INTERACTIVE CONTROL ENGINEERING
Principles and Applications

Volume 3 Feedforward, Cascade, and Ratio Control

CHAPTER 10

# Commanding an Attitude Instead of a Force in a Rocket Landing Simulator

Pages 225–254

Edited by John D. Hedengren

**How to cite this chapter.** "Commanding an Attitude Instead of a Force in a Rocket Landing 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. 10, pp. 225–254.

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

**The app for this chapter**
https://apmonitor.com/apps/rocket-landing/

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

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

CHAPTER 10

# Commanding an Attitude Instead of a Force in a Rocket Landing Simulator




**The controller moves** 1 engine tilt on its gimbal 2 engine thrust (throttle)

**The controller reads** 1 position, descent speed and tilt of the vehicle

**What disturbs it** 1 crosswind and gusts

*The only engine sits below the vehicle, so tilting it to steer toward the pad first pushes the vehicle the wrong way while the crosswind leans it downwind. Illustration, not to scale.*

**IN PLAIN TERMS** **Steering a Rocket Down by Tilting It**

**The main idea.** A rocket landing on one engine cannot push itself sideways. To reach the pad it must first lean over. A second, faster automatic controller makes it lean, and only then does the rocket start to slide.

**The challenge.** The engine hangs below the middle of the rocket. When it swings to start the lean, it first shoves the rocket the wrong way, and no setting can stop that. So the part that steers toward the pad must stay several times slower than the part that holds the lean.

**Why it matters.** A reusable rocket must land gently on a small pad, and these controllers decide if its planned path can be flown at all.

**THREE WAYS TO READ THIS CHAPTER**

| | | |
|---|---|---|
| **Explore** | 12 min | the first page, the history of the throttled descent engine (Section 10.2), the fact on the misled throttle (Section 10.3), the fact on the rate term and the two try-it segments (Section 10.7) and the key points |
| **Study** | 60 min | the whole chapter with its three pause points and the twelve exercises at its end |
| **Practice** | 20 min | the first page, the results (Section 10.7), Technology in Practice (Section 10.9), Limitations (Section 10.10) and the open challenge |

**The control question** A lander on one gimballed engine cannot push itself sideways, so its position loop commands a tilt and borrows its authority from an attitude loop. How fast may that outer loop be, when the engine hung below the center of mass puts a right-half-plane zero in its path that no attitude gain can move?

**Learning objectives**

- Derive the attitude plant gain, the right-half-plane zero and the closed-loop frequencies of all three loops from the vehicle's constants.
- Explain the wrong-way first motion of a thrust-vectored vehicle and compute its timing from the linearized step response.
- Trace a steady crosswind back through the cascade to the standing pad offset of two proportional-derivative loops.
- Diagnose a slow actuator, a saturated gimbal and an engine short of deceleration authority, and say which of them retuning can repair.
- Repair a misbehaving cascade from the inside out.

**Prerequisites** Sections 1.4.5 (inner-loop speed rule), 1.4.3 (static feedforward), 1.6.2 (velocity form) and 1.6.4 (scenarios and the challenge). Chapter 9 for a cascade against a right-half-plane zero.

**The app** https://apmonitor.com/apps/rocket-landing/

**See also** Volume 7, Chapter 5, Volume 7, Chapter 7 and Volume 7, Chapter 9 for position loops that act through an angle, and Volume 20, Chapter 3 and Volume 20, Chapter 7 for tilting thrust with a wind estimate and for attitude control in orbit.

## 10.1 Introduction

The manipulated variable of a teaching loop is normally a force or a flow that acts on the controlled variable directly. A valve admits steam and the temperature moves, a pump adds liquid and the level moves, a motor draws current and the shaft turns. What the designer chooses is how much authority to apply and how quickly to apply it. A rocket coming down on one engine has nothing of that kind available for horizontal position. The only lateral force the vehicle can generate is the horizontal component of its own thrust, so the position loop must first ask the airframe to point somewhere else and then wait for the attitude loop to deliver the tilt. The command it issues is an angle, and that angle produces no force until a second and faster loop has delivered it.

Two consequences follow and both are geometric rather than a matter of tuning. The first is that the position loop inherits the attitude loop's speed as a ceiling, because an outer loop that revises its command before the inner loop has executed the previous one is not a cascade at all. The second is less apparent. The engine hangs below the center of mass, so the deflection that pitches the nose to the right pushes the tail to the right as well, and the vehicle accelerates the wrong way for as long as the gimbal is deflected. That transient puts a real zero in the right half plane, at a frequency set by the thrust, the pitch inertia and the distance from the center of mass to the gimbal, and no amount of inner-loop gain moves it.

Powered descent is most often presented as a trajectory problem. Optimal thrust programming for a lunar soft landing was posed in that form early on [1], the Apollo lunar module flew a polynomial guidance law that computed a commanded thrust vector for an attitude autopilot to track [2], and the modern treatment casts the same problem as a second-order cone program that solves to global optimality in bounded time [3–6]. However, a guidance law produces a desired acceleration and nothing else, and the loops that turn a desired acceleration into a gimbal angle are

ordinary proportional, integral and derivative terms arranged in a cascade, the same successive loop closure used on fixed-wing and rotary-wing vehicles [7]. Those loops decide whether the trajectory is flyable, and their limits come from the airframe.

The scenarios of the application cover cascade structure and successive loop closure, bandwidth separation between an outer and an inner loop, non-minimum-phase response and the ceiling a right-half-plane zero places on achievable bandwidth, actuator rate limits and saturation inside a feedback path, feedforward against a known nonlinearity, the offset a proportional loop leaves against a constant load, and a plant that drifts during a single run [8, 9]. The physical fundamentals behind the model are Newton's second law in two translational and one rotational degree of freedom, a moment balance taken about the center of mass with the thrust acting through a lever arm below it, and the constant exhaust-speed form of the variable-mass momentum balance [10, 11]. The behavior matters industrially because a reusable booster is graded on its last thirty seconds, and the loops that track a descent trajectory decide whether the trajectory can be flown at all.

This chapter derives, from the geometry of a thrust-vectored lander, the single frequency that bounds all three loops of an interactive planar landing simulation. The right-half-plane zero of the gimbal-to-position path and the natural frequency of the attitude loop are the same quantity $\sqrt{\ell T/I}$ apart from the square root of the attitude gain, so the bandwidth separation a cascade needs and the margin a non-minimum-phase plant allows are two readings of one number. Rigid-body flight dynamics and cascade design are both established methods, and the chapter applies them to one vehicle. Section 10.2 reviews thrust vector control and powered-descent guidance. Section 10.3 writes out a nine-state planar vehicle with its gimbal, engine, wind and touchdown models, with every equation, constant and initial condition needed to rebuild the simulator. Section 10.4 linearizes it about hover and places the three loops, and Section 10.5 documents the implemented loops and the optional predictive controller. Section 10.6 lists the scenarios, which break the cascade one mechanism at a time with a slow actuator, a saturated gimbal, a constant load and a shrinking mass, and Section 10.7 reports what they measure. Section 10.8 records defects of the application's own measurements that were found and corrected. Section 10.9 describes the real sensors, actuators and flight computer of such a vehicle. Sections 10.10, 10.11 and 10.12 state the limitations, place the results in practice and summarize them, and the exercises follow in Section 10.13.

## 10.2 Background

A rocket leaving a pad has no aerodynamic authority, because control surfaces need dynamic pressure and there is none at liftoff. The earliest large vehicles answered this with graphite vanes in the exhaust, which steered the vehicle at the cost of thrust, and the arrangement that replaced them mounts the whole combustion chamber on a pair of bearings and points it. Every large launch vehicle since has been steered that way, and the fact that the deflection acts through a lever arm below the center of mass is the reason the response begins in the wrong direction. The geometry, the moment it produces and the attitude dynamics it drives are standard material in space vehicle dynamics [10].

Landing under power is the same actuator run in reverse. The optimal thrust program for a soft lunar landing was set out as a minimum-propellant control problem in 1964, with the answer that the thrust magnitude should be at one bound or the other rather than in between [1]. The lunar module flew something more practical, a guidance law that shaped the acceleration as a polynomial in time-to-go and handed the resulting thrust direction to a digital autopilot, with the throttle and the attitude treated as separate channels [2]. That split between an outer law that owns the goal and an inner autopilot that owns the vehicle survives in every propulsive lander built since.

> ⧖ HISTORY **Deep Throttling for a Lunar Descent, 1963 to 1969**
>
> The descent stage of the Apollo lunar module had to carry the vehicle from lunar orbit to the surface and then hover over the last 1000 feet so that the crew could choose the exact landing site. That variable-thrust profile called for an engine that could be throttled, and the engine was also gimbaled to provide thrust vector control. Its design was first defined in 1963. Two throttling concepts were then developed in parallel for about 18 months: a fixed-area injector with helium injection, and a variable-area injector linked to cavitating flow-control valves. The variable-area design was chosen in January 1965, because the other suffered combustion instability and gave lower performance at reduced thrust. Development and qualification for the first lunar landing ran from August 1963 to April 1969 [12].

Mars entry, descent and landing renewed the interest and changed the mathematics. The powered descent problem was shown to admit a convex relaxation, which turns a nonconvex thrust-magnitude constraint into a second-order cone program with a proof that the relaxation is tight [3], and the argument was extended to pointing constraints and to the general soft-landing problem under the name lossless convexification [4]. Related work established the structure of the propellant-optimal solution directly [5], and the whole family is surveyed alongside the successive-convexification methods that handle six degrees of freedom [6]. The practical result is that a booster can compute a fresh optimal descent trajectory on board, several times a second, on the way down.

None of that removes the inner loops. A converged trajectory is a sequence of commanded thrust vectors, and a vehicle tracks a thrust vector by holding an attitude, which is a rate-limited actuator inside a feedback path. Rate limiting in a flight control loop is not a small correction, because it makes the effective loop gain depend on the size of the correction and it is a documented mechanism behind pilot-induced oscillation in piloted aircraft [13]. Successive loop closure, in which each loop is designed against a plant that already contains the closed inner loop, is the standard way to keep those interactions manageable, and its central requirement is a bandwidth separation between the levels [7]. Section 1.4.5 states the same requirement for a process cascade as an inner-loop speed rule. The application of this chapter is arranged so that the requirement can be violated deliberately and the consequence measured.

## 10.3 Vehicle Model

A 1200 kg single-stage lander, 8 m tall, descends in a vertical plane under one throttleable engine mounted on a gimbal 3 m below its center of mass. Horizontal position $x$ and altitude $z$ are measured in a pad-fixed frame with $x$ positive to the right and $z$ positive upward, and $z$ is the height of the landing feet so that touchdown occurs exactly at $z = 0$. The tilt $\theta$ is measured from vertical toward $+x$, so the body axis is the unit vector $(\sin\theta, \cos\theta)$ and a positive tilt buys positive horizontal acceleration. The geometry and every symbol the equations refer to are drawn in Figure 10.1.

### 10.3.1 Variables and Signals

Three controlled variables are regulated by three loops, and only two of them have a setpoint an operator can type. Table 10.1 sorts every signal into its role with the value it holds at the start of the default vertical descent.

Two of the three loops are cascaded. The horizontal position loop writes $\theta_{\mathrm{SP}}$, the attitude loop reads it and writes $\delta_{\mathrm{cmd}}$, and the gimbal is the only actuator either of them ever touches. The descent-speed loop is separate and writes $T_{\mathrm{cmd}}$, and it receives its setpoint from a guidance profile rather than from an operator, because a landing is graded on the speed at contact and not on the altitude history. Figure 10.2 draws the three loops with the limits and the actuator dynamics in the positions where they are applied.

🕮 **Pause.** Try Exercise 10.1 before reading on.

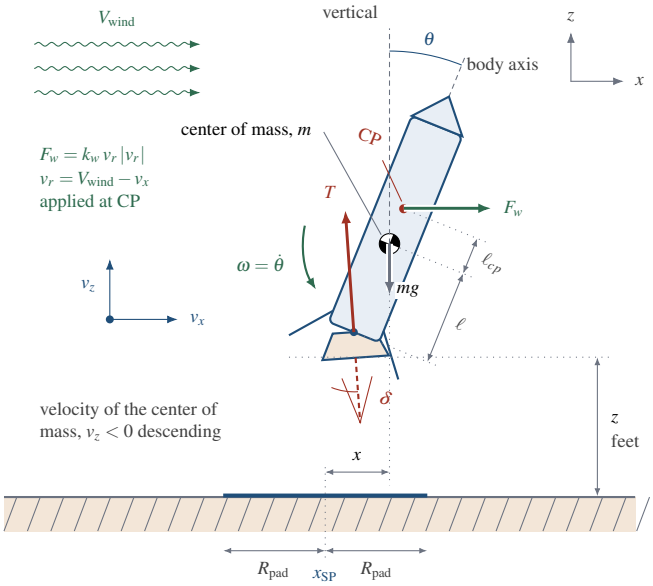



**Figure 10.1** Planar landing vehicle with one gimballed engine. The body axis leans by the tilt $\theta$ from vertical toward $+x$, and both angles are drawn exaggerated against the 12° tilt limit and the 10° gimbal travel. Thrust $T$ acts at the gimbal pivot a distance $\ell$ below the center of mass, so a deflection $\delta$ produces the moment $\ell T \sin\delta$ while the force itself points along $\theta - \delta$. The crosswind $V_{\text{wind}}$ produces the side force $F_w$ at a center of pressure $\ell_{cp}$ above the center of mass. Altitude $z$ is measured to the landing feet and horizontal position $x$ to the pad center $x_{\text{SP}}$, with the arrival graded against the pad radius $R_{\text{pad}}$.

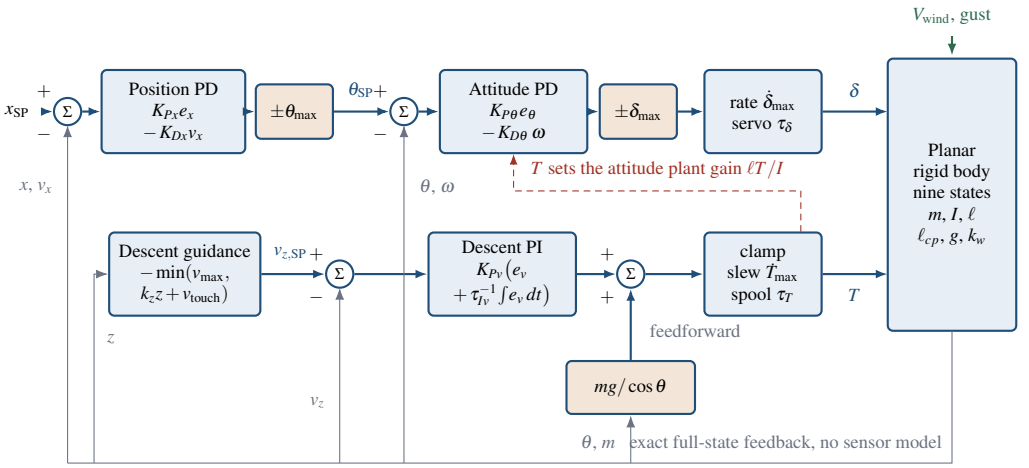



**Figure 10.2** Three loops around one gimballed engine. The position loop has no actuator of its own, so its output is the attitude setpoint $\theta_{\text{SP}}$ that the attitude loop must deliver before any horizontal force exists. Both outer stages are proportional-derivative and memoryless, and only the descent-speed loop carries an integrator, biased by the tilt-corrected hover thrust $mg/\cos\theta$. The dashed path is the coupling the cascade does not remove, because the same thrust that holds altitude also sets the gain of the attitude plant.

**Table 10.1** Classification of Every Signal. The Nominal Column Holds the Value at Ignition in the Vertical Descent Scenario.

| Symbol | Role | Nominal | Range |
|---|---|---|---|
| $x$ | state, horizontal position | 0 m | pad $\pm 25$ m |
| $z$ | state, altitude of the feet | 250 m | 30 to 600 m |
| $v_x$ | state, horizontal velocity | 0 m/s | — |
| $v_z$ | state, vertical velocity | −10 m/s | — |
| $\theta$ | CV, tilt from vertical | 0° | — |
| $\omega$ | state, pitch rate | 0 °/s | — |
| $\delta$ | MV, gimbal angle (servo state) | 0° | $\pm\delta_{\max}$ |
| $T$ | MV, engine thrust (spool state) | 11 772 N | $T_{\min}$ to $T_{\max}$ |
| $m$ | state, vehicle mass | 1200 kg | $m_{\text{dry}}$ to $m_0$ |
| $x_{\text{SP}}$ | setpoint, pad position | 0 m | −25 to 25 m |
| $\theta_{\text{SP}}$ | MV of the position loop | 0° | $\pm\theta_{\max}$ |
| $v_{z,\text{SP}}$ | guidance output, not typed | −10 m/s | — |
| $\delta_{\text{cmd}}$ | MV of the attitude loop | 0° | $\pm\delta_{\max}$ |
| $T_{\text{cmd}}$ | MV of the descent loop | 11 772 N | $T_{\min}$ to $T_{\max}$ |
| $V_{\text{wind}}$ | DV, steady crosswind | 0 m/s | $\pm 15$ m/s |
| $v_{\text{gust}}$ | DV, gust amplitude | 12 m/s | 2 to 25 m/s |
| $F_w$ | output, side force | 0 N | — |
| $T\cos(\theta-\delta)$ | output, vertical thrust | 11 772 N | — |
| $mg/\cos\theta$ | output, hover thrust | 11 772 N | — |

### 10.3.2 Translation and Rotation of the Planar Body

The vehicle is a rigid body with three degrees of freedom, two of translation and one of rotation. The thrust vector points along $\theta - \delta$ rather than $\theta + \delta$, for a reason derived in Section 10.3.3, and the crosswind adds a horizontal force. The translational balances are shown in Equation (10.1).

$$\begin{aligned} \dot{x} &= v_x, \qquad \dot{z} = v_z \\ m\dot{v}_x &= T\sin(\theta-\delta) + F_w \\ m\dot{v}_z &= T\cos(\theta-\delta) - mg \end{aligned} \tag{10.1}$$

where $m$ is the vehicle mass in kg, $T$ is the thrust in N, $\theta$ and $\delta$ are in rad, $F_w$ is the aerodynamic side force in N and $g$ is gravitational acceleration in m/s$^2$. Both lines carry units of N. The two equations share one $T$, which is the coupling the whole application is about, because $T\sin\theta$ spent on moving sideways is missing from $T\cos\theta$ holding the vehicle up.

> **+ DID YOU KNOW? A Throttle Misled by a Turning Vehicle**
>
> In the last 140 seconds of the first crewed lunar landing, the descent throttle swung about 15 percent peak to peak around a setting near 26 percent. Fast pitch maneuvers put centrifugal accelerations on accelerometers mounted high above the center of mass, and the guidance computer read them as a change in vertical acceleration [12].

Rotation is about the center of mass, driven by the gimbal moment and by the wind acting at a center of pressure above the center of mass, as shown in Equation (10.2).

$$\dot{\theta} = \omega, \qquad I\dot{\omega} = \ell T\sin\delta + \ell_{cp}F_w \tag{10.2}$$

where $I$ is the pitch inertia in kg m$^2$, $\ell$ is the distance from the center of mass down to the gimbal pivot in m and $\ell_{cp}$ is the distance from the center of mass up to the center of pressure in m. Both

terms on the right carry units of N m. There is no aerodynamic restoring moment and no damping term of any kind, so Equation (10.2) is a pure double integrator driven by the gimbal. Placing the center of pressure above the center of mass is what makes a crosswind lean the vehicle downwind instead of merely pushing it sideways, and it is the reason the attitude loop has to hold a standing gimbal deflection in a steady wind.

### 10.3.3 Sign Convention for the Gimbal Deflection

The sign convention is the one statement in the model that is easy to get backward, and every claim about the wrong-way transient inverts with it. The engine sits at $\mathbf{r} = -\ell(\sin\theta, \cos\theta)$ relative to the center of mass. Evaluate the moment $\mathbf{r} \times \mathbf{F}$ at $\theta = 0$ with the thrust deflected by an angle $\phi$ from the body axis toward $+x$, so that $\mathbf{F} = T(\sin\phi, \cos\phi)$, as shown in Equation (10.3).

$$(\mathbf{r} \times \mathbf{F})_z = r_x F_z - r_z F_x = +\ell T \sin\phi \tag{10.3}$$

That is a counter-clockwise moment on a screen whose $x$ axis points right, and counter-clockwise is the direction of decreasing $\theta$. Deflecting the exhaust toward $+x$ therefore pitches the nose toward $-x$. A controller wants one number whose sign means point the nose this way, so the application defines $\delta = -\phi$, positive for a positive pitching moment. The moment is then $+\ell T \sin\delta$ as written in Equation (10.2), and the force necessarily points along $\theta - \delta$ as written in Equation (10.1). A deflection that brings the nose to the right pushes the vehicle left until the accumulated tilt is large enough for $T \sin\theta$ to overcome $-T \sin\delta$. The application carries a built-in test that evaluates the derivative function at $\theta = 0$ and $\delta = 5^\circ$ and checks that $\dot{\omega} > 0$ and $\dot{v}_x < 0$ in the same state, so the convention is pinned numerically rather than asserted.

### 10.3.4 Wind, Actuators and Propellant

The aerodynamic side load is quadratic in the air-relative lateral speed and is applied at the center of pressure, and Equation (10.4) gives it.

$$F_w = k_w\, v_r\, |v_r|, \qquad v_r = V_{\text{wind}} + c(t) - v_x \tag{10.4}$$

where $k_w$ is the side-force coefficient in N/m$^2$ times s$^2$, $V_{\text{wind}}$ is the steady crosswind in m/s and $c$ is the decaying gust contribution in m/s. Writing $k_w = \frac{1}{2}\rho C_D A_{\text{ref}}$ makes the coefficient physical, and the shipped value of 10.6 N s$^2$/m$^2$ corresponds to $C_D A_{\text{ref}} = 17.3\,\text{m}^2$ at sea-level density. A gust event adds $v_{\text{gust}}$ to $c$ and the contribution then decays with time constant $\tau_{\text{gust}}$ through the exact discrete first-order hold $c \leftarrow c \exp(-\Delta t_c / \tau_{\text{gust}})$ evaluated once per control execution. There is no vertical drag, which for a base area consistent with the lateral coefficient above and an axial drag coefficient of 0.5 would be about 80 N at 10 m/s of descent against a 11 772 N weight.

Neither actuator can step. Both take a command that is rate limited first and then hand it to a first-order lag, which Equation (10.5) states.

$$\begin{aligned} \tau_\delta\, \dot{\delta} &= \delta_{\text{cmd}} - \delta, & \left|\dot{\delta}_{\text{cmd}}\right| &\le \dot{\delta}_{\max} \\ \tau_T\, \dot{T} &= T_{\text{cmd}} - T, & \left|\dot{T}_{\text{cmd}}\right| &\le \dot{T}_{\max} \end{aligned} \tag{10.5}$$

where $\tau_\delta$ and $\tau_T$ are the servo and spool time constants in s. The gimbal command is clipped to $\pm\delta_{\max}$ and the thrust command to $[T_{\min}, T_{\max}]$ before the rate limit acts. The lower thrust bound matters more than the upper one, because $T_{\min}$ is 68 % of a hover: the engine cannot be shut down and relit, so descending faster than the profile asks is not something the controller can arrange by closing the throttle.

Propellant depletion is an option rather than a default and is shown in Equation (10.6).

$$\dot{m} = -\frac{T}{I_{\text{sp}}\, g_0}, \qquad m \ge m_{\text{dry}} \tag{10.6}$$

in which the specific impulse $I_{sp}$ is in s and standard gravity $g_0$ in m/s$^2$. This is the constant exhaust-speed form of the variable-mass momentum balance [14], and at hover it removes 4.00 kg/s. The mass is clamped at $m_{dry}$ after the integration step and a latching propellant-exhausted annunciation is raised, after which the model keeps thrusting, which is the optimistic case and is labeled as such in the application.

### 10.3.5 Guidance, Touchdown and Numerics

The vertical channel is given a descent-speed profile rather than an altitude setpoint, as shown in Equation (10.7).

$$v_{z,\mathrm{SP}} = -\min(v_{\max},\, k_z z + v_{\mathrm{touch}}) \tag{10.7}$$

where $v_{\max}$ is the cruise descent speed in m/s, $k_z$ is the profile slope in 1/s and $v_{\mathrm{touch}}$ is the terminal descent speed in m/s. Above $z = (v_{\max} - v_{\mathrm{touch}})/k_z = 27.1$ m the command is a constant cruise and below it the descent is a first-order approach with time constant $1/k_z = 2.86$ s. The $v_{\mathrm{touch}}$ term is what makes contact happen in finite time, because with $v_{\mathrm{touch}} = 0$ the profile approaches the pad asymptotically and the vehicle never arrives.

Touchdown is a hard event applied after the integration step. At $z \le 0$ the arrival is classified once and the rigid body is then frozen, so a landed vehicle produces clean constant columns instead of numerical debris. The verdict is a tip-over if $|\theta| > \theta_{\mathrm{tip}}$, otherwise a hard landing if $|v_z| > v_{z,\max}$ or $|v_x| > v_{x,\max}$, otherwise a miss if the offset from the pad center exceeds $R_{\mathrm{pad}}$, and otherwise a landing. Attitude is checked before speed, because landing legs care about the angle at contact and not about the average angle over the descent. The actuators keep integrating after contact so the engine can be seen spooling down, and the scenario run stops at the same instant, which means the application has no settled steady state to report for any scenario other than a pure vertical descent.

Every numerical value the model uses is listed in Tables 10.2 and 10.3. Integration is classical fourth-order Runge-Kutta over the nine states at a fixed 5 ms step, the controller executes every 10 ms, and the historian logs at 20 Hz. The fastest mode in the loop is the 80 ms gimbal servo, so the integration step is one sixteenth of it and the control period one eighth. The order of work in one step and the separation of rendering from the simulation are those of Section 1.6.1, with the touchdown correction applied after the integrator, so two machines produce identical trajectories and identical exports.

Every scenario starts trimmed on the descent profile, with $T(0) = m_0 g$ and $v_z(0) = v_{z,\mathrm{SP}}(z_0)$, so the opening seconds show a controller holding an equilibrium rather than recovering from a startup transient. Scenarios override individual initial conditions, and the disturbances available in any scenario are a crosswind slider spanning $\pm 15$ m/s and three event buttons that inject a decaying gust, add 3 m/s of horizontal velocity, or add 8 °/s of pitch rate. Of the three, the velocity shove is the informative one. It moves neither the position nor the attitude, so at the instant it lands the only term anywhere in the cascade that responds is $-K_{Dx} v_x$.

### 10.3.6 Assumptions

Only mechanisms that change the tuning survive in the model. Each assumption below is given with the range in which it holds and the situation that breaks it.

*Planar rigid body.* Three degrees of freedom, fixed inertia, no yaw, no roll and no propellant slosh. The assumption holds for the small excursions the tilt limit permits. It breaks on a vehicle with a part-full tank, where slosh is a lightly damped mode close in frequency to the attitude loop, and it breaks for any out-of-plane disturbance.

*Exact full-state feedback.* The controller reads $x$, $z$, $v_x$, $v_z$, $\theta$ and $\omega$ with no noise, no bias, no quantization and no lag. This is the largest omission in the application. A real vehicle infers those states from an inertial measurement unit and a radar altimeter, and navigation error rather than tuning is what sets the achievable touchdown precision. Every result quoted here should be read as the tuning-limited case.

**Table 10.2** Vehicle, Engine and Environment Constants. Plant Parameters Are Adjustable Over the Range Given, Apart From the Three Marked Hidden, Which the Parameter Panel Does Not Show.

| Symbol | Meaning | Value | Unit | Range |
|---|---|---|---|---|
| *Vehicle* | | | | |
| $m_0$ | wet mass at ignition | 1200 | kg | 1000-1600 |
| $m_{\text{dry}}$ | dry mass (hidden) | 900 | kg | 600-1200 |
| $I$ | pitch inertia | 4500 | kg m$^2$ | 1500-12000 |
| $\ell$ | center of mass to gimbal | 3.0 | m | 1.0-6.0 |
| $\ell_{cp}$ | center of mass to CP | 1.2 | m | 0.0-3.0 |
| $L$ | body length (hidden) | 8.0 | m | 5-14 |
| $g$ | surface gravity | 9.81 | m/s$^2$ | 8.0-11.0 |
| *Engine and gimbal* | | | | |
| $T_{\max}$ | maximum thrust | 20000 | N | 13000-26000 |
| $T_{\min}$ | minimum thrust | 8000 | N | 4000-10000 |
| $\tau_T$ | spool time constant | 0.15 | s | 0.02-0.80 |
| $\dot{T}_{\max}$ | throttle slew limit | 12000 | N/s | 2000-40000 |
| $I_{\text{sp}}$ | specific impulse | 300 | s | 200-400 |
| $g_0$ | standard gravity (hidden) | 9.80665 | m/s$^2$ | fixed |
| $\delta_{\max}$ | gimbal travel limit | 10.0 | ° | 1.0-10.0 |
| $\tau_\delta$ | gimbal time constant | 0.08 | s | 0.01-0.60 |
| $\dot{\delta}_{\max}$ | gimbal rate limit | 60 | °/s | 5-120 |
| *Environment and mission* | | | | |
| $k_w$ | side-force coefficient | 10.6 | N/m$^2$ s$^2$ | 0-30 |
| $v_{\text{gust}}$ | gust amplitude | 12 | m/s | 2-25 |
| $\tau_{\text{gust}}$ | gust decay time | 2.5 | s | 0.5-10 |
| $z_0$ | ignition altitude | 250 | m | 30-600 |
| $R_{\text{pad}}$ | landing pad radius | 8 | m | 3-25 |
| $v_{z,\max}$ | touchdown speed limit | 2.0 | m/s | 0.5-6 |
| $v_{x,\max}$ | lateral speed limit | 1.0 | m/s | 0.2-4 |
| $\theta_{\text{tip}}$ | tip-over angle | 6 | ° | 1-15 |

*One thrust number with a lag and a slew limit.* Combustion dynamics, throttle-dependent specific impulse, ignition transients and the possibility of an engine that will not relight are all absent. The first-order spool is a fair representation between about 40 % and 100 % throttle and understates the difficulty at the bottom of the range.

*Constant inertia under propellant burn.* Mass falls and $I$ does not, which is the optimistic case. The rotational equations of an axisymmetric body losing mass through a nozzle carry inertia terms that change with the propellant distribution and a jet-damping contribution that this model omits [15]. A real vehicle emptying tanks near its base loses inertia faster than it loses mass, so its attitude loop gets faster as it descends rather than slower, and a gain that is well damped at ignition can be marginal at touchdown.

*One quadratic side force at a fixed center of pressure.* There is no angle-of-attack dependence, no vertical drag, no Mach dependence and no plume-ground interaction in the last few meters. Ground effect in particular is strong, unsteady and exactly where the precision requirement is tightest.

*Instantaneous inelastic touchdown.* Leg stroke, damping and crush cores are absent, so the boundary between a landing and a hard landing is a threshold rather than a structural calculation, and the pad is flat, level and stationary. On a ship deck or a slope both assumptions fail at the moment they matter most.

**Table 10.3** Controller Settings, Values Derived at the Hover Operating Point and Numerical Constants. An (adv.) Marking Means the Control Appears Only Under the Advanced Toggle.

| Symbol | Meaning | Value | Unit | Range |
|---|---|---|---|---|
| *Controller* | | | | |
| $K_{Px}$ | position gain | 0.72 | °/m | 0-3 |
| $K_{Dx}$ | position rate gain | 4.10 | ° per m/s | 0-12 |
| $K_{P\theta}$ | attitude gain | 1.20 | °/° | 0-5 |
| $K_{D\theta}$ | attitude rate gain | 0.65 | s | 0-3 |
| $K_{Pv}$ | descent-speed gain | 1500 | N per m/s | 0-6000 |
| $\tau_{Iv}$ | descent integral time | 2.5 | s | 0.2-60 |
| $\theta_{\max}$ | tilt limit | 12.0 | ° | 1-15 (adv.) |
| $v_{\max}$ | cruise descent speed | 10.0 | m/s | 2-25 (adv.) |
| $k_z$ | profile slope | 0.35 | 1/s | 0.05-1.2 (adv.) |
| $v_{\text{touch}}$ | terminal descent speed | 0.50 | m/s | 0-2 (adv.) |
| *Derived at the hover operating point* | | | | |
| $mg$ | hover thrust | 11772 | N | — |
| — | hover throttle | 31.43 | % | — |
| $a$ | attitude plant gain $\ell T/I$ | 7.848 | $1/s^2$ | — |
| $\sqrt{a}$ | right-half-plane zero | 2.801 | rad/s | — |
| $\omega_{n\theta}$ | attitude natural frequency | 3.069 | rad/s | — |
| $\zeta_\theta$ | attitude damping ratio | 0.831 | — | — |
| $\omega_{nx}$ | position natural frequency | 0.351 | rad/s | — |
| $\zeta_x$ | position damping ratio | 1.000 | — | — |
| $\omega_{nv}$ | descent natural frequency | 0.707 | rad/s | — |
| $\zeta_v$ | descent damping ratio | 0.884 | — | — |
| *Numerics* | | | | |
| $\Delta t$ | integration step (RK4) | 0.005 | s | — |
| $\Delta t_c$ | control period | 0.010 | s | — |
| — | historian rate | 20 | Hz | — |

## 10.4 Linearized Dynamics and Loop Bandwidths

Every quantitative statement in this section comes from linearizing Equations (10.1) and (10.2) about hover, with $\theta$ and $\delta$ small, $T \approx mg$ constant and no wind. That linearization is the analysis model. It is used to place the gains and to explain the measurements, and it is used nowhere inside the application.

### 10.4.1 The Attitude Plant

With $\sin\delta \approx \delta$ and the wind torque removed, Equation (10.2) becomes the pure double integrator of Equation (10.8).

$$\ddot{\theta} = a\,\delta, \qquad a \equiv \frac{\ell T}{I} = 7.848/\mathrm{s}^2 \tag{10.8}$$

where $a$ is the attitude plant gain in $1/s^2$, evaluated at $T = m_0 g = 11\,772\,\mathrm{N}$. There is no restoring moment and no damping, so proportional action supplies the entire spring and derivative action supplies the only damper the vehicle has. Substituting the attitude law $\delta = K_{P\theta}(\theta_{\mathrm{SP}} - \theta) - K_{D\theta}\omega$ gives the closed-loop pair of Equation (10.9).

$$\omega_{n\theta} = \sqrt{a\,K_{P\theta}}, \qquad \zeta_\theta = \frac{a\,K_{D\theta}}{2\,\omega_{n\theta}} \tag{10.9}$$

At $K_{P\theta} = 1.20$ and $K_{D\theta} = 0.65\,\mathrm{s}$ this places the attitude loop at 3.069 rad/s with a damping ratio of 0.831. Both gains are dimensionally clean, because $K_{P\theta}$ is an angle per angle and $K_{D\theta}$ is a time.

### 10.4.2 The Right-Half-Plane Zero of the Gimbal-to-Position Path

Horizontal acceleration comes from the direction of the thrust vector, which is $\theta - \delta$ and not $\theta$. Linearizing the first line of Equation (10.1) gives $\dot{v}_x = g(\theta - \delta)$, and eliminating $\theta$ with Equation (10.8) produces the open-loop transfer from gimbal angle to horizontal position shown in Equation (10.10).

$$\frac{x(s)}{\delta(s)} = \frac{g\left(a - s^2\right)}{s^4} \tag{10.10}$$

The path is a quadruple integrator with a zero pair at $s = \pm\sqrt{a}$, one of which sits in the right half plane at $\sqrt{a} = 2.801\,\text{rad/s}$. The physical content of that zero is the wrong-way transient, the same start in the wrong direction that a cart makes under a balanced pole in Volume 7, Chapter 7. For a step deflection $\delta$ applied to the linearized model the position response is $x(t) = g\delta(at^4/24 - t^2/2)$, so the horizontal acceleration changes sign at $t = \sqrt{2/a} = 0.505\,\text{s}$, the excursion reaches its extreme of $1.5\,g\delta/a$ at $t = \sqrt{6/a} = 0.874\,\text{s}$, and the vehicle returns through its starting point at $t = \sqrt{12/a} = 1.237\,\text{s}$. At the 10° travel limit that extreme is 0.33 m in the wrong direction.

Closing the attitude loop does not help. Substituting the same attitude law into Equation (10.10) and treating $\theta_{\text{SP}}$ as the input gives Equation (10.11).

$$\frac{x(s)}{\theta_{\text{SP}}(s)} = \frac{g\,K_{P\theta}\left(a - s^2\right)}{s^2\left(s^2 + 2\zeta_\theta\omega_{n\theta}s + \omega_{n\theta}^2\right)} \tag{10.11}$$

The numerator is unchanged. Feedback moves poles and leaves zeros where they are, so the right-half-plane zero of the vehicle survives every attitude gain that could be dialed in, and the only quantities that move it are $\ell$, $T$ and $I$. A shorter engine bay, a heavier engine or a fatter vehicle changes the achievable position bandwidth. Retuning does not. Chapter 9 meets a right-half-plane zero of the same kind in an electrical cascade, where it belongs to the power stage and moves with load and duty cycle.

🕮 **Pause.** Try Exercise 10.2 before reading on.

### 10.4.3 Bandwidth Separation and the Zero Margin

A real right-half-plane zero at $z$ limits the achievable crossover of a loop that must cross it to roughly $z/2$, because the phase the zero subtracts cannot be recovered [11]. Here that ceiling is 1.40 rad/s. The cascade imposes a second ceiling, that the position loop must be slow compared with the attitude loop, which is the inner-loop speed rule of Section 1.4.5 read in frequency. What makes this particular vehicle instructive is that the two ceilings are the same quantity in disguise, as shown in Equation (10.12).

$$\frac{\omega_{n\theta}}{\sqrt{a}} = \sqrt{K_{P\theta}} = 1.095 \tag{10.12}$$

The attitude loop sits within 10 % of the right-half-plane zero, and it does so at every attitude gain of order unity, because the ratio grows only as the square root. Making the inner loop four times faster in bandwidth costs sixteen times the attitude gain and moves the zero not at all.

At low frequency Equation (10.11) reduces to $x/\theta_{\text{SP}} \to g/s^2$, a double integrator with gain $g$, and the position law $\theta_{\text{SP}} = K_{Px}(x_{\text{SP}} - x) - K_{Dx}v_x$ closes it as shown in Equation (10.13).

$$\omega_{nx} = \sqrt{g\,K_{Px}}, \qquad \zeta_x = \frac{g\,K_{Dx}}{2\,\omega_{nx}} \tag{10.13}$$

At $K_{Px} = 0.72°/\text{m}$ (0.012 57 rad/m) and $K_{Dx} = 4.10°$ per m/s (0.071 56 rad s/m) the position loop lands at 0.351 rad/s with a damping ratio of 1.000, exactly critical. That is a factor of 8.7 below the attitude loop and a factor of 8.0 below the right-half-plane zero, and the two factors agree because Equation (10.12) says they must. The plant gain in Equation (10.13) is $g = 9.81\,\text{m/s}^2$ per rad of commanded tilt, which is not a controller gain and does not change with tuning, and the controller gain $K_{Px}$ carries units of rad/m.

### 10.4.4 The Descent Loop and Its Feedforward

The vertical channel sees $m\dot{v}_z = T\cos(\theta-\delta) - mg$, which for small angles is $m\dot{v}_z = T - mg$, an integrator driven by the excess thrust. Proportional-integral action closes it as shown in Equation (10.14).

$$\omega_{nv} = \sqrt{\frac{K_{Pv}}{m\,\tau_{Iv}}}, \qquad \zeta_v = \frac{K_{Pv}}{2\,m\,\omega_{nv}} \tag{10.14}$$

At $K_{Pv} = 1500\,\mathrm{N}$ per m/s and $\tau_{Iv} = 2.5\,\mathrm{s}$ against 1200 kg this gives 0.707 rad/s at a damping ratio of 0.884, and the proportional term alone would give a velocity time constant of $m/K_{Pv} = 0.80\,\mathrm{s}$. Both are comfortably inside the 2.86 s guidance time constant the loop has to follow.

This is the only loop carrying an integrator and the only one that needs one, because it is the only loop with a constant load. That load is the weight, and handing it to the loop as a feedforward bias, a static feedforward from a known force in the sense of Section 1.4.3, rather than making the integrator discover it is what keeps the vertical channel quiet. Tilting steals vertical thrust as $\cos\theta$, so the bias is $mg/\cos\theta$, which is 2.23 % above $mg$ at the 12° tilt limit, an extra 263 N out of 11 772 N. The same arithmetic explains why the tilt limit is where it is. At 30° the correction is 15.5 %, or 1821 N, and an engine limited to 14 kN can supply it only by giving up its deceleration margin, because 14 kN leaned 30° leaves 352 N of spare vertical force where 2228 N was available upright.

### 10.4.5 Offset in a Steady Crosswind

Neither cascaded loop has integral action, so a constant side load leaves a standing position error, and the chain that sets its size runs backward through all three levels. A 7 m/s crosswind produces $F_w = 519\,\mathrm{N}$ from Equation (10.4). Canceling that force at hover thrust requires $\sin(\theta-\delta) = -519/11772$, or −2.53°. Canceling the wind moment $\ell_{cp}F_w = 623\,\mathrm{N\,m}$ requires $\sin\delta = -623/(\ell\,mg)$, a standing gimbal deflection of −1.01°. The attitude therefore has to sit at $-2.53 - 1.01 = -3.54°$. Producing −1.01° of gimbal from a proportional gain of 1.20 requires −0.84° of attitude error, so $\theta_{\mathrm{SP}} = -4.38°$, and producing that command from $K_{Px} = 0.72°/\mathrm{m}$ requires 6.09 m of standing position error. Section 10.7 compares each link of that chain with what the application holds.

🕮 **Pause.** Try Exercise 10.3 before reading on.

## 10.5 Controller Design

All three loops execute in one function every $\Delta t_c = 10\,\mathrm{ms}$, which is a 100 Hz flight computer and twice as slow as the physics step. The original cascade carries no optimizer, horizon or internal vehicle model. The predictive option added to it is described in Section 10.5.4.

### 10.5.1 The Two Cascaded Loops

Both cascaded loops are proportional-derivative in the explicit position-plus-rate form of Equations (10.15) and (10.16), evaluated on the state vector directly.

$$\begin{aligned} \theta_{\mathrm{SP}}^{\mathrm{raw}} &= K_{Px}\,(x_{\mathrm{SP}} - x) - K_{Dx}v_x \\ \theta_{\mathrm{SP}} &= \mathrm{clamp}(\theta_{\mathrm{SP}}^{\mathrm{raw}},\, -\theta_{\max},\, \theta_{\max}) \end{aligned} \tag{10.15}$$

$$\begin{aligned} \delta_{\mathrm{cmd}}^{\mathrm{raw}} &= K_{P\theta}\,(\theta_{\mathrm{SP}} - \theta) - K_{D\theta}\,\omega \\ \delta_{\mathrm{tgt}} &= \mathrm{clamp}(\delta_{\mathrm{cmd}}^{\mathrm{raw}},\, -\delta_{\max},\, \delta_{\max}) \end{aligned} \tag{10.16}$$

Neither loop has an integral term, a derivative filter or any internal state whatsoever. The derivative terms use $v_x$ and $\omega$, which are states of the simulation rather than numerical derivatives of a measurement, so no filter time constant appears and no derivative kick is possible on a setpoint move. Both loops are therefore memoryless functions of the state, so neither needs initializing when control changes hands, although a transfer to automatic operation moves the gimbal to whatever the attitude law computes at that instant. Integral action is deliberately absent from the attitude loop, because a

reset term inside the fastest loop in the cascade spends the phase margin that loop has least of, and it is deliberately absent from the position loop, because an integrator in the outer loop of a cascade whose inner loop saturates is a windup trap. The consequence of the second choice is the steady offset derived above, and the application argues for a wind estimate biasing $x_{\mathrm{SP}}$ rather than for an outer integrator.

The gimbal command published to the actuator is the rate-limited version of $\delta_{\mathrm{tgt}}$ shown in Equation (10.17).

$$\delta_{\mathrm{cmd}} \leftarrow \mathrm{slew}\Big(\delta_{\mathrm{tgt}},\, \delta_{\mathrm{cmd}},\, \dot{\delta}_{\max},\, \Delta t_c\Big) \tag{10.17}$$

where slew moves the previous command toward the target by at most $\dot{\delta}_{\max}\Delta t_c$. The servo lag of Equation (10.5) then acts on the result, so the rate limit and the lag are separate mechanisms applied in series and can be varied independently.

### 10.5.2 The Descent Loop

The vertical loop uses the velocity-form controller of Section 1.6.2 [16] with the derivative time fixed at zero, so the proportional and integral increments are the only ones active. In this loop's symbols the recursion is Equations (10.18) to (10.24), where $b_k$, $s_k$ and $c_k$ are the $u_b$, $S_k$ and $v_k$ of Equations (1.10) and (1.11), and the bias $b_k$ changes from scan to scan.

$$e_k = v_{z,\mathrm{SP},k} - v_{z,k} \tag{10.18}$$

$$b_k = \mathrm{clamp}\left(\frac{\chi_{\mathrm{FF}}\, m_k\, g}{\max(\cos\theta_k,\, 0.25)},\, T_{\min},\, T_{\max}\right) \tag{10.19}$$

$$\Delta u_k^P = K_{Pv}\,(e_k - e_{k-1}), \qquad \Delta u_k^I = \frac{K_{Pv}}{\tau_{Iv}}\, e_k\, \Delta t_c \tag{10.20}$$

$$s_k = s_{k-1} + \Delta u_k^P \tag{10.21}$$

$$c_k = b_k + s_k + I_{k-1} \tag{10.22}$$

$$u_k = \mathrm{clamp}(c_k,\, T_{\min},\, T_{\max}) \tag{10.23}$$

$$I_k = I_{k-1} + \Delta u_k^I + \frac{\Delta t_c}{\tau_t}\,(u_k - c_k) \tag{10.24}$$

where $e_k$ is the descent-speed error in m/s, $b_k$ is the gravity feedforward bias in N, $c_k$ is the thrust the algorithm demands and $u_k$ the thrust it is allowed to command, both in N, $s_k$ accumulates the proportional increments without limit, $I_k$ holds the integral contribution in N, $\tau_t = \tau_{Iv}$ is the tracking time of the back-calculation in s, and the flag $\chi_{\mathrm{FF}} \in \{0, 1\}$ turns the bias on and off. The integral used at scan $k$ was accumulated through scan $k-1$, and the update of Equation (10.24) runs only when $K_{Pv}/\tau_{Iv} \neq 0$. While the command is clipped, the term $(\Delta t_c/\tau_t)(u_k - c_k)$ drives $I$ toward the value that reproduces the applied thrust, removing $1/250$ of the clipped excess per scan at the default $\tau_{Iv} = 2.5\,\mathrm{s}$, and the running sum $s_k$ is never truncated, as Section 1.6.2 explains. The tracking time affects no result reported here. The only shipped case that clips the thrust command, scenario 7, stays on the upper limit until contact, and a counterfactual run with a tracking time of one scan leaves its arrival speed unchanged at 3.59 m/s. The published thrust command is the rate-limited version of $u_k$, $T_{\mathrm{cmd},k} = \mathrm{slew}(u_k, T_{\mathrm{cmd},k-1}, \dot{T}_{\max}, \Delta t_c)$, so the rate limit acts once and outside the controller.

The loop is reverse acting in the ISA sense, because more thrust raises $v_z$, so a climb rate above the commanded value has to reduce thrust. The floor of 0.25 on $\cos\theta$ caps the feedforward past 75.5° of tilt, which is well outside anything the tilt limit permits and exists only to keep the expression finite if a scenario is edited into a tumble.

Equation (10.19) clamps the bias to the engine range, so switching the feedforward off leaves $b_k = T_{\min} = 8\,\mathrm{kN}$ rather than zero and the integral has to build only the remaining 3.77 kN of hover thrust. When the switch is thrown during a run, the application shifts $I_k$ by $mg/\cos\theta - T_{\min}$ (with $mg/\cos\theta$ clamped to the engine range), which is the difference between the two biases. That holds

the thrust command continuous across the switch and preloads the integral with the hover thrust the bias no longer supplies. A switch made before the run is a starting condition instead. The controller is built with $I = 0$ against whichever bias is selected, so a run started with the feedforward off begins with no preload at all, and Section 10.7 measures what that costs.

### 10.5.3 Manual Operation and the Three Models

In manual operation the throttle and gimbal commands $T_{\text{cmd}}$ and $\delta_{\text{cmd}}$ are operator-set inputs, and a hover-trim helper computes $mg/\cos\theta$ at the present mass and tilt, applies it as the throttle and centers the gimbal. All three loops keep computing throughout. The two proportional-derivative loops are memoryless and the vertical loop is fed the operator's throttle through the tracking step of Section 1.6.2, which back-solves $I_k$ so that $c_k$ of Equation (10.22) reproduces the slider exactly. The throttle slider starts every run at the trim throttle $(T_0 - T_{\min})/(T_{\max} - T_{\min})$, with $T_0 = m_0 g$ clamped to the engine range, so a vehicle taken into manual operation before the run holds its trimmed start whatever engine the scenario fits. Automatic operation drives both sliders from the commands after every controller execution, so a transfer to manual moves nothing. A transfer to automatic is bumpless for the thrust, as a by-product of the tracking arithmetic, but not for the gimbal, whose target jumps from the operator's setting to the output of Equation (10.16) and is then limited only by the slew rate. After 1 s of manual flight in scenario 2 with the gimbal centered, the switch moves the gimbal command from 0° to its 10° limit within 0.2 s.

Three models are kept apart throughout this chapter. The simulation model is Equations (10.1) to (10.6), nine nonlinear states integrated at 5 ms. The analysis model is the hover linearization of Section 10.4, used to place gains and to interpret results and present nowhere in the code. The controller's internal model is two algebraic relations only, the hover thrust $mg/\cos\theta$ and the guidance profile of Equation (10.7). The controller never sees $\ell$, $I$, $k_w$ or $\ell_{cp}$, so it knows nothing about the right-half-plane zero it is working against.

### 10.5.4 Joint Predictive Thrust and Gimbal Requests

An additional MPC, which solves the shared quadratic program of Section 1.6.3 on a model of its own, relinearizes the nominal planar force and torque balances at each decision. Its eight states are the six ideal kinematic readings and independently integrated gimbal and thrust. Geometry, inertia, gravity and servo constants use declared defaults. Nominal mass begins at 1200 kg and integrates estimated thrust over nominal $I_{sp}g_0$ when burn is enabled, with a 900 kg floor. Actual mass and servo states are not inputs. For example, an actual 1440 kg initial mass leaves this estimate 240 kg low. Measured velocity/rate differences correct additive horizontal, vertical and angular acceleration biases, with gain 0.1 and bounds 6, 4 m/s$^2$ and 2 rad/s$^2$. No wind or future event is read. This is disturbance correction [17], not inertia or mass identification. The nominal observer tracks applied requests under either selected law.

Every 0.1 s, the local affine model is exactly discretized for 60 steps, 6 s, in twelve move blocks. Future vertical-speed references integrate the existing guidance profile from current altitude. The shared convex QP [18] tracks $[x, v_x, v_z, \theta, \omega]$ with weights $[0.3, 2, 10, 100, 10]$ and terminal weights $[2, 4, 10, 300, 30]$. With thrust in kN and gimbal in rad, absolute input weights are $[0.001, 0.05]$ and move weights $[0.02, 2]$. Request boxes use current actuator limits. Block moves are at most 1.2 kN and 2°, further restricted by the physical slew allowances. Tilt is soft, priced quadratically at $10^5$ and linearly at 100. Actual command slew and touchdown idle, centering and freeze remain. The vertical PI tracks applied post-slew thrust once per scan. A rejected plan holds both requests on handover, marks fallback and returns to Cascade until recovery.

The predictor matches nominal zero-aeroload derivatives exactly. Over one 0.1 s interval at tilts up to $\pm 0.2$ rad, affine prediction errors are below 0.000 468 m/s and 13.4 µm against nonlinear integration. An actual two-input one-step optimizer check agrees with hand inversion within $3.5 \times 10^{-18}$. These are local checks. Mass and acceleration bias are held inside a forecast, command-slew transients are

**Table 10.4** Scenario Set. Every Entry Not Listed Under Change Keeps Its Default Value From Tables 10.2 and 10.3.

| # | Change | Mechanism isolated |
|---|---|---|
| 1 | none | the guidance profile and its corner |
| 2 | $x(0) = -20\,\text{m}$ | three timescales and the wrong-way transient |
| 3 | $v_x(0) = 6\,\text{m/s}$ | the rate term as the whole command |
| 4 | $\theta(0) = 10^\circ$ | cross-coupling met by feedforward |
| 5 | $V_{\text{wind}} = 7\,\text{m/s}$ | offset from proportional action |
| 6 | $z_0 = 80\,\text{m}$, 16 m/s gust at 6 s | time-to-go against loop time constant |
| 7 | $T_{\max} = 14\,\text{kN}$ | deceleration authority, not hover authority |
| 8 | $\tau_\delta = 0.18\,\text{s}$, $\dot{\delta}_{\max} = 30^\circ/\text{s}$, $x(0) = -15\,\text{m}$ | an inner loop faster than its actuator |
| 9 | $\delta_{\max} = 3^\circ$, $x(0) = -25\,\text{m}$, $v_x(0) = -2\,\text{m/s}$ | a saturated inner loop opens the cascade |
| 10 | propellant burn on, $z_0 = 320\,\text{m}$, $x(0) = -12\,\text{m}$ | three loops drifting at different rates |
| 11 | MPC, $z_0 = 100\,\text{m}$, $x(0) = -15\,\text{m}$, $v_x(0) = 2\,\text{m/s}$, $\theta(0) = 3^\circ$, wind 2 m/s, gust at 4 s | joint bounded thrust and gimbal prediction |
| 12 | graded, see Section 10.7.10 | everything at once, from poor tuning |

omitted, and there is no ground-contact, glide-cone or terminal landing guarantee. The method tracks guidance rather than minimizing propellant.

## 10.6 Scenarios

Table 10.4 lists eleven teaching scenarios and one graded challenge. Each fixes only the quantity it exists to demonstrate and inherits every other setting, so whatever a run shows is attributable to one change. All twelve open in automatic operation, and an operator is free to take manual control at any point. Each case is credited when the learner performs its listed actions and the simulation reaches its listed state, as Section 1.6.4 describes, and the challenge is judged on the case as given, so changing a model parameter or a disturbance voids the attempt.

## 10.7 Results

The numbers below come from the scorecard the application computes for itself and from the file it exports, with the physics, the controller and the historian running at the rates in Table 10.3. A run ends at touchdown, so the integrated measures cover the whole descent rather than a fixed window, and the descent lasts 32.45 s from 250 m under the default profile. Every setting named below is changed after the scenario is selected and before the run starts, and the application gives the same run when the setting is carried through Reset instead. The model has no random input, so each run repeats exactly. Table 10.5 collects the outcomes. Figure 10.3 draws the arrivals of that table against the pad radius and the touchdown speed limit.

> ✦ **DID YOU KNOW? Take Away the Rate Term and It Tips Over**
>
> From the simulator: scenario 3 starts the vehicle drifting at 6 m/s, and it lands 0.01 m from the pad center. With the position rate gain $K_{Dx}$ set to zero, the same start tips over 14.10 m from the pad, and the position integral grows from 77.7 to 348.3 m s (rows 3 and 3b of Table 10.5).

### 10.7.1 A Joint Plan and Its Model-Mismatch Limits

Case 11 starts 100 m high, 15 m left, with 2 m/s rightward drift, 3° tilt and 2 m/s wind. A gust arrives at 4 s, without preview to either controller. Figure 10.4 and Table 10.6 use the same seed and ideal measurements. MPC improves offset and precontact thrust impulse, but contacts faster vertically and moves the gimbal more. Impulse is not a fuel-optimality result. A six-point position-PD hindsight

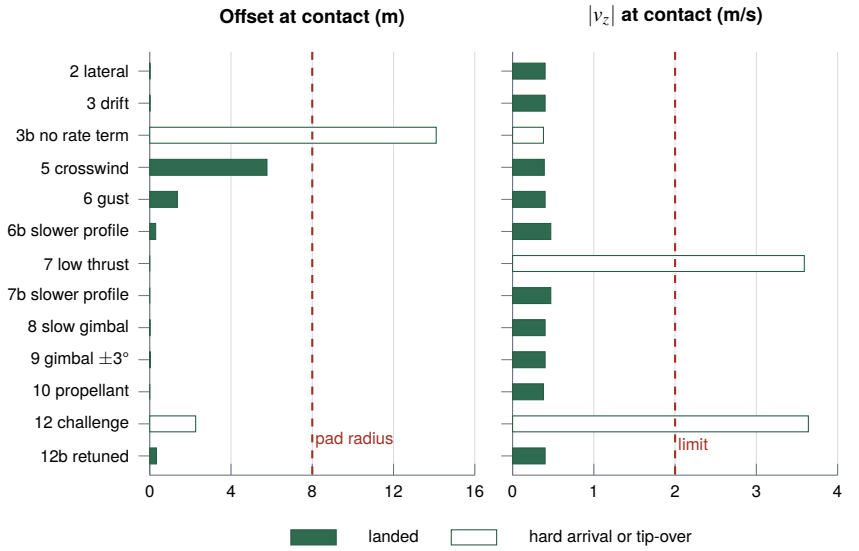



**Figure 10.3** Table 10.5 redrawn for thirteen of its rows: the vertical descent, the tilted start with and without feedforward, the two gimbal variants of scenario 8 and the 1° gimbal of row 9b are left out (row 9b arrives 42.49 m away at 57.77 m/s). Filled bars are landings and open bars are hard arrivals or tip-overs. The dashed lines are the 8 m pad radius and the 2.0 m/s touchdown limit. Notice that the failures split into two kinds: row 3b tips over far from the pad at a gentle vertical speed, while rows 7 and 12 arrive too fast, and each is repaired by a different change (rows 3, 7b and 12b).

**Table 10.5** Measured Outcomes of the Original Scenarios. Each IAE Is the Integral of the Absolute Error of That Loop Over the Whole Descent, Offset Is the Distance From the Pad Center at Contact, and the Verdict Is the Application's Own Classification of the Arrival.

| # | Set-up | Duration (s) | $x$ IAE (m s) | $\theta$ IAE (rad s) | $v_z$ IAE (m) | Peak tilt (°) | Offset (m) | $\lvert v_z \rvert$ at contact (m/s) | Verdict |
|---|---|---|---|---|---|---|---|---|---|
| 1 | vertical descent | 32.45 | 0 | 0 | 6.67 | 0.0 | 0.00 | 0.40 | landed |
| 2 | 20 m lateral | 32.45 | 123.0 | 0.264 | 6.84 | 11.4 | 0.01 | 0.40 | landed |
| 3 | 6 m/s drift | 32.45 | 77.7 | 0.255 | 6.85 | 12.2 | 0.01 | 0.40 | landed |
| 3b | the same with $K_{Dx} = 0$ | 32.46 | 348.3 | 0.612 | 6.94 | 12.0 | 14.10 | 0.38 | tip-over |
| 4 | 10° initial tilt | 32.45 | 8.7 | 0.195 | 6.83 | 10.0 | 0.00 | 0.40 | landed |
| 4b | the same, feedforward off | 31.83 | 7.9 | 0.199 | 12.94 | 10.0 | 0.00 | 0.40 | landed |
| 5 | 7 m/s crosswind | 32.46 | 152.2 | 0.475 | 6.68 | 3.6 | 5.77 | 0.39 | landed |
| 6 | gust at 20 m | 15.46 | 21.6 | 0.155 | 6.74 | 7.0 | 1.36 | 0.40 | landed |
| 6b | the same with $k_z = 0.15$ | 22.32 | 29.4 | 0.185 | 3.80 | 7.9 | 0.29 | 0.47 | landed |
| 7 | $T_{\max} = 14$ kN | 26.11 | 0 | 0 | 9.96 | 0.0 | 0.00 | 3.59 | hard |
| 7b | the same with $k_z = 0.15$ | 39.32 | 0 | 0 | 3.80 | 0.0 | 0.00 | 0.47 | landed |
| 8 | slow gimbal | 32.45 | 89.4 | 0.404 | 7.17 | 14.9 | 0.01 | 0.40 | landed |
| 8a | the same, nominal gimbal | 32.45 | 89.0 | 0.233 | 6.83 | 10.2 | 0.01 | 0.40 | landed |
| 8b | the same with $K_{P\theta} = 0.60$ | 32.45 | 89.9 | 0.406 | 6.73 | 7.9 | 0.01 | 0.40 | landed |
| 9 | $\delta_{\max} = 3°$ | 32.45 | 207.6 | 0.315 | 6.79 | 12.0 | 0.02 | 0.40 | landed |
| 9b | the same with $\delta_{\max} = 1°$ | 17.39 | 333.0 | 10.57 | 78.37 | 135.4 | 42.49 | 57.77 | tip-over |
| 10 | propellant depletion | 39.30 | 70.6 | 0.185 | 5.95 | 8.2 | 0.00 | 0.38 | landed |
| 12 | graded, as shipped | 22.82 | 148.4 | 2.081 | 11.97 | 23.8 | 2.26 | 3.64 | hard |
| 12b | graded, retuned | 29.45 | 99.6 | 0.246 | 6.76 | 5.9 | 0.33 | 0.40 | landed |

**Table 10.6** Contact and Precontact Maneuver, Cascade / MPC

| Case | Offset m | $\lvert v_z \rvert$ m/s | Time s | Impulse kN s |
|---|---|---|---|---|
| Case 11 | 0.726 / -0.079 | 0.395 / 0.519 | 17.45 / 15.52 | 217.2 / 195.0 |
| Lateral | -0.011 / -0.000 | 0.395 / 0.518 | 32.45 / 30.58 | 394.0 / 372.2 |
| Low thrust | 0.000 / 0.000 | 3.593 / 0.518 | 26.11 / 30.41 | 315.1 / 369.4 |
| Slow gimbal | -0.008 / -2.346 | 0.395 / 22.282 | 32.45 / 15.97 | 394.5 / 216.9 |

Last row MPC tips over. New-case gimbal travel: 13.7 / 87.2 degrees. MPC plans / mean solve time in microseconds: case 11 156 / 746, lateral 306 / 500, low thrust 305 / 435, slow gimbal 160 / 816. No rejected plans in these runs.

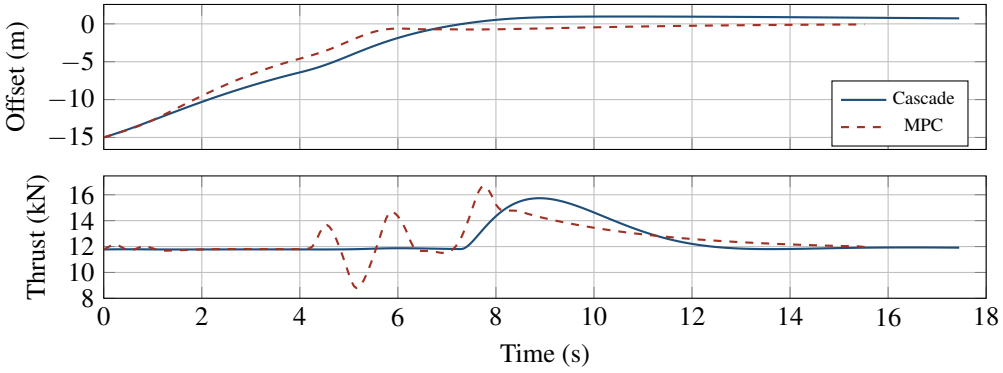



**Figure 10.4** Case 11, identical descent and unknown gust at 4 s, with the Cascade as a solid line and the MPC as a dashed one. Curves end at each controller's contact. Lower impulse and offset under MPC accompany faster vertical arrival and more gimbal travel, not a universal landing gain.

grid, three simpler control changes and two isolated bias-feedforward Cascade diagnostics complete the 43-row measurement record behind this section. The finer physics step preserves this offset and impulse ranking.

Under steady crosswind, bias-feedforward Cascade reaches 0.0308 m offset, compared with MPC –0.0363 m and original Cascade 5.771 m. Thus removing offset alone does not require a horizon. With reduced thrust, MPC lands instead of arriving hard, but uses 369.36 versus 315.06 kN s impulse. Conversely, the slow-gimbal case tips over under MPC at 22.28 m/s vertical speed while Cascade lands. Doubling inertia in the new case also tips MPC over while Cascade lands. Both suggested challenge controllers pass, MPC at −0.0248 m and Cascade at 0.3341 m, with gimbal travel 86.9 versus 29.1°. Local model error can dominate nominal benefits.

In a browser, measured over six window sizes and color themes, 100 plans average 0.96–1.54 ms. The planner then takes 0.96–1.54% of one processor core at real-time playback and 3.83–6.14% at four-times speed, excluding the nominal observer and rendering.

### 10.7.2 Lateral Displacement and the Wrong-Way Transient

Scenario 2 starts the vehicle 20 m left of the pad and puts the whole identity of the application into one 40 s frame, which is what Figure 10.5 shows. The position error of 20 m asks for $K_{Px} \cdot 20\,\mathrm{m} = 14.4°$ of tilt, which the limiter clips to 12°, so the outer loop begins saturated and stays there for 0.98 s. The gimbal is moving by the first physics step at $t = 0.005\,\mathrm{s}$, the tilt reaches 10° at $t = 0.91\,\mathrm{s}$, and the position enters the 1 m band around the pad at $t = 14.535\,\mathrm{s}$ and stays inside it. Three timescales, in the order the cascade was designed to produce them.

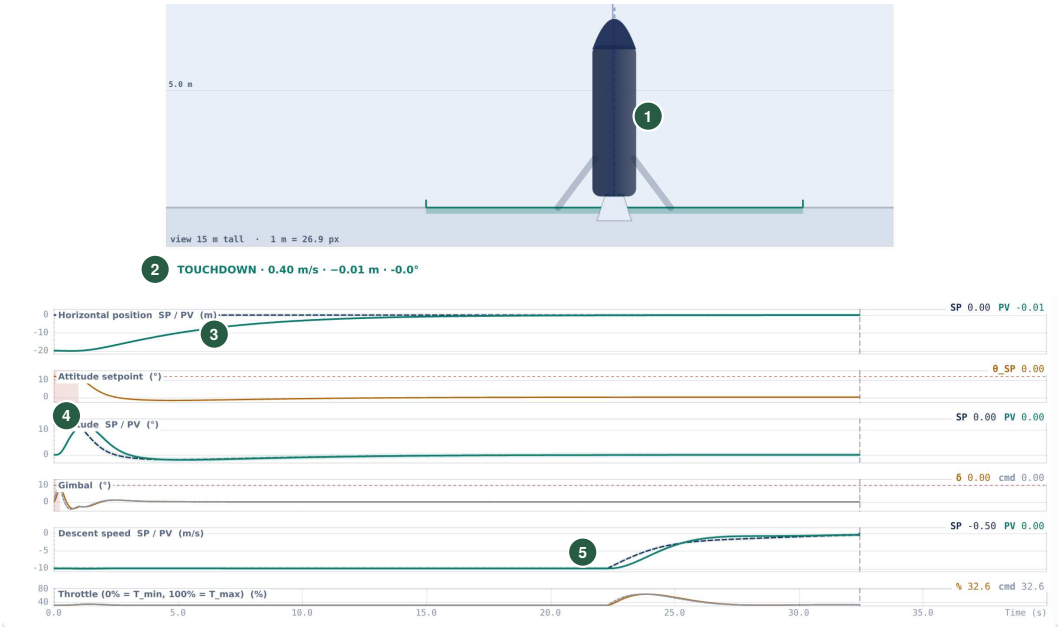



**Figure 10.5** Scenario 2 at touchdown, $t = 32.45$ s, a landing from 20 m left of the pad on a 40 s trend window that holds the whole descent. The picture is a composite of the application's window: the lower part of the plant view with its touchdown line, centered above the six trend strips, which are stacked in cascade order. The header, the value chips, the Control and Learn panels and the scorecard are cut. Callouts: (1) the vehicle standing on the pad, (2) the touchdown verdict, (3) horizontal position closing on the pad over fourteen seconds, (4) its manipulated variable, the attitude setpoint, clipped at 12° for the first 0.98 s, shown as the shaded band, because $K_{Px}$ times a 20 m error asks for 14.4°, (5) the descent strip holding −10 m/s until the profile corner at 27.1 m and then shrinking to the terminal rate. The attitude follows its command with the gimbal beneath it. In the logged run the gimbal peaks at 8.4° at $t = 0.26$ s and the attitude at 11.4° at $t = 1.25$ s. The scorecard, cut from this picture, reports a position integral of 123.0 m s, a peak tilt of 11.4° and an arrival 0.01 m from the pad center at 0.40 m/s (row 2 of Table 10.5), a peak throttle of 64 %, and saturation tiles of 2.9 % for the position loop and 0.6 % for the attitude loop, the share of the run in which each command is clipped. The position enters the 1 m band around the pad at 14.535 s.

**TRY IT IN THE APP** **2 · 20 m Lateral Displacement**

Load the scenario and press Run. Watch the horizontal position trend in the first second. The vehicle first drifts a further 0.118 m to the left, away from the pad, and only then turns toward it. It enters the 1 m band around the pad at 14.535 s and lands 0.01 m from the center.

The first thing the vehicle does, however, is travel the wrong way. Horizontal position drifts a further 0.118 m to the left, the horizontal acceleration is positive from the physics step ending at $t = 0.435$ s and the velocity reverses at $t = 0.71$ s, after which the vehicle closes on the pad and lands 0.01 m from center at 0.40 m/s. The linearized step response of Section 10.4 predicts an acceleration reversal at 0.505 s for an instantaneous full deflection, and the measured figure is earlier because the gimbal cannot step and the attitude loop has already begun to pull the deflection back before it saturates the travel limit. The excursion is 0.118 m rather than the 0.33 m of a sustained deflection for the same reason.

**Table 10.7** Scenario 2 at Four Position Bandwidths, Each Critically Damped With $K_{Dx}$ of 4.10, 5.90, 8.35 and 11.85 ° per m/s on the Slider Grid. The Attitude Loop Is Unchanged at 3.07 rad/s Throughout.

| $K_{Px}$ (°/m) | $\omega_{nx}$ (rad/s) | Ratio | $x$ IAE (m s) | $\theta$ IAE (rad s) | $\theta_{SP}$ travel (rad) | Verdict |
|---|---|---|---|---|---|---|
| 0.72 | 0.351 | 8.7 | 123.0 | 0.264 | 0.276 | landed 0.01 m |
| 1.50 | 0.507 | 6.1 | 98.5 | 0.337 | 0.407 | landed 0.00 m |
| 3.00 | 0.717 | 4.3 | 86.0 | 0.488 | 0.741 | landed 0.00 m |
| 6.00† | 1.014 | 3.0 | 82.4 | 4.607 | 10.05 | hard, 1.28 m/s lateral |

†Counterfactual. The position-gain slider stops at 3 °/m, so this row was set in a script and cannot be reproduced from the interface.

### 10.7.3 Bandwidth Separation Sweep

The bandwidth separation is not a rule of thumb in this application, it is a number that can be swept. Table 10.7 repeats scenario 2 at four position gains, each paired with the rate gain that keeps the position loop critically damped through Equation (10.13), so the only thing changing is the separation from the attitude loop.

The table shows the outer loop improving at the inner loop's expense as the position gain rises. Within the slider's range, a four-fold increase in position gain improves the position integral by 30 %, from 123.0 to 86.0 m s, while the attitude integral grows by a factor of 1.85 and the attitude setpoint travels 2.7 times as far. Doubling the gain again, beyond the slider in the counterfactual last row, improves the position integral by 33 % in all, to 82.4 m s, while the attitude integral degrades by a factor of 17.5 and the attitude setpoint travels 36 times as far. That row is not a poorly damped landing but a failed one. At a separation of 3.0 the outer loop revises $\theta_{SP}$ before the inner loop has delivered the previous command, the tilt limit is active for 16.3 s of a 32.4 s descent against 0.98 s at the default gains, and the vehicle arrives with 1.28 m/s of lateral speed against a 1.00 m/s limit. The same separation can be reached inside the slider ranges by slowing the inner loop instead, with $K_{Px} = 3.00\,°/\text{m}$, $K_{Dx} = 8.35\,°$ per m/s, $K_{P\theta} = 0.60$ and $K_{D\theta} = 0.46$ s (2.17 rad/s at a damping ratio of 0.83). That vehicle still lands, 0.00 m from the pad center, with an attitude integral of 1.134 rad s, so the ratio alone does not decide the outcome and the failure of the last row also needs an outer loop as fast as 1.01 rad/s. A metric read on the controlled variable alone would have called the last row the best of the four. The application recognizes this and prints both natural frequencies and their ratio whenever one of the three cascade gains is moved, with a warning below a ratio of four.

### 10.7.4 A Slow Gimbal Actuator

Scenario 8 slows the gimbal servo from 0.08 s to 0.18 s and halves its slew limit to 30 °/s, leaving every gain alone. A lag of 0.18 s contributes $\arctan(0.18 \times 3.07) = 28.9°$ of phase at the attitude natural frequency, which is most of the margin a proportional-derivative loop of damping ratio 0.83 has to give. Restoring the nominal actuator with everything else unchanged gives the reference run 8a. Against it, the slow gimbal raises the peak tilt from 10.2° to 14.9° on a 12° command, the attitude integral from 0.233 to 0.404 rad s and the gimbal travel from 0.491 to 1.199 rad, a factor of 2.4, all of it spent on ringing through −5.5°, +2.7° and −2.3° after the peak, which the nominal actuator does not produce.

The repair is less gain rather than more. Halving $K_{P\theta}$ to 0.60 moves the attitude loop from 3.07 rad/s to 2.17 rad/s and, because $K_{D\theta}$ is left alone, from a damping ratio of 0.83 to one of 1.18. Both changes put the loop back inside what the actuator supports, and the peak tilt falls to 7.9° with the gimbal travel down to 0.214 rad, less than half the nominal-actuator figure. The position integral is unchanged at 89.9 against 89.4 m s, so the outer loop registers neither repair. The attitude integral carries no information here, because it reads 0.404 for the ringing case and 0.406 for the well-damped one. An error integral cannot separate an overshoot that comes back from a lag that never left, while the peak and the actuator travel can. Slowing the servo one step further to $\tau_\delta = 0.20$ s at the original

gains makes the attitude loop diverge, and the vehicle tumbles through 266° and meets the pad at 64 m/s. An inner loop can never be faster than its actuator, and the failure appears first as ringing and second as a tumble.

> **TRY IT IN THE APP** **8 · Slow Gimbal Actuator**
>
> Run the scenario and watch the attitude trend. On a 12° command the tilt peaks at 14.9° and rings through −5.5°, +2.7° and −2.3°. Reset, halve Attitude gain to 0.6 and run again. The peak falls to 7.9° and the ringing is gone, while the position integral barely moves.

### 10.7.5 Gimbal Saturation

Scenario 9 reduces the gimbal travel to 3° and starts the vehicle 25 m off the pad and still moving away from it. The attitude loop asks for 14.4° of gimbal and gets 3°, so the pitching moment is frozen at $\ell T \sin 3° = 1.85\,\mathrm{kN\,m}$ no matter what the proportional-derivative law computes. While that lasts the attitude loop has no authority and the outer loop is writing commands that cannot be executed. The gimbal command is clipped for 0.91 s and the tilt command for 2.62 s, the maximum position error grows to 27.7 m against 20.1 m in scenario 2, the position integral rises from 123.0 to 207.6 m s, and the entry into the 1 m band moves from 14.535 s to 17.400 s.

The landing still succeeds, 0.02 m from the pad center, and two things save it. Neither cascaded loop has an integrator, so nothing winds up while the gimbal is pinned and the command is correct again the instant the saturation clears. The tilt limit keeps the demand bounded, so the saturation is temporary by construction. Neither protection is unlimited. At $\delta_{\max} = 1.5°$ the vehicle still lands, 0.03 m from center with a peak tilt of 13.7°, and at $\delta_{\max} = 1.0°$ it tumbles and arrives 42.5 m away at 57.8 m/s.

### 10.7.6 Steady Crosswind Offset

Scenario 5 holds a steady 7 m/s crosswind and the vehicle lands 5.77 m downwind of the pad, inside the 8 m pad radius and well away from the center. Every link of the chain derived in Section 10.4 can be read off the export, and the values below are taken at $t = 20\,\mathrm{s}$, twelve seconds before contact and with the vehicle creeping downwind at 0.04 m/s. The gimbal holds −1.00° against a predicted −1.01°, the attitude holds −3.53° against a predicted −3.54°, the attitude setpoint holds −4.37° against a predicted −4.38°, and the position sits 5.82 m downwind where the linear chain predicts 6.09 m. Matching three angles of a cascade against closed-form values is stronger evidence than matching one. Of the remaining 0.27 m of position, 0.24 m is the rate term, which is not zero because the vehicle has not stopped, and the rest is the 0.02° by which the attitude setpoint falls short of its prediction.

Doubling $K_{Px}$ to 1.44 °/m shrinks every term in the chain in proportion and the offset falls from 5.77 m to 2.96 m, a ratio of 1.95 against the predicted 2.00. It also pushes the position loop from 0.351 rad/s to 0.496 rad/s and closer to the attitude loop, which Table 10.7 shows is not free. The textbook answer of adding integral action to the position loop is the one the application declines to take, because the same scenario file that produces this offset can produce a saturated gimbal, and an outer integrator winding up against a pinned inner loop is the classic cascade failure. Volume 20, Chapter 3 takes the other route on a multirotor that also moves sideways only by tilting, and estimates the wind with a disturbance observer.

### 10.7.7 Reduced Maximum Thrust

Scenario 7 lowers the maximum thrust to 14 kN against a 11.77 kN weight, a thrust-to-weight ratio of 1.19 instead of 1.70, and starts the vehicle perfectly trimmed over the pad. Hover still exists. What has been lost is deceleration authority. 2.23 kN of spare thrust is 1.86 m/s$^2$ of retardation, while the exponential branch of the guidance profile demands $k_z|v_z| = 0.35 \times 10 = 3.5\,\mathrm{m/s^2}$ the moment the vehicle enters it. The throttle pins at 100 %, the loop stops being a loop for 3.36 s of commanded saturation, and the vehicle arrives at 3.59 m/s against a 2.0 m/s limit. Raising $T_{\max}$ in the 250 N

steps of its slider, the smallest engine that lands on the shipped profile is 14.5 kN, at 0.46 m/s, while 14.25 kN still arrives hard at 2.06 m/s.

No tuning repairs this, because the arithmetic that separates a tuning problem from a capability problem is a comparison of required against available acceleration and it does not mention the controller. Reducing $k_z$ to 0.15/s asks for 1.5 m/s$^2$ instead, which the engine has, and the same vehicle lands at 0.47 m/s with a peak throttle of 94 % and no saturation at all. The cost is 13.2 s of extra flight time, from 26.11 s to 39.32 s.

### 10.7.8 The Gravity Feedforward

Scenario 4 starts the vehicle 10° off vertical over the pad, so the attitude loop and the vertical loop are both disturbed at once. The tilt costs $1-\cos 10^\circ = 1.52\,\%$ of the vertical thrust, and the $1/\cos\theta$ term in the bias restores it before any descent-speed error develops. The descent integral is 6.83 m with the feedforward on against 6.67 m for a vertical descent that was never disturbed at all, so the vertical loop barely notices that the vehicle spent the first second lying over.

Switched off before the run, the feedforward leaves the controller with no preload, and the integral has to build the 3.77 kN between the engine's minimum thrust and hover by itself. On the vertical descent the descent integral then rises from 6.67 m to 12.94 m and the throttle travel from 67.0 % to 125.0 %. With the bias in place the descent-speed error is identically zero until the profile corner, and without it the error reaches 1.93 m/s at $t = 1.65$ s and accumulates 6.29 m of integrated error in the first ten seconds against zero. The thrust command falls to 8.29 kN while the integral builds, the vehicle falls faster than the profile asks, and the descent finishes 0.625 s early at 31.825 s. In scenario 4 the descent integral rises from 6.83 m to 12.94 m and the largest early descent-speed error from 0.090 m/s to 1.97 m/s (row 4b of Table 10.5), so the tilt correction is lost inside the hover thrust the integral is still building. Thrown at $t = 5$ s instead, the switch leaves the thrust command unchanged and the vertical descent keeps the descent integral of 6.67 m it has with the feedforward on. The feedforward is inexpensive here because the coupling is known exactly, being $\cos\theta$ with $\theta$ measured.

### 10.7.9 Propellant Depletion

Scenario 10 switches propellant depletion on and starts from 320 m. The vehicle burns 150.9 kg, from 1200 kg to 1049 kg, which is 12.6 % of its mass, and the three loops react differently because they depend on different quantities. Hover thrust falls with mass, so the peak throttle drops from 64 % to 52 %, and with it the attitude plant gain $\ell T/I$ falls from 7.85/s$^2$ to 6.86/s$^2$, moving the attitude loop from 3.07 rad/s to 2.87 rad/s and the right-half-plane zero from 2.80 rad/s to 2.62 rad/s. The horizontal authority $T/m \approx g$ barely changes, because thrust and weight fall together, so the position loop hardly notices and the vehicle lands 0.00 m from center. The vertical loop notices least of all, because the feedforward is computed from the current mass rather than the initial one.

The direction of that drift is a consequence of holding the inertia fixed. A real vehicle emptying tanks near its base loses inertia faster than it loses mass, so its attitude loop and its right-half-plane zero both move to higher frequency as the descent proceeds, and the jet damping that accompanies the mass flow is a further term this model does not carry [15]. The lesson survives the sign either way, which is that the plant a controller lands with is not the plant it was tuned for.

### 10.7.10 The Graded Challenge

The graded scenario combines a 20 m offset, 4 m/s of drift, 3° of initial tilt, a light crosswind, a 12 m/s gust at $t = 12$ s and an engine limited to 16 kN, and ships with tuning that is wrong in three separate ways. Equations (10.9) and (10.13) put the attitude loop at 1.40 rad/s with a damping ratio of 0.34 and the position loop at 0.57 rad/s with a damping ratio of 0.18, a separation of only 2.5, and $\tau_{I_v} = 20$ s leaves the vertical loop with almost no integral action. The challenge is passed by a landing within 1 m of the pad center at less than 1.5 m/s of vertical speed and less than 5° of tilt, judged at contact on the case as given. As shipped the vehicle reaches 23.8° of tilt, spends 10.96 s against the tilt limit and arrives hard at 3.64 m/s, 2.26 m off the pad and 5.22° over, missing all three limits.

Applying the tuning derived in Section 10.4, which is also the application's suggested tuning, lands it 0.33 m from the pad center at 0.40 m/s and 0.02° of tilt, inside the three limits by 0.67 m, 1.10 m/s and 4.98°, with the peak throttle at 97 % and no margin to spare. The position integral falls from 148.4 to 99.6 m s and the attitude integral from 2.081 to 0.246 rad s. Every repair made a loop slower and better damped rather than stronger, which is the usual direction of travel when a cascade misbehaves and the opposite of the first instinct most operators bring to it.

### 10.7.11 Manual Operation

Two manual experiments are worth running before any loop is closed. Trimming the engine to hover with the helper, from the vertical descent scenario, holds the thrust exactly equal to the weight and therefore holds the descent rate exactly constant. The vehicle arrives at 10.00 m/s after 25.0 s, which is a hard landing produced by a perfectly correct trim. A trim is not a landing, and the vertical loop exists to change the speed rather than to hold it.

The second experiment is to leave 0.5° of gimbal in from the same trimmed state. Equation (10.8) says the tilt then grows as $a\sin(0.5^\circ)t^2/2$, which reaches the 6° tip-over angle at $t = 1.75$ s, and the simulation crosses it at $t = 1.83$ s, the difference being the servo lag. The vehicle continues through 345° of rotation and meets the ground at a vertical speed of 61.7 m/s, 71.4 m downrange, at $t = 11.535$ s. Half a degree of a ten-degree actuator, left alone for under two seconds, is enough to lose the vehicle, which is what a plant with no restoring moment means in practice.

## 10.8 Defects in the Application Measurements

The saturation tile of each loop reports the fraction of the run during which the manipulated variable sat against a limit. When this chapter was first written, the limits handed to the metric accumulator for the two cascaded loops were constants fixed at 10° for the gimbal and 15° for the attitude setpoint, so the position loop reported 0 % saturated even in the graded scenario, where the tilt command is clipped for 10.96 s of a 22.82 s descent. The application now takes both limits from $\theta_{\max}$ and $\delta_{\max}$ as they stand, draws its limit lines there, and the position-loop tile reads 2.9 % in scenario 2 and 48.0 % in the graded scenario.

The attitude and descent tiles then still under-reported, because they scored the achieved gimbal angle and thrust, which approach a stop only through the rate limit and the servo or spool lag. Scenario 9 read 0.0 % while the gimbal command was clipped for 0.91 s, and scenario 7 read 6.3 % against 12.8 % of commanded saturation. Both tiles now count the boolean saturation flag each loop publishes on the command against the limit in force, the column the trend strips shade from, and read 2.8 % and 12.8 %. Saturated times quoted in Section 10.7 are counted from the same flags over the physics steps of the run.

The gravity feedforward switch once shifted the integral by the whole of $mg/\cos\theta$, overshooting by the 8 kN floor of Equation (10.19), so that switching it off in flight commanded a climb. The application now shifts the integral by the difference between the two biases and the transfer is continuous in both directions. A switch made before the run once preloaded the integral while the same switch carried through Reset did not, and the briefings described only the first. The application now gives the run without preload on both paths, and the briefings teach it.

## 10.9 Technology in Practice

Figure 10.6 puts the three loops of Figure 10.2 on flight hardware: navigation sensors feed one flight computer, and its outputs drive a gimbal actuator and a throttle actuator on the same engine.

#### What Measures It

The controller of this chapter reads $x$, $z$, $v_x$, $v_z$, $\theta$ and $\omega$ exactly, and Section 10.3.6 names that as the largest omission of the model, since a real vehicle infers those states from an inertial measurement unit and a radar altimeter. The lunar module's landing radar handed velocity and range to the guidance

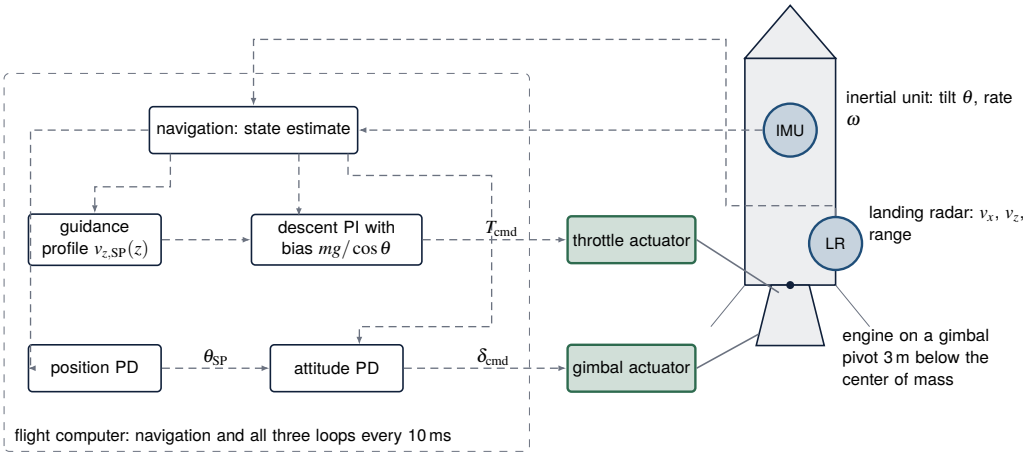



**Figure 10.6** The three loops of Figure 10.2 on flight hardware. An inertial unit and a landing radar feed a navigation estimate in the flight computer, the position loop writes the attitude setpoint for the attitude loop, and the attitude and descent loops drive a gimbal actuator and a throttle actuator on the one engine. The application replaces the sensors and the navigation block by the exact state. Solid lines are the vehicle and its devices, and dashed gray arrows are signals. Notice that both actuators act on the same engine below the center of mass, which is why the thrust that holds the vehicle up also sets the gain of the attitude loop.

computer in binary form [19]. Each of those readings carries a lag, a noise and a failure mode that the exact state vector of Section 10.5 leaves out, which is why every result in Table 10.5 is the tuning-limited case.

> **TECHNOLOGY** **Sensor: Doppler Landing Radar**
>
> The lunar module sensed its velocity and slant range relative to the surface with a landing radar that joined a three-beam Doppler velocity sensor to a frequency-modulated radar altimeter. The Doppler shift of each beam is directly proportional to the velocity along that beam, and the altimeter reads range from the frequency difference between the transmitted and the received sweep. Frequency trackers follow each return and signal lock, and the beam velocities are resolved into components for the guidance computer [19]. Ground tests found false lock-on to Doppler returns from the vehicle's own vibrating legs, engine skirt and lower structure. Less sensitivity to low-frequency vibration, an antenna turned away from the legs and a baffle against the engine bell cured it [19]. A falsely locked tracker reports the motion of the structure, not of the ground.

### What Moves It

The gimbal servo of Equation (10.5) stands for the gimbal actuator, with its 80 ms lag, its 60 °/s slew limit and its 10° travel, and the spool lag stands for the throttle actuator and the engine. The lunar descent engine was throttled over a range of ten to one, and its gimbal was a slow trim drive backed by attitude jets [12, 20]. The simulated vehicle has no jets, so its gimbal is the only source of pitching moment, and its engine spans only $T_{\min} = 8\,\mathrm{kN}$ to $T_{\max} = 20\,\mathrm{kN}$, a floor that the descent loop meets in scenario 7.

> **TECHNOLOGY** **Final Element: Throttled Engine on a Trim Gimbal**
>
> The lunar descent engine held both of the chapter's final elements. Its basic requirements were a throttling ratio of 10 to 1 and gimbaling of 6° either way in two axes. Its lowest throttle point was 10 % of rated thrust, and thrust between 65 and 92.5 % was treated as a region not to operate in. A triple-redundant electric actuator positioned the throttle [12]. The gimbal was a trim drive that moved at a constant angular rate for as long as the computer held its command on. During the descent burn the autopilot steered attitude with it to save jet propellant, and when the rather slow drive could not keep the attitude error within bounds, the reaction control jets took over [20].

### Where the Controller Runs

All three loops run in one flight computer. On the lunar module the guidance computer issued engine on and off commands, the descent-engine thrust and the gimbal commands under program control, and its digital autopilot chose which attitude jets to fire [20]. The guidance law above that autopilot and the autopilot below it are the outer and inner layers of this chapter's cascade. The reports read for this section give no execution period for the autopilot. The application executes every loop each 10 ms, which is a 100 Hz flight computer, one eighth of the gimbal servo's time constant.

> **IN PRACTICE** **Margins to Start From, Gains That Must Move**
>
> A design criteria monograph of 1970 collects what launch vehicle autopilots learned in flight. Gain margins of 6 dB and phase margins of 40° proved suitable values with which to begin a design. Because the vehicle's dynamics change rapidly in flight, gains chosen for liftoff may be inadequate later, and a preprogrammed change of gain is recommended. Propellant slosh can be a major contributor to instability, and baffles are the most common remedy. A rate gyro mounted where the structure bends senses that bending and can feed it back as control action, so where the sensor sits is itself a design choice [21].

### What Is Modernizing It

Three developments bear on this cascade. The first replaces microwave radar with light. A navigation Doppler lidar uses a frequency-modulated continuous wave to measure velocity and range together, in three fixed directions, from altitudes above 5 km, with velocity resolved to centimeters per second. Vibration of its fiber network shows up as frequency drift and costs accuracy, so the instrument's mounting is part of its design [22]. The second is terrain-relative navigation. A flash lidar records a full three-dimensional image with a single laser pulse, and a landing that must detect hazards of 30 cm needs a ground sample distance of about 15 cm [23]. Both address the navigation error that the exact state vector of this chapter hides. The third is adaptive control. An adaptive augmenting layer added to a launch vehicle autopilot extends the failures and flight anomalies the control system can accommodate while it keeps a direct link to classical gain and phase margins [24], the margins the in-practice segment above starts from.

## 10.10 Limitations

The simulator teaches a cascade and does not design a lander. The controller is given a perfect state vector, so the navigation problem that dominates real touchdown precision is absent entirely, and the numbers in Table 10.5 are tuning-limited rather than sensor-limited. Moving horizon estimation or Kalman filtering against the same model would supply that state vector, and would estimate the crosswind as a disturbance state so the standing offset of Section 10.7 could be removed by biasing $x_{SP}$ instead of by an integrator in a loop whose inner loop saturates. Planar motion excludes roll, yaw and the coupling between them that a real thrust-vector control system spends most of its authority on, and the fixed inertia excludes slosh, which is the mode most likely to sit near the attitude loop on a part-full vehicle. The controller is fixed in structure, so the convex trajectory generation that defines the modern powered-descent literature [3–6] cannot be demonstrated inside the application, although the right-half-plane zero of Equation (10.10) is exactly the property that makes a guidance

law worth having, because a trajectory that respects the vehicle's own bandwidth asks less of the loops that track it. Every disturbance is operator-triggered or constant rather than stochastic, and a fixed integration step with a fixed control period gives bit-identical repeats that no field controller delivers. A run terminates at contact, so no scenario except the vertical descent has a settled steady state to report, and every integrated measure in this chapter covers a whole descent rather than a common window. Where every modeling assumption stops holding is given in Section 10.3.6, and the three costliest are the exact state feedback, the constant inertia under propellant burn and the missing ground effect in the last few meters.

The predictive option fixes nominal geometry and lag, dead-reckons mass and corrects acceleration bias. Its soft tilt band and finite local forecast do not ensure contact feasibility, as the slow-servo and doubled-inertia failures show.

> ⚠ **OPEN CHALLENGE** **Landing in Its Own Exhaust**
>
> Every propulsive lander meets the interaction of its hot exhaust with the ground beneath it. The plume obscures cameras and sensors, threatens the vehicle's stability and tilt, raises local heating and pressure and throws debris. Apollo flight experience cannot be extrapolated to the larger landers now planned, and generalized models remain a challenge because the physics joins rarefied gas flow to granular soil, with validation work remaining for all methods and tools. Progress would be models checked against flight data gathered by instruments carried on the landers themselves [25]. The model of this chapter has no ground effect at all, as Section 10.3.6 states.

## 10.11 Discussion

An expendable booster is graded on where it puts its payload. A reusable one is graded twice, and the second grade is awarded in the last thirty seconds, on a pad a few tens of meters across, with an engine that cannot be shut off and relit and a propellant margin measured in seconds. The economics follow directly. Every kilogram of propellant reserved for the landing is a kilogram not lifted, so the descent is flown as close to the minimum as the guidance and the loops beneath it allow, which is why the powered-descent problem was worth convexifying [3, 4] and why solving it on board in bounded time is a practical rather than academic achievement [6].

What this application adds to that picture is the part the optimization does not touch. The trajectory can be optimal and the vehicle can still miss, because the loops that track it are bounded by $\sqrt{\ell T/I}$, a number fixed when the engine bay was laid out. Scenario 6 shows the same gust costing 1.36 m of pad offset on a fast approach and 0.29 m on a slow one, a factor of 4.7 for 6.86 s of extra hang time and the propellant that goes with it. The gust arrives at the same instant on both runs but not at the same altitude, 20.1 m on the shipped profile and 30.2 m at $k_z = 0.15/\mathrm{s}$, which is part of what the slower profile buys. A gust modeled as a step of side force decaying with one time constant is not a wind field, but it arrives at the same instant with the same amplitude on every run, so the comparison is a controlled experiment. That trade between how fast a vehicle comes down and how accurately it arrives is the terminal guidance design problem, and in this application it is one adjustable profile slope. The same structure, an outer loop that carries the objective and no actuator, an inner loop that carries the authority and no knowledge of the objective, and a bandwidth separation that keeps the two from contesting the same actuator, appears in quadrotors, missiles, gimballed camera mounts and ship autopilots [7], and on the bench in the ball and beam of Volume 7, Chapter 5, whose position loop asks an inner drive for an angle.

## 10.12 Conclusion

A vehicle that can only move sideways by pointing itself gives the outer loop of a cascade an angle to command instead of a force, and that single change reorders what the design has to respect. The simulation of this chapter integrates nine states under planar rigid-body dynamics with a gimbal servo, an engine spool, a quadratic side force applied above the center of mass and a graded touchdown

event, and closes two proportional-derivative loops and one velocity-form proportional-integral loop around them at 100 Hz.

Linearizing about hover puts a right-half-plane zero at $\sqrt{\ell T/I} = 2.80\,\mathrm{rad/s}$ in the path from gimbal to position, and closing the attitude loop leaves it there, so the separation the cascade needs and the margin the zero allows are the same number. The shipped tuning places the attitude loop at 3.07 rad/s and the position loop at 0.351 rad/s, a separation of 8.7. A sweep that pushes the position gain past its slider range to a separation of 3.0 improves the position integral by 33 % while degrading the attitude integral by a factor of 17.5 and turning a landing into a hard arrival.

Twelve scenarios, each with a stated completion requirement, break the cascade one mechanism at a time, with a slow gimbal costing 14.9° of peak tilt against 7.9° after the inner gain is halved, a 3° gimbal stretching the entry into the 1 m band from 14.5 s to 17.4 s, and a 7 m/s crosswind leaving a 5.77 m standing offset whose gimbal, attitude and attitude-setpoint links each match an algebraic chain to within 0.02°. The saturation tiles and the feedforward transfer found defective while this chapter was written have been corrected. The added nominal MPC improves selected arrivals but spends more gimbal motion and fails with wrong servo or inertia assumptions. The lander gives a course an accessible cascade whose outer loop must borrow its authority.

**KEY POINTS**

- The lander moves sideways only by tilting, so the position loop asks for a tilt and waits for the attitude loop to produce it.
- The engine below the center of mass puts a right-half-plane zero at $\sqrt{\ell T/I} = 2.80\,\mathrm{rad/s}$ that no attitude gain moves.
- The shipped loops are separated by a factor of 8.7, and at 3.0 the landing turns into a hard arrival.
- A steady 7 m/s crosswind leaves a 5.77 m offset, because neither cascaded loop integrates.
- A slow gimbal is repaired by less attitude gain, an engine short of deceleration only by a slower descent.

**Where else this applies.** Volume 20, Chapter 3 moves a multirotor sideways by tilting it, and Chapter 9 meets a right-half-plane zero in an electrical cascade.

## References


[1] J. Meditch, "On the problem of optimal thrust programming for a lunar soft landing," *IEEE Transactions on Automatic Control*, vol. 9, no. 4, pp. 477–484, 1964.

[2] A. R. Klumpp, "Apollo lunar descent guidance," *Automatica*, vol. 10, no. 2, pp. 133–146, 1974.

[3] B. Acikmese and S. R. Ploen, "Convex programming approach to powered descent guidance for Mars landing," *Journal of Guidance, Control, and Dynamics*, vol. 30, no. 5, pp. 1353–1366, 2007.

[4] B. Acikmese, J. M. Carson, and L. Blackmore, "Lossless convexification of nonconvex control bound and pointing constraints of the soft landing optimal control problem," *IEEE Transactions on Control Systems Technology*, vol. 21, no. 6, pp. 2104–2113, 2013.

[5] P. Lu, "Propellant-optimal powered descent guidance," *Journal of Guidance, Control, and Dynamics*, vol. 41, no. 4, pp. 813–826, 2018.

[6] D. Malyuta, T. P. Reynolds, M. Szmuk, T. Lew, R. Bonalli, M. Pavone, and B. Acikmese, "Convex optimization for trajectory generation: A tutorial on generating dynamically feasible trajectories reliably and efficiently," *IEEE Control Systems*, vol. 42, no. 5, pp. 40–113, 2022.

[7] R. W. Beard and T. W. McLain, *Small Unmanned Aircraft: Theory and Practice*. Princeton, NJ: Princeton University Press, 2012.

[8] 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.

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

[10] B. Wie, *Space Vehicle Dynamics and Control, Second Edition*. Reston, VA: American Institute of Aeronautics and Astronautics, 2008.

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

[12] W. R. Hammock, Jr., E. C. Currie, and A. E. Fisher, "Apollo experience report: Descent propulsion system," NASA Manned Spacecraft Center, Houston, TX, NASA Technical Note TN D-7143, 1973, nASA Technical Reports Server 19730011150.

[13] D. H. Klyde, D. T. McRuer, and T. T. Myers, "Pilot-induced oscillation analysis and prediction with actuator rate limiting," *Journal of Guidance, Control, and Dynamics*, vol. 20, no. 1, pp. 81–89, 1997.

[14] F. O. Eke and T. C. Mao, "On the dynamics of variable mass systems," *International Journal of Mechanical Engineering Education*, vol. 30, no. 2, pp. 123–137, 2002.

[15] S.-M. Wang and F. O. Eke, "Rotational dynamics of axisymmetric variable mass systems," *Journal of Applied Mechanics*, vol. 62, no. 4, pp. 970–974, 1995.

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

[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] D. H. Shelton, "Apollo experience report: Guidance and control systems: Lunar module stabilization and control system," NASA Lyndon B. Johnson Space Center, Houston, TX, NASA Technical Note TN D-8086, 1975, nASA Technical Reports Server 19760004104.

[21] National Aeronautics and Space Administration, "Effects of structural flexibility on launch vehicle control systems," NASA, Washington, DC, Space Vehicle Design Criteria (Guidance and Control) SP-8036, 1970, nASA Technical Reports Server 19700030458.

[22] A. Gragossian, B. Barnes, J. Heppler, J. Follman, N. Dostart, R. Amundsen, A. Ben Shabat, G. Hines, and F. Amzajerdian, "Navigation doppler lidar for lunar landers," in *Optica Advanced Photonics Congress 2022*. Optica Publishing Group, 2022, p. LsM3C.4.

[23] F. Amzajerdian, A. Gragossian, P. F. Brewster, J. M. Heppler, F. G. Wilson, G. D. Hines, S. A. Laughter, D. K. Litton, and A. Bulyshev, "Development of a compact lidar sensor for terrain relative navigation and terrain hazard avoidance," in *AIAA SciTech 2025 Forum*, 2025.

[24] J. H. Wall, J. S. Orr, and T. S. VanZwieten, "Space launch system implementation of adaptive augmenting control," in *AAS Guidance, Navigation, and Control Conference*, no. AAS 14-051, 2014, nASA Technical Reports Server 20140008747.

[25] A. M. Korzun and W. A. Chambers, "Development efforts and status of plume-surface interaction capabilities for propulsive landing systems," Presentation, AIAA SciTech Forum, NASA Technical Reports Server 20230018215, 2024.

## 10.13 Exercises

**10.1** ○ Concept check · 2 min

In the cascade of Figure 10.2, what does the horizontal position loop manipulate?

(a) The attitude setpoint $\theta_{\mathrm{SP}}$, which the attitude loop must deliver first

(b) A side thruster that pushes the vehicle toward the pad

(c) The gimbal angle $\delta$, which turns the thrust toward the pad at once

(d) The thrust magnitude $T$, whose horizontal share grows with the tilt

**10.2** ○ Concept check · 2 min

An engineer raises the attitude gain $K_{P\theta}$ from 1.20 to 4.80 at the hover operating point. What happens to the right-half-plane zero of the path from $\theta_{\mathrm{SP}}$ to $x$?

(a) It doubles in frequency, together with the attitude loop
(b) It stays at $\sqrt{\ell T/I} = 2.80\,\mathrm{rad/s}$, since feedback moves only poles
(c) It cancels against the faster attitude poles and disappears
(d) It crosses into the left half plane once the attitude loop is faster

**10.3** Concept check · 1 min

In scenario 5 the vehicle lands 5.77 m downwind of the pad center in a steady 7 m/s crosswind. What causes the offset?

(a) The 12° tilt limit stops the vehicle leaning far enough into the wind
(b) The descent integrator winds up and takes thrust from the horizontal channel
(c) Neither cascaded loop integrates, so the standing lean needs a standing position error
(d) Without a wind sensor the vehicle cannot hold any lean against a steady wind

**10.4** Theory and calculation · 6 min

Use the hover operating point of Table 10.2: $m = 1200\,\mathrm{kg}$, $g = 9.81\,\mathrm{m/s^2}$, $\ell = 3.0\,\mathrm{m}$ and $I = 4500\,\mathrm{kg\,m^2}$.

(a) Compute the attitude plant gain $a$, the zero $\sqrt{a}$, and the attitude natural frequency and damping ratio at the shipped $K_{P\theta} = 1.20$ and $K_{D\theta} = 0.65\,\mathrm{s}$. *Checkpoint:* $a = 7.848$ per $\mathrm{s^2}$, 2.801 rad/s, 3.069 rad/s and 0.831.

(b) The graded challenge ships with $K_{P\theta} = 0.25$, $K_{D\theta} = 0.12\,\mathrm{s}$, $K_{Px} = 1.90^\circ/\mathrm{m}$ and $K_{Dx} = 1.20^\circ$ per m/s. Place both loops and compute their separation. *Checkpoint:* Attitude 1.40 rad/s at 0.34, position 0.570 rad/s at 0.18, separation 2.46.

(c) Find the attitude gain that would put the attitude loop at twice the zero frequency. *Checkpoint:* $K_{P\theta} = 4$.

**10.5** Theory and calculation · 6 min

For a step deflection $\delta$ of the gimbal at hover, the linearized model gives $x(t) = g\delta\left(at^4/24 - t^2/2\right)$ with $a = 7.848\,\mathrm{s^{-2}}$.

(a) Find the times at which the horizontal acceleration changes sign, the wrong-way excursion is largest, and the vehicle passes back through its starting point. *Checkpoint:* 0.505 s, 0.874 s and 1.237 s.

(b) Compute the largest wrong-way excursion for $\delta = 5^\circ$. *Checkpoint:* 0.164 m.

(c) Scenario 10 ends its descent with $a = 6.86\,\mathrm{s^{-2}}$. Recompute the time of the acceleration reversal, and say whether a change of the attitude gains would move any of the three times. *Checkpoint:* 0.540 s.

**10.6** Theory and calculation · 7 min

Use Equation (10.14) and the descent-loop settings of Table 10.3, with $m = 1200\,\mathrm{kg}$ and hover thrust $mg = 11\,772\,\mathrm{N}$.

(a) Compute the descent loop's natural frequency, its damping ratio and the velocity time constant of the proportional term alone. *Checkpoint:* 0.707 rad/s, 0.884 and 0.80 s.

(b) Compute the gravity feedforward $mg/\cos\theta$ at 12° and at 30° of tilt, and the vertical force a 14 kN engine has to spare at 30°. *Checkpoint:* 12 035 N and 13 593 N, and 352 N against 2228 N upright.

(c) Scenario 7 limits the engine to 14 kN. Compare the deceleration the profile of Equation (10.7) demands at its corner, $k_z v_{\max}$, with what the engine can add above hover, and find the largest

profile slope on the $0.01\ \mathrm{s}^{-1}$ slider grid whose corner demand the engine can meet. *Checkpoint:* $3.5\ \mathrm{m/s}^2$ against $1.86\ \mathrm{m/s}^2$, so $k_z = 0.18\ \mathrm{s}^{-1}$.

**10.7** App challenge · 8 min

Load scenario 1, Vertical descent (start here), in Auto at the default seed. Read the clock when the run stops at touchdown and the Touchdown speed and Peak throttle tiles.

(a) Run to touchdown as loaded. Read the touchdown time and speed, and compute the altitude of the profile corner from Equation (10.7). *Checkpoint:* 32.45 s at 0.40 m/s, corner at 27.1 m.

(b) Reset, open Advanced, set Descent profile slope to 0.15 and run to touchdown. Reset again, set it to 0.80 and run once more. *Checkpoint:* 39.32 s at 0.47 m/s, then 25.38 s at 3.29 m/s with the throttle peaking at 96 %.

(c) Reload the scenario, switch to Manual before pressing Run, press Trim to hover and run to touchdown without touching anything else. *Checkpoint:* 10.00 m/s at 25.0 s.

Accept a time within 0.05 s and a speed within 0.01 m/s.

**10.8** App challenge · 8 min

Load scenario 8, Slow gimbal actuator, in Auto at the default seed. Each run goes to touchdown. Read the Peak tilt tile and the verdict.

(a) Run as loaded and read the peak tilt. Compute the phase that the 0.18 s servo lag contributes at the attitude natural frequency of 3.07 rad/s. *Checkpoint:* 14.9°, and 28.9° of phase.

(b) Reset, set Attitude gain to 0.60 and run again. Read the peak tilt, and compute the new attitude natural frequency and damping ratio with $K_{D\theta}$ left at 0.65 s. *Checkpoint:* 7.9°, 2.17 rad/s and 1.18.

(c) Reload the scenario, set Gimbal time constant to 0.20 s in the Parameters tab and run to touchdown. *Checkpoint:* A tip-over, peak tilt about 266°, at 64.2 m/s.

Accept a tilt within 0.2° in (a) and (b).

**10.9** App challenge · 8 min

Load scenario 5, Steady 7 m/s crosswind, in Auto at the default seed, and run each case to touchdown.

(a) Run as loaded and read the Pad offset tile. *Checkpoint:* 5.77 m.

(b) Export the run and read the row at `time_s` = 20: `gimbal_deg`, `theta_deg`, `theta_sp_deg`, `x_m` and `vx_mps`. Compare them with the predictions of Section 10.4, and account for the gap in position with the rate term of the position law. *Checkpoint:* −1.00, −3.53, −4.37, 5.82 m and 0.043 m/s. The rate term is worth 0.24 m.

(c) Reset, set Position gain to 1.44 °/m and run again. Compare the offset ratio with the factor of two that the chain predicts. *Checkpoint:* 2.96 m, a ratio of 1.95.

Accept an offset within 0.02 m and an angle within 0.01°.

**10.10** Application area · 4 min

Section 10.3.6 calls exact full-state feedback the largest omission of the model.

(a) Name the instruments a real lander infers its states from, and say what then sets its touchdown precision.

(b) A wind estimate is used to bias $x_{\mathrm{SP}}$ upwind by the offset the chain predicts. The estimate reads 6 m/s while 7 m/s blows. Using the chain of Section 10.4, estimate the offset that remains. *Checkpoint:* About 1.6 m.

**10.11** Application area · 4 min

The model holds the pitch inertia fixed while propellant burns.

(a) In scenario 10 the vehicle burns 150.9 kg and the model's $a$ falls from 7.85 to 6.86 $s^{-2}$. Suppose the real vehicle's inertia fell by 20 % over the same descent. Compute $a$, the attitude natural frequency at $K_{P\theta} = 1.20$ and the zero at touchdown. *Checkpoint:* 8.58 $s^{-2}$, 3.21 rad/s and 2.93 rad/s.

(b) Name the assumption that hides propellant slosh, and say why slosh matters to this loop in particular.

**10.12** Review the agent · 5 min

An AI assistant is asked how to land scenario 2 faster and replies as follows. "The ceiling on the position loop is the speed of the attitude loop beneath it. Quadruple the attitude gain to 4.80 and set the attitude rate gain to 1.30 s. The attitude loop doubles to 6.1 rad/s at the same damping of 0.83, so the ceiling doubles with it. Then run the position loop at its 3 °/m maximum with $K_{Dx} = 8.35°$ per m/s, a separation of 8.6, as comfortable as the shipped 8.7." Find the flaw and say how to check it.