:* index 2.24 at 1.57 N, 1.46 at 1.02 N, 1.06 at 0.75 N.

(b) The index is the force at firing over the fully relaxed force at the same gap. From the first run alone, find the relaxed force. Then predict the index after 15 and 45 s if its excess over one decays with the bowel's relaxation time of 15 s. *Checkpoint:* 0.70 N, then 1.46 and 1.06.

14.9 App challenge · 8 min

Load scenario 8, Partial nephrectomy: the breathing kidney. Use the default seed and a playback speed of 5, and run each case for the full 90 s. Run first with Motion compensation None, as shipped, then press Reset and run with Ventilator feedforward, then Adaptive harmonic, then MPC.

(a) Read the Tracking error tile, the RMS distance from the mark from 10 s, for each option. *Checkpoint:* 3.29, 2.19, 1.10 and 1.13 mm, within 0.01 mm.

(b) Read the During the dropout tile for None and for Adaptive harmonic. *Checkpoint:* 4.51 mm and 4.49 mm, within 0.02 mm.

Explain why feedforward falls short of the predictor, why MPC is no better than the predictor, and why the dropout costs every option the same.

14.10 Application area · 4 min

The tip force is not measured. It is estimated from the drive's effort, which carries the port's Coulomb friction of 0.3 N along the shaft, Equation (14.9). In case 10 each suture is pulled until the estimate reads 0.5 N, and the scorecard reports a true tension of 0.28 N.

(a) By how much, in newtons and as a share of the truth, does the estimate read high, and which way was the shaft sliding? *Checkpoint:* 0.22 N, 79 % of the truth, sliding outward.

(b) Name two changes a design engineer could make so that the pull reaches its intended tension.

14.11 Application area · 4 min

Case 11 asks the tip to stay within 1 mm of a coronary artery on a heart beating at 72 beats per minute, each beat drawn within ± 5 % of the mean, behind 5 Hz joint servos.

(a) Find the period of the motion and the frequency of its third harmonic, and repeat the harmonic at 100 beats per minute. *Checkpoint:* 0.83 s and 3.6 Hz, then 5.0 Hz.

(b) A beat 5 % longer than the last is read on the phase of the last interval. Find the largest phase error. *Checkpoint:* 18°.

(c) Name the surgical answer to this motion and say why it is not a control setting.

14.12 Review the agent · 5 min

A student asks an AI assistant how to improve the hold on the breathing kidney of case 8. The assistant replies: “Table 14.6 shows that the delay-times-motion term is 86 % of the mean-square error without compensation. The tracker reports 100 ms late, so set the Tracker latency to 40 ms. That removes most of the delay and with it most of the error.” Find the flaw and say how to check it in the app.

CHAPTER 15

Collaborative Robot Speed and Separation Simulator With a Truth Audit

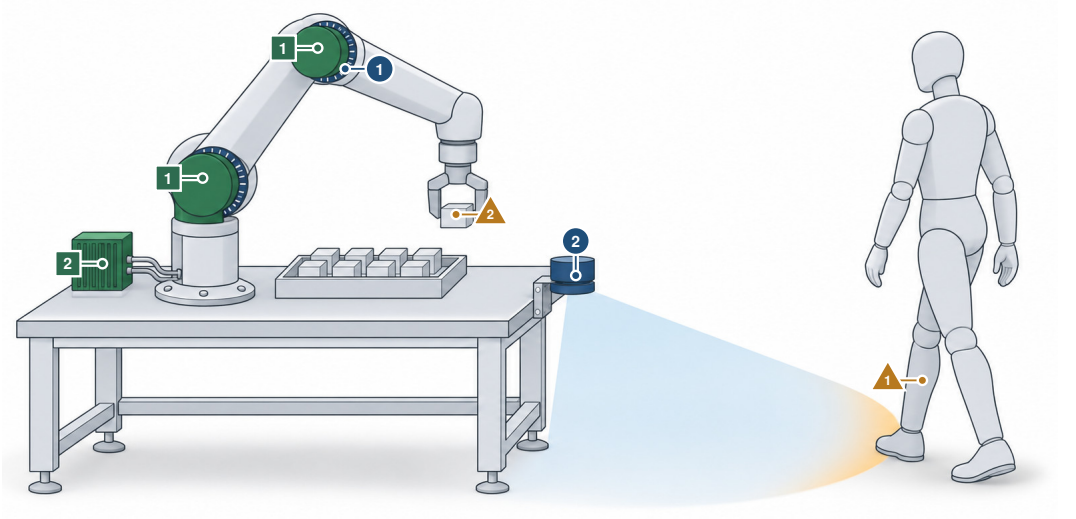

The controller moves 1 drives at the shoulder and the elbow 2 the controller that sets the arm's speed ceiling

The controller reads 1 joint encoders 2 safety laser scanner that finds the person

What disturbs it 1 the person walking in 2 a part heavier than the program believes

Notice that the scanner, not a fence, keeps the person safe: the arm slows as the person comes closer and must always be able to stop in time. Illustration, not to scale.

IN PLAIN TERMS A Robot Arm That Slows for People

The main idea. A collaborative robot works beside people with no fence between them. A scanner watches each person, and the arm slows as a person comes closer, so that it can always stop in time.

The challenge. The scan is always a little late and a little off, and a loaded arm needs more time to stop. A supervisor that trusts a hopeful braking figure acts too late, and only a replay of the truth shows how close the person came.

Why it matters. Every bit of safety distance costs parts made, so an honest stopping model keeps a shared cell both safe and busy.

THREE WAYS TO READ THIS CHAPTER

Explore	12 min	the first page, the two history segments (Section 15.2), the fact on walking speed (Section 15.4), the try-it segments of cases 2 and 5 and the fact on latency (Section 15.6) 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 measured cases (Section 15.6), Technology in Practice (Section 15.7), Limitations (Section 15.8) and the open challenges

The control question A collaborative robot without a fence must be able to stop before a person walking in can reach it, so its safe speed shrinks as the person comes closer. How large is that protective distance term by term, what does each term cost in throughput, and how can a supervisor that sees only a late, noisy scan be checked against the truth?

Learning objectives

- Compute the protective separation distance and its terms from the reaction chain, the braking model and the uncertainties.
- Solve the speed law for the fastest tool speed a measured separation allows, and explain why it needs a margin and a restart hysteresis.
- Explain why an optimistic braking figure misjudges the stop mostly in time, and read a truth audit that exposes it.
- Compute the contact speed that power and force limiting allows for a body region from the reduced mass and the region's stiffness.
- Weigh a raised collision threshold against a corrected friction model.

Prerequisites Sections 1.4.9 (stopping distance, protective separation and speed law), 1.4.2 (arm dynamics), 1.4.8 (landing force), 1.3 (loop action) and 1.6.4 (scenarios, scorecard and CSV file). Chapter 9 treats believed against true braking on a mobile robot.

The app https://apmonitor.com/apps/cobot-safety/

See also Volume 4, Chapter 2 for a protective override that rests at its clamp and is judged by its margin to a limit. Volume 8, Chapter 13 for a latched supervisor whose product history is logged apart from its instruments, and Volume 8, Chapter 2 for a safety layer that guards a hazard with no symptom in the controlled variable.

15.1 Introduction

Collaborative robots are bought to work in the same space as people, and the fence that used to separate the two was a safety function in its own right. New paradigms in industrial robotics no longer require physical separation between manipulators and people, and safety can then be posed as a hard constraint on the robot's motion, met while productivity is kept as high as possible [1]. A survey of collaborative robotics for industry relates actuation, control, safety, physical interaction, usability and productivity, and finds that the way an application is organized affects both its quality and its safety [2]. The requirements for industrial robots are set by ISO 10218 [3], and ISO/TS 15066 supplements them for collaborative operation [4, 5]. Two of its methods are feedback problems. Speed and separation monitoring, as the specification presents it, keeps a minimum protective distance between the robot and the person [6], and power and force limiting, which the application takes from the same specification, admits contact at a speed whose impact a body region can bear, the question that crash tests of robots address [7].

The protective distance is a budget of distances that all move at once. In the cell of this chapter it is 1448 mm at 1.00 m/s with a 3 kg part and a 170 ms reaction chain, and 840 mm of that is the person's own approach while the arm reacts and brakes. At rest it is still 532 mm, more than the 214 mm that the arm's path leaves a person standing at the table. The simulator is built against three beliefs. The first is that a slow enough arm is safe near a person, but an unmonitored arm held at 25 % of its speed still gives 6 violation episodes in a minute, the worst 410 mm short of the distance the truth requires. The second is that the stopping distance is the term to get right, but a brochure braking figure of 5 m/s^2 misjudges the stop of the loaded arm by 67 mm in distance and by 440 ms in time, and the time costs 704 mm of the person's walk. The third is that an optimistic model at least buys throughput. With the brochure figure the cell completes 7 cycles in a minute, the same 7 as with an honest model, and a truth audit finds 6 violation episodes, the worst 506 mm short.

The class of problem is general. A supervisor that watches a primary controller and intervenes when a safety constraint is about to be lost is a runtime assurance system, of which the automatic ground collision avoidance of military aircraft is a working example, and safety filtering of this kind

is now studied for complex control systems in general [8]. In a robot cell the filter acts on a speed. The separation can be adapted to the direction of the robot's motion as well as to the distance, and continuous adaptation of the velocity was found to beat zone-based supervision and physical barriers in productivity, by an amount that depends on how often and when an operator comes near [9]. Metrics for speed and separation monitoring accordingly quantify both the safety and the effect on production [10]. A mobile robot braking for a walker, a crane slewing near a worker and a road vehicle with automatic braking have the same shape, with a stopping model that must be honest, a measurement that is old when it is used, a threshold that needs hysteresis, and a verdict that only the true state can give. A transfer function from speed override to tool speed carries none of these.

However, the monitoring literature measures implementations on physical cells [6, 9], and the collision literature evaluates detection on physical arms [11], so a student has no plant on which the budget can be read term by term while a person walks in, one term changed, and the throughput and a verdict of safety compared side by side. This chapter itemizes the protective separation distance on a collaborative cell simulator in which every term of the budget is a setting, the stopping behavior comes from the arm's own torque limits, and each supervisory decision is audited against where the person truly was and how the arm truly stops. The protective separation distance, the speed law, power and force limiting and the momentum observer follow the standard's formulas and the literature. The model can be rebuilt from this text alone. Its controller sees nothing but its encoders, its drive torques, a hand-over button and a delayed scanner, its braking model is derived from the drives' torque limits along the program, and its audit replays the controller's own trajectory generator and servos on a copy of the true arm. Eight cases and a graded shift isolate the terms of the budget one at a time.

Section 15.2 traces the domain. Section 15.3 gives the model of the cell, and Section 15.4 the budget, the speed law built on it, the stop it assumes, the contact allowed instead and the truth audit. Section 15.5 describes the controller as implemented, and Section 15.6 measures the eight cases and the shift. Section 15.7 sets the model beside the scanners, drives and controllers of real cells, Section 15.8 states the limitations, Section 15.9 draws the results together, and the exercises of Section 15.10 work through the budget, the law and the cases.

15.2 From Fenced Cells to Shared Workspaces

For most of their history industrial robots moved where no person was meant to be, and robot accidents were being analyzed for their causes and effects by 1987 [12]. The requirements for the robot itself were collected in ISO 10218, whose 2011 edition of part 1 is one of the two parts the simulator follows [3]. BSI records that edition as withdrawn on 17 March 2025, when a new edition replaced it.

Robots built to work with a person rather than apart from one came next. The collaborative robot, or cobot, was described in 2001 as a machine for direct physical interaction with a human operator, one that creates smooth and strong virtual surfaces in a shared workspace while it steers a single mechanical degree of freedom [13]. The question of how hard a robot may strike a person was then taken to measurement, as the first history segment below tells.

> ⧖ HISTORY **Crash-Testing a Robot Arm, 2009**
>
> How hard a robot may strike a person was taken to measurement. A lightweight robot arm was impact-tested at the standard crash-test facilities used for automobiles, and industrial robots of different weights were then evaluated from the results, with the influence of the robot's mass and velocity investigated. Free impacts captured only part of the problem. The study identified the quasi-static constrained impact, in which the robot presses a body part against something, as a possible serious threat to a person even from a robot of low inertia, and it classified the injuries relevant to robotics [7].

Detecting a contact without extra sensors followed, and the second history segment, at the end of this section, traces it from 2005 to 2017. The two-link planar robot of its first study is the plant that carries the observer in this simulator.

Standards then gave the methods a common form. ISO/TS 15066 was published in February 2016 as a supplement to ISO 10218 [5], and it specifies safety requirements for collaborative industrial robot systems and their work environment [4]. Its minimum protective distance for speed and separation monitoring was examined analytically and in tests, with its input values, implementation specifications and performance expectations [6]. The approach speed of the person belongs to an older standard, which positions safeguards, laser scanners among them, by the approach speeds of parts of the body, walking and upper-limb movement, and gives a method for the minimum distance from a detection zone to a hazard [14]. Compliance does not remove every hazard, and a residual hazard potential in case of an incident has been illustrated with an experiment on a standard collaboration scenario [5]. What the simulator of this chapter adds to these pieces is a record that a physical cell cannot keep, the true separation and the true stop behind every decision of its supervisor.

> ⧖ HISTORY **A Collision Sensed Without a Sensor, 2005 to 2017**
>
> In 2005 a collision between a robot arm and an obstacle of unknown shape and place was treated as a fault of the arm's actuating system. A dynamic fault detection and isolation technique, which needs neither acceleration nor force measurements, found a collision at any point along the arm with the joint encoders alone and named the link that was struck. Its results were simulations of a two-link planar robot [15]. A survey of 2017 reviewed, extended and evaluated experimentally such model-based detection, isolation and identification from proprioceptive sensors only. It reported that the physically motivated solution had been applied to manipulators, humanoids, flying robots and commercial products [11].

15.3 The Shared Cell

Figure 15.1 shows the cell from above. The origin is at the arm's shoulder, $+y$ points from the robot toward the aisle, and the front edge of the robot's table lies at $y = 0.90\,\mathrm{m}$. A feeder P at $(-0.56, 0.26)$ m supplies parts, a fixture Q at $(0.56, 0.26)$ m receives them, and a hand-over point H at $(0.06, 0.72)$ m, with a lamp, is where a part can be handed to the person. The twelve integrated states are the two joint angles q_1 and q_2, their rates ω_1 and ω_2, and for the person's torso and hand a two-dimensional displacement from a scripted anchor with its rate. The person's script, the gripper, contact episodes, the reflex timing and the truth audit are discrete and are updated after each integration step, never inside the derivative function. The signals are classified in Table 15.1.

15.3.1 The Arm

The two joints have vertical axes, so the links sweep the plane of the table, gravity does no work, and the arm obeys Equation (15.1), the arm dynamics of Equation (1.2) with a friction scale and a contact torque added.

$$M(q)\,\ddot{q} + C(q,\dot{q})\,\dot{q} + \varphi F(\dot{q}) = \tau + J_c^{\mathsf{T}} f_c, \qquad |\tau_i| \le \tau_{i,max} \tag{15.1}$$

where $q = [q_1, q_2]^{\mathsf{T}}$, both angles counterclockwise seen from above, F is the joint friction with its true scale φ, and $J_c^{\mathsf{T}} f_c$ is the torque of any contact force. The links are uniform rods of lengths l_1, l_2

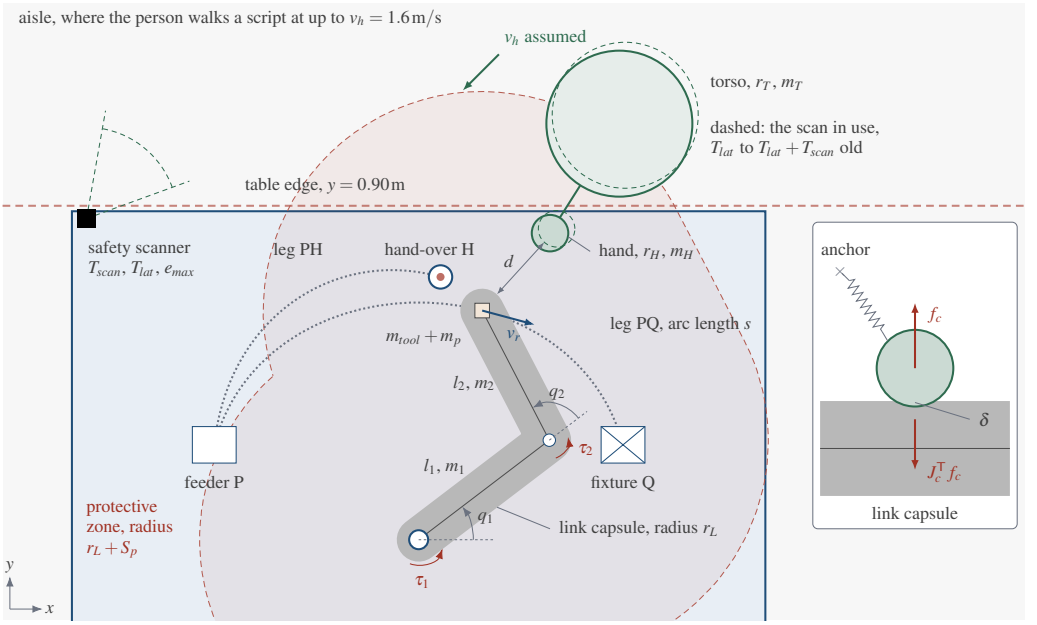

Figure 15.1 Plan view of the shared cell, to scale apart from the magnified contact on the right. The upper arm of length l_1 turns through q_1 from the x axis and the forearm of length l_2 through q_2 from the upper arm, driven by τ_1 and τ_2, and each link is a capsule of radius r_L. The tool carries m_{tool} and, while a part is held, its mass m_p along leg PQ at v_r. The person's torso and hand are discs of radii r_T and r_H carrying the effective masses m_T and m_H on supports to a scripted anchor, and the dashed discs are the scan the supervisor is using, from which it measures the separation d. The shaded capsule of radius $r_L + S_p$ is the protective zone at $S_p = 0.54\,\mathrm{m}$, a little above its value at rest. A contact overlap δ pushes the body with f_c and the arm with $-f_c$ through the contact Jacobian J_c. The scanner is drawn at the table's corner but has no geometry in the model.

Table 15.1 Signals of the Cell by Role, With Their Defaults

Symbol	Role	Default	Range
q_1, q_2	states, CVs of the servos	114.8°, 87.0°	the paths
ω_1, ω_2	states, joint rates	0	—
ξ_T, $\dot{\xi}_T$	states, torso on its support	0	—
ξ_H, $\dot{\xi}_H$	states, hand on its support	0	—
τ_1, τ_2	MVs, joint torques	held 10 ms	± 20, $\pm 14\,\mathrm{N\,m}$
v_{cap}, α	MV of the supervisor, ceiling	v_{nom}, 100 %	0–100 %
—	DV, the person's script	per case	scripted legs
λ_v	DV, walking speed scale	100 %	30–100
λ_m	DV, true part mass	100 % of m_L	50–200
φ	DV, true gearbox friction	100 %	50–300
—	DV, scanner dropout	off	1.5 s event
q_m	measured, joint encoders	17 bit	—
τ	measured, drive torques	as applied	—
$\hat{p}_T$, $\hat{p}_H$	measured, scan of the person	40 ms, 100 ms late	$\leq 40\,\mathrm{mm}$ off
—	measured, hand-over button	off	on, off
d, S_p	supervisor's separation, budget	—	—
d_{true}	truth, the audit's separation	—	—

and masses m_1, m_2, and a point mass $m_{tip} = m_{tool} + m_p$ at the tip carries the wrist, the gripper and any part held. With $l_{ci} = l_i/2$ and $I_i = m_i l_i^2/12$ the inertia coefficients and the remaining terms are given by Equation (15.2).

$$\begin{aligned} a_1 &= m_1 l_{c1}^2 + I_1 + (m_2 + m_{tip})\, l_1^2 + J_{1m} \\ a_2 &= m_2 l_{c2}^2 + I_2 + m_{tip}\, l_2^2, \qquad a_3 = (m_2 l_{c2} + m_{tip} l_2)\, l_1 \\ M &= \begin{bmatrix} a_1 + a_2 + 2a_3 \cos q_2 & a_2 + a_3 \cos q_2 \\ a_2 + a_3 \cos q_2 & a_2 + J_{2m} \end{bmatrix} \\ C\dot{q} &= a_3 \sin q_2 \begin{bmatrix} -(2\dot{q}_1\dot{q}_2 + \dot{q}_2^2) \\ \dot{q}_1^2 \end{bmatrix} \\ F_i &= b_i \dot{q}_i + f_{ci} \tanh(\dot{q}_i / 0.05) \end{aligned} \tag{15.2}$$

where J_{1m} and J_{2m} are the motors' rotor inertias reflected through their gears, b_i and f_{ci} are the viscous and Coulomb coefficients of each joint, and the hyperbolic tangent smooths the Coulomb term over 0.05 rad/s. At the feeder pose the empty arm has $M_{11} = 2.273\,\mathrm{kg\,m^2}$ and $M_{22} = 0.483\,\mathrm{kg\,m^2}$, and the 3 kg part raises them to 3.416 and $0.963\,\mathrm{kg\,m^2}$. The generalized momentum $p = M(q)\dot{q}$ obeys $\dot{p} = \tau + C^{\mathsf{T}}\dot{q} - \varphi F + J_c^{\mathsf{T}} f_c$, because $\dot{M} = C + C^{\mathsf{T}}$, and for this arm $C^{\mathsf{T}}\dot{q} = [0,\ -a_3 \sin q_2\, (\dot{q}_1^2 + \dot{q}_1\dot{q}_2)]^{\mathsf{T}}$. That identity is what the collision observer of Section 15.5 integrates.

Velocity is affine along each rigid link, so the fastest point of the arm is the elbow or the tip, and its speed $v_r = \max(l_1|\dot{q}_1|,\ \|J(q)\dot{q}\|)$ is the robot speed in the protective distance. On every path of this program the tip is the fastest point. A body struck at a point of the arm along a direction $\hat{n}$ meets the operational-space mass $m_{eff} = 1/(\hat{n}^{\mathsf{T}} J_p M^{-1} J_p^{\mathsf{T}} \hat{n})$ of that point [16]. Along its direction of travel the gripper presents 3.3 to 5.4 kg empty and 6.5 to 9.0 kg with the part.

15.3.2 The Person

The person is a torso disc of radius $r_T = 0.20\,\mathrm{m}$ and a hand disc of radius $r_H = 0.05\,\mathrm{m}$ in the plane of the links. A script moves an anchor through legs that walk to a point, accelerating and stopping at $2\,\mathrm{m/s^2}$ and never faster than 1.6 m/s, pause, reach the hand out to at most 0.62 m from the torso's center, hold it, withdraw it, move it across, lean over the table, press the hand-over button, wait for the lamp, wait for a part with a given patience, or repeat. A walking-speed disturbance λ_v scales every walk. The torso and the hand are each a mass on a soft support to the anchor, so a strike pushes a body part of the right effective mass, as shown in Equation (15.3).

$$\begin{aligned} m_B \ddot{\xi}_B &= f_{c,B} - m_B \omega_B^2\, \xi_B - 2\zeta_B m_B \omega_B\, \dot{\xi}_B \\ p_B &= a_B(t) + \xi_B \end{aligned} \tag{15.3}$$

where $B \in \{T, H\}$, $a_B(t)$ is the scripted anchor, ξ_B the displacement from it, $m_T = 40\,\mathrm{kg}$ and $m_H = 0.6\,\mathrm{kg}$ are the effective masses of the chest and hand regions, $\omega_T = 2\pi \cdot 2\,\mathrm{Hz}$, $\omega_H = 2\pi \cdot 5\,\mathrm{Hz}$ and $\zeta_B = 0.7$. A hand that is not reached out lies inside the torso and cannot be struck on its own. A struck person reacts after 0.15 s, straightening up from a lean or pulling a held hand back.

15.3.3 Contact

Where a body disc overlaps a link capsule, a unilateral spring pushes the two apart along the line between their closest points, as shown in Equation (15.4).

$$\begin{aligned} \delta &= r_L + r_B - \|p_B - p_L\| \\ F &= \max\big(0,\ k_B\,\delta + c_B\,\dot{\delta}\,\min(1,\ \delta/1\,\mathrm{mm})\big), \qquad f_c = F\,\hat{n} \end{aligned} \tag{15.4}$$

where p_L is the closest point on the link's centerline, $\hat{n}$ the unit vector from it to the body, k_B the body region's effective stiffness and $c_B = 2\zeta_c\sqrt{k_B m_B}$ with $\zeta_c = 0$, an elastic contact. The force enters the

arm through the Jacobian of the contact point. The body-region values of Table 15.2 are illustrative teaching values, in the manner of the body model of ISO/TS 15066 as the application describes them, and not the standard's tables. For two bodies meeting through such a spring the peak force is $v\sqrt{\mu k}$, Equation (1.15) with the reduced mass μ in place of the moving mass, and the application's check drives the tip at 0.6 m/s into a hand and a chest and finds peaks of 121.0 and 186.5 N against predictions of 119.2 and 187.0 N, with contacts lasting 9 ms and 40 ms. An episode starts above 0.5 N and its peak is compared with the region's limit.

15.3.4 Scanner and Separation

Every T_{scan} the safety scanner captures the torso and, once the hand is out of the torso's outline, the hand, as shown in Equation (15.5).

$$\hat{p}_B(t_k) = p_B(t_k) + \mathrm{sat}_{e_{max}}\big(\sigma\, \nu_{B,k}\big), \qquad \sigma = e_{max}/3 \tag{15.5}$$

where $\nu_{B,k}$ is a pair of independent standard normal deviates from the scenario's seeded generator and the saturation shortens any error longer than $e_{max} = 40\,\mathrm{mm}$ to that length, so $\sigma = 13.3\,\mathrm{mm}$ per axis. The scan reaches the supervisor T_{lat} later, and none is taken during a dropout. The scanner sees the person anywhere, and the arm never hides the person from it. The supervisor's separation is the distance from its own links to the scanned body, as shown in Equation (15.6).

$$d = \min_{B} \Big[\mathrm{dist}\left(\hat{p}_B,\ \mathscr{L}(q_m)\right) - r_L - r_B \Big] \tag{15.6}$$

where $\mathscr{L}(q)$ is the pair of link segments from the shoulder to the elbow and from the elbow to the tip, evaluated at the encoder angles q_m, and the minimum runs over the torso and a scanned hand. The fixed base is not part of either segment. The true separation d_{true} is the same expression with the true angles and positions, and a check against 2000 sampled link points agrees with the segment distance to $4.4 \times 10^{-8}\,\mathrm{m}$.

15.3.5 Program, Paths and Braking Capability

The program picks a part at P, carries it along leg PQ, places it at Q and returns along QP, dwelling 0.5 s at each station, and a cycle is a part placed or handed over. Asked for a part, it takes leg PH to the hand-over point and HP back. Each leg is a cubic Bézier curve, PQ with control points $(-0.42, 0.78)$ and $(0.42, 0.78)$ m, 1.470 m long, and PH with $(-0.52, 0.62)$ and $(-0.22, 0.80)$ m, 0.885 m long. It is parameterized by its arc length s through a table of 240 intervals, each integrated by three-point Gauss-Legendre quadrature. A reference $(s, \dot{s}, \ddot{s})$ gives the joint reference through the inverse kinematics with the elbow bent counterclockwise, $q = \mathrm{IK}(B(s))$, $\dot{q} = J^{-1} t\, \dot{s}$ and $\ddot{q} = J^{-1}(t\ddot{s} + \kappa \dot{s}^2 - \dot{J}\dot{q})$, with t the unit tangent and κ the curvature vector. The tool reaches $y = 0.74\,\mathrm{m}$ at most, 0.80 m with the capsule.

The torque that follows a path at speed $\dot{s}$ while braking at a is affine in a, $\tau = \tau_0 + a\,\tau_a$, so each drive's limit bounds the deceleration as shown in Equation (15.7).

$$\begin{aligned} a_{cap}(s, \dot{s}) &= \min_{j} \ \frac{\mathrm{sgn}(\tau_{a,j})\, \tau_{j,max} - \tau_{0,j}}{\tau_{a,j}} \\ a_{wc} &= \min_{\text{program}} \ a_{cap}(s, \dot{s}) \end{aligned} \tag{15.7}$$

where the worst case a_{wc} runs over the four legs, 49 positions on each and speeds of 0 to 1.2 m/s in five steps, and a point whose constant-speed torque already exceeds a limit counts as zero. With the arm's nominal friction the worst case is $4.74\,\mathrm{m/s^2}$ empty, 3.13 with the 3 kg part and 2.04 with a 6 kg one. The correct braking model plans on $a = 0.9\, a_{wc}$ for the payload the program believes it holds, 4.26, 2.81 or $1.84\,\mathrm{m/s^2}$, and the optimistic one on a single brochure figure of $5\,\mathrm{m/s^2}$ whatever the payload.

Table 15.2 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
Arm and drives			
l_1, l_2	link lengths	0.45, 0.40 m	V
m_1, m_2	link masses	8, 4 kg	V
m_{tool}	wrist and gripper	1.5 kg	V
r_L	link capsule radius	60 mm	V
$\tau_{1,max}$, $\tau_{2,max}$	drive torque limits	20, 14 N m	V
J_{1m}, J_{2m}	rotor inertias, reflected	0.10, 0.03 kg m^2	H
b_1, b_2	viscous friction	1.0, 0.5 N m s/rad	H
f_{c1}, f_{c2}	Coulomb friction	1.0, 0.5 N m	H
—	encoder count	$2\pi/131072$ rad	H
Program			
a_{nom}	path acceleration	1.5 m/s^2	V
—	dwell at P and Q	0.5 s	V
—	longest hand-over wait	8 s	H
Person and contact			
r_T, r_H	torso and hand radii	0.20, 0.05 m	C
—	walking acceleration, top speed	2 m/s^2, 1.6 m/s	C
—	reach, reaction time	0.62 m, 0.15 s	C
f_T, f_H, ζ_B	body supports	2, 5 Hz, 0.7	H
ζ_c	contact damping ratio	0	H
Body regions, m_B, k_B and F_{max} (illustrative)			
—	hand	0.6 kg, 75 N/mm, 280 N	C
—	lower arm	2.0 kg, 40 N/mm, 320 N	C
—	chest	40 kg, 25 N/mm, 280 N	C
—	head	no contact permitted	C
Scanner and budget			
e_{max}	true scan error bound	40 mm	V
v_h, Z_r	approach speed, robot uncertainty	1.6 m/s, 10 mm	C
T_{sup}, T_{act}	supervisor and servo cycles	20, 10 ms	C
—	braking share, brochure figure	0.9, 5 m/s^2	C

15.3.6 Initial State and Numerics

Table 15.2 lists the constants, and Table 15.3 the supervisor and controller settings, which are the student's controls. Every case starts with the arm at rest at the feeder, holding nothing, and the person at the case's own starting point, and the scanner's queue is seeded with error-free samples of the person taken before $t = 0$ so that the supervisor has a reading at once. The twelve states advance by the fourth-order Runge-Kutta method on a fixed 2 ms step, in the engine of Section 1.6.1. Whenever a body part lies within $(v_r + 2.5\,\mathrm{m/s}) \times 6\,\mathrm{ms} + 4\,\mathrm{mm}$ of a link, the reach of the arm and a fast body part in three steps, the step is split into eight sub-steps of 0.25 ms, because a hand strike rings at 360 to 390 rad/s, and the peak force is taken at every sub-step. The servos and the observer run every 10 ms from $t = 0$, the supervisor every 20 ms, and the historian logs at 20 Hz. Inside a step the controller acts first, the integrator second, and then come the discrete events, the outputs and the log, as in Figure 1.11. The scan errors come from a generator seeded at 20261103, so every reset replays the same scans. Halving the step moves the worst violation of case 4 from 506 to 517 mm, and a second halving leaves it at 517 mm, and moves the first hand strike of case 8, with its friction model corrected, from 68.13 to 68.18 N. Sixteen sub-steps move a hand strike's peak by 0.008 % from eight, and without torque or friction the arm keeps its kinetic energy over 2 s to 2×10^{-13}. The application passes its 43 verification checks.

15.3.7 Assumptions

A planar arm at table height. The two-link arm stands in for a six-axis collaborative robot working in one horizontal layer, with no wrist, no vertical motion and rigid links. Its drive limits are teaching values chosen so that its braking, 2 to 5 m/s^2 at the tool, is of the order that collaborative robots' stopping data show, as the application states.

A person of two discs on a script. The legs, the head and the arm between shoulder and hand have no geometry, the head appears only as a body region a risk assessment could name, and the person never changes their mind.

An ideal scanner. It has a period, a latency and a bounded error, and nothing else. It has no detection capability, response-time tables, muting or multiple fields, the arm never hides the person from it, and its error is independent from scan to scan.

Elastic transient contact. Only free impacts are modeled. Clamping against a fixture, the quasi-static constrained impact that the crash tests found dangerous [7], pressure limits and edges are outside the model.

A path-following stop. The protective stop decelerates the reference along the path at the braking model's deceleration, which the application calls a category 2 stop in the terms of IEC 60204-1, and the arm then holds position under power. Brakes and stops of other categories are not modeled.

Friction as a scale factor. The true gearbox friction is the nominal model times φ, the same shape at every temperature and load, although the static friction of a real robot joint changes with both load torque and temperature [17].

Teaching values. Every limit and threshold is a teaching value. The method follows ISO 10218 and ISO/TS 15066, and real values come from those standards, ISO 13855, the robot maker's stopping data and a risk assessment [3, 4, 14].

15.4 The Budget and the Law

The protective separation distance is where the cell stops being a trajectory tracking problem. It turns the scanner's delay, the arm's speed and braking and the uncertainties into one moving limit, and the supervisor schedules the speed against it.

15.4.1 The Budget

The supervisor writes the protective separation distance in the form of ISO/TS 15066, as shown in Equation (15.8).

$$S_p = \underbrace{v_h(T_r + T_s)}_{S_h} + \underbrace{v_r T_r}_{S_r} + \underbrace{\frac{v_r^2}{2a}}_{S_s} + C + Z_d + Z_r$$
$$T_s = \frac{v_r}{a}, \qquad T_r = T_{scan} + T_{lat} + T_{sup} + T_{act} \tag{15.8}$$

where $v_h = 1.6\,\mathrm{m/s}$ is the approach speed the application takes from ISO 13855, v_r the speed of the arm's fastest point as the encoders report it, a the braking model's deceleration, C the intrusion distance that fingers reach past what the sensor resolves, and Z_d and Z_r the position uncertainties the supervisor assumes for the scanner and the robot. It is Equation (1.18) of Chapter 1 with the two uncertainties added, written with T_r for t_r, T_s for t_s and a for A_b. The reaction time budgets the whole chain from a movement of the person to the arm starting to brake. A movement just after a scan waits up to one period to be seen, the scan then takes T_{lat} to arrive, and the supervisor and the servos are each allowed one cycle, so $T_r = 40 + 100 + 20 + 10 = 170\,\mathrm{ms}$ by default. The intrusion distance is 200 mm, which the application sizes after the $8(d_s - 14)$ mm that its model panel quotes from ISO 13855 for a sensor resolving $d_s = 40\,\mathrm{mm}$.

> **+ DID YOU KNOW? Walking Faster Than the Budget Assumes**
>
> ISO 13855 takes the approach of a person as 2000 mm/s, or 1600 mm/s where the separation exceeds 500 mm. Studies of gait put average normal walking on level ground at 1250 to 1510 mm/s, and average fast walking at 1710 to 1840 mm/s [6].

With the 3 kg part and the correct model at 1.00 m/s, the terms are $840 + 170 + 178 + 200 + 50 + 10 = 1448\,\mathrm{mm}$. The empty arm, which brakes at $4.26\,\mathrm{m/s^2}$, needs 1195 mm, a 6 kg part 1845 mm, the

brochure figure 1122 mm, and a 300 ms latency 1968 mm. At rest only $v_h T_r + C + Z_d + Z_r = 532$ mm is left, the person's approach through the reaction chain, the intrusion and the uncertainties. The stopping time enters twice, once through the arm's braking distance and again as time during which the person keeps walking. With round inputs, 1 m/s of tool speed, $T_r = 0.2$ s, 5 m/s^2, $C = 0.2$ m and 0.06 m of uncertainty, the formula gives 1.200 m, and 0.716 m at 0.25 m/s, and a shortcut that leaves out the person's approach during the stop gives 0.880 m, short by the 0.32 m a person walks while the arm brakes. The application's own check reproduces the first two.

15.4.2 The Speed Law

S_p is a quadratic in v_r, so the fastest speed whose budget fits in a distance D is its positive root, Equation (15.9), the speed law of Section 1.4.9 with these terms.

$$v_{allow}(D) = a\left[\sqrt{\beta^2 + \frac{2(D - D_0)}{a}} - \beta\right], \qquad \beta = \frac{v_h}{a} + T_r \tag{15.9}$$

where $D_0 = v_h T_r + C + Z_d + Z_r$ and $v_{allow} = 0$ for $D \le D_0$. The supervisor applies it to the measured separation less a margin $m = 50$ mm. With the 3 kg part it allows 1.00 m/s at 1.5 m, 0.50 m/s at 1.0 m, 0.28 m/s at 0.8 m and nothing at or below 582 mm, as shown in Figure 15.2(a). Full speed needs 1498 mm with the part, 1245 mm empty, 1895 mm with 6 kg and 1172 mm under the brochure figure, which is how the optimistic model buys its extra speed near a person.

15.4.3 Sampling and the Margin

The speed law acts only through scans, and a scan arrives every 40 ms. A person walking in at 1.6 m/s closes 64 mm between scans (Section 1.4.9), and in case 2 the measured separation falls by 64 to 96 mm per scan in the 0.3 s before the first stop, the arm's own motion and the scan error included. The band in which the law can slow the arm before the stop test $d < S_p(v_r)$ trips is only the margin wide, so a head-on approach at full speed can cross it between two scans. In case 2 the scan used at 3.96 s puts the person 1271 mm from the empty arm, where 1.00 m/s is still allowed, the next, at 3.98 s, reads 1207 mm and trims the ceiling to 0.95 m/s, and the one after, at 4.02 s, reads 1112 mm against an S_p of 1149 mm at 0.94 m/s, so the arm stops from nearly full speed. The margin decides how the approach ends. With no margin the first stop latches at 1.00 m/s, at 50 mm at 0.94 m/s, at 100 mm at 0.72 m/s, and at 150 mm and 200 mm the law has brought the arm to rest when the stop latches at 4.42 s. The cell completes 7 cycles in the minute with every margin tried from 0 to 200 mm and 6 with 300 mm, while the protective stops fall from 8 with no margin to 2 at 50 mm and 1 at 150 mm.

15.4.4 Hysteresis, Stop and Restart

The ceiling follows the speed law with backlash, as shown in Equation (15.10).

$$\begin{aligned} v^s_{cap} &\leftarrow \mathrm{clamp}\big(v^s_{cap},\ v_{allow}(d - m - h),\ v_{allow}(d - m)\big) \\ \mathrm{stop} &\Leftarrow d < S_p(v_r) \\ \mathrm{restart} &\Leftarrow v_r < 0.01\,\mathrm{m/s} \ \wedge\ d \ge S_p(0) + m + h \end{aligned} \tag{15.10}$$

where $h = 100$ mm is the restart hysteresis. The ceiling comes down at once when the separation shrinks and goes back up only when the separation would allow it with h to spare. A protective stop latches, and the arm restarts on its own once it is at rest and the separation exceeds $S_p(0) + m + h = 682$ mm. Without h every scan that reads a little farther lets a slowed arm creep, and the next one stops it again, which is the chattering of case 5.

15.4.5 The Stop the Budget Assumes

The budget is only as good as its deceleration. The protective stop commands a reference that slows along the path at the braking model's a, and Figure 15.2(b) shows what the loaded arm does with

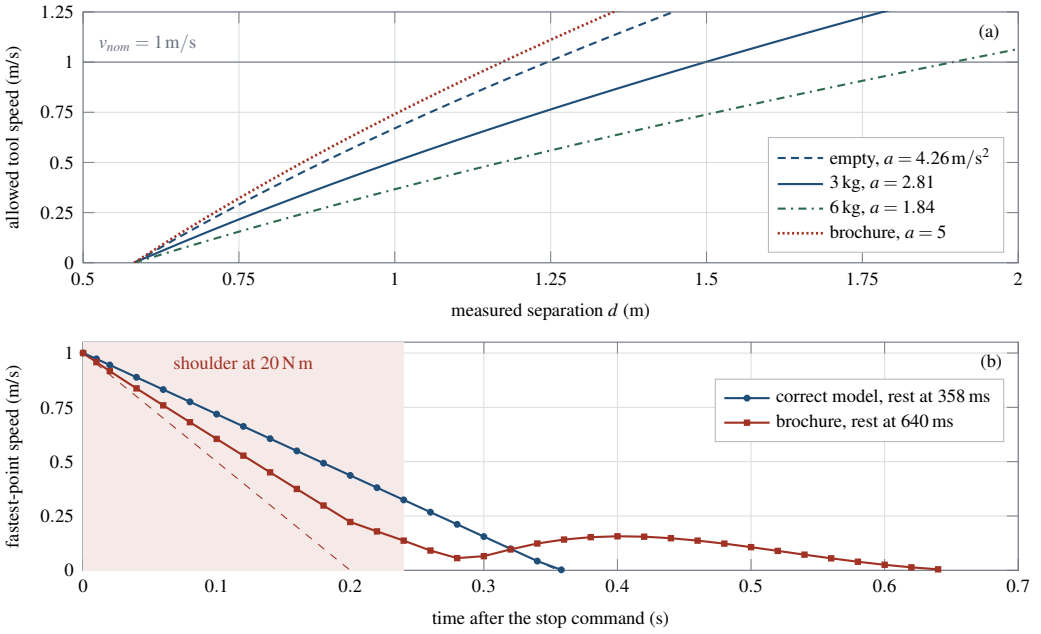

Figure 15.2 The speed law and the stop it relies on, both from the application. (a) The fastest tool speed whose protective distance fits in the measured separation less the 50 mm margin, at the default 170 ms reaction chain, empty (densely dashed), with the 3 kg part (solid), with a 6 kg part (dash-dotted) and under the brochure figure (dotted). Every curve starts at 582 mm, and full speed needs 1245, 1498, 1895 and 1172 mm. (b) The speed of the arm's fastest point after a protective stop commanded from 1.00 m/s with the 3 kg part, at a servo cycle, under the correct model (circles) and the brochure figure (squares). The thin dashed references reach rest at 355 and 200 ms. The shaded interval is where the shoulder drive sits at its limit, and the brochure's second hump is the arm moving back to a reference that has already stopped.

it from 1.00 m/s. Under the correct model the reference stops after the model's 178 mm, the drives never reach a limit, and the arm comes to rest after 178 mm in 358 ms against the model's 355 ms, within 0.1 mm of the model's distance. Under the brochure figure the reference stops after 100 mm in 200 ms. The shoulder drive cannot supply the torque, sits at its 20 N m limit for the first 240 ms, and lets the tool run 30 mm past the stopped reference before the servo pulls it back. The arm comes to rest 640 ms after the command, having moved 167 mm at its fastest point. The distance is 67 mm longer than the brochure promises, but the time is 440 ms longer, and at 1.6 m/s that time is 704 mm of the person's walk. An optimistic stop is wrong mostly in time. Chapter 9 meets the same error on a pallet robot that plans its stop on the braking of the empty robot.

🕮 **Pause.** Try Exercise 15.2 before reading on.

15.4.6 Contact Instead of Separation

On the hand-over legs the person must come close, and speed and separation monitoring would only stop the arm. Power and force limiting admits contact instead, at the speed at which the energy of a transient impact is stored in the body's spring at its force limit, as shown in Equation (15.11).

$$\frac{1}{2}\,\mu\, v^2 = \frac{F_{max}^2}{2k_B} \;\Rightarrow\; v_{PFL} = \frac{F_{max}}{\sqrt{\mu k_B}} \qquad \mu = \left(\frac{1}{m_B} + \frac{1}{m_R}\right)^{-1} \tag{15.11}$$

where $m_R = (m_1 + m_2)/2 + m_{tool} + m_L$ is the standard's estimate of the robot's effective mass as the application implements it, half the moving mass of the arm plus the tool and the programmed payload m_L. With the 3 kg part $m_R = 10.5$ kg, the chest allows 0.61 m/s, the lower arm 1.23 m/s and the hand 1.36 m/s, and the head permits no contact. Empty, $m_R = 7.5$ kg and the chest allows 0.70 m/s. The limit has to be taken for the most vulnerable region the motion could reach, which here is the chest. The estimate is conservative for this arm, whose own operational-space mass at the gripper is below it. Struck head-on at the empty-arm limit of 0.70 m/s in the pose of the application's contact check, the chest takes 219 N against its 280 N, because the arm presents only 4.3 kg there.

15.4.7 The Truth Audit

The supervisor decides from an old, noisy scan and a braking model. At every supervisor cycle in which the arm moves or is about to, the audit replaces each measured or assumed term by its true value and checks the inequality of Equation (15.12).

$$d_{true}(t_c) \;\geq\; v_h\,(\text{age} + T_{rest}) + D + C + Z_r \tag{15.12}$$

where $d_{true}(t_c)$ is the true separation between the arm now and the person where they truly were when the scan in use was taken at t_c, age is that scan's age, T_{rest} is how long the arm truly takes, from now, to come to rest, and D is how far its fastest point travels meanwhile. T_{rest} and D come from replaying the controller's own trajectory generator and servos on a copy of the true arm, with the true payload, the true friction and the drives' limits, for the T_r − age left of the reaction chain at the ceiling just set and then through the stop the braking model commands. The replay feeds the servos the true angles and rates, without the encoder quantization or the half-period lag of the backward difference, and it holds no bodies, so it is a slightly idealized copy of the controller. Rest means within 2 mm of where the arm finally stops. There is no Z_d term, because the truth has no sensor error. A cheap bound skips the replay when the person is far enough that no stop of this arm could matter. A violation episode starts when the inequality fails and ends when it holds with 20 mm to spare, and it means that a person walking straight in at 1.6 m/s from where they really were could have reached the arm while it still moved. The audit writes nothing the controller reads, and it stands aside on the legs that power and force limiting governs, where the contact force against its limit is the test. Volume 8, Chapter 13 keeps a record of the same kind, the treatment of every parcel of milk logged apart from the instruments its diversion valve acts on.

15.5 The Controller as Implemented

The controller works from its 17-bit joint encoders, the torques its drives delivered, the scans that have arrived and the hand-over button, and from nothing else. It is built in layers, a servo and a collision observer every $\Delta t = 10$ ms, a supervisor every 20 ms, and the program above them. The framework's velocity-form PID of Section 1.6.2 is not used. No loop integrates its error, so the servos carry no integral term for a torque limit to wind up and need no anti-windup. The settings are listed in Table 15.3.

15.5.1 Trajectory Generator and Servos

The reference is a point $(s, \dot{s}, \ddot{s})$ on the current leg, advanced once per control period with $\ddot{s}$ held over it, a trapezoid profile in the sense of Section 1.4.4. Its target speed is $\dot{s}^*$ of

$$\dot{s}^* = \max\big(0,\ \min(v_{prof}(s),\ v_{cap})\big), \qquad v_{prof}(s) = \min\Big(v_{nom},\ \sqrt{2\min(a_{nom}, a)\,(L - s)}\Big)$$

where v_{prof} is the leg's own profile into its station at the end of the leg's length L, and v_{cap} is the supervisor's ceiling. The reference speeds up at a_{nom} and slows at $\max(a_{nom}, a)$, the same deceleration that the budget assumes for S_s, and it advances by the trapezoid rule. The servo law of each joint is Equation (15.13).

$$\tau_i = \text{clip}_{\tau_{i,max}}\Big(\tau_{ff,i} + K_{p,i}\,e_i + K_{d,i}\big(\dot{q}_{ref,i} - \tfrac{\Delta t}{2}\ddot{q}_{ref,i} - \hat{\omega}_i\big)\Big) \tag{15.13}$$

Table 15.3 Supervisor and Controller Settings, Defaults and Ranges

Symbol	Setting	Default	Range
—	speed and separation monitoring	on	on, off
v_{nom}	tool speed at 100 %	1.00 m/s	0.2–1.2
m, h	margin, restart hysteresis	50, 100 mm	0–300
α_{man}	manual speed override	100 %	0–100
T_{lat}	scanner latency	100 ms	20–400
T_{scan}	scan period	40 ms	40–100
Z_d	declared scan uncertainty	50 mm	0–300
C	intrusion distance	200 mm	0–600
—	braking model	correct	or optimistic
m_L	programmed payload	3 kg	0–6
—	PFL on hand-over legs, region	on, chest	4 regions
—	collision detection	on	on, off
r_{th}	residual threshold	3 N m	0.5–20
T_{reflex}	declared reflex time	50 ms	10–200
K_O	observer gain	40 s^{-1}	5–100
$\hat{\varphi}$	friction in the model	100 %	50–300
K_{p1}, K_{d1}	shoulder gains	400, 40	to 1500, 120
K_{p2}, K_{d2}	elbow gains	200, 16	to 800, 60
—	computed-torque feedforward	on	on, off

Gains in N m/rad and N m s/rad, lower bounds 50, 0, 20 and 0.

where $\mathrm{clip}_c(x) = \max(-c, \min(c, x))$, $e_i = q_{ref,i} - q_{m,i}$ and $\hat{\omega}_i = (q_{m,i,k} - q_{m,i,k-1})/\Delta t$ is the unfiltered backward-difference rate, half a period old, which the shift $\frac{\Delta t}{2}\ddot{q}_{ref}$ matches. The feedforward $\tau_{ff} = \hat{M}\ddot{q}^* + \hat{C}\dot{q}^* + \hat{\varphi}\hat{F}$ is the torque that the controller's arm model assigns to the reference half a period ahead, since the torque computed at the start of a period is applied until its end, with the programmed payload m_L while a part is carried and the friction scaled by $\hat{\varphi}$. The clip at each drive's limit is the only output limit, and the torque has no rate limit. Against the diagonal inertias along the paths, the shipped gains give the shoulder 11.7 to 13.3 rad/s at damping ratios of 0.59 to 0.66 empty and 9.3 to 10.8 rad/s at 0.47 to 0.54 with the part, and the elbow 20.3 rad/s at 0.81 empty and 14.4 rad/s at 0.58 loaded. The loops are reverse acting in the convention of Section 1.3, $a = +1$, because the torque rises when a joint falls short of its reference, as the application's sign check confirms. With no gravity on the links and the friction in the feedforward, no steady torque is left for an integral term to find.

15.5.2 Collision Observer and Reflex

The generalized-momentum residual estimates the external joint torque without differentiating the joint rates [11, 15], and the implementation is the discrete form of Equation (15.14).

$$\begin{aligned} p_k &= \hat{M}(q_{m,k})\,\hat{\omega}_k, \qquad \hat{\beta}_k = \hat{C}^{\mathsf{T}}\hat{\omega}_k - \hat{\varphi}\hat{F}(\hat{\omega}_k) \\ I_k &= I_{k-1} + \left[\tfrac{1}{2}(\tau_{k-1} + \tau_{k-2}) + \hat{\beta}_k + r_{k-1}\right]\Delta t \\ r_k &= K_O\,(p_k - p_0 - I_k) \end{aligned} \tag{15.14}$$

where τ_{k-1} and τ_{k-2} are the torques the drives held over the last two periods, averaged because the backward-difference rate makes p_k the momentum of half a period ago, and the model uses the programmed payload and the modeled friction. For a constant external torque the residual approaches it through a first-order lag of rate $K_O = 40\,s^{-1}$. When either joint's residual exceeds r_{th} the reflex stops the arm and holds it for 2 s, and then until the scanned person is at least 200 mm from the arm. A reflex with no contact under way is booked as a nuisance stop, and one that comes later than the declared T_{reflex} after the first touch as late. With the model right the residual stays within about 1 N m in free motion.

Table 15.4 The Eight Cases and the Shift

#	Change from the reference line	Teaching point
1	manual override, no monitoring, walk-in at 1.2 m/s	the budget by hand
2	walk-in at 1.6 m/s, stand 5 s	the law and its price
3	$T_{lat} = 300$ ms, aisle walk 1.9 m away	time and uncertainty are distance
4	brochure 5 m/s^2, walk-in at 1.6 m/s	an optimistic stop
5	$h = 0$, person stands 1.47 m away	chattering
6	hand-over asked for, eager reach across	contact allowed, and the head
7	$m_L = 6$ kg, aisle walk	payload in stop and contact
8	no monitoring, 0.5 m/s, friction 250 %, hand in the path	threshold against model error
9	600 s of mixed traffic, $T_{lat} = 150$ ms, brochure, $h = 0$, PFL off	graded shift

15.5.3 Supervisor and Program

Every 20 ms the supervisor computes the reaction chain, the braking model's deceleration for the payload the program believes it holds, the measured separation of Equation (15.6) from the latest scan and the budget of Equation (15.8) at the measured speed, and it updates the ceiling and the latches of Equation (15.10). The chain T_r is a budget. At the default timing each scan arrives on a supervisor cycle, and the generator and the servos begin a stop in the same execution as the decision, so the code applies at most 140 ms of the 170 ms. A scan older than $T_{lat} + 3T_{scan} = 220$ ms trips a watchdog that stops the arm like a protective stop. Volume 4, Chapter 2 has a protective layer of the same kind, an override that rests at its clamp in healthy operation and is judged by its margin to a limit. On legs PH and HP and at the hand-over point, with power and force limiting on and a body region that permits contact, the ceiling is $\min(v_{nom}, v_{PFL})$ instead, and any latched stop is released. A stop, a fault or a reflex sets the ceiling to zero, and the speed override the scene shows is $\alpha = v_{cap}/v_{nom}$. A halt is counted each time the ceiling reaches zero while the arm is on a leg, whether a stop or the speed law brought it there. Asked for a part, the program waits at the feeder with the lamp lit until the scanner sees a hand within 0.30 m of the hand-over point, carries the part there, and opens the gripper for 0.4 s once it sees the hand within 0.15 m of the gripper. It gives up after 8 s at either wait, and a person whose patience runs out leaves.

In manual mode the operator's override α_{man} replaces the speed law, as the override knob of a teach pendant does, while the stop latch, the power and force limit, the watchdog and the reflex keep acting and the servos keep running. In automatic mode the manual override tracks the supervisor's, so a switch of mode moves nothing.

15.5.4 Simulation, Controller and Audit Models

Equations (15.1) to (15.6) make up the simulation model. The controller holds four internal models, the arm model of the feedforward and the observer with the programmed payload and the modeled friction, the braking model, the budget with its declared Z_d, and the estimate m_R for contact. The audit's model is the controller's generator and servos on a copy of the true arm, without contact. No control model is identified from data, and a wrong payload or friction is a wrong model in every safety function at once.

15.6 Eight Cases and a Shift

Table 15.4 lists the cases. Each opens on the reference line of Tables 15.2 and 15.3, with monitoring on, the correct braking model, the 3 kg part, the 50 mm margin, 100 mm of hysteresis, power and force limiting at the chest and detection at 3 N m, and alters that line in the ways its row gives. The numbers below come from headless runs at the default seed, over the first simulated minute unless another span is named. Every variant is set at $t = 0$ before Run, the same run as one carried through Reset, and a run that pauses at its first violation is continued. The results are collected in Table 15.5.

Table 15.5 Truth-Audit Scorecard of Each Case and Variant. Runs Last 60 s Except Where Named

Case and configuration	Result
1, override 100 %	11 cycles, 13 violations, worst 1054 mm, first at 4.18 s
1, override 50 and 25 %	8 cycles, 9 violations (598 mm); 5 cycles, 6 (410 mm)
1, override 0	no cycle, no violation
2, as shipped	7 cycles, no violation, closest 118 mm, 2 stops
2, monitoring off	11 cycles, 12 violations, worst 1054 mm
3, T_{lat} 100, 200, 300 ms, 240 s	660, 600, 540 h^{-1}, no violation
3, T_{lat} 50, 150, 400 ms, 240 s	675, 630, 495 h^{-1}, no violation
3, $Z_d = 200$ mm or $T_{scan} = 100$ ms	615 h^{-1} either way, no violation
3, monitoring off, 240 s	690 h^{-1}, 44 violations, worst 604 mm
4, brochure	7 cycles, 6 violations, worst 506 mm, first at 6.22 s
4, correct model	7 cycles, no violation, closest 61 mm
5, $h = 0$	13 stops, 178 halts, no cycle
5, $h = 100$ mm	1 stop, 2 halts, no cycle
5, h 0 and 100 mm, 90 s	17 stops, 4 cycles; 2 stops, 3 cycles
6, PFL at the chest	3 hand-overs, 8 cycles, peak 76 N, reflex 12.5 ms
6, PFL at the hand or lower arm	3 hand-overs, 8 cycles, peak 76 N
6, PFL off, or the head named	no hand-over, 1 cycle
7, 6 kg and 3 kg, 240 s	615 and 660 h^{-1}, no violation
7, 6 kg undeclared, 240 s	110 nuisance stops, 1 cycle
7, undeclared, reflex off, 240 s	630 h^{-1}, 83 violations, worst 567 mm
8, $r_{th} = 3$ N m	19 nuisance stops, 1 cycle
8, $r_{th} = 5$ N m	no nuisance stop, 7 reflexes within 16.5 ms, 6 cycles
8, $r_{th} = 12$ N m	one reflex 117 ms late, 6 cycles
8, model at 250 %	no nuisance stop, within 18.5 ms, 6 cycles
8, model right, $r_{th} = 12$ N m	no detection, 32 contacts, 8 cycles

15.6.1 The Budget in the Operator's Hands

In case 1 nothing monitors the cell and the operator holds the override while a person walks in from 2.85 m at 1.2 m/s, stands at the table for 5 s and walks back. At 100 % the arm completes 11 cycles in the minute and the audit finds 13 violation episodes, the first 4.18 s after the run starts and the worst 1054 mm short of the true protective distance, while the moving arm comes within 214 mm of the person. Held at 50 % for the minute it makes 8 cycles with 9 episodes, the worst 598 mm, and at 25 % 5 cycles with 6, the worst 410 mm. Only a standstill avoids them all, and then nothing is made. The reason is Section 15.4's 532 mm at rest. At any speed above zero the person's approach through the 170 ms chain, the intrusion distance and the uncertainties are due, and they exceed the 214 mm the path leaves.

🕮 **Pause.** Try Exercise 15.1 before reading on.

15.6.2 The Law Against a Person Walking In

Case 2 turns monitoring on and sends the person in at 1.6 m/s. The first protective stop comes at 4.02 s from 0.94 m/s, for the sampling reason of Section 15.4. The arm restarts at 4.26 s, at rest with the person still coming and 754 mm away, creeps, stops again at 4.38 s at 579 mm against an S_p of 583 mm, and restarts at 10.70 s once the person, walking away, is 709 mm off. In the minute the cell completes 7 cycles against 11 unmonitored, the audit never finds it late, its closest call leaving 118 mm, and the arm spends 30.8 s at full speed, 22.7 s slowed by the law and 6.6 s stopped. Unmonitored, the same walk gives 12 violation episodes, the worst 1054 mm. The price of the law is 4 cycles in a minute of a person coming and going.

TRY IT IN THE APP **2 · The Separation Law at Walking Speed**

Load the case and press Run. Watch the override dial as the person walks in: it falls before they reach the table, and the arm stops while they stand there. Over the minute the scorecard counts 7 cycles and no separation violated. Set Speed and separation monitoring Off on the supervisor card and run again: 11 cycles and 12 violation episodes.

15.6.3 Latency and Uncertainty Are Distance

Case 3 walks the person along the aisle 1.9 m from the robot, never within 0.99 m of the arm, with the scanner's latency at 300 ms. The chain grows to 370 ms and S_p at 1 m/s with the part to 1968 mm, so a person the arm used to pass at full speed now slows it. Over four minutes the cell makes 540 cycles an hour, against 600 at 200 ms, 660 at 100 ms and 675 at 50 ms, and it spends 158 s of the four minutes slowed at 300 ms against 52 s at 100 ms. At 100 ms a declared uncertainty of 200 mm instead of 50 mm costs 45 cycles an hour, as does a 100 ms scan period, while 25 mm gains 15. None of the monitored runs violates, with closest calls of 103 to 301 mm, because the terms are honest. The same walk unmonitored makes 690 cycles an hour with 44 violation episodes at 300 ms, the worst 604 mm, and 23 at 100 ms, the worst 276 mm, because the audit charges the person's walk through the whole reaction chain.

DID YOU KNOW? **Safety Nearly Free on a Fast Scanner**

From the simulator: along the aisle walk of case 3, the unmonitored cell makes 690 cycles an hour, and the audit, charging the person's walk through a 300 ms latency, finds 44 violation episodes. Monitored with 50 ms of latency, the cell makes 675 cycles an hour with none (Table 15.5).

15.6.4 An Optimistic Braking Model

Case 4 has the supervisor plan with the brochure's 5 m/s^2 and sends the person straight at the table at 1.6 m/s as the arm starts its loaded leg. The audit finds 6 violation episodes in the minute, the first at 6.22 s and the worst 506 mm short. At the first one the supervisor measured 1486 mm against an S_p of 970 mm at 0.77 m/s and ran at full override, while the truth asked for 1558 mm and had 1478 mm. Figure 15.3 shows the Separation view at that moment, with the measured separation, the budget, the override and the audit's true distances. The truth places the person where they really were when the 100 ms-old scan was taken, and it has the arm come to rest 620 ms after the decision, 70 ms of reaction chain and the rest braking, after 196 mm of travel. The cell completes 7 cycles with either model, and the brochure's 7 protective stops keep the arm stopped for 20.6 s of the minute against 16.5 s for the correct model's 9. With the correct model the audit's closest call leaves 61 mm. Unmonitored, the same walk gives 11 cycles and 11 episodes, the worst 1468 mm.

15.6.5 Chattering at the Boundary

Case 5 parks the person 1.47 m from the robot for a minute with no restart hysteresis, then steps them back to 1.72 m. At that distance the speed the law allows is small. The restart threshold and the speed ceiling keep the arm on the boundary, and the scanner's 40 mm of error adds variation. The run makes 13 protective stops and 178 halts of the ceiling in the first minute with no cycle completed. The halts do not need the scan error. With the true error bound at the smallest value the Show Model dialog allows, the protective stops disappear and the ceiling still halts nearly as often, because the person stands where the law allows almost nothing and the ceiling flickers on and off zero. The scan error turns that flicker into latched protective stops. With 100 mm of hysteresis the arm stops once, halts twice and waits. The hysteresis is not free. After the step back, which begins at 61.4 s, the arm speeds up only with 100 mm to spare, its first cycle ends at 65.7 s instead of 64.5 s, and each cycle takes about 8.9 s instead of 8.1 s, so the 90 s run completes 3 cycles against 4. For that it ends the chatter, 2 protective stops in the 90 s instead of 17. Smaller values bracket the choice, since 20 mm still halts the ceiling 75 times in the first minute and 50 mm 9 times. Collaborative applications are judged on



Figure 15.3 The cell in case 4, An optimistic braking model, at $t = 6.22$ s, where the application paused itself at the first separation violation, with no control changed, in the Separation view, cropped to the plant view. (1) The badge row with the separation, the state, the override and S_p. (2) The stacked bars give S_p of 970 mm as S_h 519, S_r 131, S_s (60 mm, too narrow for its value), C 200 and Z (60 mm), with the measured separation of 1486 mm below them and the margin and hysteresis marked beyond S_p. (3) The truth-audit bar reads a true d of 1478 mm at the scan against a true S_p of 1558 mm, the person walking $1.6\,\text{m/s} \times (100 + 620)$ ms while the arm truly travels 196 mm, and the chain below it reads $40 + 100 + 20 + 10 = 170$ ms with $a = 5.00\,\text{m/s}^2$ and $T_s = 155$ ms. (4) The speed law draws S_p against the tool speed up to 1.00 m/s, entirely below $d - m = 1436$ mm. (5) The status line reads Separation violated, the truth audit finds the person too close.

human factors as well as on production figures [2], and the application's briefing adds that an arm which stops and starts on every scan wears its drives and is hard for the person beside it to read.

> **TRY IT IN THE APP** **5 · A Person at the Boundary**
>
> Press Run and watch the protective stops climb while the person stands still: 13 in the first minute, with 178 halts of the speed ceiling and no cycle. Press Reset, set Restart hysteresis to 100 mm on the supervisor card and run again. The arm stops once, halts twice and waits for the person to step back.

15.6.6 Contact Allowed at the Hand-Over

In case 6 the person presses the hand-over button, comes to the table when the lamp shows that the robot is ready at the feeder, and holds out a hand. Under power and force limiting at the chest, leg PH runs at 0.61 m/s with the part and HP at 0.70 m/s empty, and the hand-overs complete at 12.14, 33.06 and 54.26 s, 3 in the minute with 8 cycles. Each time, with the part in hand, the person reaches across for the next one and meets the returning gripper 65.5 ms after the hand-over, at 76 N, 27 % of the hand's 280 N. The reflex stops the arm 12.5 ms after the first touch and has it at rest 163 ms after, while the hand bounces off it four more times, at 62, 52, 45 and 39 N. With power and force limiting off, the person at the table is inside the protective distance of any moving arm. A protective stop latches at 9.46 s while the arm waits at the feeder, the person gives up at 30.1 s, and the arm, released once they have gone, reaches an empty hand-over point at 35.2 s, waits 8 s and carries the part back, so the minute has no hand-over and one cycle. Naming the head as the region the motion

could reach forbids contact and gives the same result, an honest infeasibility rather than a slower success. Naming the hand or the lower arm lets the legs run at the full 1.00 m/s, because their limits of 1.36 and 1.23 m/s do not bind, and gives the same 76 N strike, because the gripper is still gathering speed out of the hand-over point, at 0.11 m/s, when the hand meets it. The chest is why the legs run at 0.61 m/s.

15.6.7 A Heavier Part

Case 7 programs a 6 kg part and walks the person along the aisle. The drives' worst case falls from 3.13 to 2.04 m/s^2 and the model's plan to 1.84 m/s^2, S_p at 1 m/s grows by 397 mm to 1845 mm, and the arm slows for the person sooner, 615 cycles an hour over four minutes against 660 with 3 kg, with no violation. Contact tightens as well, since m_R rises to 13.5 kg and the chest allows 0.56 m/s. With 3 kg programmed and the part truly twice as heavy, the observer, whose model expects 3 kg, reads every acceleration of the loaded arm as a contact, and the arm makes 110 nuisance stops and one cycle in the four minutes. Chapter 10 shows the torque budget that a payload left out of the controller's model breaks on an arm of the same form. Those stops hide a worse fault. With the reflex off, the same undeclared part lets the cell make 630 cycles an hour while the audit finds 83 violation episodes, the worst 567 mm short, because the braking model still plans for 3 kg.

15.6.8 A Threshold Against a Wrong Model

Case 8 switches monitoring off and runs the arm at 0.5 m/s, relying on detecting contact, while a person holds a hand in the path after the feeder and the gearboxes run at 250 % of the friction the model expects, a change of the kind that temperature causes in the friction of a robot joint [17]. With the 3 N m threshold the residual reaches 3.3 N m in free motion and trips, 19 nuisance stops and one cycle in the minute. At 5 N m the nuisance stops end and every strike on the hand, 100 N at the worst, is caught within 16.5 ms, and at 12 N m the first reflex comes 117 ms after the touch, past the declared 50 ms. Re-identifying the friction, the model at 250 %, lets 3 N m work. The first strike, 68 N at 1.312 s, is caught 18.5 ms later, the rest within 7.5 ms, the free-motion residual stays below 1.1 N m, and the arm makes 6 cycles with no nuisance stop. With the model right and the threshold at 12 N m no strike is detected at all, and the arm makes 8 cycles through 32 contacts. Monitoring is off, so the audit flags separation violations throughout, 14 in the shipped minute. A raised threshold trades false alarms for late ones, and identifying the model is the fix.

Pause. Try Exercise 15.3 before reading on.

OPEN CHALLENGE **Choosing a Collision Threshold**

Selecting the threshold on the monitoring signal, so that it avoids false positives and stays sensitive at the same time, is a major practical problem of collision detection. In an observer, unknown friction torques, quantization, measurement noise and delay all enter the estimate of the external torque, so the residual is not zero even without contact [11]. The threshold has to clear a model error that changes with load and temperature. Progress would be a threshold that adapts to what the model is known to get wrong, catching a light strike in time without stopping on every acceleration.

15.6.9 The Shift

The challenge runs 600 s of mixed traffic at 150 ms of latency and 50 mm of declared uncertainty. Its loop, a walk along the aisle, a 1.6 m/s approach toward the fixture side, 20 s standing at 1.47 m, a request for a part and a reach into the fixture, repeats every 76 to 94 s, depending on how quickly the hand-over goes. The target asks for at least 55 cycles, no violation, no contact above its region's limit and at most 30 protective stops. Changing a model parameter or a disturbance voids an attempt. Table 15.6 compares the configurations. As shipped, with the brochure figure, no hysteresis and power and force limiting off, the cell makes 39 cycles with 47 violation episodes, the worst 466 mm, 46 stops and 552 halts, and no hand-over completes. Each shipped setting fails a different part of the

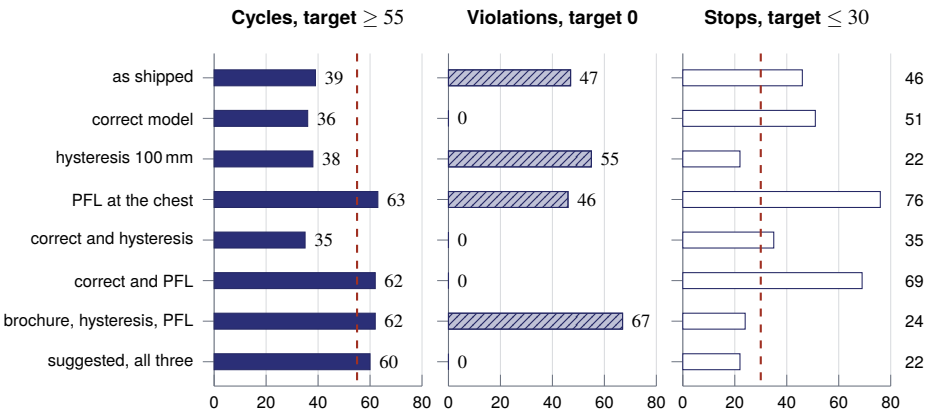

Figure 15.4 The shift of 600 s in its eight configurations, from Table 15.6: cycles (solid bars), violation episodes (hatched bars) and protective stops (open bars), with the targets as dashed lines. Notice that each single change fails a different part of the target, and only the suggested configuration, with the correct model, hysteresis and power and force limiting together, passes all three.

Table 15.6 The Shift of 600 s in Eight Configurations

Configuration	Cycles	Hand-overs	Violations	Stops
as shipped	39	0	47	46
correct model	36	0	0	51
hysteresis 100 mm	38	0	55	22
PFL at the chest	63	8	46	76
correct and hysteresis	35	0	0	35
correct and PFL	62	8	0	69
brochure, hysteresis and PFL	62	8	67	24
suggested, all three	60	7	0	22

target. The correct model alone removes the violations and loses 3 cycles, hysteresis alone removes the chatter, and power and force limiting alone finds the throughput in the hand-overs. The suggested configuration makes 60 cycles, 7 of them hand-overs, with no violation, no contact and 22 stops, and the audit's closest call leaves 120 mm. Figure 15.4 sets the cycles, violations and stops of the eight configurations against their targets.

15.7 Technology in Practice

The cell of Section 15.3 reads joint encoders, drive torques, a hand-over button and a delayed safety scanner, and it moves the arm through two torque-limited joint drives. Figure 15.5 draws its loops with real devices in the place of the schematic's symbols.

What Measures It

A fenceless cell detects people with safety-rated sensing equipment. The accepted technologies are light curtains, laser scanners that respond to diffuse reflection, standardized in IEC 61496-3, and pressure-sensitive mats. A sensor without a performance standard, such as a thermal camera, may only support a primary safety-rated one. Mats respond quickly but resolve poorly, while cameras and lasers resolve well but depend on lighting, processing time, capture rate and occlusion [6]. The scanner of Section 15.3, 40 ms between scans, 100 ms late and up to 40 mm off, stands for such a device with its occlusion left out. Contact has no sensor of its own. The drive torque follows from the motor current, and an observer compares it with the model, but unknown friction, quantization, noise

and delay leave a residual even with no contact [11]. The residual of about 1 N m in free motion in Section 15.5 stands for that floor.

> **TECHNOLOGY Sensor: The Safety Laser Scanner**
>
> A laser scanner sweeps a laser rangefinder over a plane and finds the range to what it strikes from the time of flight divided by the speed of light [18]. In one testbed for speed and separation monitoring, two scanners about 0.4 m above the floor watched the work area. Software removed a recorded background, grouped the returns into legs and people, and a Kalman filter estimated each person's position and velocity [19]. The range depends on the target's reflectivity and the beam's angle, and a beam that falls across an edge can report phantom points [18]. A person hidden behind another is not seen at all [19]. In the model a lost scan leaves the supervisor an ever older position of the person until its watchdog stops the arm.

What Moves It

Each joint has a servo drive whose torque saturates, and a protective stop is a deceleration commanded through those drives. Arms built for power and force limiting are typically lighter, slower and of shorter reach, while a conventional arm converted to collaborative work has no power and force limiting and needs speed and separation monitoring instead [6]. The torque limits of Table 15.2 stand for the drives, and the worst-case braking of Equation (15.7) for the stopping data a maker would measure.

> **TECHNOLOGY Final Element: Joint Drives and the Protective Stop**
>
> A robot is stopped by removing power from its drives, by applying brakes or by controlling the motors against the motion. IEC 60204-1 sorts stops into three categories: category 0 removes power and applies the brakes at once, category 1 brakes under control and then removes power, and category 2 brakes under control and keeps the power on [6]. The stopping time and distance measured under ISO 10218-1 are functions of the initial speed and the payload. Pose, path, brake wear and a shifted center of mass of the payload change them, so they should be measured again from time to time [6].

Where the Controller Runs

A discrete-part plant may give each cell a supervisor above a machine controller and a robot controller, and such loops run with cycle times from milliseconds to minutes [20]. In one testbed the monitoring ran on an external computer, which sent slow, pause and resume commands over a network socket to tasks written in the robot's own language, and the slow command set the controller's general speed override [19]. A delay of about 220 ms was measured on one arm between being at a place and reporting it [6]. In this chapter the servos and the observer act every 10 ms and the supervisor every 20 ms, and the ceiling reaches the arm as the speed override α.

> **IN PRACTICE Measuring the Reaction Time and Proving the Stop**
>
> No standard method measures the reaction time of a sensor that detects people, so the integrator supplies it. One way drives the robot along a line past a laser beam wired to its inputs, records where it stops at several speeds, and fits the stop positions to a quadratic in speed whose coefficients give the reaction time and the worst-case deceleration. Repeated trials of a rail-mounted six-axis arm gave 0.113 s [6]. A testbed proved its stops with a mannequin on a rolling platform, pushed toward the moving robot while an independent measurement system tracked both, and the robot had to stop before the contact threshold [19].

What Is Modernizing It

Sensing is the first development. Safety-rated cameras and sensors on the robot are expected to be critical for cells that are not structured, and people will have to be told apart from parts [6]. Thermal cameras have been used to tell a person from a machine, so that the protective distance can shrink

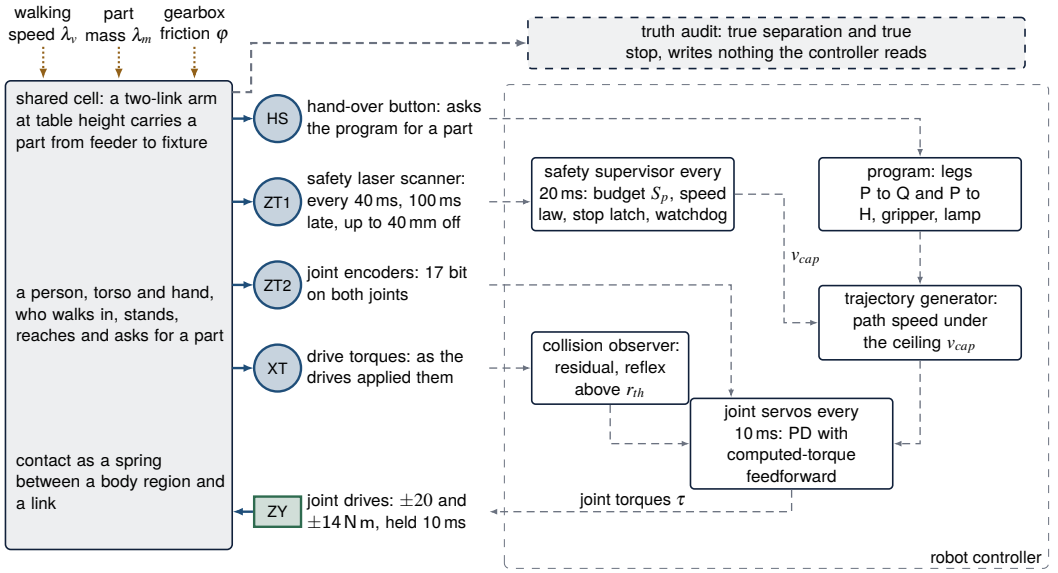

Figure 15.5 The loops of the chapter with real devices in place of the schematic's symbols. A safety supervisor acting every 20 ms reads the safety laser scanner ZT1 and the encoders ZT2, computes the budget and sets the speed ceiling. The trajectory generator follows the program under that ceiling, the servos turn the reference into torques for the joint drives ZY, and the collision observer compares the drive torques XT with the arm model to trigger the reflex. Notice that the truth audit (dashed) reads the cell itself and feeds nothing back, and that no sensor measures contact.

when what approaches is not a person [21]. Detection without sensors is the second: momentum-observer methods for collisions have reached commercial products [11]. Wireless networking is the third. A collaborative work-cell testbed links programmable logic controllers, robot arms and machine tools over wired and wireless industrial networks and measures how the wireless links perform [22].

15.8 Limitations

The arm is planar, rigid and at table height, the person two discs on a script, and the scanner an ideal sensor with a bounded error, no occlusion and no detection capability. Contact is elastic and transient, so clamping, pressure limits and the constrained impacts the crash tests found dangerous [7] are absent. The protective stop follows the path, and brakes, other stop categories and the safety-rated monitoring of the stop itself are not modeled. The truth audit replays the stop with the controller's own servos on the true arm, an idealized validation rather than a measurement on a robot.

Every limit is a teaching value. The body-region masses, stiffnesses and force limits are illustrative, and the standard's estimate m_R stands for the robot's effective mass. The method follows the 2011 editions of ISO 10218, now replaced, and ISO/TS 15066 of 2016, and a real cell takes its values from the current standards, ISO 13855, the robot maker's stopping data, measured transient forces and a risk assessment [3, 4, 14]. The targets of the shift were chosen for teaching, and none of the numbers is a measure of industrial performance.

OPEN CHALLENGE **Proving a Fenceless Cell Safe**

Speed and separation monitoring still lacks methods to verify and validate that it works and keeps working. No method captures the measurement and performance uncertainty of the robot system as a whole, no standard exists for measuring where a person is, and the minimum distance has not been validated as safe for every robot configuration. People vary, sensors are uncertain, and it is hard to tell when a robot actually began its stop [6]. Progress would be a repeatable test, with a human surrogate and an independent measurement synchronized to the robot, that grades a whole cell as the truth audit grades the simulator.

15.9 Conclusion

A collaborative cell earns its missing fence only if the arm can always stop in time, and this simulator makes the budget behind that guarantee, and its truth, measurable. At 1.00 m/s with a 3 kg part the protective separation distance is 1448 mm, and at rest it is still 532 mm, more than the 214 mm the arm's path leaves a person at the table, so no override lets the arm work safely beside them. The speed law costs 4 cycles in a minute of a person coming and going, each 100 ms of latency from 100 to 400 ms costs 45 to 60 cycles an hour, and a brochure braking figure buys no cycle while it produces 6 violation episodes, the worst 506 mm, because the loaded arm takes 640 ms rather than 200 ms to stop. Hysteresis turns 17 protective stops into 2 for one cycle in 90 s, power and force limiting at the chest makes the hand-over possible at 0.61 m/s, and a threshold raised against a wrong friction model detects a strike 117 ms late. The shift improves from 39 cycles with 47 violations to 60 cycles with none. The model is planar with an ideal scanner and elastic contact. A fuller model would add a vertical axis, clamping contact and pressure limits, a scanner with occlusion and detection capability, and a person who changes course. The simulator treats safety as a feedback loop whose price is paid in throughput and whose verdict belongs to the truth.

KEY POINTS

- A collaborative arm without a fence is safe only if it can always stop before a person walking in could reach it.
- At rest the protective separation distance is still 532 mm, more than the 214 mm the path leaves a person at the table.
- A brochure braking figure buys no cycle and gives 6 violation episodes, because the loaded arm takes 640 ms rather than 200 ms to stop.
- Hysteresis turns 17 protective stops into 2 for one cycle in 90 s.
- A threshold raised against a wrong friction model detects a strike 117 ms late, and identifying the model is the fix.

Where else this applies. A record kept apart from the instruments a supervisor acts on also judges the diversion valve of Volume 8, Chapter 13.

References

[1] A. M. Zanchettin, N. M. Ceriani, P. Rocco, H. Ding, and B. Matthias, "Safety in human-robot collaborative manufacturing environments: Metrics and control," *IEEE Transactions on Automation Science and Engineering*, vol. 13, no. 2, pp. 882–893, 2016.

[2] F. Vicentini, "Collaborative robotics: A survey," *Journal of Mechanical Design*, vol. 143, no. 4, p. 040802, 2021.

[3] *Robots and robotic devices. Safety requirements for industrial robots. Robots*, BSI British Standards Standard BS EN ISO 10 218-1:2011, Sep. 2011.

[4] *Robots and robotic devices. Collaborative robots*, BSI British Standards Published Document PD ISO/TS 15 066:2016, Mar. 2016.

[5] M. J. Rosenstrauch and J. Krüger, "Safe human-robot-collaboration-introduction and experiment using ISO/TS 15066," in *2017 3rd International Conference on Control, Automation and Robotics (ICCAR)*, 2017, pp. 740–744.

[6] J. A. Marvel and R. Norcross, “Implementing speed and separation monitoring in collaborative robot workcells,” *Robotics and Computer-Integrated Manufacturing*, vol. 44, pp. 144–155, 2017.

[7] S. Haddadin, A. Albu-Schäffer, and G. Hirzinger, “Requirements for safe robots: Measurements, analysis and new insights,” *The International Journal of Robotics Research*, vol. 28, no. 11-12, pp. 1507–1527, 2009.

[8] K. L. Hobbs, M. L. Mote, M. C. L. Abate, S. D. Coogan, and E. M. Feron, “Runtime assurance for safety-critical systems: An introduction to safety filtering approaches for complex control systems,” *IEEE Control Systems*, vol. 43, no. 2, pp. 28–65, 2023.

[9] C. Byner, B. Matthias, and H. Ding, “Dynamic speed and separation monitoring for collaborative robot applications – concepts and performance,” *Robotics and Computer-Integrated Manufacturing*, vol. 58, pp. 239–252, 2019.

[10] J. A. Marvel, “Performance metrics of speed and separation monitoring in shared workspaces,” *IEEE Transactions on Automation Science and Engineering*, vol. 10, no. 2, pp. 405–414, 2013.

[11] S. Haddadin, A. De Luca, and A. Albu-Schäffer, “Robot collisions: A survey on detection, isolation, and identification,” *IEEE Transactions on Robotics*, vol. 33, no. 6, pp. 1292–1312, 2017.

[12] B. C. Jiang and C. A. Gainer, “A cause-and-effect analysis of robot accidents,” *Journal of Occupational Accidents*, vol. 9, no. 1, pp. 27–45, 1987.

[13] M. A. Peshkin, J. E. Colgate, W. Wannasuphoprasit, C. A. Moore, R. B. Gillespie, and P. Akella, “Cobot architecture,” *IEEE Transactions on Robotics and Automation*, vol. 17, no. 4, pp. 377–390, 2001.

[14] *Safety of machinery. Positioning of safeguards with respect to the approach speeds of parts of the human body*, BSI British Standards Standard BS EN ISO 13 855:2010, May 2010.

[15] A. De Luca and R. Mattone, “Sensorless robot collision detection and hybrid force/motion control,” in *Proceedings of the 2005 IEEE International Conference on Robotics and Automation*, 2005, pp. 999–1004.

[16] O. Khatib, “Inertial properties in robotic manipulation: An object-level framework,” *The International Journal of Robotics Research*, vol. 14, no. 1, pp. 19–36, 1995.

[17] A. C. Bittencourt and S. Gunnarsson, “Static friction in a robot joint—modeling and identification of load and temperature effects,” *Journal of Dynamic Systems, Measurement, and Control*, vol. 134, no. 5, p. 051013, 2012.

[18] 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.

[19] S. Szabo, W. Shackleford, R. Norcross, and J. Marvel, “A testbed for evaluation of speed and separation monitoring in a human robot collaborative environment,” National Institute of Standards and Technology, Gaithersburg, MD, NIST Interagency/Internal Report 7851, 2012.

[20] 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.

[21] U. B. Himmelsbach, T. M. Wendt, N. Hangst, P. Gawron, and L. Stiglmeier, “Human–machine differentiation in speed and separation monitoring for improved efficiency in human–robot collaboration,” *Sensors*, vol. 21, no. 21, p. 7144, 2021.

[22] Y. Liu, M. Hany, K. Montgomery, and R. Candell, “A collaborative robot work-cell testbed for industrial wireless communications,” National Institute of Standards and Technology, Gaithersburg, MD, NIST Interagency/Internal Report 8365, 2021.

15.10 Exercises

15.1 ⌕ Concept check · 2 min

In scenario 1 nothing monitors the cell, and a person stands at the table for 5 s, where the moving arm passes within 214 mm of them. Why does no speed override above zero keep the arm working safely while they stand there?

(a) Below about half speed the drives brake so gently that the arm's stopping distance stays longer than the gap

(b) The scanner cannot resolve a person who stands inside the intrusion distance of the edge of the table

(c) A lower override lengthens the reaction chain, because the supervisor then waits longer between its cycles

(d) Even a crawling arm owes the person's approach through the reaction chain, the intrusion and the uncertainties

15.2 Concept check · 2 min

In scenario 4 the supervisor plans the stop of the loaded arm on a brochure braking figure of $5\,\mathrm{m/s^2}$, while the drives can supply only about $3.13\,\mathrm{m/s^2}$ with the part. How does the true stop from $1.00\,\mathrm{m/s}$ compare with the one the model promises?

(a) It is somewhat longer in distance and much longer in time, and the person keeps walking through that time

(b) It is much longer in distance, because both drives saturate and the arm coasts at nearly full speed

(c) It matches the promise, because the servos follow the stopping reference to within a millimeter

(d) It is shorter than promised, because the joint friction adds its torque to the braking of the drives

15.3 Concept check · 2 min

In scenario 8 the gearboxes run at 250 % of the friction that the observer's model expects, and the 3 N m residual threshold gives 19 nuisance stops in a minute. What does raising the threshold to 12 N m, with the model unchanged, do?

(a) The nuisance stops end and every strike is caught sooner, because the residual is larger

(b) The nuisance stops end, but a strike can be caught later than the declared reflex time

(c) Nothing changes, because the observer gain and not the threshold sets the reflex timing

(d) The nuisance stops grow, because a higher threshold admits more of the friction error

15.4 Theory and calculation · 5 min

The arm moves at $v_r = 1.00\,\mathrm{m/s}$ with the 3 kg part. The correct braking model plans on $a = 0.9 \times 3.13\,\mathrm{m/s^2}$. The reaction chain is $T_r = 170\,\mathrm{ms}$, the approach speed $v_h = 1.6\,\mathrm{m/s}$, the intrusion distance $C = 200\,\mathrm{mm}$ and the uncertainties $Z_d = 50\,\mathrm{mm}$ and $Z_r = 10\,\mathrm{mm}$.

(a) Compute a, the stopping time T_s and the terms S_h, S_r and S_s of Equation (15.8). *Checkpoint:* $a = 2.817\,\mathrm{m/s^2}$, $T_s = 355\,\mathrm{ms}$, $S_h = 840\,\mathrm{mm}$, $S_r = 170\,\mathrm{mm}$, $S_s = 178\,\mathrm{mm}$, each within 1 mm.

(b) Add the terms to S_p, and give S_p with the arm at rest. *Checkpoint:* 1448 mm within 2 mm, and 532 mm.

(c) How much of S_h is the person's walk during the stop, and how does it compare with the arm's own braking distance S_s? *Checkpoint:* 568 mm, more than three times S_s.

15.5 Theory and calculation · 6 min

Use the speed law of Equation (15.9) with the values of the previous exercise, $a = 2.817\,\mathrm{m/s^2}$ and $T_r = 170\,\mathrm{ms}$, and the margin $m = 50\,\mathrm{mm}$ that the supervisor subtracts from the measured separation.

(a) Compute β and D_0. *Checkpoint:* $\beta = 0.738\,\mathrm{s}$ and $D_0 = 532\,\mathrm{mm}$.

(b) Compute the speed allowed at a measured separation of 1.0 m. *Checkpoint:* 0.50 m/s, within 0.01 m/s.

(c) Find the measured separation at which full speed, 1.00 m/s, is first allowed, the separation at or below which nothing is allowed, and the separation a stopped arm needs before it restarts with $h = 100$ mm. *Checkpoint:* 1498 mm within 2 mm, 582 mm and 682 mm.

15.6 Theory and calculation · 5 min

Power and force limiting allows the speed of Equation (15.11), with $m_R = (m_1 + m_2)/2 + m_{tool} + m_L$. The links weigh 8 and 4 kg and the wrist and gripper 1.5 kg. The chest has $m_B = 40$ kg, $k_B = 25$ N/mm and $F_{max} = 280$ N, and the hand $m_B = 0.6$ kg, $k_B = 75$ N/mm and the same 280 N.

(a) With the 3 kg part, compute m_R, the reduced mass μ for the chest and the speed the chest allows. *Checkpoint:* 10.5 kg, 8.32 kg and 0.61 m/s.

(b) Compute the speed the hand allows, and say which limit the hand-over legs must obey and why. *Checkpoint:* 1.36 m/s; the chest's.

(c) Repeat (a) for a 6 kg part. *Checkpoint:* 13.5 kg and 0.56 m/s.

15.7 App challenge · 9 min

Load scenario 4, An optimistic braking model, and change no setting. Press Run at 1x with the default seed. The run pauses itself at the first separation violation.

(a) On the Separation view, read the measured separation d and S_p on the supervisor card, and the true separation and true S_p on the truth audit line, with the clock reading. *Checkpoint:* 6.22 s; $d = 1486$ mm against $S_p = 970$ mm; true 1478 mm against 1558 mm.

(b) Press Run again, at 2x if you wish, and let the run reach 60 s, pressing Run after any pause. Read the Separation violated and Cycles tiles. *Checkpoint:* 6 episodes, worst 506 mm; 7 cycles.

(c) Press Reset, switch to Advanced and set Braking model to Correct on the scanner card before Run. Run 60 s and read the same two tiles. What did the optimistic model buy? *Checkpoint:* never violated; 7 cycles; nothing.

15.8 App challenge · 9 min

Load scenario 3, Latency and uncertainty are distance, which ships with a scanner latency of 300 ms, and use the default seed and a speed of 5x. Each run lasts 240 s.

(a) Run with no change. Read the Cycles tile, the Separation violated tile, and T_r on the scanner card. *Checkpoint:* 36 cycles, 540 h^{-1}; never; 370 ms.

(b) Press Reset, set Scanner latency to 100 ms before Run and run again. Read the Cycles tile and compute the cost of each 100 ms of latency between the two runs. *Checkpoint:* 44 cycles, 660 h^{-1}; 60 cycles an hour per 100 ms.

(c) Press Reset, keep 100 ms and set Position uncertainty the supervisor assumes to 200 mm before Run. Read the Cycles tile and say why none of the three runs violates. *Checkpoint:* 41 cycles, 615 h^{-1}; the terms are honest.

15.9 App challenge · 8 min

Load scenario 5, A person at the boundary, which ships with no restart hysteresis. Use the default seed and a speed of 5x, and read the Protective stops and Cycles tiles at the end of each run.

(a) Run 60 s with no change. *Checkpoint:* 13 protective stops, 0 cycles.

(b) Press Reset, set Restart hysteresis to 100 mm before Run, and run 90 s, through the step back at about 61 s. *Checkpoint:* 2 protective stops, 3 cycles.

(c) Press Reset, set Restart hysteresis back to 0 before Run and run 90 s. What does the hysteresis cost and what does it buy? *Checkpoint:* 17 protective stops, 4 cycles.

15.10 Application area · 3 min

A cell integrator proposes to copy the chapter's 1448 mm into the protective field of a safety scanner on a real collaborative cell that runs at 1.00 m/s with a 3 kg part. Give two reasons the number cannot be used, and name where the real values must come from.

15.11 Application area · 4 min

A person reaches a hand into fixture Q while the arm places a part there.

(a) Which contact hazard does the model leave out, and why does the speed of Equation (15.11) not cover it?

(b) Name two idealizations of the scanner in the model, and say what a real scanner's occlusion by the arm would do to the separation the supervisor uses.

15.12 Review the agent · 5 min

An AI assistant is asked how to pass the challenge, Ten minutes of mixed traffic, and replies as follows. "Keep the brochure braking figure, because it lets the arm run faster near the person. Add 100 mm of restart hysteresis and power and force limiting at the chest on the hand-over legs. The shift then makes 62 cycles with 8 hand-overs and 24 protective stops, so it beats the target of 55 cycles and 30 stops with more cycles than the configuration the app suggests." Find the flaw and say how you would check it.

CHAPTER 16

Summary and Outlook

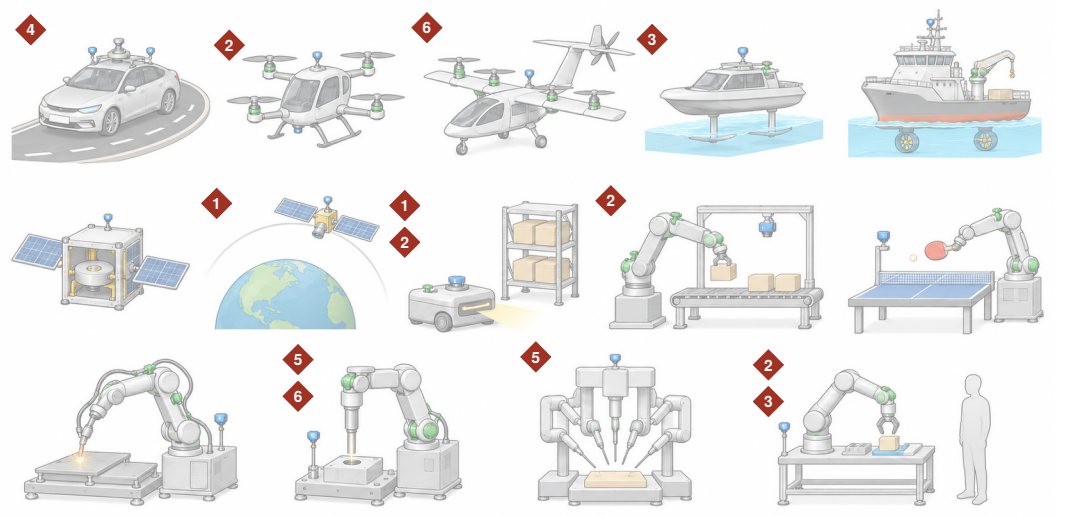

Open gaps 1 trusted uncertainty 2 verified told values 3 proof of infeasibility 4 model mismatch 5 soft contact 6 event-keyed benchmarks

Notice that the pallet robot, the peg and the cobot carry two gaps each, and that no gap marks the vessel, the small satellite, the table tennis arm or the welding arm. Illustration, not to scale.

IN PLAIN TERMS **What Fourteen Moving Machines Taught**

The main idea. Each chapter of this book let a machine move on its own: a car, an aircraft, a boat, a satellite, a robot in a warehouse and arms that pick, weld and operate. None could see all it needed, so each acted on what it believed.

The challenge. A number written down in advance can be wrong. A guess takes time to improve, and some choices must be made before the result is known. When a machine cannot stop in the space it can see, the safe answer is to slow down or not go.

Why it matters. The machines that worked were told what was checked, kept a margin and were judged against the truth.

Fourteen chapters put a moving machine to work on a task it could not see in full: a car, an air taxi, an aircraft handing its weight to a wing, a foiling boat, a vessel on station, two satellites, a pallet robot and six arms. Section 1.4.10 named four ways a controller closes the gap between what it believes and what is true. It can be told, it can estimate, it can bound the unknown and keep a margin, or it can commit to a prediction before the outcome is known. This chapter adds the measured results to that reading and compares how each system completed its task under uncertainty. It then gives a short guide for choosing a structure, looks at where practice is heading, reports how the volume was made with AI agents and what they got wrong, and lists the problems that remain open.

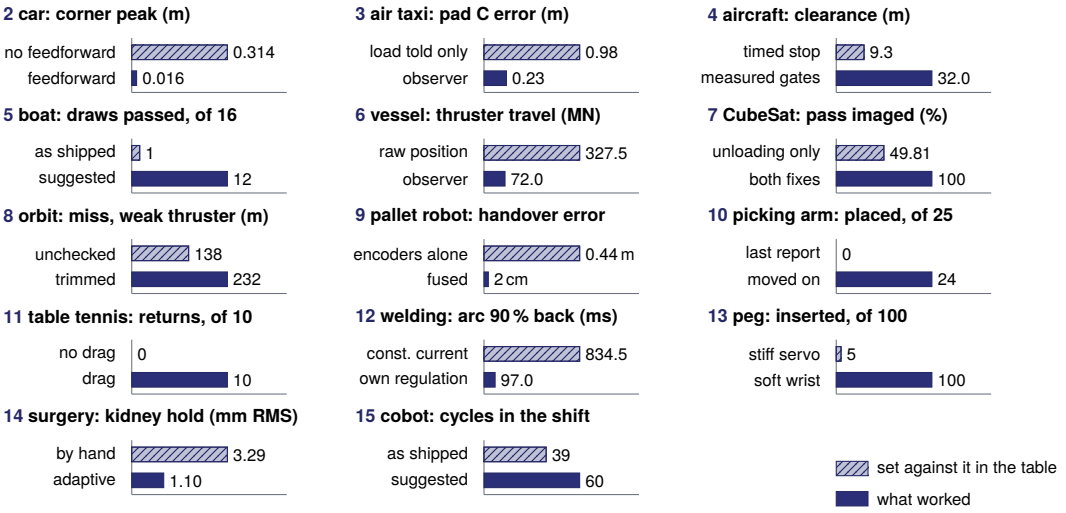

Figure 16.1 One pair of results from each chapter, with the numbers as the "Measured result" column of Tables 16.1 and 16.2 prints them. Each panel has its own scale from zero. A larger value is better in the panels of Chapters 4, 5, 7, 8, 10, 11, 13 and 15, and a smaller one in the others. The pallet robot's pair is printed in two units and drawn on one axis in meters, and a result of zero draws no bar. The cobot's suggested shift also ends with no violation against 47. Notice that in most rows the better result came from a change of structure, of what the controller was told or of what it estimated, not from tuning alone.

16.1 What the Chapters Showed

Tables 16.1 and 16.2 extend Table 1.2. Each row names what was uncertain, a structure that completed the task, one that did not, and the measured result with the table or section of the chapter that reports it. Every number is the chapter's own, and every result is scored against the true state that the app logs beside the controller's estimate. Figure 16.1 draws one pair of results from each row.

The columns of the two tables make five points. The first four follow the four ways of Section 1.4.10.

Telling worked when the sheet was right. A value given to the controller beforehand acted at once. The air taxi told its load climbed 0.05 m behind where fixed tuning fell 1.78 m behind, the picking arm told its payload made no hits where the arm model alone made 11, and curvature feedforward cut the corner peak of Chapter 2 nearly twentyfold. A told value carries no error signal of its own, so a wrong one is believed in full. The pallet robot planned with its empty braking hit a cart at 0.74 m/s despite an emergency stop. The cobot planned with a brochure figure gained no cycle and produced six violations, because the loaded arm took 640 ms rather than 200 ms to stop (Section 15.4).

Estimating worked when the estimator had a state for the unknown. The vessel's observer separated wave motion from drift, and a vessel model with a bias estimate but no wave states already captured 92 % of the reduction in thruster travel (Section 6.6). The satellite's filter with a bias state held the camera to 0.054° through the Sun's blind window, where the same filter without it drifted 2.35°. Only an estimator with a spin state read a 40 rps backspin before the bounce in Chapter 11 (Table 11.4). An estimate takes time to form, and the chapters measured it. After the squall of Chapter 6 the bias estimate covered 63 % of the change in 237 s, and a gain lowered for a heavy sea took 17.35 instead of 5.05 s to notice a real step.

Table 16.1 How Each Machine Completed Its Task Under Uncertainty, Chapters 2 to 8

Chapter	Uncertain	What worked	What failed	Measured result
2, lane keeping	Speed, curvature, a misaligned linkage, the rack's lag.	A lateral gain divided by speed, with curvature feedforward. A predictive controller against the misalignment.	No curvature feedforward. The predictive controller on a 0.45 s rack that its 0.10 s model misjudges.	Corner peak 0.314 against 0.016 m (Table 2.5). Recovery to $1/e$ in 0.84 and 0.85 s at 13.0 and 30.6 m/s (Section 2.6). IAE 0.326 against 8.850 m s on the misalignment (Table 2.7), unstable above 26.4 m/s on the slow rack (Section 2.6.7).
3, air taxi	Load, balance, wind among buildings.	Load-sheet feedforward. A disturbance observer with a 0.67 s time constant.	Fixed tuning with a heavy load. Feedforward without the observer at the gap's jet.	Heavy climb 1.78 m behind with fixed tuning, 0.05 m when told (Table 3.7). Pad C error 0.99, 0.98 and 0.23 m, only the observer passes (Table 3.8), on 20 of 20 seeds against 5 for each other mode (Section 3.6).
4, lift-plus-cruise aircraft	Airspeed, lift margin, mass, a failed motor or pitot.	Rotor stop gated on measured airspeed and lift margin. An airspeed derived from lift beside the pitot.	A timed schedule tuned for 250 kg in calm air. The pitot alone.	Stop at 37.9 against 55.7 m/s, clearance 9.3 against 32.0 m (Table 4.8). Unchecked pitot: into the wood at 38.5 s. A failed motor leaves 4.8 % of margin: divert (Table 4.7).
5, foiling boat	Waves and sensor loss, inside a band 0.80 m tall between slam and ventilation.	Contouring. A cascade on the measured heave rate. Slowing off the foils in a sea too large.	Platforming in a 0.5 m following sea. Every setting tried in a 1.4 m head sea.	Cascade contouring clean on 17 of 17 draws, the PID's on 7 (Section 5.6). Challenge passed on 1 of 16 draws as shipped, 12 as suggested (Table 5.6). None of 33 795 runs flew the head sea past 12.49 s.
6, vessel on thrusters	Wave motion inside the measurement, the environmental force, a heading the job leaves free.	An observer with bias and wave states, bias feedforward, the minimum-thrust heading.	The raw position handed to the loops. A slow integral in place of the feedforward.	Thruster travel 327.5 against 72.0 MN for a mean offset of 0.55 against 0.51 m (Table 6.4). Heading: 2297 to 738 kW. Squall without feedforward: 7.7 m off, past the disconnect radius with the waves (Section 6.6).
7, CubeSat attitude	Gyro bias, stored momentum, a failed wheel.	A filter with a bias state together with magnetic unloading. The eigenaxis slew.	Either fix alone. Any tuning tried after a wheel fails.	Pass imaged 24.12, 49.81, 20.64 and 100 %: shipped, unloading only, bias state only, both (Table 7.6). Blind-window error 2.35 against 0.054°. Failed wheel 81.56 to 84.75 % (Table 7.5).
8, satellite in low orbit	Drag, the debris covariance, thruster output.	A predictive planner on the mean altitude. Collision probability, not miss distance. A burn checked and trimmed.	A deadband deciding every 6 h. A plan that trusts a weak thruster.	Storm: 6.08 h below the band against none. A pass predicted 1.6 km away carries $P_c = 1.8 \times 10^{-4}$. Weak thruster: 138 m unchecked, 232 m trimmed (Table 8.7). Campaign: 3 unresolved passes as shipped, 0 suggested (Table 8.8).

Table 16.2 How Each Machine Completed Its Task Under Uncertainty, Chapters 9 to 15

Chapter	Uncertain	What worked	What failed	Measured result
9, pallet robot	Position, braking with a load, space it cannot see.	A fused estimate, braking planned for the payload, slowing for unseen space.	The empty robot's braking with a pallet on. Encoders alone. Unseen space treated as clear.	Cart hit at 0.74 m/s against a stop 0.25 m short. Handover 0.44 m past against 2 cm. Worker hit at 0.72 m/s against 0.92 m kept. Shift fails at 11.1 s, passes 3 of 3 in 120.6 s (Table 9.6).
10, picking arm	Report age, belt speed, payload.	Each report moved on by the encoder travel since its stamp. Meeting the product at belt speed. The payload in the arm model.	Aiming at the last report, at any latency. Waiting at rest. The arm model alone.	Placed 24 of 25 against 0. At rest 0 placed and 60 failed, moving with it 30 placed. Arm model alone: 11 hits, torques 47.1 and 24.3 N m. Challenge 0 to 98 placed (Table 10.4).
11, table tennis arm	Drag, spin, the time left to swing.	A flight model with drag, a commit 200 to 275 ms after the serve, a minimum-jerk swing, a filter with a spin state.	A predictor without drag. Commits too early or too late. Stepped commands. A slow camera.	Returns 0 and 10 of 10 without and with drag. Commits at 50, 200 and 450 ms: 16, 19 and 11 of 20. Slow camera: 0 of 10 in all 144 designs (Table 11.4). Match 0–11 shipped, 11–9 suggested (Table 11.5).
12, arc welding cell	Torch height, a seam that wanders ahead of the arc.	The arc's own regulation. The current read as a height. A preview filed by distance.	A constant-current arc. A look-ahead configured 20 mm short.	After a 3 mm step the arc is 90 % regrown in 97.0 ms, at constant current in 834.5 ms. Short look-ahead: 1.451 mm against 1.455 mm predicted, 372 mm unfused (Section 12.6). Graded seam 1198 to 0 mm defective (Table 12.6).
13, peg in a hole	Hole position, friction, a late force signal.	Compliance: a 3 N/mm spring, a soft wrist, contact detection and a search.	A 200 N/mm servo against a 0.4 mm offset. A force loop behind 80 ms of delay.	Damaged at 136.1 N against inserted at 13.7 N. At 80 ms: 42.2 N true, 13.4 N read. Of 100 parts: stiff 5, impedance 66, sequence 87, sequence with a soft wrist 100 at 30.2 N (Table 13.5).
14, surgical instrument	Tremor, delay, organ motion, tissue creep.	Motion scaled 3:1 with a 6 Hz filter. An adaptive harmonic predictor.	A 2 Hz filter. A force hold at 2.5 N on creeping tissue.	Artery contacts 6, 0 and 3 for 1:1 at 30 Hz, 3:1 at 6 and 3:1 at 2 Hz. Kidney 3.29 by hand, 1.10 mm RMS adaptive. Tear at 13.3 s (Table 14.5). Delay times motion 86 % of the error, 91 % of it removed (Table 14.6).
15, collaborative arm	The person's motion, latency, braking with a load.	A separation budget on the true braking. Restart hysteresis. Power and force limiting at the chest.	A brochure braking figure. No hysteresis. Any speed at the table.	Brochure: 6 violations, the worst 506 mm, against none. Stops 17 against 2 in 90 s (Table 15.5). Shift 39 cycles with 47 violations against 60 with none (Table 15.6).

Bounding cost time, and sometimes the bound said do not go. Slowing for space the robot cannot see kept 0.92 m from a worker at a cost of 8.7 s. The cobot's speed law cost 4 cycles in a minute of a person coming and going, and every 100 ms of latency from 100 to 400 ms cost 45 to 60 cycles an hour. Where the margin could not be met, the right answer was to decline the task. The aircraft with a failed motor diverted. The foiling boat in a 1.4 m head sea was safest kept off its foils. At the cobot's table the protective distance at rest, 532 mm, exceeds the 214 mm the path leaves a person, so no speed override works there.

Committing traded certainty for time. The table tennis arm returned 16 or 17 of 20 serves when it committed early, 19 at 200 to 275 ms, and 10 to 12 from 350 ms on, when the swing no longer fit although the prediction was sharper (Section 11.4). The satellite in low orbit paid 1.20 cm/s to clear a pass with 8.8 h of lead, 2.10 with 5.8 h and 5.08 with less than an orbit (Section 8.7). In Chapter 14 the interaction of the surgeon's delay with the organ's motion held 86 % of the error of a hold, and an adaptive predictor of the breathing removed 91 % of it.

Structure decided before tuning, and the verdict belonged to the truth. In Chapter 6 the three PID loops accounted for almost none of the difference between a good and a poor system, and the heading cut the thruster power by 68 % with nothing in the control system changed. The cascade on heave rate flew 17 draws of a sea that the PID flew 7 times, and the satellite reached a full pass only with both fixes. Every app scored against a truth the controller never saw, and the gap was often silent. The pallet robot handed its load 0.44 m past the station with the truth outside its own ellipse. The vessel at a low observer gain wandered to 4.14 m while the trace reported that all was well.

16.2 Choosing a Method

The chapters suggest an order of questions for a new moving machine. Figure 16.2 draws the same questions as a flowchart.

1. *What does the task score, and against which truth?* Score the true state, and log it beside the estimate, as the audit of Chapter 15 does.
2. *What can the controller be told, and how is the sheet checked?* Feed forward the load, the payload and the curvature (Chapter 3, Chapter 10, Chapter 2). Take a braking figure from a test of the loaded machine, not from a brochure (Chapter 9, Chapter 15).
3. *Is part of the measured motion real and not an error?* Give the estimator a state for it, as the wave states of Chapter 6 and the platform-or-contour choice of Chapter 5 do, and a bias state for any drifting sensor (Chapter 7).
4. *How old is each piece of information?* Move a report forward by the motion since its stamp (Chapter 10), key a preview by distance rather than time (Chapter 12), and predict across a delay that multiplies motion (Chapter 14).
5. *Can the machine stop or abort inside what it can see?* Set the speed from believed braking and visible space, gate a mode change on a measured margin (Chapter 4), and add hysteresis where a bound flickers (Chapter 15). Decide in advance what the machine does when the margin is not there.
6. *Must it commit before it knows?* Find the window in which the prediction is good enough and the action still fits (Chapter 11), and plan early when lead time is cheap (Chapter 8).
7. *Does contact or a shared actuator turn an error into a force or a limit?* Make the contact compliant before raising any gain (Chapter 13), and allocate among actuators with their stored momentum and free variables in view (Chapter 7, Chapter 6).

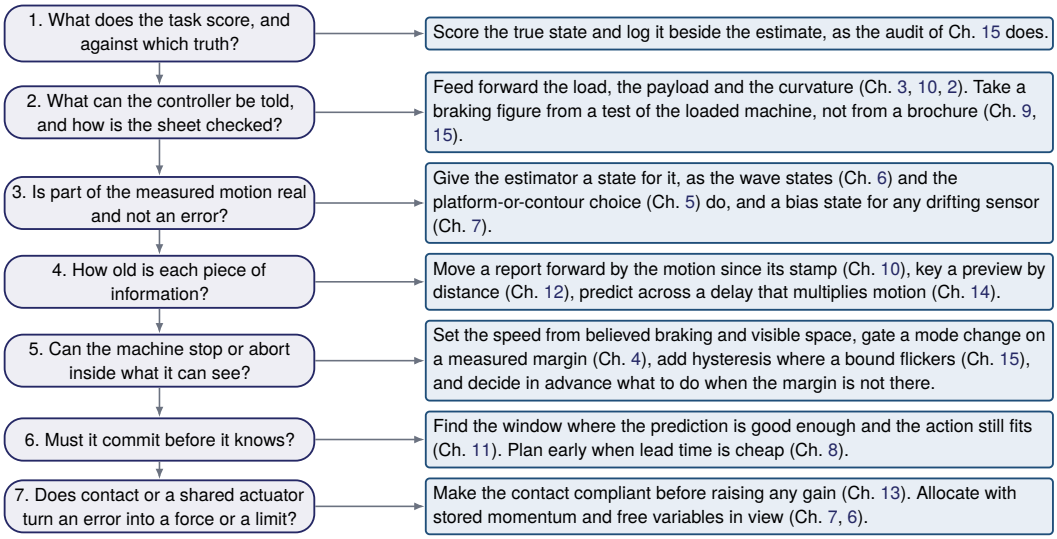

Figure 16.2 The seven questions of this section in order, with the list's answer and the chapters where it decided the result. Notice that the truth and the told values come first, and that the questions of timing and commitment follow the questions of what the machine knows.

A predictive controller earned its place where it carried a constraint or a model the simpler law lacked. In Chapter 2 it removed the standing error of a misaligned linkage, and in Chapter 8 it never let the mean altitude reach the band through a storm. Its gains were often smaller than they first looked. On the foiling boat it rode case 6's sea at 0.74 m/s^2 against the shipped cascade's 1.30, but a cascade retuned by the same rule rode at 0.90 and rode the 0.65 to 1.0 m following seas more smoothly (Table 5.5). The satellite's predictive law reached the next site in 39.6 against 52.1 s, a gain that came from staying inside the wheels' limits rather than from looking ahead (Section 7.6.11). Volume 9 treats that choice in full.

16.3 Where Practice Is Heading

Evidence that cannot come from driving alone. A statistical study found that demonstrating the safety of autonomous vehicles by test driving would take hundreds of millions, and in some cases hundreds of billions, of miles, and concluded that driving alone is not a practical way to assess it [1]. An analysis of autonomous vehicle safety names validating inductive learning against novel inputs and reaching the dependability of a full fleet as open challenges, and suggests that the largest may be a single design and deployment process that joins every specialty's safety concerns [2]. In the editor's view, evidence for moving machines will come increasingly from simulators scored against a logged truth, with seeds and disturbances published so that a result can be checked.

Learned controllers trained in simulation. An autonomous racing drone trained by deep reinforcement learning in simulation, combined with data collected in the physical world, won several races against each of three human champions and set the fastest recorded race time [3]. A survey of foundation models in robotics finds that models pretrained on internet-scale data may improve perception, decision-making and control, and names the scarcity of robot data, safety guarantees, uncertainty quantification and real-time execution as open challenges [4]. A review of safety filters, which monitor a controller's commands and intervene to keep the system safe, proposes one framework for them and argues that deployment needs model-based methods together with data-driven

ones, whose guarantees are less well understood [5]. The speed law of Chapter 15, the envelope of Chapter 9 and the gates of Chapter 4 are filters of that kind.

> ⧖ HISTORY **1999 to 2025: Gyros and Stars in a Small Cube**
>
> The first history box of Chapter 1 left the lunar spacecraft's gyros on a stable member, realigned in flight by sighting stars. A NASA survey of small spacecraft technology records what followed. In 1999 two universities developed the CubeSat, an educational platform built on a cube 10 cm on a side, and CubeSats are now a common form of small spacecraft that can weigh only a few kilograms. Three-axis control with satellite navigation had flown for decades on spacecraft of about 100 kg, and only in the past few years has it reached the smallest classes. Fiber-optic and micromachined gyros carry the attitude between star tracker updates, which typically arrive a few times a second [6]. In the editor's reading the arc ends where Chapter 7 begins: the same stars and the same gyro bias, in a cube.

Machines among people. A survey of human-robot collaboration in industry describes collaborative robots entering the market and names guaranteed safe interaction as the first challenge, ahead of intuitive interfaces for the operators who program them [7]. A review of autonomous mobile robots in intralogistics reports their introduction in manufacturing, warehousing, cross-docks, terminals and hospitals, with decisions moving from a central unit to the robots themselves [8]. In the editor's view the believed braking and the truth audit of Chapter 15 will matter more as such fleets share aisles with people.

Autonomy in the operating room and in orbit. A framework for regulatory, ethical and legal discussion identifies six levels of autonomy for medical robots [9]. A review of surgical autonomy places the bulk of commercial platforms at level 0, with no decision autonomy [10]. In orbit, collision avoidance has become an essential part of satellite operations. The European Space Agency released conjunction messages collected from 2015 to 2019 for a competition to predict the final collision risk of an encounter from its earlier messages [11]. The planner of Chapter 8 flies on that kind of prediction, and the editor's view is that both domains will keep a person in charge of the decision while software proposes it.

Weight measured on board. United States guidance on aircraft weight and balance lets an operator use an onboard system that measures the actual weight and the location of the center of gravity. Such an operator may not need to narrow the loading envelope for variation in passenger weight and seating, although it must still allow for the system's own tolerances and keep a load manifest [12]. The load sheet of Chapter 3 is the told value that such a system checks. In the editor's view the braking figures of Chapter 9 and Chapter 15 still wait for the same step, from an assumed value to one measured in service.

Learned models with physics inside. A framework for robotic welding joins physical laws to a model learned from data. It predicted weld instability with an accuracy of up to 98 % over the next 50 ms and 86 % up to 1 s ahead, and it reduced the dependence on costly data, which is scarce for unstable welds and for low-volume work on complex paths [13]. The cell of Chapter 12 carries a physical model of the arc and no learned one. In the editor's view a learned model that keeps the physics is easier to bound under the mismatch of gap 4 of Section 16.5 than one that does not.

Mobile robots tested with their loads. The published performance standards for autonomous industrial vehicles address navigation, obstacle description and detection, environmental conditions, vehicle configuration and terminology. In NIST experiments heavy payloads affected the vehicles when they accelerated, on slippery floors and at the transition from floor to ramp, and the report recommends that the capabilities standard combine ramps and other conditions with the payload the

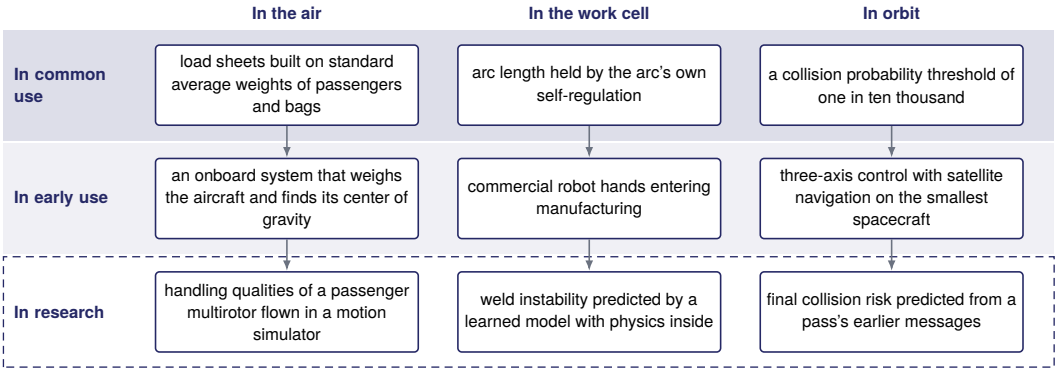

Figure 16.3 Developments in three threads. The top lane holds what the cited source calls standard, most common or widely accepted, the middle lane what it reports in service or entering use, and the bottom lane what it reports from a simulator, a laboratory or a competition on recorded data, in the editor's reading. Sources by thread: in the air [12, 15], in the work cell [13, 16, 17] and in orbit [6, 11, 18]. Notice that each thread moves from a value assumed or held by a rule toward one measured or predicted in service.

vehicle carries [14]. The braking that the pallet robot of Chapter 9 planned without its load is the same combination seen from the controller's side. Figure 16.3 sets developments of this section in three threads.

16.4 Engineering With AI Agents

An AI agent is a language model that works toward a goal in a loop. It plans, calls tools such as a compiler, a simulator or a browser, reads the results and continues [19]. In the editor's view, agents can now carry out much of the work that used to fill the hours of a control project. They derive a model from balances or fit one to plant data. They write the simulator and its tests, code the controller, tune it against stated criteria, script commissioning checks, write the documentation and watch running loops for degradation. The apps and chapters of this series were produced that way under the editor's direction, and the account that follows this section reports what the agents got wrong in this volume and how it was caught.

In the editor's view, this changes where an engineer's value lies. When a model, a controller and a report can each be produced in minutes, producing them is no longer the scarce skill. The work moves to five things.

Quality control. An agent's output arrives as a claim: this model fits, this tuning settles in four minutes, this code implements the equation. The engineer checks claims against the plant or the simulator. A reported number is measured again on a fresh run at the stated seed. The code is read for what it clamps, what it initializes and what it reads, because a display that shows the controller's model in place of the plant looks correct until the two differ. A tuning is tried on a disturbance the agent did not test. The most useful checks are independent of the author, made by a second agent or a person who has not seen the first answer, and they are checks that can fail.

Supervision. The engineer sets the scope of each task and decides what an agent may touch. Reading a historian and proposing a tuning is one level of access. Writing a setpoint, downloading a configuration or changing an interlock is another, and it passes through the plant's management of change with a named approver. Supervision also fixes the order of work: a baseline before a change, one change at a time and a comparison that shows what moved.

Creativity. An agent answers the question it is given. Choosing the question remains the engineer's task: which loop is worth improving, whether a disturbance calls for feedforward, a cascade or a change to the process, which step test the operation can tolerate, and which comparison would show that a predictive controller earns its place over a retuned PID loop. The experiment that separates two explanations is designed by someone who can imagine both.

Process insight. A model contains what its author put in. The equipment does more. A valve sticks, the dead time of an analyzer grows when its sample flow drops, a thermowell lags, a heat exchanger fouls and an operator intervenes at the same hour each night. The engineer who knows the unit recognizes a result that is too good, asks what the model left out and knows which measurement to distrust first.

Accountability. Responsibility does not transfer to a tool. Functional safety standards assign the activities of the safety lifecycle to persons and organizations of stated competence, and they require an assessment that includes someone independent of the design team [20]. An agent may draft the hazard analysis or the test procedure. A person owns the decision to accept it, to start up and to keep running.

These five are the professional form of the habit introduced in Section 1.7: check before trusting, and know enough to check.

16.4.1 How This Volume Was Made

Each of the fourteen apps was built and tested by agents. An agent then wrote the paper on it, keeping verification notes that trace each number to a run. Six of the fourteen referee reports call themselves self-reviews, and the other eight are written in a referee's voice. None names a reviewer independent of the writing. The defect logs hold 133 entries with a severity. Of these, 31 are wrong output or behavior, one of them marked a blocker, 85 are documentation or content, 15 cosmetic and 2 something else, a test and a check allowance. Some logs remove repaired entries or list a round's fixes without a severity, so the counts are a floor. The editor set the scope of each round and what each agent could change. Four cases show what went wrong and what caught it. In each, a controller or a comparison met uncertainty in a way its text did not say.

Controllers that were told the truth. Two logs record a controller that read the run's hidden configuration. The satellite's pointing laws took the inertia of the run and its filters the tracker noise of the run. With the moments of inertia raised two and a half to three times, a slew still settled as the declared spacecraft does, and with a tracker twenty times noisier than declared the camera still imaged the whole pass. The foiling boat's flap map was trimmed on the run's mass, passenger and weed, so its trim flap moved from 0.496° to −0.027° with a passenger moved aft and to 2.097° at 1900 kg, before any measurement had changed. In a volume about uncertainty, a controller that reads the true plant has none to handle. The boat's was found by the paper agent of a later round and the satellite's by an audit of what each controller reads. Both were measured the same way: move a hidden value before Run and ask whether the controller moves while the measurements do not. No shipped run moved when the leaks were closed, 19 of 19 and 17 of 17 hashed sets, so only a learner who changed a value could have met them. With the leak closed, the satellite images 78.0 % of the pass with the noisier tracker.

An advantage with two causes. The vessel's texts credited the whole of a 1.34 against 0.63 m gap in mean offset to the raw measurement chasing the sea. Reading the controller against its own comments showed that the raw setting also switched the environment feedforward off. A copy of the controller with that condition removed put about 0.4 m of the 0.71 m gap on the missing feedforward, with the travel unaffected. The same consistency pass found that every automatic scenario started

with its bias estimate and integrals at zero, so the vessel drove off station and left the watch circle at 91.6 s in scenario 4. After both repairs the offsets are 0.55 and 0.51 m, and the chapter rests the case for the wave filter on the travel, 327.5 against 72.0 MN, where it belongs.

Comparisons that were not like for like. The peg-in-hole challenge drew each part's offset and tilt from the same random stream as the sensor noise, so the number of noise samples before part k depended on the controller. Between the stiff and the suggested runs 1 of 100 parts was the same. The paper agent traced the draw and gave the parts a generator of their own, and ten further seeds confirmed the ranking. In the picking cell the draft compared grasp errors of 84.7 and 55.0 mm at latencies of 300 and 0 ms, which would make the latency worth 29.7 mm where the belt predicts 45 mm. That mismatch led the self-review to the code, where the figure was the distance to the nearest product, not to the one the arm was sent for.

Mode changes that stranded a plan. In Chapter 8 a reboost committed before a switch to manual, whose window passed in manual, never fired and blocked every later reboost. Left alone for three days the satellite spent 61.16 h outside its band, against 0.53 h in automatic. The fix closes such a burn as expired and plans again. The book round then found a second path to the same state: pressed a third time while earlier burns are pending, the manual burn button queues a burn that never starts and never expires, so ten presses fly 10 cm/s instead of 50. That defect has since been repaired, and pending manual burns now fly in full.

The book round. The fourteen chapter records list 202 rows of numeric verification, each a stated value or group of values set beside its reproduction, and all fourteen report no drift. The writer of Chapter 1 read each app's action setting and found eleven sentences in eight chapters that named a loop's action against the convention of Section 1.3. Two causal claims of the papers were corrected by run. The surgical force loop holds its estimate, not the force, and the true force dips below what the dissection needs. The cobot's chatter persists with the scan error at its smallest setting. The pass on Chapter 4 found that its graded challenge passes on the timed schedule once the altitude integral time is 60 s, because the climb is slower and the gust, fixed at 64 s, passes before the timed stop. The app now times the gust from the start of the transition, and the chapter's review exercise shows the timed schedule failing again with that integral time. Four chapter passes led to five corrections of Chapter 1. Ten of the fourteen passes logged an app defect, most of them stale or rounded numbers in briefings or code comments.

Every error above was caught by a check that could fail: a hidden value moved with the measurements held, a counterfactual copy of the code, a difference that did not match the belt speed, a mode switched in the middle of a plan, a causal claim tested on the variable it names. None was caught by rereading the text that contained it. Four working rules follow. Give the controller the declared plant and never the run's. Apportion every advantage by counterfactual before explaining it. Compare designs on the same draws, with disturbances timed to the event they test. Exercise every mode change a user can make, in every order.

16.5 Gaps in Motion Control and Autonomy

Six gaps belong to this volume's topic, and the gaps common to agent-era practice follow them in Section 16.5.1. In the editor's view none is closed. Each item gives the chapters' evidence, then the state of the literature with its sources. Why each gap is hard and what progress would look like are the editor's view. Figure 16.4 places the gaps by the chapters that met them and by the maturity of the method available, and the same numbers mark the machines in the picture that opens this chapter.

1. *Uncertainty a safety layer can trust.* The pallet robot handed its load over with the truth outside its own ellipse, and the satellite in low orbit carried a probability set by an ellipse 971 by 150 m.

Speed laws and collision thresholds consume these bounds as if they were right. A handbook of conjunction assessment notes that probability methods almost always need a covariance for each object. The catalog's covariances have had twenty years of work on their realism, those in operators' own predicted orbits are much more uneven, and its appendix gives test statistics that check predicted covariances against precise positions from onboard satellite navigation [18]. The gap is hard because the truth that would check a stated uncertainty is seldom measured in operation. Practice needs estimators whose stated uncertainty is checked against the truth in operation, with a defined response when it is not.

2. *Told values that are verified.* The cobot's brochure braking, the pallet robot's empty braking and the air taxi's load sheet were each believed in full. The picking arm given an assumed belt speed placed 10 products and failed 40 after the speed changed (Table 10.4). Online estimates of mass, braking and friction that a safety monitor may use are not routine. The Open challenge box after this list takes the gap further.

3. *Proof that a task cannot be done.* The foiling boat's head sea defeated 33 795 settings, and a search tuned to one draw flew 59.07 s on it and the minute on none of fifteen others (Section 5.6). The cobot's table shows what such a proof looks like: at rest the protective distance of 532 mm exceeds the 214 mm the path leaves a person, so no speed override works there (Chapter 15). Practice sets limits from capacity and geometry. Guidelines for vessels on dynamic positioning ask for a consequence analysis that raises an alarm when the thrusters left after the worst failure could not hold the vessel in the prevailing weather [21], and a hydrofoil specification of 1977 makes the effective strut length at least 1.4 times the significant wave height of the 90th-percentile sea [22]. Neither bounds what a controller could achieve. The gap is hard because a negative result must hold for every setting, not only those tried. Methods that certify what no controller in a class can achieve, and so set an operating limit, remain thin.

4. *Predictive and learned controllers under model mismatch.* The lane-keeping predictive controller lost stability on a rack its model misjudged, and its limit fell from 32.3 to 16.9 m/s as the rack lag grew from 0.40 to 0.60 s (Section 2.6.7). Learned policies share the exposure, and neither comes with a stated margin against mismatch. A learned table tennis player trained in simulation and moved to the real robot without further training kept a gap between the two: its return rates on the real table stayed below the more than 80 % seen in simulation, and its builders corrected its choices online from real play [23]. The gap is hard because the mismatch that matters is the one the model leaves out, so the model cannot measure it. Progress would be a margin against mismatch, stated and tested on the real plant, beside every predictive or learned design.

5. *Contact with soft and moving bodies.* The surgical model is lumped and rigid, the peg-in-hole model omits the arm's dynamics and a damage criterion that depends on duration, and no setting tracked the beating heart within the vessel's width. An autonomous robot that joined small bowel in pigs detected the breathing motion and the deformation of the tissue, notified its operator to start a replanning step, and synchronized its tool with the breathing [24]. The gap is hard because tissue changes its shape and position under the tool, and neither change is measured directly. Models of deforming tissue and of damage that a controller can be designed against are still research.

6. *Benchmarks whose disturbances follow the event.* Two chapters had to build their tests this way. The transition challenge meets its gust 36.9 s into the transition (Chapter 4), and the peg-in-hole challenge draws its parts from a random stream of their own, so the sensor noise an attempt consumes cannot change the next part (Chapter 13). Section 16.4.1 tells how both needs were found. A NIST roadmap for robot dexterity lists grasping objects while in motion among the capabilities that need test methods, and its first tests were to assume static conditions, with no picking from a conveyor belt [17]. The gap is hard because a disturbance fixed in time rewards

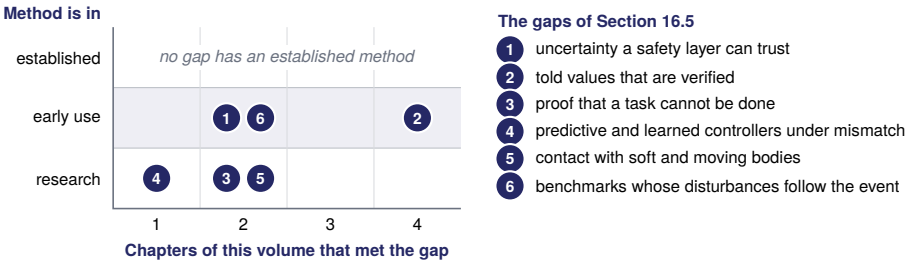

Figure 16.4 The six gaps of this section. Across is the number of chapters that each item names as evidence: two for gap 1 (Chapters 8 and 9), four for gap 2 (3, 9, 10 and 15), two for gap 3 (5 and 15), one for gap 4 (2), two for gap 5 (13 and 14) and two for gap 6 (4 and 13). Up is the maturity of the method available, the editor's reading of the sources cited in each item and in the Open challenge box. Notice that the gaps about stated uncertainty, told values and benchmarks have methods in early use, while proof of infeasibility, model mismatch and soft contact are still research.

a design that happens to arrive early or late. Shared benchmarks for moving machines need disturbances keyed to the maneuver, common random draws for every design and several seeds per verdict.

OPEN CHALLENGE **A Braking Figure That Stays True**

Four chapters name gap 2, and three of them believed a braking figure or a weight that nothing checked. A NIST study of collaborative cells notes that many installations treat the stopping time and distance as constants, though they should be measured again from time to time as the brakes wear. Not every integrator or user maintains a robot regularly, and the study asks for simple, inexpensive tests of joint brake wear [25]. In the editor's view the gap is hard because braking changes with load, pose and wear while the monitor needs one number it can trust. Progress would be a braking figure measured in service and checked against each stop the machine makes.

16.5.1 Gaps Common to Agent-Era Practice

The final chapter of each volume lists gaps specific to its topic. Seven more are common to every topic in the series. In the editor's view none of them is closed.

1. *Verifying agent-written control code and models before they reach a plant.* Tests show that a controller behaves on the cases tested. For controllers learned from data, guarantees of stability and constraint satisfaction rest on assumptions about the model and its uncertainty, and open problems remain [26]. In the editor's view, code written by an agent has no comparable body of method yet. Practice needs a routine path from an agent's draft to code a plant can trust: independent derivation, tests of what the code reads and clamps, and trials with hardware in the loop.

2. *Trustworthy models from sparse or poor plant data.* Plant records are short, taken in closed loop, noisy and marked by unmeasured disturbances. A fitting routine returns a model from any record. What is missing is a dependable statement, delivered with the model, of where it can be trusted, and a method that says so when the data cannot identify a parameter.

3. *Safety cases and standards for agents in commissioning and operation.* Functional safety standards define a lifecycle of specified requirements, verification and assigned human roles [20]. In the editor's view, how an agent that writes configuration, runs commissioning checks or advises an operator is qualified, limited and audited under them is not yet settled.

4. *Security of agent access to control systems.* An agent with credentials to a control system is a new path into it. Industrial security standards state the requirements a control system must meet at each security level [27]. An agent that reads documents, messages or web pages can also be steered by text placed there. In the editor's view, least privilege, read-only access by default and human approval of every write are a starting point and not a solution.

5. *Benchmarks and reproducibility.* An agent asked twice may deliver two different controllers, and a result reported without its plant, seed and scoring cannot be checked. Control needs shared benchmark plants with fixed disturbances and published measures, on which agent-built and conventional designs are compared like for like. The apps of this series, with a default seed and an exported record of every run, are a start.

6. *Transfer from simulation to hardware.* A controller developed on a simulator inherits what the simulator omits. The robotics literature documents the gap between simulated and real performance and the methods used to narrow it, such as randomizing the simulated plant and adapting on the real one [28]. Process plants add slow dynamics, rare upsets and little freedom to experiment, so each transfer still needs a staged plan from the simulator to a bench rig to the unit.

7. *Keeping engineers' judgment sharp when agents do the routine work.* Automation removes the routine work through which skill is built and leaves the person the hard cases, an irony described for process operators in 1983 [29]. The same now applies to engineers. An engineer who has never tuned a loop by hand or closed a balance on paper is poorly placed to review an agent that has. Deliberate practice, of the kind the exercises in this book provide, has to be designed into education and into work.

KEY POINTS

- A machine coped with what it could not see when it was told what had been checked, or estimated what drifts.
- A told value acted at once, and a wrong one was believed in full.
- An estimator needed a state for each unknown, and time to form it.
- Where no margin could be met, the right answer was to decline the task.
- Structure decided before tuning, and every verdict belonged to the truth.

References

[1] N. Kalra and S. M. Paddock, "Driving to safety: How many miles of driving would it take to demonstrate autonomous vehicle reliability?" *Transportation Research Part A: Policy and Practice*, vol. 94, pp. 182–193, 2016.

[2] P. Koopman and M. Wagner, "Autonomous vehicle safety: An interdisciplinary challenge," *IEEE Intelligent Transportation Systems Magazine*, vol. 9, no. 1, pp. 90–96, 2017.

[3] E. Kaufmann, L. Bauersfeld, A. Loquercio, M. Müller, V. Koltun, and D. Scaramuzza, "Champion-level drone racing using deep reinforcement learning," *Nature*, vol. 620, no. 7976, pp. 982–987, 2023.

[4] R. Firoozi, J. Tucker, S. Tian, A. Majumdar, J. Sun, W. Liu, Y. Zhu, S. Song, A. Kapoor, K. Hausman, B. Ichter, D. Driess, J. Wu, C. Lu, and M. Schwager, "Foundation models in robotics: Applications, challenges, and the future," *The International Journal of Robotics Research*, vol. 44, no. 5, pp. 701–739, 2025.

[5] K.-C. Hsu, H. Hu, and J. F. Fisac, "The safety filter: A unified view of safety-critical control in autonomous systems," *Annual Review of Control, Robotics, and Autonomous Systems*, vol. 7, pp. 47–72, 2024.

[6] S. V. Weston, C. D. Burkhard, J. M. Stupl, R. L. Ticknor, B. D. Yost, R. A. Austin, P. Galchenko, L. K. Newman, and L. Santos Soto, "State-of-the-art small spacecraft technology," National Aeronautics and Space Administration, Ames Research Center, Moffett Field, CA, Technical Publication NASA/TP-20250000142, Jan. 2025. [Online]. Available: https://ntrs.nasa.gov/citations/20250000142

[7] V. Villani, F. Pini, F. Leali, and C. Secchi, "Survey on human–robot collaboration in industrial settings: Safety, intuitive interfaces and applications," *Mechatronics*, vol. 55, pp. 248–266, 2018.

[8] 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.

[9] G.-Z. Yang, J. Cambias, K. Cleary, E. Daimler, J. Drake, P. E. Dupont, N. Hata, P. Kazanzides, S. Martel, R. V. Patel, V. J. Santos, and R. H. Taylor, "Medical robotics—regulatory, ethical, and legal considerations for increasing levels of autonomy," *Science Robotics*, vol. 2, no. 4, p. eaam8638, 2017.

[10] A. Attanasio, B. Scaglioni, E. De Momi, P. Fiorini, and P. Valdastri, "Autonomy in surgical robotics," *Annual Review of Control, Robotics, and Autonomous Systems*, vol. 4, pp. 651–679, 2021.

[11] T. Uriot, D. Izzo, L. F. Simões, R. Abay, N. Einecke, S. Rebhan, J. Martinez-Heras, F. Letizia, J. Siminski, and K. Merz, "Spacecraft collision avoidance challenge: Design and results of a machine learning competition," *Astrodynamics*, vol. 6, no. 2, pp. 121–140, 2022.

[12] Federal Aviation Administration, "Aircraft weight and balance control," Advisory Circular 120-27F, 6 May 2019, https://www.faa.gov/documentLibrary/media/Advisory_Circular/AC_120-27F.pdf, 2019.

[13] J. Liu, F. Jiang, S. Tashiro, S. Chen, and M. Tanaka, "A physics-informed and data-driven framework for robotic welding in manufacturing," *Nature Communications*, vol. 16, p. 4807, 2025.

[14] 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.

[15] S. Withrow-Maser, J. Aires, A. Ruan, C. Malpica, and S. Schuet, "Evaluation of heave disturbance rejection and control response criteria on the handling qualities evaluation of urban air mobility (UAM) eVTOL quadrotors using the vertical motion simulator," in *Vertical Flight Society Aeromechanics for Advanced Vertical Flight Technical Meeting*, San Jose, CA, 2022, nASA Technical Reports Server 20210026423, https://ntrs.nasa.gov/citations/20210026423.

[16] R. B. Madigan, T. P. Quinn, and T. A. Siewert, "Control of gas-metal-arc welding using arc-light sensing," National Institute of Standards and Technology, Boulder, CO, NIST Interagency/Internal Report 5037, 1995.

[17] J. Falco, J. Marvel, and E. Messina, "A roadmap to progress measurement science in robot dexterity and manipulation," National Institute of Standards and Technology, Gaithersburg, MD, NIST Interagency/Internal Report 7993, 2014.

[18] National Aeronautics and Space Administration, "NASA spacecraft conjunction assessment and collision avoidance best practices handbook," National Aeronautics and Space Administration, Washington, DC, Special Publication NASA/SP-20205011318, Dec. 2020. [Online]. Available: https://ntrs.nasa.gov/citations/20205011318

[19] L. Wang, C. Ma, X. Feng, Z. Zhang, H. Yang, J. Zhang, Z. Chen, J. Tang, X. Chen, Y. Lin, W. X. Zhao, Z. Wei, and J. Wen, "A survey on large language model based autonomous agents," *Frontiers of Computer Science*, vol. 18, no. 6, p. 186345, 2024.

[20] International Electrotechnical Commission, "Functional safety – Safety instrumented systems for the process industry sector – Part 1: Framework, definitions, system, hardware and application programming requirements," IEC 61511-1:2016, Edition 2.0, 2016.

[21] International Maritime Organization, "Guidelines for vessels with dynamic positioning systems," MSC/Circ.645, Maritime Safety Committee, London, 1994, adopted 6 June 1994.

[22] D. R. Stark, W. E. Farris, and A. O. Harang, "Hydrofoil ship control and dynamics specifications and criteria, volume III," Boeing Marine Systems, for the David Taylor Naval Ship Research and Development Center, Document D321-51313-1, Mar. 1977, contract N00600-75-C-1107. Copy in the library of the International Hydrofoil Society, https://foils.org/control/.

[23] D. B. D'Ambrosio, S. Abeyruwan, L. Graesser, A. Iscen, H. Ben Amor, A. Bewley, B. J. Reed, K. Reymann, L. Takayama, Y. Tassa *et al.*, "Achieving human level competitive robot table tennis," in *2025 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2025, pp. 74–82.

[24] H. Saeidi, J. D. Opfermann, M. Kam, S. Wei, S. Leonard, M. H. Hsieh, J. U. Kang, and A. Krieger, "Autonomous robotic laparoscopic surgery for intestinal anastomosis," *Science Robotics*, vol. 7, no. 62, p. eabj2908, 2022.

[25] J. A. Marvel and R. Norcross, "Implementing speed and separation monitoring in collaborative robot workcells," *Robotics and Computer-Integrated Manufacturing*, vol. 44, pp. 144–155, 2017.

[26] L. Brunke, M. Greeff, A. W. Hall, Z. Yuan, S. Zhou, J. Panerati, and A. P. Schoellig, "Safe learning in robotics: From learning-based control to safe reinforcement learning," *Annual Review of Control, Robotics, and Autonomous Systems*, vol. 5, pp. 411–444, 2022.

[27] International Electrotechnical Commission, "Industrial communication networks – Network and system security – Part 3-3: System security requirements and security levels," IEC 62443-3-3:2013, Edition 1.0, 2013.

[28] W. Zhao, J. P. Queralta, and T. Westerlund, "Sim-to-real transfer in deep reinforcement learning for robotics: A survey," in *Proc. 2020 IEEE Symposium Series on Computational Intelligence (SSCI)*, 2020, pp. 737–744.

[29] L. Bainbridge, "Ironies of automation," *Automatica*, vol. 19, no. 6, pp. 775–779, 1983.

16.6 Exercises

16.1 Discussion · 8 min

An agent proposes to delete the load sheet of Chapter 3 and instead raise the altitude loop's integral action "so that the loop learns each passenger's mass by itself and a wrong sheet can never mislead it". Using Table 3.5 and the chapter's case 4, say what the proposal would gain and what it would lose, and what you would build instead.

16.2 Discussion · 8 min

Section 16.4.1 reports a loophole in the challenge of Chapter 4 and parts drawn differently for each controller in Chapter 13. Write the grading rules for a benchmark that compares agent-built and conventional controllers of moving machines, and say which failure each rule prevents.

16.3 Discussion · 8 min

An earlier version of the vessel app of Chapter 6 reported mean offsets of 1.34 and 0.63 m with the raw measurement and with the wave filter, and its texts credited the whole difference to the controller chasing the sea. Today Table 6.4 gives 0.55 and 0.51 m. Explain how the gap was apportioned, and write the business case for the wave filter that the chapter's numbers support.

16.4 Discussion · 8 min

A plant manager asks for two changes to raise throughput: switch off the pallet robot's slowing for unseen space in Chapter 9, which costs 8.7 s a run, and replace the collaborative arm's measured braking in Chapter 15 with the maker's brochure figure, which is faster. Discuss both requests with the chapters' numbers, and say what you would offer instead.

16.5 Discussion · 7 min

An agent reports a setting that flies the 1.4 m head sea of Chapter 5, where a search tuned to one draw flew 59.07 s of the minute and its best settings failed fifteen others. Say what evidence you would ask for before accepting the claim, and what the chapter's result means for an operating limit on the boat.

16.6 Discussion · 9 min

An agent hands you the code of a new attitude controller for a small satellite and a passing test report. List the checks you would run before the code reaches a hardware test, and say for each which defect of this volume it would have caught.

Answers to Selected Exercises

Short answers are given for the odd-numbered exercises. Numeric answers from an app are for its default seed and settings unless the exercise says otherwise.

Chapter 1. Introduction

1.1 (b). The angle rises with the torque, so $a = +1$, and an output that falls as the PV rises is reverse acting.

1.3 Quintic 3.0 s, set by the speed limit. Trapezoid 2.1 s.

1.5 2.04 m/s, and 1.52 m/s at 0.75 m/s^2.

1.7 It asks 50 N and 10 N. Moment first: 40 N and 0 N, so 40 N and 12 N m. Clipping: 50 N and 9 N m.

1.9 (a) About 31.6 N, 1.2 N and 80 N. (b) $g = 4\,\mathrm{s}^{-1}$, 0.25 s, delay margin 0.39 s.

Chapter 2. Speed Dependence of Stanley Lane Keeping

2.1 (a). The correction falls as $1/(v + k_{soft})$ while the yaw rate per radian rises as v, so Equation (2.9) leaves a rate $k_e v/(v + k_{soft})$ close to k_e.

2.3 (c). The corner needs 6.617° and only 6° is available, so no gain closes the gap.

2.5 (a) 0.994 m. (b) 0.503 m. (c) 0.88/s, and 0.085 m at the suggested gain, about a third of the band.

2.7 (a) IAE 0.028 m s, 0.831°, 0.000 m. (b) 8.797 m s, 0.830°, 0.314 m. (c) The prediction is 0.314 m. The corner needs $\arctan(L/R)$ either way, and only its source changes.

2.9 (a) 0.09 m, 0.326 m s, 0.499°. (b) 8.850 m s, –0.225 m, as predicted. (c) –0.080 m. The ratio $31/11 = 2.82$ is the ratio of $v + k_{soft}$.

2.11 (a) 1.384°. (b) 0.92 °/s and a lag of about 0.09°. (c) So that the driver can override the system at any time. An angle limit reproduces saturation but not the handover.

Chapter 3. Drone Taxi Feedforward and Wind Estimation

3.1 (a). The tilt loop is proportional-derivative, so the trim moment mgc/ℓ can come only from $K_{c,\theta}$ times a standing error.

3.3 (c). The error comes from 1.37 kN of weight the feedforward leaves out, and correcting the feedforward removes it at the shipped gain.

3.5 (a) 1.37 kN, 1.96 m. (b) 496 kg, 1.22 kN, 723 N/m, 1.68 m. (c) 0.67 s and 2.0 s.

3.7 (a) 0.08 m and 1.75 m. (b) 0.05 m. (c) The heavy climb reaches about 89 % of the 1.96 m budget, and the feedforward removes the lag.

3.9 (a) One touchdown of two in specification, pad C at 0.99 m, 22.5 %. (b) The same failure at 0.98 m, 22.4 %. (c) Both touchdowns in specification, worst 0.23 m at 0.41 m/s, 22.6 %, a pass.

3.11 (a) 124 kg and 1.22 kN. (b) The feedforward falls short, so the vehicle lifts off late and lags its climb by 1.54 m. The supervisor computes the margin on the believed weight, so it reads too high and may dispatch a load it should hold.

Chapter 4. Lift-Plus-Cruise VTOL Control Allocation

4.1 (d). Equation (4.3) builds the lift on $\mathbf{v}_a = (\dot{x} - u, \dot{z} - w)$, and the gust raises u by up to 10 m/s.

4.3 (b). The lamp reports demand the allocation cannot meet, and with a pitot 12 m/s high the allocation believes the wing makes 1.8 times its real lift.

4.5 (a) 40.4 m/s and 43.8 m/s. (b) 50.9 m/s and 47.0 m/s. (c) 75 % and 88 %.

4.7 (a) 54.9 m/s, 1.7 m, 3.8°. (b) 39.7 m/s, 6.7 m, 2.4°. (c) 49.6 m/s over the ground and 39.8 m/s through the air, 91 %.

4.9 (a) 9.3 m of clearance against 25 m, the only failing part, with the stop at 37.9 m/s, 31.7 m below the reference and a 2.2° margin. (b) 32.0 m, 3.6°, 73 s, 0.43 m/s and 0.13 m, 25.6 %. Every part passes.

4.11 (a) At 48.3 m/s, where the wing can lift 122 % of the true weight. (b) Below 30 m/s the wing carries too little of the weight for its lift to measure the airspeed, so a second, independent source is needed there.

Chapter 5. Foiling Boat Ride Height

5.1 (c). The waves arrive at 5.47 rad/s and are far smaller than the 0.80 m band, so platforming rides at 0.87 m/s^2 against 2.31 for contouring.

5.3 (a). The process gain from rear flap to pitch is negative, so the loop is direct acting.

5.5 (a) 3.54 s, 1.78 rad/s, 0.322 rad/m. (b) 1.54 and 5.09 rad/s, 4.09 s. (c) 5.47 rad/s. All three lie above the platform crossover and below the contour crossover.

5.7 (a) 0.750 and 0.709 m. (b) 0.434, 0.245 and −0.004 m. (c) Near the surface the stiffness settles the boat in about a second, while from 0.50 m the time constant is 7.9 s and grows as the boat sinks.

5.9 (a) 0.74 m/s^2. (b) 0.98 m/s^2. (c) 1.30 m/s^2. (d) 43 %: the preview is worth 0.24 of the 0.56 m/s^2.

5.11 The strut, the speed, the cavitation number and the history of the flow, and the partially ventilated regime. The foil must sink to a full chord (0.35 m for the main foil) to recover, so the boat drops well below where the lift was lost.

Chapter 6. Dynamic Positioning and Wave Filtering

6.1 (b). The hull has no restoring force, so the loop needs derivative action, and derivative action needs the phase that the filter takes away.

6.3 (d). The heading leaves a 3° band at 17 s, before any scheduled change, while the power stays between 54 and 68 %.

6.5 (a) 0.571 rad/s and 1.37 m/s. (b) 0.19 and about 405 kN. (c) About 2.12 MN, so the demand runs from clamp to clamp.

6.7 (a) 327.5 MN and 72.0 MN. (b) 452 kW h and 412 kW h, ratios 4.5 for travel and 1.10 for energy.

6.9 (a) Minimum norm: mean 0.70 m, largest 1.87 m. Fixed split: 2.86 m and 5.40 m. (b) 128 kN against 105 kN: the solved allocation hands the lost share to the remaining azimuth.

6.11 The bias estimate must converge on the environmental load. The chapter's estimator covers 63 % of a load step in 237 s and comes within 2 % about ten and a half minutes after it.

Chapter 7. CubeSat Attitude and Wheel Momentum

7.1 (b). A wheel moves momentum but cannot remove it, and only an external torque can empty it.

7.3 (d). Limiting the length of the rate command keeps the eigenaxis, and limiting its parts does not.

7.5 (a) 15.24 min along the orbit normal, 29.59 and 29.47 min in the plane. (b) 35.1, 68.1 and 67.9 min predicted. (c) The field lies near the orbit plane, always nearly across the orbit normal, and it passes along an excess in the plane four times an orbit.

7.7 (a) 37.84 %. (b) 245.0 min (4.083 h) and 250.8 min (4.180 h). (c) 40.0 %. The wheel fills at the predicted rate and saturates at the predicted time.

7.9 (a) 24.1, 100.0 and 100.0 %. (b) 49.8, 40.0 and 33.0 % with unloading only, 20.6, 100.0 and 100.0 % with the bias state only. (c) 100.0, 40.0 and 33.0 %, target met. Each remedy fixes one of two independent failures.

7.11 (a) 0.71 to 1.43 µN m, 2100 to 4200 times less than a wheel. (b) Image only where the field lies well away from the failed axis, choose attitudes that ask little of that axis, or use another external torque, as Kepler used the pressure of sunlight in K2.

Chapter 8. Low Earth Orbit Maintenance and Avoidance

8.1 (a). The secular term $-3\Delta v\,\tau$ of Equation (8.8) grows with the time from the burn to the pass.

8.3 (c). The burn point stays where it was, and the far side rises by $2\Delta a$, so the mean perigee stays at 499.770 km and the apogee reaches 500.575 km.

8.5 (a) $K = 7.72 \times 10^{-4}$ and $P_c = 1.79 \times 10^{-4}$. (b) $s = d/\sqrt{2} = 1.21$, at most 1.94×10^{-4}, of which the actual covariance gives 92 %. (c) $d = \sqrt{2\ln(K/P^*)} = 2.51$.

8.7 (a) 6.08 h outside, 2.5 g (2.52 g in column `propellant_used_g`), lowest 499.677 km at 24.3 h; the drag scale estimate passes 2.5 at 8.73 h. (b) 0.00 h, 1.7 g (1.68), lowest 499.848 km. The deadband waits for a decision after the crossing.

8.9 (a) 3 of 3 unresolved, 1.7×10^{-3}, 5.53 h, 2.8 g: failed. (b) 0 of 3, 4.9×10^{-5}, 5.51 h, 4.6 g: failed on the band. (c) 0 of 3, 7.7×10^{-5}, 0.00 h, 6.0 g: passed. The avoidance planner resolves the passes and the predictive planner keeps the band.

8.11 (a) One prograde burn leaves the mean in the band and the apogee 325 m above it. (b) Hold the mean perigee and apogee inside the band, columns `h_perigee_km` and `h_apogee_km` and the Mean-orbit perigee and apogee trend.

Chapter 9. Warehouse Robot Speed and Stopping Distance

9.1 (d). The stop tests the gap against the believed 2.26 m/s^2, and the motors can stop the loaded robot at only 0.76 m/s^2.

9.3 (b). The cart entered 0.40 m ahead of a robot that needed 0.59 m, while a walk out of unseen space is avoidable by definition.

9.5 (a) 2.01 m/s and 1.30 m/s; the cruise limit binds and the robot drives at 1.5 m/s. (b) 0.42 m. (c) 0.12 m; below that unseen gap the floor allows more speed than the visible space can stop.

9.7 (a) 0.74 m/s at 18.6 s, one E-stop. (b) 0.25 m, 40.3 s. (c) The empty-robot plan reserves 0.95 m at 1.5 m/s where 1.92 m is needed.

9.9 (a) The worker is hit at 1.47 m/s at 11.1 s. (b) 3 of 3, 0.27 m, 120.6 s, 48.4 kJ, a pass. (c) 3 of 3 but 0.11 m and one violation, a fail.

9.11 (a) 0.81 m/s^2 and 1.39 m, 0.08 m less than needed. (b) Under-reading the mass is unsafe. Plan with an upper bound on the mass, or a lower bound on the braking.

Chapter 10. Robot Conveyor Picking on Late Camera Reports

10.1 (c). The report is still 0.42 s old at the grasp, so the gripper closes 63.0 mm behind the product it was sent for.

10.3 (a). The torque check sees only the arm, so it approves carries that need 47.1 N m at a 40 N m shoulder.

10.5 (a) 20.0 rad/s, $\zeta = 0.63$. (b) 0.641 m and 1.617 kg m^2. (c) 15.7 rad/s and $\zeta = 0.49$. The frequency falls by 21 % and the damping ratio by 21 %, both as $1/\sqrt{M_{11}}$.

10.7 (a) 0 placed, 42 failed, 23 missed. (b) The same counts: the report is still 0.42 s old at the grasp. (c) 24 placed, 0 failed, 0 missed.

10.9 (a) 0 placed, 35 dropped, 158 failed. (b) 35 placed, 0 dropped, 158 failed. (c) 98 of 100, no drop, no failed grasp, 1.97 mm rms. (d) 98 of 100, 0.25 mm rms.

10.11 (a) $60/0.857 = 70.0$ picks per minute. (b) Two arms, which need a shared assignment of the stream.

Chapter 11. Table Tennis Robot Prediction and Commitment

11.1 (c). Without drag the predicted ball covers the rest of the table too fast, so the error is in time, not in place.

11.3 (a). An estimator can estimate only what its model has a state for, and the spin shows only through the Magnus lift and the bounce.

11.5 (a) 3.41 m/s. (b) Grip needs an impulse of 2.44 m/s against an available 4.21 m/s. The ball leaves at 3.66 m/s with 29.1 rev/s of topspin. (c) Grip would need 4.47 m/s against 3.96 m/s, so it slides, and leaves at 2.18 m/s.

11.7 (a) 16 of 20. (b) 19 of 20. (c) 10 of 20. Early commits freeze a prediction with a poor hit time, and late ones leave too little time to swing.

11.9 (a) 0–11 at 28.2 s. (b) 11–9 at 66.3 s. (c) 4–11 at 44.3 s. A commit fixed by the clock, which is the tuned arm's own blind spot.

11.11 240 ms with about 110 ms left, 120 ms with 231 ms left, and 45 ms with 306 ms left. Only the 200 Hz, 10 ms camera leaves time for the swing.

Chapter 12. Robotic Arc Welding With Seam Preview

12.1 (b). A shorter arc needs less voltage, so the source drives more current, and the faster melting burns the extension back until the arc is long again.

12.3 (d). The preview is early but old, the arc current is fresh but late, and the combination uses each where it is right.

12.5 (a) 1.455, 0.732, 0.732, 1.455 mm. (b) 1.22 mm; 20 and 60 mm lose fusion. (c) About 16.7 mm either way.

12.7 (a) 4.29 mm, 252 A; 3.81 mm, 272 A; about 0.48 mm not recovered. (b) 3.99 mm, then 0.99 mm. (c) 4.44 mm at 6.88 m/min.

12.9 (a) 0 mm defective, 0.13 mm, 12.3 mm/s, 18.1 m. (b) 331 mm undercut. (c) $H \approx 545\,\mathrm{J/mm}$ gives $v_{max} = 10.9\,\mathrm{mm/s}$, below the 12.5 mm/s travel.

12.11 (a) None on a straight seam, since a shift along it changes nothing; 1.455 mm on a 3.5 mm bend, as in case 4. (b) Check the filed readings against the laser line, or teach a feature and see where the reading puts it.

Chapter 13. Robotic Peg-in-Hole Force Control

13.1 (a). The correction of $x_o - c$ against K_x, amplified by the chamfer factor λ, is 361 N/mm times the offset for the servo and about 1 N for the spring.

13.3 (c). The side forces that squeeze the peg grow with $K_\theta - mgh_G$, and their friction is what stops it.

13.5 (a) 0.47 kg. (b) $m_n = 0.83\,\mathrm{kg}$ and 1.44 N per mm/s. (c) 6.9 mm/s by the estimate, 7.4 mm/s by the fit.

13.7 (a) 136.1 N, damaged at 0.223 s. (b) 13.7 N, in at 1.14 s. (c) 13.7 N, in at 0.63 s, because the sequence regulates the insertion force after contact.

13.9 (a) 42.2 N true against a largest reading of about 13.3 N, in at 1.72 s with no retry. (b) 16.8 N on the tile, 8.3 N while searching. (c) 1.59 s.

13.11 The peg jams at 4.96 mm with 56.3 N. The compliance center belongs at the tip, as on a remote center compliance. The planar model has one tilt axis, and a round peg adds a second.

Chapter 14. Surgical Robot Instrument Control

14.1 (d). The surgeon steers by a forward model that expects no lag from the robot, and the 2 Hz filter adds about 120 ms at 1 Hz.

14.3 (b). No incursion, but 8.6 s of safe motion blocked, with the tip held 11.79 mm from the vessels' axis.

14.5 (a) 0.339 and 0.040, 34 µm and 4.0 µm. (b) 37.9 ms and 120.3 ms. (c) 0.25 s, so the loop would be unstable without the forward model. (d) 82.4 ms.

14.7 (a) 1.64 N, 0.64 N, slip at 11.5 s. (b) 3.61 mm and 4.83 mm. (c) The flap tears at 13.3 s.

14.9 (a) 3.29, 2.19, 1.10 and 1.13 mm. (b) 4.51 and 4.49 mm. Feedforward has the wrong breath shape, MPC tracks the same forecast with limits that never bind, and every option freezes and decays its offset while the tracker is blind.

14.11 (a) 0.83 s, 3.6 Hz, and 5.0 Hz at 100 beats per minute. (b) 18°. (c) A mechanical stabilizer holds the patch of heart still, or the heart is stopped. Both change the plant.

Chapter 15. Collaborative Robot Speed and Separation

15.1 (d). At rest the protective distance is still $v_h T_r + C + Z_d + Z_r = 532\,\text{mm}$, more than the 214 mm the path leaves.

15.3 (b). At 12 N m the first reflex comes 117 ms after the touch, past the declared 50 ms.

15.5 (a) 0.738 s and 532 mm. (b) 0.50 m/s. (c) 1498 mm, 582 mm and 682 mm.

15.7 (a) 6.22 s, 1486 mm against 970 mm, true 1478 mm against 1558 mm. (b) 6 episodes, worst 506 mm, 7 cycles. (c) No violation and 7 cycles, so the optimism bought no cycle.

15.9 (a) 13 stops, no cycle. (b) 2 stops, 3 cycles. (c) 17 stops, 4 cycles. The hysteresis costs one cycle in 90 s and turns 17 protective stops into 2.

15.11 (a) Clamping, the quasi-static constrained impact, with its pressure limits and edges. Equation (15.11) assumes a free body that stores the impact energy in its spring and recoils. (b) No occlusion, no detection capability, errors independent between scans. A hidden person makes the measured separation wrong in the unsafe direction.

Chapter 16. Summary and Outlook

16.1 A good answer uses the error budget: a loop not told about a force supplies it only through its error, about $\Delta m g / K^{e}_{c,z}$, so the heavy climb fell 1.78 m behind with fixed tuning against 0.05 m when told. A faster integral shortens how long that error lasts but cannot remove the need for it, and it repays the excursion with an overshoot of equal area. The proposal is right that a sheet can be wrong: in case 4 a sheet 20 % light left the vehicle on the pad 1.17 s into the climb and 1.54 m behind. The design to build keeps the sheet for the immediate force, adds the disturbance observer, which in case 4 brought the lag to 0.02 m three seconds after liftoff and estimated −1.21 kN of the 1.22 kN missing, and compares the two to flag a wrong sheet.

16.3 A good answer explains the counterfactual: the raw setting also switched the environment feedforward off, so the scenario pair differed in two things. A copy of the code with only that condition removed put about 0.4 m of the 0.71 m gap on the missing feedforward and left the travel unchanged. The case for the filter therefore rests on thruster travel, 327.5 against 72.0 MN over the same window, with energy only about 9 % lower and position no better. The model-only column, 91.1 MN, shows that a vessel model with a bias estimate earns most of the saving, and the wave states the remaining 19.1 MN.

16.5 A good answer separates passing one realization from flying a sea state. It asks for the result on draws the search never saw, with the number passed and the spread, and for the margin to slam and ventilation on each. It notes that the chapter's own search found a passage through particular waves, not a way to fly the sea. It treats the case as an operating limit: slow below the takeoff speed, which at 12 knots kept the hull from slamming for the minute, and it states that the limit rests on a search of 33 795 settings, not on a proof.

Glossary

Each term is given as this volume uses it, in one plain sentence, followed by the chapter where it is first used in earnest. The two indexes give every page.

A* search A route-finding method that searches a map of free cells for the shortest path from start to goal. *Chapter 9.*

actuator saturation The state in which an actuator is at its limit, so a larger command produces no more force, torque or turn. *Chapter 2.*

analytical redundancy A second value of a measured quantity computed from a model and other readings, used to detect a faulty sensor. *Chapter 4.*

angle of attack The angle between a wing or foil and the flow that meets it, which sets how much lift it makes. *Chapter 4.*

anti-windup Logic that stops a controller's integral term from growing while the actuator is held at a limit. *Chapter 1.*

azimuth thruster A propeller unit that can turn about a vertical axis, so it can push a vessel in any horizontal direction. *Chapter 6.*

B-dot detumbling A law that drives magnetic coils against the change of the measured magnetic field to slow a tumbling satellite. *Chapter 7.*

blind corner A corner past which a robot's scanner cannot see, so a person or a cart may appear there without warning. *Chapter 9.*

cascade control Nested loops in which an outer loop sets the reference of a faster inner loop, such as position over speed. *Chapter 1.*

chance constraint A limit that must hold with a stated probability rather than always, such as a collision probability kept below a threshold. *Chapter 8.*

collaborative robot A robot arm built or set up to work in the same space as people, without a fence between them. *Chapter 15.*

collision probability The chance that two objects in orbit pass within a set distance of each other, computed from their position uncertainties. *Chapter 8.*

complementary filter An estimator that takes a slow, steady sensor at low frequencies and a fast, drifting one at high frequencies. *Chapter 1.*

computed-torque control A law for an arm that computes the joint torques from its dynamic model and adds feedback on the remaining error. *Chapter 10.*

conjunction assessment The routine check of a satellite's predicted orbit against tracked objects to find close approaches and judge their risk. *Chapter 8.*

constant-voltage source A welding power supply that holds its output voltage, so the current changes as the arc length changes. *Chapter 12.*

contact bumper A soft strip on a guided vehicle that triggers a stop when it is pressed against an obstacle. *Chapter 9.*

control allocation Sharing the total force and moment the loops ask for among several actuators, within each one's limits. *Chapter 1.*

conveyor tracking Following a product on a moving belt, using the belt's measured travel to predict where the product is. *Chapter 10.*

cross-track error The sideways distance between a vehicle and the path it is meant to follow. *Chapter 2.*

CubeSat A small satellite built to a standard outside size and interface, so that many can share one deployer into orbit. *Chapter 7.*

dead reckoning Estimating position by adding up measured motion, such as wheel turns, from a known start, with error that grows. *Chapter 1.*

delay margin The largest added time delay a feedback loop can tolerate before it becomes unstable. *Chapter 1.*

differential drive A vehicle steered by driving its left and right wheels at different speeds. *Chapter 1.*

disturbance observer An estimator that infers an unmeasured force from the motion it causes, so the controller can cancel it. *Chapter 1.*

dynamic positioning Holding a vessel at a fixed position and heading with its own thrusters, without anchors. *Chapter 6.*

eigenaxis rotation Turning a body from one attitude to another about the single axis that makes the turn shortest. *Chapter 1.*

electrode extension The length of welding wire between the contact tip and the arc, which the current heats as it passes. *Chapter 12.*

encoder A sensor that counts pulses as a shaft or wheel turns, giving its angle or the distance traveled. *Chapter 1.*

error budget A breakdown of a final error into the shares caused by each source, such as noise, delay

or model error. *Chapter 1*.

extended Kalman filter (EKF) A Kalman filter for a nonlinear system, which uses the model linearized about the current estimate. *Chapter 1*.

feedforward Control action computed from a known reference or disturbance before any error appears. *Chapter 1*.

flow scheme The chain from mission to planner, motion loops, allocation, actuators, body, sensors and navigation that every machine here shares. *Chapter 1*.

foiling boat A boat that rises on underwater wings, called foils, so its hull flies clear of the water. *Chapter 5*.

forgetting factor A weight below one that makes a recursive estimator count recent data more than old data. *Chapter 14*.

fulcrum effect The reversal of motion through a fixed entry port, where moving the handle one way moves the tip the other way. *Chapter 14*.

gain scheduling Changing a controller's gains with a measured condition, such as speed or mass, so it stays well tuned. *Chapter 1*.

gas metal arc welding Welding in which a wire fed through a torch melts in an electric arc under a shielding gas. *Chapter 12*.

guidance, navigation and control The three tasks of a moving vehicle: deciding where to go, finding where it is and making it move there. *Chapter 1*.

gyro bias A small steady offset in a gyroscope's rate reading, which makes an integrated angle drift. *Chapter 7*.

Hall thruster An electric thruster that ionizes a gas and accelerates the ions with an electric field, giving small thrust for long times. *Chapter 8*.

hand-eye calibration Finding where a camera sits relative to a robot arm, so the camera's coordinates can be turned into the arm's. *Chapter 10*.

hysteresis Switching on and off at two different thresholds, so a decision does not chatter when a signal sits near one value. *Chapter 1*.

impedance control Control that makes a robot behave like a chosen spring and damper when it touches something. *Chapter 1*.

inertia matrix The table of masses and inertias that relates an arm's joint torques to its joint accelerations. *Chapter 1*.

inertial measurement unit (IMU) A package of gyroscopes and accelerometers that measures a body's turning rates and accelerations. *Chapter 3*.

interception Planning a move that meets a moving target at the place where it will be when the arm arrives. *Chapter 10*.

jamming A peg stuck in a hole because the side contact forces and their friction balance the push. *Chapter 13*.

jerk The rate of change of acceleration, limited in a trajectory so that a move starts and stops smoothly. *Chapter 1*.

Kalman filter An estimator that blends a model's prediction with each measurement, weighting them by their uncertainties. *Chapter 1*.

kinematic bicycle model A vehicle model that lumps each axle into one wheel and assumes the tires do not slip sideways. *Chapter 2*.

kinematics The geometry of motion, relating positions and speeds without regard to the forces that cause them. *Chapter 1*.

laser scanner A sensor that sweeps a laser beam and measures the distance to surfaces from the time the light takes to return. *Chapter 1*.

latency The time between an event and the moment its measurement reaches the controller. *Chapter 11*.

lift-plus-cruise aircraft An aircraft with separate rotors for vertical lift and a propeller and wing for forward flight. *Chapter 4*.

low Earth orbit An orbit a few hundred kilometers above the Earth, where thin air still drags on a satellite. *Chapter 8*.

magnetorquer A coil on a satellite that makes a torque by pushing against the Earth's magnetic field. *Chapter 7*.

Magnus effect The sideways force on a spinning ball moving through air, which curves its flight. *Chapter 11*.

mean elements Orbit values averaged over the short wobbles caused by the Earth's flattening, used to plan station keeping. *Chapter 8*.

minimum-jerk trajectory The smooth move between two rests that keeps the total squared jerk as small as possible. *Chapter 1*.

model predictive control (MPC) Control that plans a sequence of future moves on a model, applies the first and plans again at the next step. *Chapter 2*.

multirotor An aircraft lifted and steered by several fixed rotors whose speeds or thrusts are varied. *Chapter 3*.

observer An algorithm that estimates states that are not measured from a model and the measurements that are. *Chapter 6*.

path following Steering a vehicle onto a path and along it, with the error measured as a distance to the path. *Chapter 1*.

peg-in-hole The assembly task of inserting a cylindrical part into a closely fitting hole. *Chapter 1*.

pitot tube A probe facing the airflow whose pressure gives the airspeed. *Chapter 4.*

platforming and contouring Two ways to ride waves: holding the boat level while waves pass beneath, or following the water surface. *Chapter 5.*

power and force limiting A way of working beside people that allows contact but keeps its force and energy below set limits. *Chapter 15.*

protective separation distance The distance a robot must keep from a person so it can stop before contact, allowing for both their motions. *Chapter 1.*

protective stop A controlled stop that a safety function commands, after which the robot may resume. *Chapter 15.*

pure pursuit A path-following law that steers toward a point a set distance ahead on the route. *Chapter 2.*

reaction wheel A spinning wheel inside a spacecraft whose speed changes turn the craft the other way by exchanging momentum. *Chapter 7.*

remote center compliance A passive wrist that lets a held part shift and tilt about a point near its tip, so it slides into a hole. *Chapter 13.*

remote center of motion A fixed point, at the entry port, about which a surgical instrument pivots. *Chapter 14.*

repetitive control Control that learns a disturbance of known period, such as a heartbeat, and cancels it in later periods. *Chapter 14.*

ride control The control system that holds a foiling boat's height and pitch above the water. *Chapter 5.*

rigid-body dynamics The balances of force and moment that set how a solid body accelerates and turns. *Chapter 1.*

runtime assurance A monitor that checks a controller's commands while it runs and replaces them when they would break a safety limit. *Chapter 15.*

SCARA arm An assembly arm whose joints turn about vertical axes, so it moves easily sideways but is stiff vertically. *Chapter 10.*

seam tracking Steering a welding torch along a joint line that a sensor finds as the weld proceeds. *Chapter 12.*

self-regulation The property of a welding arc that partly restores its own length after a disturbance, with no controller. *Chapter 12.*

sensor fusion Combining several sensors so that each covers the others' weaknesses in one estimate. *Chapter 1.*

speed and separation monitoring Slowing or stopping a robot so that it always keeps at least the protective separation distance from a person. *Chapter 1.*

Stanley controller A path-following law that turns the wheels to the path heading plus a correction that shrinks as speed grows. *Chapter 1.*

star tracker A camera that recognizes star patterns to give a spacecraft's attitude. *Chapter 7.*

stopping distance The distance a machine travels from the moment it must stop until it is at rest, reaction included. *Chapter 1.*

teleoperation Operating a robot from a distance, with a person moving a console and the robot copying the motion. *Chapter 14.*

through-the-arc sensing Finding torch height and position from the welding current and voltage, with no separate sensor at the torch. *Chapter 12.*

time stamp The recorded time at which a measurement was taken, used to correct the measurement for its age. *Chapter 10.*

trajectory generation Planning a reference of position against time that respects limits on speed, acceleration and jerk. *Chapter 1.*

transition flight The phase in which an aircraft hands its weight from its rotors to its wing, or back. *Chapter 4.*

tunnel thruster A propeller in a tube across a vessel's hull that pushes it sideways. *Chapter 6.*

ventilation Air drawn down from the surface onto a foil, which suddenly destroys much of its lift. *Chapter 5.*

virtual fixture A software boundary that stops or guides a robot's tool so it cannot enter a forbidden region. *Chapter 14.*

wave filtering Removing the wave-frequency motion from a vessel's measured position, so the thrusters do not chase each wave. *Chapter 1.*

weathervaning The turning of a hull by wind and current, whose yaw moment swings the vessel like a weathervane. *Chapter 6.*

wheel slip The difference between a wheel's rolling speed and the ground speed, which grows on a slippery floor. *Chapter 9.*

wire feeder The drive that pushes welding wire through the torch at a set speed. *Chapter 12.*

wrist force sensor A sensor between an arm and its tool that measures the force and torque of a contact. *Chapter 13.*

Index of Control Topics

A* search, 221
active constraint, 338
actuator
    fault, 198
    lag, 29, 58, 84
    rate limit, 29, 84, 334
actuator saturation, 42
AI agents
    engineering practice, 387
    open problems, 391
    reproduction, 389
    truth audit, 388
    verification record, 388
    while studying, 19
analytical redundancy, 89
autonomy
    validation, 385

back-calculation, 36, 90, 221, 314, 339
backlash, 335
bias estimation, 130, 168
braking
    verification, 391
bumpless transfer, 17, 36, 63, 90, 222, 247, 339
Butterworth filter, 338

cascade control, 104
    motion, 6, 54
chance constraint, 188
collaborative robot
    history, 357
collision detection
    history, 357
    momentum observer, 366
    open problems, 371
collision probability, 188
    threshold, 206
commit rule, 269
commit time, 269
commitment
    under a deadline, 384
complementary filter, 8, 104
    bias state, 113
compliance center, 313
    passive or active, 321
computed-torque control, 238
control allocation, 10, 64, 79, 130, 221
    daisy chain, 81
    directional actuators, 170
    weighted pseudo-inverse, 142
    wing first, 90
control education, 187
control layers, 15
counterfactual run, 388
covariance realism, 206
creep, 335

data association, 245
data export, 18
dead reckoning, 8, 219
    calibration, 230
deadband, 197
decision against a deadline, 258
delay margin, 317
disturbance observer, 6, 54, 89, 170
    history, 65
dynamic positioning
    equipment class, 153

effective mass, 310, 359
eigenaxis rotation, 169
emergency stop, 222
encoder, 13
encounter frequency, 110
error budget, 12, 66, 329, 384
    factorial, 344
extended Kalman filter, 195, 211
    spin state, 267

fault detection, 89
fault tolerance
    reaction wheels, 182
feasibility check, 91
feedforward
    acceleration, 6
    computed torque, 366
    estimated disturbance, 148
    from a load sheet, 381
    inverse servo model, 269
    measured disturbance, 293
feedforward control

curvature, 33
inverse dynamics, 246
model error, 68
final control element, 13
flow scheme, 4
force control, 11, 306, 339
history, 308
performance tests, 323
force estimation, 336
force sensor, 14
forecast
drag, 197
forgetting factor, 339
forward kinematics, 332
forward model, 336
friction
Coulomb, 311

gain scheduling
load, 64
speed, 7, 25
gravity feedforward, 66
guidance, navigation and control, 4

hand-eye calibration, 253
handling qualities
open problems, 74
haptic feedback
open problems, 349
harmonic predictor, 339
hybrid position and force control, 307
hysteresis, 93, 109, 363

impact force, 316
impedance control, 11, 306
inertia matrix
payload term, 240
inertial navigation
history, 3
small spacecraft, 386
infeasibility, 175
inferential control, 282
integrator windup, 96
interception, 235
speed matching, 243
under a sensing delay, 259
inverse kinematics, 334
damped least squares, 268

jamming, 317

Kalman filter, 18
hand-placed gains, 139
input disturbance, 38
kinematics, 5

lane detection, 46
open problems, 48
lane keeping assist
regulation, 47
laser scanner, 14
latency
measured before tuning, 274
least-squares fit
of a flight, 267
linear relative motion, 194
linearization, 66
loop action, 315

machine learning
physics-informed, 386
Mahalanobis distance, 197, 220
measurement delay, 8, 220, 235, 336
measurement noise, 30, 88
and actuator travel, 42
minimum-jerk trajectory, 6
peak acceleration, 266
mode switching, 79
mode transition
defects, 389
model error
payload, 250
model predictive control, 104, 340
against a geometric law, 385
attitude, 170
in the apps, 17
steering, 37
momentum management, 166
design practice, 181
motion control, 1
design procedure, 384
open problems, 389
motion scaling, 338
move blocking, 171, 340
moving average, 192
moving horizon estimation, 18
multiplicative extended Kalman filter, 168

navigation
satellite and inertial, 72

observer
bias state, 381
wave-frequency, 139

occlusion, 220
offset, 34
  proportional control, 66
orbit averaging, 167
organ motion
  open problems, 344
over-actuated systems, 130

path following, 7, 24, 221
PD control, 64, 246
penalty contact, 310
PID control
  anti-windup, 17
  derivative filter, 141
  loop action, 113, 221, 291, 339
  velocity form, 17, 63, 89, 291, 314, 339
platforming and contouring, 112
preview control, 115, 280, 282
probability dilution, 190
proportional control
  mechanical, 113
protective separation distance, 354
pure pursuit, 27, 221

quaternion feedback, 169

recursive least squares, 339
redundancy
  common-mode failure, 153
reinforcement learning
  sim-to-real, 385
repetitive control, 330
repetitive disturbance, 330
ride control, 104
rigid-body dynamics, 5
robot table tennis
  open problems, 275
robotic assembly
  open problems, 324
robotic welding
  open problems, 300
rotational stiffness, 313
rotor speed control, 73
Routh stability criterion, 35
runtime assurance, 355

safety monitoring, 11, 211
  truth audit, 354, 365
  truth check, 222
sampled-data control, 35
satellite navigation, 13
scenarios, 18
  completion rules, 18
scheduling
  earliest deadline, 245
  pick order, 237
  throughput ceiling, 250
scorecard, 18
seam tracking
  narrow gap, 289
self-regulation, 280
  welding arc, 281
sensor fallback, 114
sensor fusion, 8
servo drive
  safety stop, 274
  speed and acceleration limits, 263, 334
simulation engine, 16
simulator framework, 15
soft sensor, 282
spacecraft attitude control, 158
speed and separation monitoring, 11, 354
  commissioning, 373
  open problems, 375
speed envelope, 384
speed law, 363
stability margins
  hydrofoil, 124
standard linear solid, 335
Stanley controller, 24, 32
state estimation, 130
stopping distance, 11, 211, 363
  speed envelope, 220
stress relaxation, 335
supervised autonomy, 332
  history, 332
supervisory control, 65, 79, 91, 130, 340
  heading, 143
  mode sequencing, 169
surgical robotics
  history, 331
  regulation, 347
switched system, 80

teleoperation, 329
through-the-arc sensing
  contact tube wear, 300
time stamp, 243
trajectory generation, 6, 65, 90
  protective stop, 365
  quintic, 245, 340
  step, trapezoid and quintic, 268
trajectory prediction, 267

transition flight
open problems, 99
tremor filtering, 338

uncertainty, 12, 187

ventilation
prediction, 124
virtual fixture, 338
visual servoing, 237

wave filtering, 130
history, 133

Index of Applications

adaptive fill, 283
air taxi, 54
angle of attack, 84
approach gate, 65
atmospheric density, 188
atmospheric drag, 188
automated guided vehicles, 211
autonomous mobile robot, 386
    performance standards, 386
autonomous vehicle
    history, 4
avoidance maneuver, 197
azimuth thruster, 133

B-dot detumbling, 167
back-transition, 93
battery
    power limit, 61, 84
belt encoder, 240
blind corner, 220
bounce
    gripping and sliding, 263
building wake, 57
burn-through, 287
butt joint, 283

cable-driven wrist, 335
    in practice, 347
camera
    frame rate and latency, 264
    high-speed stereo pair, 274
capability plot, 150
center of mass, 58
chamfer, 309
clothoid, 29
collaborative robot, 354, 355, 386
collision avoidance, 187
conjunction assessment, 190
    history, 190
conjunction screening, 195
constant-current source, 281
constant-voltage source, 281
contact bumper, 213
contact tip, 281
conveyor tracking, 235
    cell model, 238
cross-track error, 30
CubeSat, 158
    deployer, 159

DARPA Grand Challenge, 27
Delta robot, 237
differential drive, 5, 215
diversion, 91
drag torque, 164
drive wheel, 228
dynamic positioning, 130
    history, 132

Earth-observation satellite, 158
electric power steering, 46
electric vertical takeoff and landing, 57, 79
electrode extension, 282
encounter plane, 195

figure of merit, 61, 84
floor friction, 216
foil lift, 109
foiling boat, 104
free height stiffness, 112
free-surface effect, 109
    history, 106
fulcrum effect, 331

gas metal arc welding, 280
gearbox friction, 371
geomagnetic field
    dipole, 163
GNSS navigation, 195
    receiver, 204
go-around, 65, 91
grasp performance testing, 253
grasp tolerance, 240
gravity gradient, 164
    libration, 166
guide wire, 213
gust, 84
gust front, 70
gyro bias, 164

Hall thruster, 188
    history, 189

in service, 205
hand-over, 370
heading error, 30
heat input, 287
hover, 82
hull resistance hump, 110
humping, 283
hydroacoustic position reference, 151
hydrofoil, 104
mechanical height control, 107

instrument, 13
ISO/TS 15066, 362

joint limits, 263

kinematic bicycle model, 24, 28

lack of fusion, 287
lane departure, 31
lane departure warning, 27
lane keeping, 7, 24
laparoscopy, 330
laser scanner, 228
laser seam sensor, 281
structured light, 298
lift capacity, 88
lift rotor
power, 84
speed control, 97
lift-plus-cruise aircraft, 79
load sheet, 64, 91
low Earth orbit, 187

machine vision, 251
magnetic attitude control
history, 161
magnetometer, 165
magnetorquers, 159
Magnus effect, 260
mean elements, 188
melting rate, 285
mobile robots, 212
momentum theory, 61, 81
momentum unloading, 170
motion control
milestones, 3
motor failure, 94
multirotor, 54
attitude by thrust difference, 56

needle insertion, 335

nodal regression, 194

oblateness of the Earth, 188
offshore construction vessel, 130
orbit determination, 195
orbit maintenance, 187
organ motion
cardiac, 336
respiratory, 336

paddle
impact law, 263
partial nephrectomy, 332
passenger comfort, 57
payload, 83
peg-in-hole, 11
peg-in-hole insertion, 306
performance standards, 230
physiological tremor, 332
pick-and-place, 235
Pierson–Moskowitz spectrum, 110
pitot tube, 88
blockage, 97
porosity, 287
position reference
bias and redundancy, 150
power and force limiting, 364
propellant budget, 192
protective stop
in practice, 373

quadrotor
history, 56

radar altimeter, 88
range scanner, 217
reaction wheels, 159
back-EMF fade, 164
failure, 161
in service, 180
loss of two, 162
reflector navigation, 217
remote center compliance, 308
remote center of motion, 331
replanning, 221
robot arm
joint drives, 323
two-link planar, 354
robot table tennis, 258
history, 260
robotic assembly, 306
robotic welding, 280

robotics, 1
root gap, 286

safety scanner, 360
  in practice, 373
satellite
  Earth-observation, 187
SCARA arm, 235
  history, 237
seam tracking, 281
  history, 283
selective compliance, 237
six-axis arm, 263
slamming, 110
space debris, 190
spacecraft attitude, 158
spin stabilization, 160
spray transfer, 285
stall, 84
stall speed, 87
  approach margin, 99
star tracker, 161
  in service, 179
steering misalignment, 34
steering rack, 29
stroke
  inverted from the landing, 268
sun sensor, 165
sun-synchronous orbit, 162, 188
surgical robot
  autonomy, 386
surgical robotics, 329
suturing, 335

table-tennis ball
  aerodynamics, 260
  flight model, 261
takeoff sequencer, 114
through-the-arc sensing, 281
thrust margin, 65
thruster wear, 146
tilt rotor, 81
  conversion, 82
tissue tracking
  endoscopic, 347
trailing-edge flap
  servo, 123
transition flight, 80
  incident, 96
transition spiral, 29
tuned opponent, 269
tunnel thruster, 133
  in service, 152

ultrasonic height sensor, 122
undercut, 287
underfill, 287
urban air mobility, 57, 80
urban wind, 57

V-groove, 283
vacuum gripper, 238
  cup, 251
  holding force, 243
ventilation, 109
viscoelastic tissue, 335

wand, 113
warehouse robot, 211
warm ischemia, 332
watch circle, 133
wave drift force, 136
wave-frequency motion, 136
weathervaning, 136
weight and balance, 58
  onboard measurement, 386
  standard average weights, 74
weld pool, 281
welding distortion, 283
welding power source, 299
welding procedure qualification, 287
welding robot, 281
wheel slip, 215
wire feeder, 281
  dynamics, 299
wrist force sensor, 311
  six-axis, 322

About the Editor

John D. Hedengren is a professor in the Department of Chemical and Biological Engineering at Brigham Young University, where he leads the Process Research and Intelligent Systems Modeling (PRISM) group. He received the B.S. and M.S. degrees in chemical engineering from Brigham Young University and the Ph.D. degree from the University of Texas at Austin in 2005. He developed advanced process control applications for polymer production at ExxonMobil Chemical Company and commissioned nonlinear model predictive controllers with PAS, Inc. before joining the faculty at Brigham Young University in 2011.

His teaching covers process dynamics and control, engineering computing, data science, machine learning, and dynamic optimization, with course material published online for use inside and outside the classroom. He developed the Temperature Control Lab, a portable heat transfer device now used by many universities for process control and data science instruction, and he maintains the APMonitor and GEKKO optimization software used in those courses. He serves as Chair of Control Education for the IEEE Control Systems Society and as a trustee of the CACHE Corporation. He received the AIChE David Himmelblau Award in 2014, the AIChE Computing Practice Award in 2018, and the John R. Ragazzini Award in 2024, and he was inducted into the Process Automation Hall of Fame in 2025.

Volumes in This Series

Interactive Control Engineering: Principles and Applications has 20 volumes. Nine are organized by control technology and eleven by industry or application area. Every chapter in every volume has a simulator at https://apmonitor.com/apps/.

1 **Process Dynamics, Modeling, and Identification.** How plants respond, and how to find the model.
2 **Practical PID Control.** Tuning, anti-windup and implementation that works on real loops.
3 **Feedforward, Cascade, and Ratio Control.** Choosing a structure from where the disturbance enters.
4 **Override, Split-Range, and Constraint Control.** When limits decide what the controller may do.
5 **Delays, Slow Measurements, and State Estimation.** Controlling what arrives late or cannot be measured.
6 **Multivariable and Plantwide Control.** Pairing loops, managing inventories, closing the recycle.
7 **Nonlinear Dynamics, Oscillations, and Stabilization.** Plants that will not stay put without control.
8 **Batch, Run-to-Run, and Hybrid Control.** Recipes, endpoints and learning from the last run.
9 **Model Predictive Control and Economic Optimization.** Planning moves ahead, and choosing where to operate.
10 **Control in Upstream and Midstream Oil and Gas.** From the wellbore to the terminal.
11 **Control in Petroleum Refining.** Throughput, quality and limits from crude unit to blending.
12 **Control in Petrochemicals and Polymers.** Severity, selectivity and grade from cracker to pellet.
13 **Control in Chemical Plants.** Recycle, equilibrium and energy in continuous chemical processes.
14 **Control in Semiconductor Manufacturing.** Run-to-run control where little can be measured in the chamber.
15 **Control in Mining and Mineral Processing.** Variable ore, late assays and equipment at its limit.
16 **Control in Iron, Steel, and Metals Production.** Slow furnaces, fast strip, and measurements that arrive once.
17 **Control of Energy Conversion and Storage.** Fast regulation and slow dispatch on a changing grid.
18 **Control of Nuclear and Fusion Systems.** Reactor physics, protection limits and coordinated control.
19 **Control of Biological Systems and Human Performance.** Variable subjects, sparse samples and limits that cannot be crossed.
20 **Motion Control, Robotics, and Autonomous Systems.** Tracking, navigation, contact and allocation.