INTERACTIVE CONTROL ENGINEERING
Principles and Applications

Volume 3 Feedforward, Cascade, and Ratio Control

CHAPTER 11

Canceling Coriolis Coupling by Feedforward in an Orbital Rendezvous Simulator

Pages 255–285

Edited by John D. Hedengren

How to cite this chapter. "Canceling Coriolis Coupling by Feedforward in an Orbital Rendezvous Simulator," in *Feedforward, Cascade, and Ratio Control*, John D. Hedengren, Ed., vol. 3 of *Interactive Control Engineering: Principles and Applications*. Advanced Process Solutions, LLC, 2026, ch. 11, pp. 255–285.

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

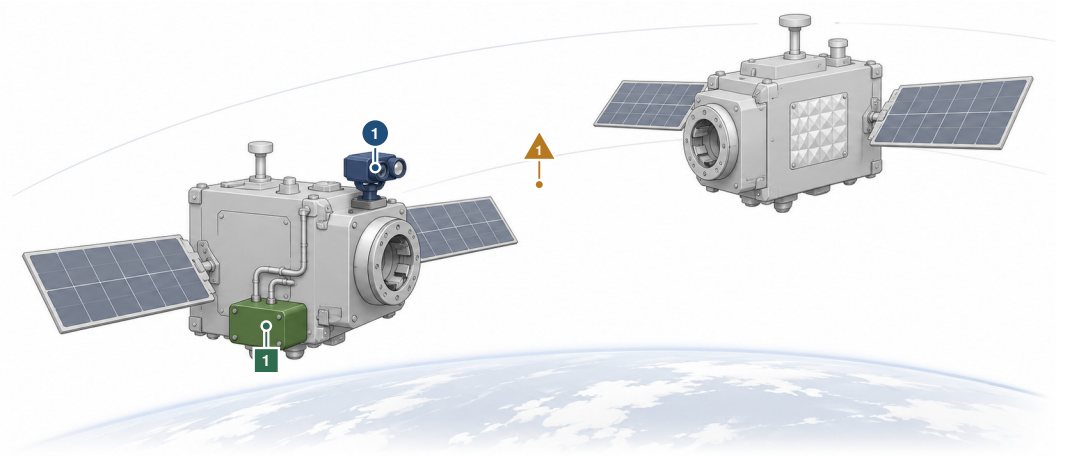
The app for this chapter
https://apmonitor.com/apps/orbital-rendezvous/

The volume online: contents, other chapters and citation records
https://apmonitor.com/apps/books/v03-feedforward-cascade/#orbital-rendezvous

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

CHAPTER 11

Canceling Coriolis Coupling by Feedforward in an Orbital Rendezvous Simulator



The controller moves 1 thrusters on three axes

The controller reads 1 relative position and velocity sensor

What disturbs it 1 differential drag along the track

The approaching craft reads its position relative to the target and fires its thrusters, but because both circle the Earth, a push aimed straight at the target also lifts it or lowers it. Illustration, not to scale.

IN PLAIN TERMS **Steering Up Close to a Craft in Orbit**

The main idea. The view from the target turns once each orbit, so near another craft a push in one direction also moves the arriving craft in another. The controller already knows these side effects and adds small extra pushes that cancel them before any error appears.

The challenge. A faster controller hides the side effects but never removes them. The extra pushes also burn fuel, so there is a best speed that uses the least.

Why it matters. Every supply flight to a space station ends with this slow approach, flown on fuel that cannot be refilled.

THREE WAYS TO READ THIS CHAPTER

Explore	12 min	the first page, the history of proximity operations and the fact on manual braking (Section 11.2), the first try-it segment (Section 11.3.3), the second try-it segment and the fact on station keeping (Section 11.6) and the key points
Study	55 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 11.6), Technology in Practice (Section 11.7), Limitations (Section 11.8) and the open challenge

The control question A chaser near a target in orbit is flown in axes that turn once per orbit, and the gravity-gradient and Coriolis terms make three thrusters on three axes a coupled plant. What does canceling that known coupling by feedforward buy that more feedback gain cannot, and what does it cost in propellant?

Learning objectives

- Derive the free-drift paths of the Clohessy-Wiltshire equations and explain why a hold behind the target is free.
- Show that the model feedforward leaves three damped second-order loops, and compute the stiffness an uncompensated loop loses.
- Compute the propellant-optimal bandwidth and the bandwidth the thrusters can support.
- Measure what a wrong mean motion, a thrust limit and a degraded thruster cost at two bandwidths.
- Compare a propellant-pricing MPC with the compensated loop on propellant and docking time.

Prerequisites Sections 1.4.2 (where a disturbance enters), 1.4.3 (static feedforward and the ideal compensator) and 1.6.4 (scenarios and exported files). Section 1.6.3 for the MPC only. Chapter 10 for a gravity feedforward on another vehicle.

The app https://apmonitor.com/apps/orbital-rendezvous/

See also Volume 6, Chapter 2 for a static feedforward that removes loop interaction. Volume 20, Chapter 7 for the attitude loop this chapter leaves out, Volume 20, Chapter 8 for orbit maintenance under drag and oblateness, Volume 20, Chapter 6 for thruster allocation and a failed thruster.

11.1 Introduction

Three thrusters on three orthogonal axes present as three independent actuators, and the natural pairing puts each one with the position it appears to move [1]. That pairing is correct almost everywhere a spacecraft flies. It stops being correct inside a few kilometers of another vehicle, because the frame in which the relative position is measured is attached to the target and the target is following a curved path. The axes rotate once per orbit, at 1.131 mrad/s in a 400 km circular orbit, and every consequence of writing Newton's second law in rotating axes appears as a term in the relative equations of motion.

Two of those terms make the pairing false. A Coriolis term of magnitude $2n$ times a velocity appears with the axes exchanged, so an along-track velocity produces a radial acceleration and a radial velocity produces an along-track acceleration. A gravity-gradient term of magnitude $3n^2$ pushes a radial displacement further out rather than pulling it back. Neither is a force acting on the chaser. Both are exact once the target's mean motion is known, which is the property the simulator of this chapter is built around. A proportional-derivative loop on each axis, tuned as though the axes were independent, is fighting a positive feedback path on the radial axis and is being pushed sideways by a path it cannot see. Adding four algebraic terms that the flight computer can already evaluate removes both effects and turns the plant into three uncoupled second-order systems.

Relative motion of this kind has a settled analytical description. The linearized equations for a satellite near a circular reference orbit were published as a terminal guidance system for satellite rendezvous [2], and they are the same equations that appear in the lunar theory of the previous century [3]. Rendezvous guidance built on them ranges from piloted braking gates and glideslope profiles [4] to trajectory families chosen so that a failed burn leaves a passively safe relative orbit [5], and to model predictive formulations that carry the thrust bound and the approach corridor as explicit constraints [6, 7]. However, the step that makes any of those methods work is elementary and is usually left implicit. The coupling terms are known, so they are subtracted rather than rejected, and what remains for feedback is only the unmodeled part of the plant.

The plant carries a specific set of process dynamics and control principles: the difference between a pairing that is wrong and a plant that is coupled, an unstable open-loop mode and a marginally stable one on the same vehicle, feedforward from a model rather than from a measured disturbance, the sensitivity of a feedforward term to the parameter it is evaluated with, actuator saturation as a

limit on achievable bandwidth, and proportional offset under a constant load [1, 8, 9]. The physical fundamentals are Newton's second law written in a rotating frame, a first-order expansion of the gravitational difference between two nearby bodies, and the separation between a bounded periodic mode and a secular one. The behavior matters industrially because station resupply, crew rotation, on-orbit servicing and debris removal all reduce to the same terminal approach, flown against a propellant budget that cannot be replenished in flight.

This chapter measures what one control decision is worth on a plant whose coupling is known exactly. It works with an interactive rendezvous simulator, built on the framework of Section 1.6, in which the Clohessy-Wiltshire terms are canceled or left in place by a single switch. The central result is a propellant argument. Because the feedforward term $3n^2x$ must itself be produced by the thrusters, the cumulative velocity change of a compensated regulator has an interior minimum in closed-loop bandwidth at $\omega_n = n\sqrt{3\mathrm{e}} = 3.23\,\mathrm{mrad/s}$, which a fine sweep of the simulation places at 3.14 mrad/s. The six-state simulator's thruster bound, docking test, disturbance set and bumpless-transfer arrangement are set down completely enough that another implementation reproduces the same trajectories. Its set-ups present the plant before the controller and break the coupling, the actuator bound and the model error one at a time. The Clohessy-Wiltshire model and the algebraic cancellation of known terms are both established methods, and the chapter uses them as they stand.

Section 11.2 places the plant in rendezvous practice and guidance. Section 11.3 gives the relative motion model, its free-drift solutions, its constants and its assumptions, with the equations and initial states a reimplementation needs. Section 11.4 derives the model feedforward, its sensitivity to the assumed altitude, the bandwidth the thrusters allow and the predictive controller. Section 11.5 lists the set-ups, and Section 11.6 measures the compensation at fixed gains, the bandwidth sweep, saturation, model error, station keeping, a degraded thruster, the graded challenge and the predictive controller. Section 11.7 sets the loops on flight hardware, Sections 11.8 and 11.9 state the limits of the model and what transfers to other plants, and the exercises of Section 11.11 close the chapter.

11.2 Background

Rendezvous was the enabling technique of the lunar program and it was flown by hand first. The Gemini program established that two crewed vehicles could be brought together and docked, and the piloted procedures worked out there, with braking gates, a defined approach corridor and a station-keeping phase before contact, carried forward into Apollo and then into the Space Shuttle [10]. What the crews were taught is the lesson the simulator of this chapter reproduces. Pointing the nose at the target and thrusting does not close the range, and the correct action at long range is frequently a burn along the velocity vector rather than toward the object being approached.

> ⧖ HISTORY **Proximity Operations Gets a Name, 1973 to 1977**
>
> Gemini and Apollo flew an inertial final approach and braked at the end. By mid-1973 the reusable orbiter raised a new concern: its primary reaction control jets, of 870 pounds thrust, could contaminate a small target with their exhaust, induce attitude rates on it or even push it away. Piloted simulations in 1976 showed that the inertial approach did not remove the concern, while approaches along the target's velocity vector or along its radius vector were acceptable. By March 1977 the term proximity operations had been coined for flight within 2000 feet of the target, characterized by nearly continuous trajectory control, whereas rendezvous maneuvers come at intervals of hours or tens of minutes [11].

Automated rendezvous developed on a separate track and for a different reason. A station in permanent occupation has to be resupplied whether or not a pilot is aboard the arriving vehicle, and Progress freighters, the European Automated Transfer Vehicle and the commercial cargo and crew vehicles that followed all perform the terminal approach without a hand on a controller [12]. The transition from piloted to automated operation is not a matter of replacing the pilot with a controller, because the pilot was also supplying the navigation, the fault detection and the decision to abort.

A review of what automation has to take over, and of the flight demonstrations that attempted it, is given in [13].

> ✦ DID YOU KNOW? **Braking Left to the Crew**
>
> A 1966 memo on Apollo terminal guidance deleted automatic braking from the guidance computer. Based on Gemini experience, crews could brake as well given range and range rate, and the program would have used 50 to 100 words of scarce storage. The first automated rendezvous and docking was flown in October 1967 by two uncrewed spacecraft [11].

The guidance literature splits along the same line. Impulsive targeting solves a two-point boundary value problem for the two burns that move a vehicle from one relative state to another, which is a purely open-loop construction and requires that the relative motion between the burns is predicted correctly. Glideslope approaches replace it with a closed-loop rule in which the closing speed is commanded as a linear function of the remaining range, executed as a sequence of small burns, and the relative navigation needed to support it is treated together with the guidance in [4]. Both families raise a question no amount of tuning answers, which is what the vehicle does if the next burn does not happen. Trajectories chosen so that the uncorrected relative motion never intersects the target, known as safety ellipses and passive-abort trajectories, are designed by adding that requirement as a constraint on the free-drift solution [5]. Model predictive control is the natural home for the whole set, because the thrust bound, the approach cone, the plume-impingement keep-out region and the terminal speed limit are all constraints, and formulations that carry them explicitly for proximity operations are given in [6, 7].

The simulator of this chapter sits underneath that literature rather than beside it. An optimizer that plans a constrained approach still has to be compared against the simplest thing that could work, and the simplest thing that could work on this plant is a proportional-derivative loop with the known terms fed forward. The simulator also carries such a planner as an option, and Section 11.6.9 weighs it against that loop on propellant and on time to dock.

11.3 Relative Motion Model

Two vehicles fly near one another in a circular orbit of radius $R = R_E + h_{orb}$ about a central body of gravitational parameter μ. The target occupies the origin of a right-handed frame whose axes are carried around with it, and the chaser is described by its position and velocity in that frame. Figure 11.1 defines the frame, the sense of each axis and every symbol the equations use.

11.3.1 Classification of Variables

Three loops are declared, one per axis, and each one pairs a relative position with the thruster acceleration on the same axis. Table 11.1 names every signal, the part it plays in the loop and the range over which it can be set.

Nothing in the vehicle is measured imperfectly. The controller reads the true position and velocity on all three axes with no noise, no bias, no lag and no navigation filter, which is the largest single idealization in the model and is revisited in Section 11.3.5. The only imperfection available is in the actuator, because a de-rated thruster still consumes the command it was given and does not produce the acceleration.

11.3.2 The Clohessy-Wiltshire Equations as Implemented

Both vehicles are point masses in the gravity field of a spherical central body. Expanding the difference of their gravitational accelerations to first order in the separation, and writing the result in axes that

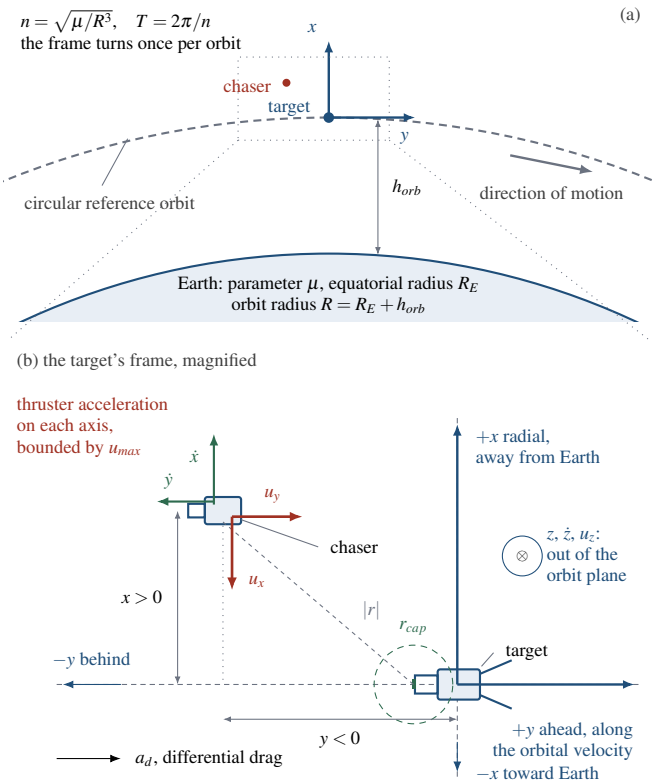


Figure 11.1 Relative motion geometry and the local-vertical local-horizontal frame. The frame origin rides on the target, x points radially away from the central body, y points along the orbital velocity, and $z = \hat{x} \times \hat{y}$ completes the right-handed triad along the orbit angular momentum, which is into the page in this view. A chaser behind and above the target therefore has $y < 0$ and $x > 0$. The whole frame turns at the mean motion n, and the chaser is driven only by its own bounded thruster acceleration (u_x, u_y, u_z) and by the optional along-track differential drag a_d. Panel (b) is magnified by roughly four orders of magnitude relative to panel (a).

rotate at the constant rate n about the orbit normal, gives the system in Equation (11.1).

$$\begin{aligned} \ddot{x} &= 3n^2 x + 2n\dot{y} + u_x \\ \ddot{y} &= \quad -2n\dot{x} + u_y + a_d \\ \ddot{z} &= -n^2 z \quad\quad + u_z \end{aligned} \tag{11.1}$$

where x, y and z are the relative position components in m, u_x, u_y and u_z are the commanded thruster accelerations in m/s^2, a_d is the along-track differential drag in m/s^2 and n is the mean motion of the reference orbit in rad/s. Every line carries units of m/s^2. The terms $3n^2x$ and $-n^2z$ are the net of the gravity gradient and the centrifugal term of the rotating frame, and the two terms $\pm 2n$ times a velocity are Coriolis. None of the four is a force on the chaser, and all four vanish if the same motion is written in an inertial frame. All four are consequences of choosing axes in which the target is stationary.

The mean motion follows from the circular-orbit condition and is given by Equation (11.2).

$$n = \sqrt{\frac{\mu}{R^3}}, \qquad R = R_E + h_{orb}, \qquad T = \frac{2\pi}{n} \tag{11.2}$$

Table 11.1 Classification of Every Signal. Entries Marked (adv.) Are Available Only in the Advanced Set.

Symbol	Role	Default	Range
x	CV, radial position	0 m	—
y	CV, along-track position	0 m	—
z	CV, cross-track position	0 m	—
$\dot{x},\dot{y},\dot{z}$	states, relative velocity	0 m/s	—
x_{SP}	setpoint, radial	0 m	± 300 m
y_{SP}	setpoint, along-track	0 m	± 500 m
z_{SP}	setpoint, cross-track (adv.)	0 m	± 200 m
u_x	MV, radial acceleration	0	$\pm u_{\max}$
u_y	MV, along-track acceleration	0	$\pm u_{\max}$
u_z	MV, cross-track acceleration	0	$\pm u_{\max}$
a_d	DV, differential drag	0	$\pm 20\,\mu\mathrm{m/s^2}$
—	DV, degraded thruster axis	none	x, y or z
$1-\eta$	DV, thrust lost on that axis	0 %	0 to 100 %
ω_n	tuning, in-plane bandwidth	10.0	1 to 40 mrad/s
ζ	tuning, damping ratio	1.00	0.20 to 2.00
$\omega_{n,z}$	tuning, cross-track (adv.)	8.0	1 to 40 mrad/s
$\hat{h}_{orb}$	tuning, assumed altitude (adv.)	400 km	200 to 1000 km
—	option, control architecture	PD + comp.	three, see text
T_{arr}	tuning, MPC arrival time (adv.)	26 min	5 to 40 min
$\lvert r\rvert$	output, range to the port	0 m	—
$\lvert v\rvert$	output, relative speed	0 m/s	—
$\dot{r}$	output, closing speed	0 m/s	—
Δv	output, cumulative burn	0 m/s	—
n	output, true mean motion	1.131	mrad/s
T	output, orbital period	5553.6	s

where μ is the gravitational parameter in $\mathrm{m^3/s^2}$, R_E is the equatorial radius of the central body in m and h_{orb} is the orbital altitude in m. At the default 400 km this evaluates to $n = 1.131\,367\,\mathrm{mrad/s}$ and $T = 5553.6\,\mathrm{s}$, which is 92.56 min. Every coupling term in Equation (11.1) scales as n or n^2, so the coupling weakens with altitude. Between 200 km and 1000 km the mean motion falls by 15.8 % and the gravity-gradient coefficient $3n^2$ by 29.1 %.

Equation (11.1) is linear and time invariant, and the only nonlinearities anywhere in the simulation are the thruster clamp, the de-rating of a faulted thruster, the washout of the bumpless-transfer bias and the three latched proximity events. The state vector is exactly the six relative coordinates, and the propellant total is accumulated after the integration step rather than carried as a seventh state.

11.3.3 Free-Drift Behavior of the Three Axes

Setting $u_x = u_y = u_z = a_d = 0$ in Equation (11.1) leaves the free-drift solution, and the three axes behave completely differently. For a chaser released from rest at a pure radial offset x_0 the in-plane solution is Equation (11.3).

$$x(t) = (4 - 3\cos nt)\,x_0, \qquad y(t) = 6\,(\sin nt - nt)\,x_0 \tag{11.3}$$

The radial component is bounded and periodic and reaches $7x_0$ at half a period. The along-track component is not bounded, because the secular term $-6ntx_0$ never stops growing, and its slope $6nx_0$ is 0.679 m/s for $x_0 = 100\,\mathrm{m}$ at the default altitude, or 40.7 m of drift per minute from an initial condition in which nothing was moving. A radial offset is an unstable mode of this plant.

An along-track offset is the opposite case. Substituting $x = z = 0$ and $\dot{x} = \dot{y} = \dot{z} = 0$ into Equation (11.1) makes every right-hand side vanish, because every term contains x, z or a velocity. A

vehicle parked anywhere along the velocity vector with zero relative velocity is at an exact equilibrium of the plant and stays there at no propellant cost. This asymmetry is why hold points on a real approach are placed along the velocity vector rather than above or below the target.

The cross-track axis is an undamped harmonic oscillator at exactly the orbital frequency, and its solution for a pure out-of-plane rate $\dot{z}_0$ is Equation (11.4).

$$z(t) = \frac{\dot{z}_0}{n} \sin nt \tag{11.4}$$

The amplitude is $\dot{z}_0/n$, which is 88.4 m for 0.10 m/s, and the error returns to zero once per orbit whatever the controller does. A plane error is bounded and self-correcting, so the cross-track loop is given its own and slower bandwidth.

The case that matters most for the control problem is a pure closing velocity. For a chaser starting at y_0 with an along-track rate $\dot{y}_0$ and nothing else, the radial solution is Equation (11.5).

$$x(t) = \frac{2\dot{y}_0}{n} (1 - \cos nt) \tag{11.5}$$

The Coriolis term has converted a velocity that pointed exactly at the target into a radial excursion whose amplitude is $2\dot{y}_0/n$, and for 0.5 m/s that amplitude is 883.9 m. Nothing in a coast removes it. Figure 11.2 draws all three cases in the plane of the orbit, computed from Equations (11.3) and (11.5) with the same mean motion the simulator uses.

Pause. Try Exercise 11.1 before reading on.

> **TRY IT IN THE APP** **3 · Aimed Straight at the Target**
>
> Load the scenario, which starts in Manual with every thruster at zero, and press Run. Watch the range to the port. It falls to a closest approach of 115.22 m at 363.5 s, with the chaser 73.7 m above the port, and then grows as the chaser climbs away, although its velocity pointed straight at the port.

11.3.4 Numerical Constants, Limits and Latched Events

Table 11.2 collects every constant in the model. A fixed 0.5 s step of fourth-order Runge-Kutta advances the six states, the control law is evaluated on the same 0.5 s period, and one sample is logged per simulated second. Both rates are set against the fastest thing in the problem, which is the closed loop at a few mrad/s rather than the orbit itself. A control period of 0.5 s is 8×10^{-4} of one period of a closed loop at 10 mrad/s. A step does its work in the order of Section 1.6.1, with the propellant accounting and the proximity tests placed after the integration. The right-hand side of Equation (11.1) writes nothing back into the state, so the four Runge-Kutta stages cannot disagree.

Three events are latched after the integration step and are the only discontinuities in the run. Inside the capture radius r_{cap} a relative speed at or below v_{cap} declares a dock and the run pauses once so the propellant total can be read, except in the graded challenge, which runs on to its verdict at 1500 s. A relative speed above v_{cap} declares a collision. Beyond $r_{\max}$ a loss-of-proximity flag is raised, with 5 % hysteresis, marking the range at which the linearization of Equation (11.1) stops being a fair description. Propellant is accumulated as the time integral of the commanded acceleration magnitude given by Equation (11.6).

$$\Delta v = \int_0^t \sqrt{u_x^2 + u_y^2 + u_z^2}\, \mathrm{d}\tau \tag{11.6}$$

A second total is accumulated the same way from the accelerations actually delivered, so a de-rated thruster makes the two diverge. Neither is an impulsive budget, because the propulsion model is a continuous bounded acceleration.

Two continuous disturbances and two one-shot disturbances are available in every set-up. A differential-drag bias of $\pm 20\,\mu\mathrm{m/s^2}$ enters the along-track equation directly. A selected axis can

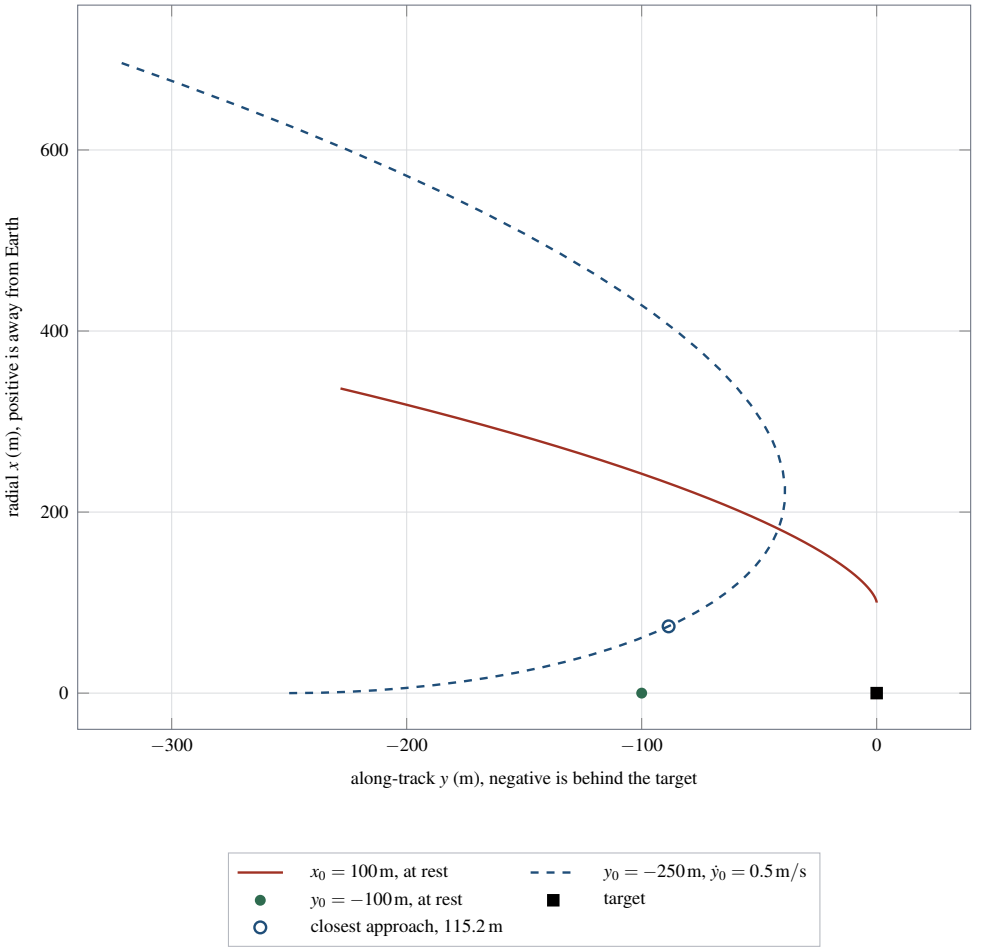


Figure 11.2 Twenty minutes of free drift from three initial conditions, evaluated from the closed-form solutions of Section 11.3.3 at $n = 1.131\,367\,\mathrm{mrad/s}$. The radial offset (solid) curls away and falls behind, the along-track offset (filled circle) does not move at all, and the coast whose relative velocity points exactly at the target (dashed) rises above it and passes over the top with a miss distance of 115.2 m at $t = 364\,\mathrm{s}$ (open circle). The target is the filled square. Every thruster is off in all three cases.

have between 0 % and 100 % of its thrust removed, and the commanded value stays the manipulated variable so the failure appears as a gap between the command and the applied trace. A venting impulse adds 0.03 m/s to $\dot{y}$ and 0.01 m/s to $\dot{x}$, and a plume-impingement event adds 0.02 m/s to $\dot{x}$ away from the target. In manual operation a 0.05 m/s translation pulse is available in each of the six directions as an instantaneous change in the corresponding velocity state, charged to the propellant budget.

11.3.5 Assumptions

First-order expansion about a circular orbit. Equation (11.1) drops every term of second order and higher in the separation divided by the orbit radius, and it assumes the reference orbit is exactly circular. Eccentricity, the J_2 oblateness term and third-body effects are all absent. The model declares its own validity boundary at $r_{\max} = 2000\,\mathrm{m}$ against an orbit radius of 6778 km, a ratio of 3×10^{-4}. The oblateness term is the first thing to matter over long horizons, because it makes the relative

Table 11.2 Plant, Controller and Numerical Constants. An Adjustable Parameter Shows the Range Over Which It Can Be Set. A Hidden Row Is a Constant the Model Uses and Does Not Report.

Symbol	Meaning	Value	Unit	Range
Orbit				
μ	gravitational parameter	3.986004418	10^{14} m^3/s^2	fixed
R_E	equatorial radius (hidden)	6378137	m	fixed
h_{orb}	true orbital altitude	400	km	200-1000
n	mean motion, derived	1.131367	mrad/s	—
T	orbital period, derived	5553.6	s	—
$3n^2$	gravity-gradient term	3.840×10^{-6}	1/s^2	—
$2n$	Coriolis coefficient	2.263×10^{-3}	1/s	—
Propulsion and proximity				
$u_{\max}$	thrust acceleration limit	0.050	m/s^2	0.002-0.050
Δv_p	translation pulse (hidden)	0.05	m/s	fixed
—	duty threshold (hidden)	1	% of $u_{\max}$	fixed
r_{cap}	capture radius	3.0	m	1.0-10.0
v_{cap}	maximum contact speed	0.10	m/s	0.01-0.50
$r_{\max}$	loss-of-proximity range	2000	m	500-10000
Controller				
ω_n	in-plane bandwidth	10.0	mrad/s	1-40
ζ	damping ratio	1.00	—	0.20-2.00
$\omega_{n,z}$	cross-track bandwidth	8.0	mrad/s	1-40
$\hat{h}_{orb}$	assumed orbital altitude	400	km	200-1000
K_P	in-plane position gain	1.000×10^{-4}	1/s^2	—
K_D	in-plane rate gain	2.000×10^{-2}	1/s	—
τ_b	bias washout (hidden)	60	s	fixed
T_{arr}	MPC arrival time	26	min	5-40
Metrics and numerics				
—	settled band, each loop	1.0	m	—
—	equilibrium tolerance	0.01	m	—
Δt	integration step (RK4)	0.5	s	—
Δt_c	control period	0.5	s	—
—	historian rate	1	Hz	—

motion secular in a way the circular model does not capture [14], and the drift it produces is a formation flying problem rather than a rendezvous problem only because a rendezvous lasts an hour. Volume 20, Chapter 8 flies a single satellite with the oblateness term and drag in the model, over days rather than an hour.

Point masses with no attitude. Neither vehicle has an orientation, so thrust is available instantly along any axis, the docking port is always pointing the right way and there is no attitude loop competing for the same thrusters. A real chaser holds an attitude to aim both its sensors and its docking mechanism, and the attitude loop is the faster and more delicate of the two. Volume 20, Chapter 7 treats such a loop on a small satellite, with reaction wheels and magnetic torquers.

Continuous bounded acceleration. Reaction control thrusters are on and off devices with a minimum impulse bit, arranged in a geometry that rarely delivers a pure single-axis force, and the vehicle loses mass as they fire. The model has none of that, so a command of 0.001 mm/s^2 is executed exactly.

Exact relative navigation. The controller reads the true state, so the rate feedback that supplies all of the damping costs nothing here and is never free on orbit. Close-range relative navigation is a filtered combination of lidar, camera and radio ranging, and its error rather than its tuning is what sets achievable docking precision [13].

One environmental acceleration. Differential drag is the only non-gravitational effect besides thrust, it acts along-track only, and the adjustable range of $\pm 20\,\mu\mathrm{m/s^2}$ is a teaching stress range, not a representative drag prediction at either altitude. Solar radiation pressure and plume impingement as a continuous effect are absent.

Geometric docking. Contact is declared by a radius and a speed. There is no soft-capture mechanism, no lateral or angular misalignment tolerance and no contact dynamics, so a dock and a collision are separated by a number rather than by a structural calculation.

Two deliberately unphysical set-ups. Mean motion is one of the best-known numbers on any real mission, since it falls out of the tracking solution to many digits, so the 15.8 % ephemeris error of set-up 8 is not a credible failure. It stands in for the quantities that genuinely are uncertain, such as a mass property, a thrust coefficient or an inertia, and the arithmetic of Equation (11.12) is identical for all of them. A thruster that delivers a smooth 30 % of its command is not how reaction control thrusters fail either, but it isolates the one question set-up 10 exists to ask, which is what a feedforward term does when the actuator it assumes to be exact is not.

11.4 Controller Design

One switch selects among three architectures, the proportional-derivative loops alone, the same loops with the model feedforward described below, which is the default, and the model predictive controller of Section 11.4.6. The default controller has two parts. The feedback part is a proportional-derivative law on each axis, written as a gain on the error minus a gain on the measured rate and evaluated on the state vector directly, as shown in Equation (11.7).

$$\begin{aligned} a_x &= K_P\,(x_{\mathrm{SP}} - x) - K_D\,\dot{x} \\ a_y &= K_P\,(y_{\mathrm{SP}} - y) - K_D\,\dot{y} \\ a_z &= K_{P,z}\,(z_{\mathrm{SP}} - z) - K_{D,z}\,\dot{z} \end{aligned} \tag{11.7}$$

where a_x, a_y and a_z are the feedback demands in $\mathrm{m/s^2}$, K_P and $K_{P,z}$ are position gains in $1/\mathrm{s^2}$ and K_D and $K_{D,z}$ are rate gains in 1/s. There is no integral term anywhere, no derivative filter and no internal state, so the feedback law is a memoryless function of the state. It is therefore not the velocity-form PID of Section 1.6.2, and its bumpless transfer is arranged separately, as Section 11.4.4 describes.

11.4.1 Gain Parameterization

Six raw gains obscure the quantity a designer actually chooses, so the loops are parameterized the way a second-order closed loop is specified, by a natural frequency and a damping ratio [1], as shown in Equation (11.8).

$$K_P = \omega_n^2, \quad K_D = 2\zeta\,\omega_n, \quad K_{P,z} = \omega_{n,z}^2, \quad K_{D,z} = 2\zeta\,\omega_{n,z} \tag{11.8}$$

The radial and along-track axes share one ω_n because they are dynamically coupled through the Coriolis terms, and giving them different bandwidths only means that whatever coupling survives arrives at a frequency the other loop is not prepared for. The cross-track axis has no cross terms at all and is given its own and usually slower ω_n, because Equation (11.4) says its error is bounded and returns to zero once per orbit without help. The damping ratio is shared across all three, because it sets the shape of the response rather than its speed. At the reference tuning $\omega_n = 10\,\mathrm{mrad/s}$ the in-plane loop is 8.84 times faster than the orbit itself.

11.4.2 Model Feedforward and Exact Cancellation

The feedforward block is an algebraic copy of the coupling terms of Equation (11.1), evaluated with the mean motion the controller believes it is flying at, and the total command is Equation (11.9).

$$\begin{aligned} u_x &= -3\hat{n}^2 x - 2\hat{n}\dot{y} + a_x \\ u_y &= \quad\quad + 2\hat{n}\dot{x} + a_y \\ u_z &= \hat{n}^2 z \quad\quad + a_z \end{aligned} \tag{11.9}$$

where $\hat{n} = \sqrt{\mu/(R_E + \hat{h}_{orb})^3}$ is computed from the controller's own assumed altitude, which is a separate control from the plant altitude so that the mismatched case can be set up deliberately. Substituting Equation (11.9) into Equation (11.1) with $\hat{n} = n$ leaves Equation (11.10).

$$\ddot{e} + K_D\,\dot{e} + K_P\,e = 0 \tag{11.10}$$

on each of the two in-plane axes, with $K_{P,z}$ and $K_{D,z}$ in place of K_P and K_D on the cross-track axis, where e is the position error on the axis concerned. Three coupled axes, one of them unstable, have become three independent damped second-order systems with their poles exactly where ω_n and ζ place them. The cancellation requires only the state the controller already measures and one number it already knows, with no prediction horizon and no online solve. The coupling terms enter each axis at the same point as the thrust, so the ideal compensator of Equation (1.4) reduces to a pure gain on measured states, with no lead–lag and no delay. The app's own check on this identity drives the compensated loop from the start of set-up 6, $x_0 = 100\,\text{m}$ at rest, at $\omega_n = 12\,\text{mrad/s}$ and $\zeta = 0.7$ for 600 s and compares it against the analytic second-order response. The worst deviation is 0.148 m against a tolerance of 0.165 m, and that tolerance is not arbitrary. It is the half-sample of travel that a zero-order hold at Δt_c introduces against a continuous law, and the deviation halves whenever Δt_c halves, to 0.074 m at 0.25 s.

Pause. Try Exercise 11.2 before reading on.

Switching the compensation off sets all three feedforward terms to zero and leaves Equation (11.7) alone against Equation (11.1). The radial axis then closes as Equation (11.11).

$$\ddot{x} + K_D\dot{x} + \left(K_P - 3n^2\right)x = 2n\dot{y} \tag{11.11}$$

The gravity-gradient term subtracts directly from the position gain. At $\omega_n = 3\,\text{mrad/s}$ the loss is $3n^2/K_P = 42.7\,\%$, at $\omega_n = 10\,\text{mrad/s}$ it is 3.84 %, and the whole of the radial stiffness is gone at $\omega_n = \sqrt{3}\,n = 1.9596\,\text{mrad/s}$. The right-hand side is the term that has no analogue in a single-axis problem at all, because it drives an axis whose own error is zero.

11.4.3 Sensitivity to an Error in the Assumed Altitude

Feedforward cancels the difference between the plant terms and the model terms, so a mismatch in $\hat{n}$ leaves the residuals of Equation (11.12).

$$\ddot{x} + K_D\dot{x} + \left[K_P - 3\left(n^2 - \hat{n}^2\right)\right]x = 2\left(n - \hat{n}\right)\dot{y} \tag{11.12}$$

The useful comparison is against the gains rather than against n itself. A target at 200 km flown with an ephemeris for 1000 km gives $n = 1.1834\,\text{mrad/s}$ and $\hat{n} = 0.9962\,\text{mrad/s}$, an error of 15.8 %. The surviving radial term is $3(n^2 - \hat{n}^2) = 1.224 \times 10^{-6}\,/\text{s}^2$, which is 30.6 % of K_P at $\omega_n = 2\,\text{mrad/s}$ and only 7.6 % of it at 4 mrad/s. The surviving Coriolis term $2(n - \hat{n}) = 3.743 \times 10^{-4}\,/\text{s}$ is 9.4 % of K_D at the lower bandwidth and 4.7 % at the higher one. The same model error is therefore worth four times as much to a slow loop as to a fast one, which is the trade every feedforward design turns on. As in Section 1.4.3, the feedback kept beside the feedforward removes what a wrong model term leaves behind.

11.4.4 Actuator Bound and Bumpless Transfer

Each axis is clamped independently at $\pm u_{\max}$ and there is no rate limit, no minimum impulse bit and no mass depletion. Because Equation (11.7) carries no integrator there is nothing to wind up while the clamp is active, so recovery is immediate once the error shrinks. The demand at the start of a maneuver is K_P times the initial error, so the bandwidth an actuator can support is bounded by Equation (11.13).

$$\omega_n \lesssim \sqrt{\frac{u_{\max}}{e_0}} \tag{11.13}$$

where e_0 is the initial error on the axis in m. A degraded thruster multiplies the clamped command by its efficiency η on the selected axis only, and the commanded value remains the manipulated variable that the trends and the export report.

A proportional-derivative controller has no integrator for the tracking of Section 1.6.2 to write into, so bumpless transfer is arranged explicitly. While the operator holds the actuator, each axis stores the bias $b_i = u_{i,\mathrm{man}} - (ff_i + a_i)$, which is the difference between what the operator is commanding and what the law would command. On transfer to automatic operation the command is $ff_i + a_i + b_i$, which reproduces the operator's value exactly, and the bias then decays as $b_i \leftarrow b_i \exp(-\Delta t_c/\tau_b)$ with $\tau_b = 60\,\mathrm{s}$. Authority moves to the control law over about a minute rather than in one step, and only after the bias has washed out is the loop the pure system of Equation (11.10). In the other direction the three manual commands follow the controller output, so returning to manual operation moves nothing either. What they follow is the law's request before the clamp, which the clamp then limits exactly as it limited the automatic command, so in set-up 7 the stations read −20.23 and 19.15 mm/s^2 at 10 s while the thrusters still receive $\pm 4\,\mathrm{mm/s^2}$.

11.4.5 Three Models Kept Apart

The simulation model is Equation (11.1) with the true n from Equation (11.2) at the true altitude, integrated over six states. The analysis model is Equations (11.3) to (11.5) and (11.10) to (11.12), which this chapter uses to choose gains and to read the measurements and which appears nowhere in the app's code. What the proportional-derivative controller itself holds is the four algebraic terms of Equation (11.9) and the single number $\hat{n}$, and nothing further. That controller does not know μ, R_E, $u_{\max}$ or the docking criteria, and it never predicts forward. The predictive controller of Section 11.4.6 holds more, namely Equation (11.1) at $\hat{n}$ without the drag term, the declared μ and R_E, the limit $u_{\max}$ and the contact speed v_{cap}, and it predicts five minutes ahead. The instruments are the same for both, the six true relative states read without error.

11.4.6 Model Predictive Control on the Exact Model

The third architecture is a model predictive controller (MPC), an option that only set-up 11 selects. Every $T_s = 2\,\mathrm{s}$, once in four executions of the control law, it solves one quadratic program for the thrust over the next $N = 150$ steps and applies only its opening block, which it keeps until the following decision replans from fresh measurements. This is the receding-horizon law reviewed in [15] and applied to proximity operations in [6, 7]. It differs from the framework's MPC option of Section 1.6.3 in its objective, which prices propellant rather than tracking error. The first plan is made when the run is initialized, so the MPC flies from $t = 0$.

Its model is not fitted. Equation (11.1) is linear, so the plan predicts with the equations the plant integrates, without the drag term and at the $\hat{n}$ of Equation (11.9), which comes from $\hat{h}_{orb}$ on the declared R_E and μ and never from the true altitude. Under an exact zero-order hold it is Equation (11.14).

$$\xi_{k+1} = \Phi\xi_k + \Gamma\left(u_k^+ - u_k^-\right),\ \Phi = \mathrm{e}^{AT_s},\ \Gamma = \int_0^{T_s} \mathrm{e}^{A\tau}\mathrm{d}\tau B \tag{11.14}$$

where $\xi = (x, \dot{x}, y, \dot{y}, z, \dot{z})^T$, A and B are the matrices of Equation (11.1) at $\hat{n}$, and $u_k^{\pm} \geq 0$ are the three positive and three negative thrust accelerations held over step k. With $\hat{h}_{orb} = h_{orb}$ the free prediction

equals the closed-form Clohessy-Wiltshire solution to better than 10^{-7} m over the horizon, so there is no residual to report and the lesson is propellant rather than model error, except in set-up 8, where $\hat{n}$ is 15.8 % low for this controller as for the feedforward. Splitting each thrust into two nonnegative parts makes propellant linear, $T_s \sum_i (u_{i,k}^{+} + u_{i,k}^{-})$ per step, which is what the opposed thruster pairs spend. The six thrusts are held over blocks of 1, 2, 4, 8, 16, 32 and 87 steps, 42 decisions.

The 300 s horizon is 5.4 % of the orbit and a fifth to a sixth of an approach, so the rest of a transfer is charged as a terminal cost. Each transfer, begun at a new setpoint or when the MPC takes over outside its box, has an arrival time T_{arr}, and the plan's end state is charged with the propellant of the two-burn Clohessy-Wiltshire transfer from that state to rest on the setpoint at T_{arr}. For the coast τ left after the last judged step, 298 s ahead, both burns are fixed linear maps of the predicted position error $\tilde{r}$ and velocity $\tilde{v}$ there, as shown in Equation (11.15).

$$\begin{gathered} \Delta v_1 = -M\tilde{r} - \tilde{v}, \qquad \Delta v_2 = -\left(\Phi_{vr} - \Phi_{vv} M\right)\tilde{r}, \\ M = \Phi_{rv}^{-1}\Phi_{rr} \end{gathered} \tag{11.15}$$

where Φ_{rr}, Φ_{rv}, Φ_{vr} and Φ_{vv} are the position and velocity blocks of $e^{A\tau}$. The six components are soft rows pinned to zero, each slack $\varepsilon_j \geq |\Delta v_j|$ priced like propellant, and the objective is Equation (11.16).

$$J = \sum_{k=0}^{N-1} \sum_{i} \left[T_s \left(u_{i,k}^{+} + u_{i,k}^{-} \right) + R \left(u_{i,k}^{+2} + u_{i,k}^{-2} \right) \right] + \sum_{j} \left(\sigma_j \varepsilon_j + \rho_j \varepsilon_j^2 \right) \tag{11.16}$$

where $R = 10^{-6}$ per $(\mathrm{mm/s^2})^2$ keeps the program strictly convex, and the slack prices in m/s of propellant are $\sigma = 1$, $\rho = 10^{-4}$ on the terminal rows and $\sigma = 10$, $\rho = 100$ on the envelope and the box. No position error is weighted during a transfer. On 16 plans of set-ups 3, 7, 10 and 11 every terminal slack equals the hand-solved two-burn $|\Delta v|$ to 1.5×10^{-13} m/s, and the applied command is the program's first block exactly, so nothing outside the solve schedules a burn. Once less than 10 s of coast remains beyond the horizon, the terminal rows give way to a box of ± 0.5 m per axis from the arrival step, or a minute ahead if later, with the final velocity priced as the propellant to stop it, and after T_{arr} the plan holds that box with a centering weight of 10^{-3} per $\mathrm{m^2}$ s. Because an end state at rest on the target costs nothing, T_{arr} is a deadline, and the plan arrives early whenever the coast allows.

The thrusts are bounded hard, $0 \leq u_{i,k}^{\pm} \leq u_{\max}$, with $u_{\max}$ read as the proportional-derivative clamp reads it. With no corridor or keep-out zone in the app, the docking rule supplies a soft envelope on the closing speed of each axis, Equation (11.17).

$$\left| \dot{r}_{i,k} + \left(r_{i,k} - r_{\mathrm{SP},i} \right) / 200\,\mathrm{s} \right| \leq v_{cap}/2 + \varepsilon_{e,i} \tag{11.17}$$

where r_i is the position on axis i. Only the side the chaser occupies is applied, both sides within 1 m, so the closing speed allowed is half the contact speed plus 5 mm/s per meter still to go, judged every second step. Each plan has 42 thrusts, at most nine slacks and 84 bound rows, and is solved by the dual active-set method [16] with a ridge of $10^{-9}(1 + 2RN)$.

The arrival time was chosen once, on set-up 11 alone, from 5 to 40 min, as the least propellant among settings that dock no later than set-up 6 (1785.0 s), enter the 1 m box no later (2212.5 s) and reject no plan. The 25 and 26 min settings lie within 1 %, at 0.3552 and 0.3567 m/s, and the longer is shipped. The usual tracking form, a position weight at every step with a propellant price and a terminal weight from the discrete Riccati solution, was flown on the same start and solver at 12 settings (prices 3×10^5 to 10^7, terminal bandwidths 1, 3 and 10 mrad/s). Each spent more than 0.4191 m/s, docked after 1785 s, collided, or stalled 21 to 63 m off, because a linear propellant price outweighs a quadratic incentive that vanishes near the target.

Table 11.3 Set-Up List. Every Entry Not Named Under Change Keeps Its Default From Table 11.2. The First Four Run in Manual Operation With Every Thruster Command at Zero.

#	Change	Mechanism isolated
1	$x(0) = 100\,\mathrm{m}$	an unstable free-drift mode
2	$y(0) = -100\,\mathrm{m}$	an exact equilibrium of the plant
3	$y(0) = -250\,\mathrm{m}$, $\dot{y}(0) = 0.5\,\mathrm{m/s}$	Coriolis turns closing speed into altitude
4	$y(0) = -60\,\mathrm{m}$, $\dot{z}(0) = 0.10\,\mathrm{m/s}$	a bounded, uncoupled third axis
5	slow set, compensation off, $x(0) = 100\,\mathrm{m}$	diagonal feedback against a coupled plant
6	slow set, compensation on, $x(0) = 100\,\mathrm{m}$	exact cancellation of the same coupling
7	$u_{\max} = 4\,\mathrm{mm/s^2}$, $x(0) = 200\,\mathrm{m}$, $y(0) = -200\,\mathrm{m}$	bandwidth the actuator cannot support
8	$h_{orb} = 200\,\mathrm{km}$, $\hat{h}_{orb} = 1000\,\mathrm{km}$, $\omega_n = 2\,\mathrm{mrad/s}$	feedforward on a wrong parameter
9	$y_{\mathrm{SP}} = y(0) = -20\,\mathrm{m}$	station keeping and offset under load
10	along-track thruster at 30 %, $x(0) = 80\,\mathrm{m}$, $y(0) = -80\,\mathrm{m}$	feedforward through a lying actuator
11	slow set, MPC, $T_{arr} = 26\,\mathrm{min}$, $x(0) = 100\,\mathrm{m}$	a plan that pays only for propellant
12	graded, see Section 11.6.7	all of the above, from poor tuning

Offset-free predictive controllers add an integrating disturbance estimated by a Kalman filter [17], as about half of the framework's MPC options do (Section 1.6.3). This one departs from that default and starts each plan from the six measured states the proportional-derivative loops read. Such a filter leaves set-ups 3, 7 and 11 unchanged, but it reads set-up 10's degraded thruster as an outside force proportional to the command and chases it, spending 11.54, 1.92, 1.14 and 1.00 m/s at filter time constants of 106, 1004, 9992 and 150 246 s, against 0.99 m/s without one. Without it, drag of 10 and 20 µm/s^2 leaves a steady hold error of 0.024 and 0.031 m, against the proportional a_d/K_P of 0.100 and 0.200 m. None of the three architectures integrates.

Any plan whose program ends short of optimal, or whose thrust is undefined, infinite, beyond $\pm u_{\max}$ or a negative split, is rejected, and the compensated law of Equation (11.9) takes over without a bump, because under the MPC its bias b_i tracks the applied command at every execution. The record is marked, and the next accepted plan takes control back. A rejection forced at 602 s of set-up 11 moves the command by one execution of the compensated law to within 10^{-12} m/s^2, and no run in this chapter rejects a plan. Under the MPC the saturation counter counts the command at the limit. The program minimizes the pairs' propellant, $\int (|u_x| + |u_y| + |u_z|)\,\mathrm{d}t$, while Equation (11.6) counts the vector's length, so both are reported.

11.5 Scenario Definitions

Table 11.3 lists eleven teaching set-ups and one graded challenge. The app numbers them as its scenarios 1 to 11 and the challenge, each with the briefing and completion rule of Section 1.6.4. The first four run with every thruster off, so the plant is met before the controller is, and set-ups 5 and 6 differ by exactly one boolean. Set-up 11 is the start and tuning of set-up 6 with the model predictive controller selected, so set-ups 5, 6 and 11 fly one transfer under three architectures. Two tunings recur. The reference set is $\omega_n = 10\,\mathrm{mrad/s}$, $\zeta = 1.00$ and $\omega_{n,z} = 8\,\mathrm{mrad/s}$ with compensation on. The slow set used for the central comparison is $\omega_n = \omega_{n,z} = 3\,\mathrm{mrad/s}$ and $\zeta = 1.00$, which puts the loop only 2.65 times faster than the orbit and is where the coupling terms stop being a correction and become the dominant dynamics.

11.6 Results

All of the values below come from the app's own scorecard, tiles and exported samples (Section 1.6.4), taken with the physics step, the control period and the log rate set as in Table 11.2. The exported file,

Table 11.4 Free Drift With Every Thruster Off. The Closed-Form Column Is Evaluated From Section 11.3.3 at the Same Mean Motion.

#	Quantity	Simulated	Closed form
1	x at 1200 s (m)	336.536	336.536
1	y at 1200 s (m)	−228.163	−228.163
1	x at 2778 s (m)	700.000	700.000
1	y at 2778 s (m)	−1886.568	−1886.568
2	y at 3600 s (m)	−100.000	−100
2	Δv over 3600 s (m/s)	0.0000	0
3	closest approach (m)	115.22	115.22
3	time of closest approach (s)	363.5	363.6
3	x at closest approach (m)	73.7	73.75
3	x at 1200 s (m)	696.904	696.904
4	peak $\lvert z\rvert$ (m)	88.389	88.389
4	z at 5556 s (m)	0.238	0.238

at 1 Hz, carries the feedforward and feedback contributions separately on each axis, both propellant totals and the true and assumed mean motions, from which every tabulated result can be recomputed. An integrated measure depends on how long the run lasted, so every one is quoted beside its horizon. A setting that departs from a set-up is chosen before the run starts, and the app gives the same run whether it is chosen then or carried through a Reset. Docking times and closest approaches are read at the 0.5 s physics step, and a run continued past its docking pause is resumed unchanged. The saturation fraction is measured against the $u_{\max}$ in force on each axis, while overshoot and settling time are reported only after a setpoint move, so for these transfers from an initial offset the results quote instead the time from which the range to the port stays below 1 m. Table 11.4 collects the four free-drift set-ups and Table 11.5 the closed-loop ones.

11.6.1 Effect of the Compensation at Fixed Gains

Set-ups 5 and 6 differ by one boolean and are the controlled experiment the simulator exists to run. Both start 100 m above the target with the loops at $\omega_n = 3\,\mathrm{mrad/s}$ and $\zeta = 1.00$. The compensated run docks at 1785 s having spent 0.4191 m/s, and its along-track axis never moves more than 0.0074 m from the setpoint it started on. The uncompensated run with the same gains takes 3516 s to dock, which is 1.97 times as long, spends 0.5402 m/s over the full 6000 s against 0.4195 m/s for the compensated run on the same horizon, which is 28.8 % more, and drags the along-track axis 11.475 m away from a setpoint nothing asked it to leave. That last measurement carries the argument, and it is the trace shown in Figure 11.3.

> **TRY IT IN THE APP** **5 · PD Only, No Orbital Compensation**
>
> Run the scenario and watch the along-track trend, which starts on its setpoint. While the radial loop brings the chaser down, the along-track axis is pulled 11.475 m away. Load 6 · PD + orbital compensation and run again. The along-track error stays within 0.0074 m, and the chaser docks at 1785.0 s.

Two mechanisms produce it and Equation (11.11) names both. The gravity-gradient term removes 42.7 % of the radial position gain at this bandwidth, which is why the radial approach is slow, and the Coriolis term on the right-hand side drives the along-track axis with the radial rate the controller is deliberately generating. Neither is a disturbance in the usual sense, because both are consequences of the controller's own action on a plant whose known terms it is not using.

The uncompensated loop degrades gracefully into instability rather than abruptly. At $\omega_n = 2\,\mathrm{mrad/s}$ the net radial stiffness $K_P - 3n^2$ is $1.6 \times 10^{-7}\,\mathrm{/s^2}$, four percent of K_P, and the radial error

Table 11.5 Closed-Loop Outcomes. Each Error Integral Accumulates the Absolute Error of Its Own Loop Across the Horizon Named in the Third Column, Δv Is the Commanded Total of Equation (11.6), the Docking Time Is the Physics Step on Which the Range First Fell Inside $r_{cap} = 3\,\mathrm{m}$ at or Below $v_{cap} = 0.10\,\mathrm{m/s}$, and "Inside 1 m" Gives the Time From Which the Range Stayed Below 1 m. Row 5e Uses a Bandwidth Between the 0.5 mrad/s Settings of the App and Is a Model Run.

#	Set-up	Horizon	Δv	IAE		Peak	Docked	Note
				radial	along-track	$\lvert y - y_{SP}\rvert$	at	
		(s)	(m/s)	(m s)	(m s)	(m)	(s)	
5	slow set, compensation off	3000	0.5294	113888	21896	11.475	—	not inside 1 m
6	slow set, compensation on	3000	0.4191	66536	8.17	0.0074	1785.0	inside 1 m at 2213 s
5b	set-up 5 run to 6000 s	6000	0.5402	116506	25145	11.475	3516.0	inside 1 m at 4282 s
6b	set-up 6 run to 6000 s	6000	0.4195	66581	8.45	0.0074	1785.0	—
5c	compensation off, $\omega_n = 8$	3000	0.6547	26573	3537	6.080	789.0	—
6c	compensation on, $\omega_n = 8$	3000	0.6607	24924	3.2	0.0074	669.5	—
5d	compensation off, $\omega_n = 2$	6000	2.0917	542598	9310	1.979	—	radial still at 80.19 m
5e	compensation off, $\omega_n = 1.96$	6000	2.3017	599372	108	0.024	—	radial still at 99.78 m
6d	compensation on, $\omega_n = 2$	6000	0.4775	99897	12.6	0.0074	2677.0	—
7	$u_{\max} = 4\,\mathrm{mm/s^2}$, $\omega_n = 10$	2400	2.2313	61475	51357	200.0	766.0	233.5 s clipped
7b	the same at $\omega_n = 4.5$	2400	1.4082	89518	88729	200.0	1462.5	38.5 s clipped
7c	the same at $\omega_n = 4$	2400	1.3504	99799	99810	200.0	1643.0	never clipped
7d	the same at $\omega_n = 3$	2400	1.3022	132707	132723	200.0	2190.0	never clipped
7e	$u_{\max} = 50\,\mathrm{mm/s^2}$, $\omega_n = 10$	2400	2.3015	39848	39850	200.0	657.5	never clipped
8	$\hat{n}$ 15.8 % low, $\omega_n = 2$	3000	0.9130	143325	92568	100.0	—	radial still at 12.48 m
8b	the same with $\hat{n} = n$	3000	0.7444	98914	98931	100.0	2882.0	—
8c	$\hat{n}$ 15.8 % low, $\omega_n = 4$	3000	0.7188	56572	47713	100.0	1602.5	—
8d	the same with $\hat{n} = n$	3000	0.6995	49916	49923	100.0	1441.0	—
9	V-bar hold at –20 m	3600	0.0000	0.0	0.0	0.000	—	exact equilibrium
9b	the same with $a_d = 10\,\mu\mathrm{m/s^2}$	3600	0.0360	0.0	340	0.100	—	offset a_d/K_P
9c	the same with $a_d = 20\,\mu\mathrm{m/s^2}$	3600	0.0720	0.0	680	0.200	—	offset a_d/K_P
9d	the same hold at $x = +20\,\mathrm{m}$	3600	0.2765	0.0	0.0	0.000	—	radial offset, not free
10	along-track thruster at 30 %	1800	2.1882	15939	24554	80.0	986.5	delivered 0.8888 m/s
10b	the same with compensation off	1800	2.3175	19045	25415	80.0	1014.0	delivered 0.9125 m/s
10c	the same with a healthy thruster	1800	0.9206	15939	15940	80.0	550.5	—
12	graded, as shipped	1500	0.4646	91918	42159	60.0	—	radial grew to 87.86 m
12b	graded, retuned	1500	0.5750	7969	11756	60.0	494.5	inside 1 m at 624 s

is still 80.19 m after 6000 s with 2.0917 m/s already spent. A model run at 1.96 mrad/s, between the 0.5 mrad/s settings the app offers and just above $\sqrt{3}\,n = 1.9596\,\mathrm{mrad/s}$, has a net stiffness of

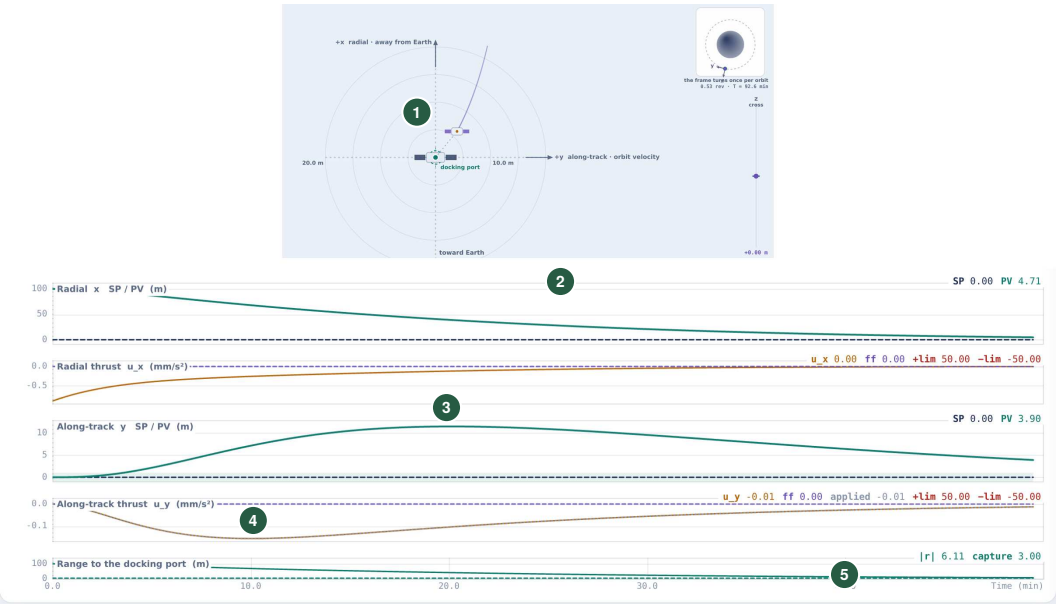


Figure 11.3 A 100 m radial transfer with the orbital compensation switched off. Set-up 5 is shown at $t = 2968$ s, read on the app's clock, which the crop leaves out, flown by three independent proportional-derivative loops at $\omega_n = 3$ mrad/s. The picture is a composite of two parts of one screenshot: the plant view (top, reduced) and the trend panel (bottom). (1) The chaser just above the docking port. (2) The radial measurement, fallen from 100 m to 4.71 m in fifty minutes. (3) The along-track measurement, the diagnostic. It began on its setpoint, nothing disturbed it, and the Coriolis term $-2n\dot{x}$ produced by the radial correction has carried it to 11.5 m and back to 3.90 m, for an along-track error integral of 21 773 m s on the scorecard, which the crop leaves out. (4) The along-track thrust strip. Each thrust strip is scaled to its own command, at most 0.90 mm/s^2 radially and 0.16 mm/s^2 along-track, and the ± 50 mm/s^2 limits lie far off both axes. (5) The range to the port, still 6.11 m.

1.6×10^{-9}/s^2 and leaves the radial error at 99.78 m, essentially where it started, after 2.3017 m/s of propellant spent holding it there. The compensated loop at the same $\omega_n = 2$ mrad/s docks at 2677.0 s for 0.4775 m/s.

11.6.2 Bandwidth Sweep and the Propellant Minimum

A fair comparison requires sweeping the bandwidth, because enough feedback authority masks any plant. Table 11.6 repeats the 100 m radial transfer at thirteen bandwidths with and without the compensation, each critically damped, over a 9000 s horizon. Figure 11.4 draws the same columns on logarithmic axes.

The compensated column has an interior minimum, and it is there for a reason worth deriving. With the cancellation exact, the critically damped error from rest is $x(t) = x_0(1 + \omega_n t)\mathrm{e}^{-\omega_n t}$, the feedback part of the command is $\ddot{x}$, and its absolute integral is $2\omega_n x_0/\mathrm{e} = 0.736\,\omega_n x_0$, which grows with bandwidth. The feedforward term $-3\hat{n}^2 x$ is a thrust command in its own right, and $\int_0^\infty x\,\mathrm{d}t = 2x_0/\omega_n$, so it contributes $6n^2 x_0/\omega_n$, which shrinks with bandwidth. Setting the two equal gives Equation (11.18).

$$\omega_n^\star \approx n\sqrt{3\mathrm{e}} = 2.856\,n = 3.23\,\text{mrad/s} \tag{11.18}$$

A model sweep in steps of 0.01 mrad/s puts the measured minimum at 3.14 mrad/s with 0.4189 m/s, within 3 % of the prediction. The bandwidth setting of the app moves in 0.5 mrad/s steps, and its

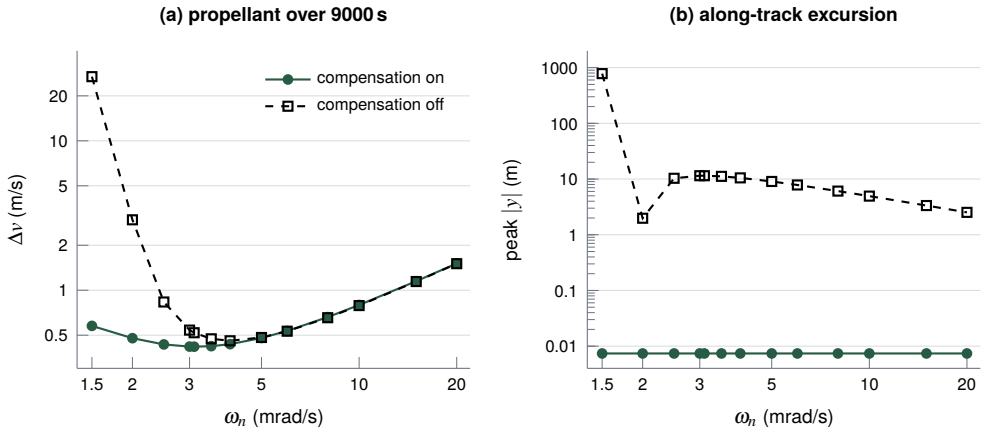


Figure 11.4 Table 11.6 drawn on logarithmic axes, the 100 m radial transfer at thirteen bandwidths with the compensation on (filled circles, solid) and off (open squares, dashed). The row at 3.1 mrad/s is the model run of the table. Notice in (a) that the two propellant curves merge above about 6 mrad/s, the compensated one lowest near 3 mrad/s and the other at 4 mrad/s, and in (b) that the gap in excursion never closes, from 0.0074 m against 2.523 m even at 20 mrad/s.

Table 11.6 The 100 m Radial Transfer at Thirteen Bandwidths, Each Critically Damped, Over a 9000 s Horizon. Peak $|y|$ Is the Largest Along-Track Excursion From a Setpoint the Vehicle Started On. A Dash in a Docked Column Means the Range Never Fell Inside r_{cap} Within the Horizon. The Row Marked † Lies Between the 0.5 mrad/s Settings of the App and Is a Model Run.

ω_n	compensation on			compensation off		
(mrad/s)	Δv (m/s)	docked (s)	peak $\|y\|$ (m)	Δv (m/s)	docked (s)	peak $\|y\|$ (m)
1.5	0.5772	3568.5	0.0074	26.8822	—	780.995
2.0	0.4775	2677.0	0.0074	2.9605	—	1.979
2.5	0.4337	2142.0	0.0074	0.8341	6239.5	10.319
3.0	0.4195	1785.0	0.0074	0.5409	3516.0	11.475
3.1†	**0.4189**	1727.5	0.0074	0.5193	3264.5	11.476
3.5	0.4223	1530.0	0.0074	0.4730	2569.5	11.175
4.0	0.4354	1339.0	0.0074	**0.4589**	2050.5	10.499
5.0	0.4788	1071.0	0.0074	0.4827	1468.5	9.050
6.0	0.5343	892.5	0.0074	0.5318	1143.0	7.828
8.0	0.6607	669.5	0.0074	0.6547	789.0	6.080
10.0	0.7962	535.5	0.0074	0.7896	601.0	4.941
15.0	1.1495	357.0	0.0074	1.1433	378.0	3.345
20.0	1.5110	268.0	0.0074	1.5055	277.0	2.523

cheapest setting is 3.0 mrad/s at 0.4195 m/s. The reference tuning of 10 mrad/s costs 0.7962 m/s, which is 90 % more propellant than the minimum for a docking at 535.5 s instead of 1727.5 s. Rendezvous is flown slowly because that trade is close to linear at the top of the range.

The uncompensated column carries three separate results. It has a floor that the compensated column does not, because Equation (11.11) is unstable whenever $K_P < 3n^2$, and no damping ratio repairs that. At 1.5 mrad/s the radial error grows to 3111.9 m in 9000 s with 26.8822 m/s spent, far outside the 2 km range at which the model is declared valid, while the compensated loop at the same bandwidth docks at 3568.5 s for 0.5772 m/s. Its own propellant minimum is then at 4.0 mrad/s rather than 3.1 mrad/s, displaced upward because the gravity gradient eats part of the nominal gain

and a larger ω_n is needed to achieve the same closed-loop speed. At its own best point it spends 0.4589 m/s against 0.4189 m/s, a penalty of only 9.6 %, and above 6 mrad/s the two propellant columns are indistinguishable to within 1 %. What never converges is the along-track excursion. Even at 20 mrad/s, eighteen times the orbital rate, the uncompensated loop still pulls the along-track axis 2.523 m off a setpoint it started on, against 0.0074 m compensated, a factor of 340. The value of canceling known coupling is not primarily propellant. It is the removal of the cross-axis interference, and that benefit does not diminish with bandwidth. The general rule is that raising loop gain attenuates a disturbance but cannot remove a structural coupling, because the coupling is inside the loop rather than outside it [1].

Pause. Try Exercise 11.3 before reading on.

11.6.3 Thruster Saturation

Set-up 7 starts 200 m above and 200 m behind with the thruster limit cut to 4 mm/s^2, about one twelfth of the default, and leaves the reference tuning in place. The initial demand on each axis is $\omega_n^2 e_0 = 20\,\mathrm{mm/s^2}$, five times what is available, so both in-plane axes pin at the clamp. The command is clipped for 233.5 s of the 2400 s run, and the saturation fraction reads 9.71 % on the radial loop and 5.46 % on the along-track loop. The vehicle still docks at 766.0 s, and the price is 2.2313 m/s of propellant against 2.3015 m/s for the same tuning with the full 50 mm/s^2 available. While saturated the loop is not a proportional-derivative controller at all. It is a constant-acceleration burn, and the poles the gains were chosen to place do not exist.

The remedy is a slower loop rather than a stronger one. Equation (11.13) gives $\sqrt{u_{\max}/e_0} = 4.47\,\mathrm{mrad/s}$, and the measurements bracket it closely. At 4.5 mrad/s the clamp is still active for 38.5 s, and at 4.0 mrad/s it never acts. Slower is also cheaper here, because 4.0 mrad/s spends 1.3504 m/s against 2.2313 m/s at 10 mrad/s, a saving of 39.5 % bought with 877.0 s of extra approach time. Pushing further down to 3.0 mrad/s saves only another 3.6 % while adding 547.0 s, which locates the knee of the trade at about 4 mrad/s for this maneuver.

11.6.4 Feedforward With an Incorrect Mean Motion

Set-up 8 puts the target in a 200 km orbit and loads the flight computer with an ephemeris for 1000 km, so $\hat{n}$ is 15.8 % low, and runs the loop deliberately slowly at $\omega_n = 2\,\mathrm{mrad/s}$. The residuals of Equation (11.12) hand back 30.6 % of the position gain and 9.4 % of the rate gain. Measured against the same case with $\hat{n} = n$, the propellant rises from 0.7444 m/s to 0.9130 m/s, a penalty of 22.6 %, and the chaser is still 12.48 m out after fifty minutes where the matched model has it docked at 2882 s.

The second half of the experiment is the one that generalizes. Leaving the model wrong and raising ω_n to 4 mrad/s drops the propellant penalty from 22.6 % to 2.7 %, from 0.7188 m/s against 0.6995 m/s, and both cases dock. A quantity that is wrong by a fixed fraction stops mattering once the loop closed around it is fast enough, so the question a feedforward model has to answer is not whether it is right but whether it is right enough at the bandwidth being paid for.

The error integrals give a cleaner diagnostic than the propellant does. Exact cancellation makes the two in-plane axes behave identically, and they start from identical 100 m errors in this set-up, so their error integrals should match. With $\hat{n} = n$ they read 98914 and 98 931 m s, equal to within 0.02 %. With $\hat{n}$ wrong they read 143325 and 92 568 m s, a difference of 55 %. Asymmetry between two axes that were set up symmetrically is the signature of residual coupling, and it appears in the reported metrics without any need to plot a trajectory.

11.6.5 Station Keeping on the Two In-Plane Axes

Set-up 9 parks the chaser 20 m behind the target on the velocity vector, on the setpoint, at rest. Over a full hour the propellant counter does not move from 0.0000 m/s and both error integrals stay at zero, because the state is an exact equilibrium of Equation (11.1) and a proportional-derivative controller sitting on an equilibrium commands nothing. Moving the radial setpoint to +20 m before the run

holds the chaser 20 m above that point instead, and it costs 0.2765 m/s per hour, because $3n^2x$ must be canceled continuously while the along-track offset still costs nothing. The 20 m offset and the tuning are the same in both cases, and the choice of axis is the whole difference between free and 6.6 m/s per day.

> **+ DID YOU KNOW? Free Behind the Target, Paid For Above It**
>
> From the simulator: holding 20 m behind the target on the velocity vector spends 0.0000 m/s of propellant in an hour, while holding 20 m above it with the same tuning spends 0.2765 m/s (rows 9 and 9d of Table 11.5). The gravity-gradient term must be canceled for as long as the hold lasts.

Adding differential drag converts the hold into the textbook demonstration of proportional offset. Under a constant along-track acceleration a_d the loop settles where $K_P e = a_d$, and at $a_d = 10\,\mu\text{m/s}^2$ against $K_P = 10^{-4}/\text{s}^2$ the measured offset is 0.100 m, matching a_d/K_P to the resolution of the export. Doubling the drag doubles the offset to 0.200 m. On a real vehicle that offset is either accepted, removed with an integral term, or removed by feeding the estimated drag forward, which is the same mechanism as the orbital compensation one level up.

11.6.6 Feedforward Through a Degraded Thruster

Set-up 10 removes 70 % of the along-track thrust without telling the controller. The commanded propellant total rises from 0.9206 m/s to 2.1882 m/s while only 0.8888 m/s is delivered, and the docking time stretches from 550.5 s to 986.5 s. The mechanism is the feedforward itself. The radial channel is healthy and generates $\dot{x}$ exactly as designed, and every part of that $\dot{x}$ demands $2\hat{n}\dot{x}$ of along-track thrust to cancel the Coriolis term it produces. Only 30 % of that arrives, and the remainder appears as an along-track error that the proportional-derivative loop must then chase with the same crippled thruster. The along-track error integral rises from 15940 to 24 554 m s, a 54 % degradation.

The radial axis is untouched, and that is worth stating precisely because it is easy to expect otherwise. Its error integral reads 15 939 m s in both the healthy and the faulted case, and its error stays inside the 1 m band from 638 s in both. The cancellation $-3\hat{n}^2x - 2\hat{n}\dot{y}$ is exact whatever value $\dot{y}$ takes, so the radial axis stays decoupled from an along-track channel that is failing. Turning the compensation off makes both axes worse rather than better, with the commanded total rising to 2.3175 m/s and the docking time to 1014.0 s, so a degraded actuator is not a reason to abandon the model term. Sweeping the fault from 0 to 95 % of thrust lost over a 3000 s horizon finds no crossover at all. The compensated loop is cheaper at every point, by 1.1 % of commanded propellant with a healthy thruster, 6.0 % at 70 % lost and 2.1 % at 95 % lost, where neither architecture reaches the port. Exact cancellation degrades gracefully with actuator authority, because the term it cancels is proportional to the state rather than to the command. It is a reason to detect the fault, because the mismatch between the commanded acceleration and the acceleration a navigation filter observes is the signal that identifies a failed thruster, and control allocation across a redundant set is what recovers from it. Neither is modeled here. Volume 20, Chapter 6 allocates the demands of its position loops over a set of thrusters and runs a case with one of them failed.

11.6.7 Graded Challenge

The graded set-up starts the chaser 40 m above and 60 m behind the port with 0.02 m/s of residual closing rate and a thruster limit of 25 mm/s^2, and hands over tuning that is wrong in three ways at once. The compensation is off, $\omega_n = 2\,\text{mrad/s}$ puts $K_P = 4 \times 10^{-6}/\text{s}^2$ within 4 % of $3n^2 = 3.84 \times 10^{-6}/\text{s}^2$, and $\zeta = 0.4$ would produce overshoot at a docking port if the loop reached one. At the delivered settings the closest approach is 60.43 m at 518.5 s, after which the range only increases, and the radial error grows from 40 m to 87.86 m by 1500 s with 0.4646 m/s spent. The run is judged at 1500 s. It passes when the chaser has docked inside 3 m below 0.10 m/s without a collision, both in-plane errors have stayed within ± 1 m over the last third of the run, from 990 s, and no more than 0.9 m/s of

propellant has been spent, with the model parameters and disturbances as given. Changing a model parameter or a disturbance voids the attempt.

One successful configuration makes three changes. Compensation is not required for a pass. PD only at 10 mrad/s and damping ratio 1.00 docks near 549 s with a worst scored error of 0.13 m. Turning compensation on removes the terms eating the radial stiffness and costs nothing. Setting $\omega_n = 10\,\mathrm{mrad/s}$ demands $\omega_n^2|e_0| = 7.21\,\mathrm{mm/s^2}$ at the start from $|e_0| = 72.1\,\mathrm{m}$, against a $25\,\mathrm{mm/s^2}$ limit, and the largest thrust magnitude measured over the whole run is $7.00\,\mathrm{mm/s^2}$. Setting $\zeta = 1.00$ removes the overshoot. The retuned run docks at 494.5 s for 0.5750 m/s, inside the budget, its range stays below 1 m from 624 s, and its worst in-plane error over the judged window is 0.032 m. That leaves 0.968 m inside the band and 0.325 m/s under the budget, and the same settings are the suggested tuning, with $\omega_{n,z} = 8\,\mathrm{mrad/s}$ added. The radial error integral falls from 91918 to 7969 m s, a factor of 11.5. The division of the work between the two changes is the result worth recording. The gains moved the response from divergent to a slower pass, and the model term reduced the remaining error within the judged window.

11.6.8 Free Drift in Manual Operation

The first four set-ups are flown by hand with every thruster at zero, and two of them contradict the expected behavior. Releasing the chaser 100 m directly above the target produces 40.7 m per minute of backward drift and a radial excursion to 700 m at half an orbit, from an initial condition in which nothing was moving. Releasing it 100 m behind the target produces nothing at all, to the resolution of the export, for as long as the run is left going. The third carries the Coriolis lesson of the chapter's title. A coast with 0.5 m/s of pure closing speed from 250 m behind, with the relative velocity vector pointing exactly at the docking port, reaches a closest approach of 115.22 m at 363.5 s and is 73.7 m above the port when it gets there, after which it climbs away. Aiming at a target is correct in a fixed frame and incorrect in this one.

11.6.9 Propellant Against Time Under Predictive Control

Table 11.7 sets the MPC beside both proportional-derivative architectures on five approaches, each from the same state over the same horizon, the loops at their set-up's tuning and every MPC row at the 26 min arrival of set-up 11 with no retuning. No plan was rejected, and Table 11.8 reduces each case to a saving and a delay.

On set-up 11, the start of set-up 6 at 3 mrad/s, the MPC docks at 1322.5 s at 0.045 m/s and uses 0.3567 m/s by 3000 s, against 1785.0 s and 0.4191 m/s for the compensated loop, 14.9 % less and 462.5 s sooner. It spends 47 % of its total in the first 10 s. Counted over the pairs it is 4 % above the two-burn transfer arriving at the same 1322.5 s, 0.401 against 0.387 m/s, and on the tile 26 % above, 0.357 against 0.284 m/s, because a diagonal burn costs the pairs more than its length. The price is the along-track axis, which starts on its setpoint and swings 20.85 m behind the target at $t = 626\,\mathrm{s}$, where the compensated loop holds it within 0.007 m, as shown in Figure 11.5. Released above the target, the chaser drifts backward as Equation (11.3) describes, and the planned transfer rides that drift, whereas the compensated loop pays for $3n^2x$ and the Coriolis terms at every step to fly a straight radial line.

The 26 min arrival was tuned on set-up 11 alone, and against loops at 10 mrad/s the trade turns over. On set-ups 7, 3 and 10 the MPC spends 33 to 55 % less than the compensated loop and docks 530 to 905 s later (33 to 57 % and 502 to 820 s against the loop alone), because it keeps its own schedule. On set-up 7 its $4\,\mathrm{mm/s^2}$ thrusters cannot brake the cheaper coast in time, so it passes 23.62 m beyond the target with an envelope slack of 0.116 m/s. In every set-up in which the MPC acts there is no collision, and every contact lies between 0.045 and 0.067 m/s. On set-up 8, the one mismatch case, the MPC docks at 1516.0 s on 0.5894 m/s while neither loop docks within 3000 s, because replanning every 2 s from the measured state absorbs the error the feedforward carries.

Moving the arrival time on set-up 11 traces the MPC's own trade. Over 3000 s, 10 min costs 0.645 m/s for a docking at 562.5 s, 20 min costs 0.439 at 1144.5 s and 40 min costs 0.382 at 2294.0 s, against 0.3567 m/s at 1322.5 s for the shipped 26 min. Set beside the compensated column of Ta-

Table 11.7 The Three Architectures on Five Approaches. Δv Is the Tile of Equation (11.6) at the End of the Horizon, Pairs Is the Propellant of the Opposed Thruster Pairs That the MPC Minimizes, Contact Is the Relative Speed at Docking, Past Port Is the Largest Distance Beyond the Setpoint on an Axis That Started More Than 1 m Off, and Excursion Is the Largest Error of an Axis That Started on Its Setpoint. Set-Ups 5 and 6 Supply the Proportional-Derivative Rows of Set-Up 11.

					Docked		Inside	Past		IAE		
#	Architecture	Horizon (s)	Δv (m/s)	Pairs (m/s)	at (s)	Contact (m/s)	1 m (s)	port (m)	Excursion (m)	radial (m s)	along-track (m s)	Plans
11	PD only (set-up 5)	3000	0.5294	0.6723	—	—	—	0.00	11.475	113888	21896	—
11	PD + comp. (set-up 6)	3000	0.4191	0.5103	1785.0	0.008	2212.5	0.00	0.007	66536	8	—
11	MPC	3000	**0.3567**	0.4014	1322.5	0.045	1367.0	0.00	20.847	44839	18908	1500
7	PD only	2400	2.2408	2.9446	863.5	0.021	1029.5	1.95	—	64809	48875	—
7	PD + comp.	2400	2.2313	3.0036	766.0	0.026	892.0	0.00	—	61475	51357	—
7	MPC	2400	**1.4901**	1.7837	1545.0	0.064	1585.5	23.62	—	82401	254033	1200
3	PD only	1800	1.5891	2.0137	706.0	0.020	880.0	0.13	12.935	5883	44939	—
3	PD + comp.	1800	1.5508	1.9581	621.5	0.026	747.0	0.00	0.010	5	44838	—
3	MPC	1800	**0.9134**	1.0648	1526.0	0.057	1586.5	0.00	62.858	61467	191652	900
8	PD only	3000	1.9458	2.0930	—	—	—	0.00	—	412324	128933	—
8	PD + comp.	3000	0.9130	1.1528	—	—	—	0.32	—	143325	92568	—
8	MPC	3000	**0.5894**	0.7281	1516.0	0.059	1580.5	14.08	—	39814	118641	1500
10	PD only	1800	2.3175	2.6513	1014.0	0.035	1601.0	18.36	—	19045	25415	—
10	PD + comp.	1800	2.1882	2.5354	986.5	0.035	1551.5	15.91	—	15939	24554	—
10	MPC	1800	**0.9875**	1.1434	1516.0	0.067	1572.0	14.95	—	32770	99709	900

Table 11.8 The Saving and the Delay of the MPC at the 26 min Arrival Time, Against Each Proportional-Derivative Architecture. A Negative Delay Is an Earlier Docking, and a Dash Means That Architecture Never Docked.

	MPC		against PD + comp.		against PD only	
#	Δv (m/s)	docked (s)	less Δv (%)	delay (s)	less Δv (%)	delay (s)
11	0.3567	1322.5	14.9	−462.5	32.6	—
7	1.4901	1545.0	33.2	779.0	33.5	681.5
3	0.9134	1526.0	41.1	904.5	42.5	820.0
10	0.9875	1516.0	54.9	529.5	57.4	502.0
8	0.5894	1516.0	35.4	—	69.7	—

ble 11.6, these points compare the two controllers at equal docking time, and each lies below that column. Industrial predictive controllers usually carry their economics as a linear objective in a steady-state target calculation [18], while this one prices propellant linearly inside the dynamic plan itself.

On every hold of Table 11.9 the MPC spends what the compensated loop spends, and under drag it holds 4.2 and 6.5 times closer, with no integral action in either. The graded challenge is unchanged, and the MPC fails it at the shipped arrival with either tuning, because it reads neither. At the 1500 s verdict it has not docked, its worst in-plane error from 990 s is 35.35 m and it has spent 0.217 m/s, since 26 min lies at 1560 s. With $T_{arr} = 10$ min, not shipped, it docks at 542.5 s on 0.424 m/s and passes.

Halving the physics step to 0.25 s leaves the Δv of every architecture on set-ups 3, 7, 8, 10 and 11

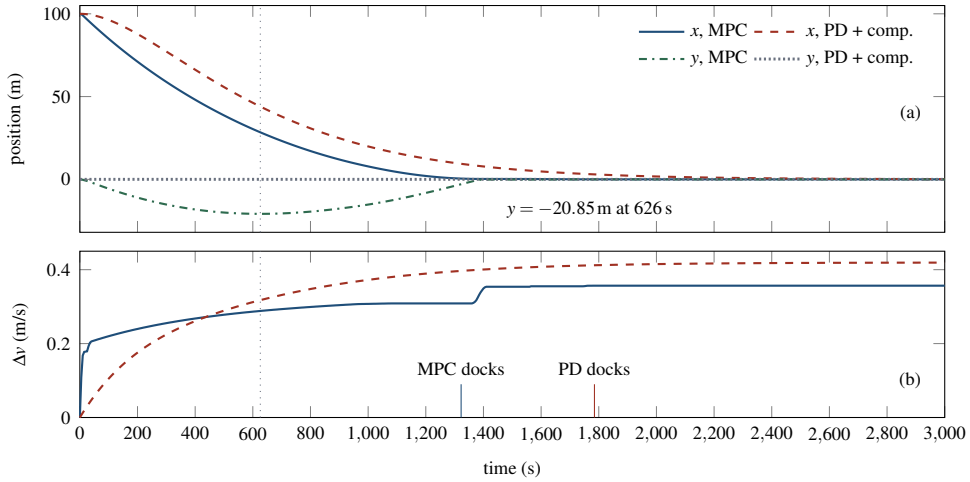


Figure 11.5 Set-up 11 under the MPC and the same start under PD with compensation, both at 3 mrad/s. (a) Radial position x under the MPC (solid) and PD (dashed), and along-track position y under the MPC (dash-dotted) and PD (heavy dotted). (b) Propellant used, the tile of Equation (11.6), under the MPC (solid) and PD (dashed). At 626 s (thin vertical dotted line) the MPC has spent 0.288 m/s against 0.318 m/s and its along-track axis is 20.85 m behind the target, where the compensated loop holds it within 0.007 m. The MPC docks at 1322.5 s and the loop at 1785.0 s, and by 3000 s they have used 0.3567 and 0.4191 m/s.

Table 11.9 Station Keeping Over 3600 s and the Graded Challenge at Its 1500 s Verdict. Each Hold Entry Gives Δv in m/s and the Largest In-Plane Error in m. The Challenge's Worst Error Is Judged From 990 s, and Its 10 min Arrival Is Not Shipped.

Hold	PD only	PD + comp.	MPC
9, at $y = -20$ m	0.0000, 0.000	0.0000, 0.000	0.0000, 0.000
9b, $a_d = 10\,\mu\text{m/s}^2$	0.0360, 0.100	0.0360, 0.100	0.0360, 0.024
9c, $a_d = 20\,\mu\text{m/s}^2$	0.0720, 0.200	0.0720, 0.200	0.0720, 0.031
9d, $x_{SP} = +20$ m	0.2869, 0.799	0.2765, 0.000	0.2765, 0.001

Graded challenge	Δv	docked (s)	worst (m)
delivered, PD, fails	0.4646	—	87.86
retuned, comp., meets	0.5750	494.5	0.03
MPC, 26 min, fails	0.2170	—	35.35
MPC, 10 min, meets	0.4243	542.5	0.00

unchanged to four decimals and every docking time within 0.2 s. In Chrome on the built app a plan takes 2.8 to 6.2 ms on average and at most 30.5 ms across set-ups 7, 9, 11 and the graded challenge, which is 0.14 to 0.31 % of a core at real time and 4.3 to 9.2 % at the top playback speed of 30 times, measured in two runs at load averages of 1.3 to 2.8 on the host. At that speed the page keeps a mean frame of 16.7 ms, with a 99th percentile of 16.8 ms, and the graded challenge takes 50.0 s of wall time under both the MPC and the loop alone.

11.7 Technology in Practice

Figure 11.6 puts the compensated loops of Equation (11.9) on flight hardware: relative navigation sensors feed a filter in the flight computer, the feedforward and feedback demands are summed there, and the sum is turned into firings of a set of reaction control thrusters.

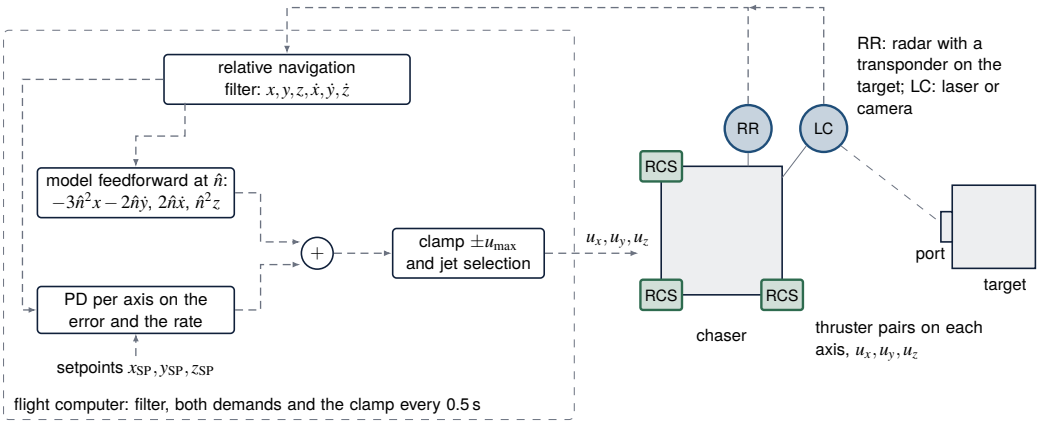


Figure 11.6 The compensated loops of Equation (11.9) on flight hardware. A radar working with a transponder on the target and a laser or camera sensor feed a relative navigation filter in the flight computer. The model feedforward and the PD demand are summed, clamped and turned into firings of thruster pairs on the chaser. The application replaces the sensors and the filter by the exact state and the jet selection by a clamp on each axis. Solid outlines are the vehicles and their devices, and dashed gray arrows are signals. Notice that the feedforward block reads the same state as the feedback block and one number, $\hat{n}$, and nothing else.

What Measures It

The controller of this chapter reads the six relative states exactly, which Section 11.3.5 names as the largest single idealization of the model. The orbiter's rendezvous radar reached the end of its tracking envelope at about 80 feet even on small payloads, and visual ranging with a sight or with camera tilt angles was imprecise, loaded the crew and gave no direct range rate [11]. Studies of docking with a space station later found the radar inadequate in minimum range and accuracy, which led to hand-held and payload-bay laser sensors [11]. In the model each reading is the true state, so none lags, drifts or jumps.

> **TECHNOLOGY Sensor: Rendezvous Radar With a Transponder**
>
> The Apollo lunar module measured range, range rate, line-of-sight angle and angle rate to the command module with a continuous-wave tracking radar that worked with a transponder on the other vehicle. Range came from the time delay of a three-tone phase modulation and range rate from the two-way Doppler shift. Its listed minimum range was 80 feet. In flight its range agreed with a separate radio ranging system to a mean bias of less than 0.04 percent of the median range. In flight tests the range tone tracking loop slipped cycles at low signal-to-noise ratio, and each slip moved the range reading by a multiple of 2400 feet [19], a step on a range trend that the vehicle never made.

What Moves It

The bounded acceleration u_{max} on each axis of Table 11.2 stands for a set of reaction control thrusters, which Section 11.3.5 describes as on and off devices with a minimum impulse bit. The orbiter carried primary jets of 870 pounds thrust and vernier jets of 25 pounds [11], and the plumes of the large jets shaped how it approached a target. A real vehicle commands opposed pairs of thrusters, which is what the propellant of the predictive controller counts in Section 11.4.6, while the per-axis clamp of the model delivers any command up to u_{max} exactly.

TECHNOLOGY Final Element: Reaction Control Thrusters

Thrusters give a spacecraft both translation and attitude control. A cold gas thruster has no combustion and expels a stored gas through a nozzle, which gives very small minimum impulse bits for precise maneuvering, while its low specific impulse rules out large maneuvers. The minimum impulse bit sets how precisely the vehicle translates and the thrust sets its authority. Depending on the thruster layout and the center of gravity, the thrusters also impart torques that attitude control must absorb. One recent small proximity mission lost thrust on several axes, likely after a controller failure, and others met propellant leaks [20]. On a trend a lost thruster shows as a command that no longer produces its acceleration, the gap of set-up 10.

Where the Controller Runs

On the orbiter, guidance, navigation and control software ran in two flight computers in parallel during rendezvous and proximity operations, and in three from the station flights onward. Clohessy-Wiltshire targeting was first preferred for proximity operations because it took about 5 seconds on the flight computer, against 30 to 60 seconds expected for Lambert targeting, which in practice later took less than 15 seconds [11]. The parts of the report read for this section give no execution period for the orbiter's digital autopilot. The application evaluates the law of Equation (11.9) every 0.5 s, and the four feedforward terms are added to the feedback demand before the clamp, so feedback and feedforward share one layer of the controller.

IN PRACTICE Braking Gates and the Approach Axes

Piloted approaches were flown against braking gates, a closing rate assigned to each range. A shuttle profile of the 1970s gave 45 feet per second at 9000 feet, falling in steps to 5 at 1000 feet and to zero at 100 feet, with coast periods long enough for longer burns caused by dispersions or failed jets. The pilot controlled range, range rate and line-of-sight angles directly, and final approaches ran along the velocity vector or the radius vector. An approach from below along the radius met what the source calls the natural braking effect of orbital mechanics [11], which is the term $3n^2x$ of Equation (11.1) pushing a vehicle below the target further down.

What Is Modernizing It

Three developments bear on these loops. The first is light detection and ranging, an emerging spacecraft sensor that matured in terrestrial use and already flies on larger spacecraft capable of proximity operations, with rendezvous and docking among its uses [20]. The second is camera-based navigation. A NASA demonstration launched in 2023 navigated relative to an object from angles alone, with no prior knowledge of it, and reached a relative position knowledge with 0.5 percent error relative to range [20]. Both address the navigation error that the exact state of this chapter hides. The third is autonomous guidance, navigation and control software for operations when ground communication is limited or unavailable [20]. The feedforward of this chapter is the least demanding part of such software.

11.8 Limitations

The simulator teaches a coupling structure and does not design an approach. The controller is handed a perfect state vector, so the relative navigation that dominates real terminal guidance is absent and every number in Table 11.5 is tuning-limited rather than sensor-limited. There is no attitude, so the loop that actually consumes most of a chaser's control authority and most of its engineering effort does not appear, and thrust is available along any axis at any instant. Continuous bounded acceleration replaces a set of on and off thrusters with a minimum impulse bit and a geometry that couples translation into rotation, which is a substantial softening of the terminal phase. The control structures are fixed, and the one constrained formulation among them, the predictive controller of Section 11.4.6, carries only the thrust bound and a soft closing-speed envelope, which puts the corridor and keep-out constraints of current proximity-operations work [6, 7] and the passively safe

trajectory families [5] outside anything a run can show, even though Equation (11.1) is precisely the prediction model each of them carries. The linearization itself is declared valid only inside 2 km, and it ignores the oblateness term that dominates relative motion over many orbits [14]. Disturbances are either constant or fired by the operator and never stochastic, and one fixed step against one fixed control period makes every repeat bit-identical, which no flight computer achieves. The individual failure modes of every modeling assumption are given in Section 11.3.5, and the expensive ones are the exact navigation, the absent attitude dynamics and the continuous thruster.

The predictive controller's model leaves out the drag, a degraded thruster, an error in $\hat{h}_{orb}$ and the impulsive events, which it meets only by replanning every 2 s without an estimator, and its one setting, the 26 min arrival, was chosen on one slow case, so a different value moves every delay in Table 11.8.

> ⚠ **OPEN CHALLENGE** **Proximity Operations on Small Spacecraft**
>
> A small spacecraft that is to approach another needs a suite of sensors for relative position and attitude, sometimes of an object it has no prior knowledge of, and much of its navigation must run without commands from the ground. Many recent demonstrations had trouble with their control systems and actuators, and one ran out of propellant before it could attempt its docking. The source expects the field to grow, with more demand for software that has proven flight heritage [20]. Progress would be such software flying routinely, with navigation error rather than tuning setting the docking precision, as Section 11.3.5 says it does.

11.9 Discussion

A spacecraft that cannot rendezvous is a spacecraft that must carry everything it will ever need. Every station resupply flight, every crew rotation and every proposal to refuel or repair a satellite already in orbit reduces to the same terminal problem of arriving slowly at a moving object with a propellant budget that cannot be topped up on the way. On-orbit servicing has been studied for decades and its enabling technologies are surveyed in [21], and the same approach and capture phase is the hard part of every proposal to remove a piece of large debris from a congested orbit [22]. What separates those missions from a station resupply is that the object being approached is uncooperative, tumbling and sometimes not designed to be grappled at all.

The simulator adds one statement to that picture, about where the difficulty does not lie. The terminal phase is not hard because the dynamics are unknown. Equation (11.1) is one of the best-conditioned models in engineering, its single parameter is known to many digits, and a controller that uses it collapses a coupled unstable plant to three independent second-order systems with four multiplications. The difficulty is in the navigation, the attitude, the discrete actuators and the consequences of an abort, none of which the model removes. Recognizing which part of a problem the model solves, and refusing to spend feedback authority on the part it already solved, is the transferable skill, and it is the same decomposition that appears as computed-torque control in robotics, as gravity and Coriolis compensation on gantry machines, and as static feedforward in process control [9]. In Volume 6, Chapter 2 one static feedforward from the pump command to the concentrate valve removes most of the interaction between the two loops of a reverse osmosis train, which is the same move on a process plant.

11.10 Conclusion

Relative motion near a target in orbit is written in axes that turn once per orbit, and the terms this produces make three thrusters something other than three independent actuators. The simulator of this chapter integrates the six Clohessy-Wiltshire states about a 400 km circular orbit at $n = 1.131\,\mathrm{mrad/s}$, with a bounded thruster acceleration per axis, a differential-drag disturbance, a de-ratable thruster and a geometric docking test, and closes one proportional-derivative loop per axis around them every 0.5 s.

Adding four algebraic terms evaluated at the controller's own mean motion cancels the plant's coupling identically and reduces each axis to a damped second-order system, verified against the analytic response to 0.148 m on a 100 m transfer. On the central comparison the compensated loop docks in 1785 s for 0.4191 m/s and holds its along-track axis within 0.0074 m, where the same gains without the model terms take 3516 s, spend 28.8 % more propellant and pull the along-track axis 11.475 m off a setpoint it started on.

A bandwidth sweep locates a propellant minimum at 3.14 mrad/s against a predicted $n\sqrt{3\mathrm{e}} = 3.23\,\mathrm{mrad/s}$, and shows that enough bandwidth buys away the propellant penalty but never the cross-axis excursion, which is still a factor of 340 apart at 20 mrad/s. Twelve set-ups separate the unstable radial mode, the free along-track equilibrium, a 115.22 m miss from a velocity vector pointed at the port, thruster saturation, a 15.8 % model error and a thruster delivering 30 % of its command, and each carries its own completion rule, with the graded challenge passed only by a docking that holds within 1 m on a 0.9 m/s budget. An added model predictive controller on the same exact model docks the start of set-up 6 462.5 s sooner on 14.9 % less propellant by letting the along-track axis swing 20.85 m, and against loops tuned at 10 mrad/s it saves 33 to 55 % while docking 530 to 905 s later. Throughout, a known coupling can be switched off and the cost of doing so read from the same measurements.

KEY POINTS

- The target's frame turns once per orbit, so three thrusters are a coupled plant: a radial offset drifts away, a hold behind is free.
- Four algebraic terms at the controller's own mean motion cancel the coupling and leave three damped second-order loops.
- With the same gains, the compensated loop docks in 1785 s for 0.4191 m/s, and the uncompensated one takes 3516 s and pulls the along-track axis 11.475 m off.
- More bandwidth buys away the propellant penalty but never the cross-axis excursion, still a factor of 340 apart at 20 mrad/s.
- The feedforward is itself thrust, so propellant is least near $\omega_n = n\sqrt{3\mathrm{e}}$.

Where else this applies. Volume 6, Chapter 2 removes loop interaction with one static feedforward, and Chapter 10 feeds gravity forward on another vehicle.

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] 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.

[3] G. W. Hill, "Researches in the lunar theory," *American Journal of Mathematics*, vol. 1, no. 1, p. 5, 1878.

[4] H. B. Hablani, M. L. Tapper, and D. J. Dana-Bashian, "Guidance and relative navigation for autonomous rendezvous in a circular orbit," *Journal of Guidance, Control, and Dynamics*, vol. 25, no. 3, pp. 553–562, 2002.

[5] L. Breger and J. P. How, "Safe trajectories for autonomous rendezvous of spacecraft," *Journal of Guidance, Control, and Dynamics*, vol. 31, no. 5, pp. 1478–1489, 2008.

[6] S. Di Cairano, H. Park, and I. Kolmanovsky, "Model predictive control approach for guidance of spacecraft rendezvous and proximity maneuvering," *International Journal of Robust and Nonlinear Control*, vol. 22, no. 12, pp. 1398–1427, 2012.

[7] A. Weiss, M. Baldwin, R. S. Erwin, and I. Kolmanovsky, "Model predictive control for spacecraft rendezvous and docking: Strategies for handling constraints and case studies," *IEEE Transactions on Control Systems Technology*, vol. 23, no. 4, pp. 1638–1647, 2015.

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

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

[10] J. L. Goodman, "History of space shuttle rendezvous and proximity operations," *Journal of Spacecraft and Rockets*, vol. 43, no. 5, pp. 944–959, 2006.

[11] ——, "History of space shuttle rendezvous," NASA Johnson Space Center, Mission Operations Directorate, Flight Dynamics Division, Houston, TX, Report JSC-63400, Revision 3, Oct. 2011, nASA Technical Reports Server 20110023479.

[12] W. Fehse, *Automated Rendezvous and Docking of Spacecraft*. Cambridge, UK: Cambridge University Press, 2003.

[13] D. C. Woffinden and D. K. Geller, "Navigating the road to autonomous orbital rendezvous," *Journal of Spacecraft and Rockets*, vol. 44, no. 4, pp. 898–909, 2007.

[14] S. A. Schweighart and R. J. Sedwick, "High-fidelity linearized J_2 model for satellite formation flight," *Journal of Guidance, Control, and Dynamics*, vol. 25, no. 6, pp. 1073–1080, 2002.

[15] 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.

[16] 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.

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

[18] 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.

[19] P. Rozas and A. R. Cunningham, "Apollo experience report: Lunar module landing radar and rendezvous radar," NASA Manned Spacecraft Center, Houston, TX, NASA Technical Note TN D-6849, 1972, nASA Technical Reports Server 19720016521.

[20] 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

[21] A. Flores-Abad, O. Ma, K. Pham, and S. Ulrich, "A review of space robotics technologies for on-orbit servicing," *Progress in Aerospace Sciences*, vol. 68, pp. 1–26, 2014.

[22] M. Shan, J. Guo, and E. Gill, "Review and comparison of active space debris capturing and removal methods," *Progress in Aerospace Sciences*, vol. 80, pp. 18–32, 2016.

11.11 Exercises

11.1 ○ Concept check · 1 min

In scenario 3 the chaser coasts from 250 m behind the target with 0.5 m/s of closing speed aimed exactly at the port, and it passes 73.7 m above the port. Which term of Equation (11.1) starts the rise?

(a) The Coriolis term $2n\dot{y}$ in the radial equation
(b) The gravity-gradient term $3n^2x$ in the radial equation
(c) The Coriolis term $-2n\dot{x}$ in the along-track equation
(d) The differential-drag term a_d in the along-track equation

11.2 ○ Concept check · 2 min

Which structure does the default controller of the rendezvous simulator use?

(a) A cascade with an inner thrust-acceleration loop on each axis
(b) A PD loop per axis plus a feedforward that cancels the gravity-gradient and Coriolis terms
(c) A PID loop per axis whose integral action removes the coupling as a steady load
(d) A steady-state decoupler computed from the plant's matrix of process gains

11.3 ○ Concept check · 2 min

In Table 11.6, above 6 mrad/s the propellant used with and without the compensation agrees within 1 %. At 20 mrad/s, what does the compensation still buy?

(a) A docking time about half as long
(b) A propellant saving that grows with the bandwidth
(c) An along-track excursion of 0.0074 m in place of 2.523 m
(d) Nothing that the scorecard or the exported file can show

11.4 Theory and calculation · 6 min

Use $\mu = 3.986\,004\,418 \times 10^{14}\,\mathrm{m^3/s^2}$ and $R_E = 6\,378\,137\,\mathrm{m}$.

(a) Compute the mean motion n and the period T of a 400 km circular orbit from Equation (11.2). *Checkpoint:* $n = 1.131\,367\,\mathrm{mrad/s}$, $T = 5553.6\,\mathrm{s} = 92.56\,\mathrm{min}$.

(b) A chaser is released at rest 100 m above the target. From Equation (11.3), give the slope of the secular along-track drift in m/min, and x and y at half a period. *Checkpoint:* 40.7 m/min backward. $x = 700\,\mathrm{m}$, $y = -6\pi x_0 = -1885\,\mathrm{m}$.

(c) Give the amplitude of the radial arc for a pure closing speed of 0.5 m/s, and of the cross-track swing for $\dot{z}_0 = 0.10\,\mathrm{m/s}$. *Checkpoint:* 883.9 m and 88.4 m.

11.5 Theory and calculation · 6 min

Use Equations (11.11) and (11.12) with $K_P = \omega_n^2$, $K_D = 2\zeta\omega_n$ and $\zeta = 1$.

(a) At 400 km, what fraction of the radial position gain does the gravity gradient remove from the uncompensated loop at $\omega_n = 3$ and 10 mrad/s, and at what ω_n is the radial stiffness gone? *Checkpoint:* 42.7 %, 3.84 %, 1.9596 mrad/s.

(b) Scenario 8 flies a 200 km orbit with a model for 1000 km. Compute n, $\hat{n}$, the residual $3(n^2 - \hat{n}^2)$ and the residual $2(n - \hat{n})$. *Checkpoint:* $n = 1.1834\,\mathrm{mrad/s}$, $\hat{n} = 0.9962\,\mathrm{mrad/s}$, $1.224 \times 10^{-6}\,\mathrm{/s^2}$, $3.743 \times 10^{-4}\,\mathrm{/s}$.

(c) Express the two residuals as fractions of K_P and K_D at $\omega_n = 2$ and 4 mrad/s, and say why doubling the bandwidth shrinks one fraction faster than the other. *Checkpoint:* 30.6 % and 9.4 % at 2, 7.6 % and 4.7 % at 4.

11.6 Theory and calculation · 7 min

With the cancellation exact and $\zeta = 1$, the radial error from rest at x_0 is $x(t) = x_0(1 + \omega_n t)\,\mathrm{e}^{-\omega_n t}$ and the feedback command is $a_x = \ddot{x}$. Take $x_0 = 100\,\mathrm{m}$ at 400 km.

(a) Show that $\int_0^\infty |\ddot{x}|\,\mathrm{d}t = 2\omega_n x_0/\mathrm{e}$ and $\int_0^\infty x\,\mathrm{d}t = 2x_0/\omega_n$, so that the feedforward $-3n^2x$ spends $6n^2x_0/\omega_n$. *Checkpoint:* $0.736\,\omega_n x_0$ and $2x_0/\omega_n$.

(b) Minimize the sum of the two parts over ω_n and evaluate the minimum and the sum there. *Checkpoint:* $\omega_n^\star = n\sqrt{3\mathrm{e}} = 3.23\,\mathrm{mrad/s}$, about 0.475 m/s.

(c) Table 11.6 gives 0.4189 m/s near this bandwidth. Why is the estimate higher, and why is its location still close? *Checkpoint:* The estimate is about 13 % high.

11.7 App challenge · 8 min

Use the default seed and leave every other setting as the scenario loads it.

(a) Load scenario 5, PD only, no orbital compensation, and run to 3000 s. Read the Δv used tile, the radial and along-track IAE on the scorecard, and the largest $|$`along_track_error_m`$|$ in the exported file. *Checkpoint:* 0.529 m/s, 113888 and 21 896 m s, 11.475 m, not docked.

(b) Load scenario 6, PD + orbital compensation, and run to 3000 s, pressing Run again after the pause at docking. Read the same four values and the first row of the `docked` column that reads 1. *Checkpoint:* 0.419 m/s, 66536 and 8.17 m s, 0.0074 m, docked in the 1785 s row.

(c) Explain where the along-track error of (a) comes from, and why the radial approach of (a) is slower.

Integrals within 1 %.

11.8 App challenge · 8 min

Load scenario 7, Limited thruster authority, and use the default seed. The thrust limit is 4 mm/s^2, the start 200 m above and 200 m behind. Each run lasts 2400 s. Press Run again after the pause at docking.

(a) Change nothing and run. Read the Δv used tile, the first exported row with `docked` at 1, and the time at a limit for the radial and along-track loops on the scorecard. *Checkpoint:* 2.231 m/s, 766 s, 9.71 % and 5.46 %.

(b) Compute the bound of Equation (11.13). Press Reset, set In-plane bandwidth to 4.5 mrad/s before Run and repeat. Then Reset, set 4.0 mrad/s and repeat. *Checkpoint:* 4.47 mrad/s. 1.408 m/s, 1463 s, 1.58 % and 0.04 %. 1.350 m/s, 1643 s, no time at the limit.

(c) What fraction of the propellant does the move from 10 to 4 mrad/s save, and what does it cost in time? *Checkpoint:* 39.5 %, 877 s later.

11.9 App challenge · 9 min

Load the challenge, Dock on a Δv budget, and use the default seed. Each run lasts 1500 s. Start every part from the reload button beside the scenario list, which restores the shipped settings, and make the stated settings before Run.

(a) Change nothing. Report the verdict line. *Checkpoint:* Not docked, worst in-plane error 87.86 m, Δv 0.465 m/s.

(b) Set only Control architecture to PD + orbital comp. *Checkpoint:* Docked at 998 s, worst 14.96 m, 0.376 m/s, fails.

(c) Keep PD only and set In-plane bandwidth to 10 mrad/s and Damping ratio to 1.00. *Checkpoint:* Docked at 549 s, worst 0.13 m, 0.583 m/s, meets the target.

(d) Make all three changes. *Checkpoint:* Docked at 495 s, worst 0.03 m, 0.575 m/s, meets the target.

Which changes are needed to pass, and what does the compensation add once they are made?

11.10 Application area · 3 min

Set-up 9 measures that holding 20 m above the target costs 0.2765 m/s per hour, while holding 20 m behind it costs nothing.

(a) Show that the hourly cost is $3n^2x \times 3600\,\mathrm{s}$ and give the cost per day. *Checkpoint:* 6.64 m/s per day.

(b) A 10 000 kg chaser burns propellant at an effective exhaust velocity of 2157 m/s (a specific impulse of 220 s). How much propellant does a day at the radial hold cost? *Checkpoint:* About 31 kg.

11.11 Application area · 3 min

Take a differential drag of 1 µm/s^2, the order the app's briefing for scenario 9 gives for two dissimilar vehicles at 400 km and a tenth of the 10 µm/s^2 of set-up 9b.

(a) What along-track offset does the PD loop hold at the reference tuning, 10 mrad/s, and at the slow tuning, 3 mrad/s? *Checkpoint:* 0.010 m and 0.111 m.

(b) What does holding against it cost per orbit and per hour? *Checkpoint:* 5.6 mm/s per orbit, 3.6 mm/s per hour.

Name the three remedies for the offset and say which one an operator would choose at a 20 m hold.

11.12 Review the agent · 5 min

An engineer asks an AI assistant whether the orbital feedforward is worth keeping. It replies: "Table 11.6 shows that above 6 mrad/s the two architectures spend the same propellant within 1 %. Drop the feedforward, tune PD only at 10 mrad/s, and the chaser docks at 601 s instead of 535.5 s for the same fuel, with no exposure to an error in the mean motion." Find the flaw, and say how you would check it in the app.